

ESP32-P4

ESP-IDF Programming Guide



Release v5.2.7
Espressif Systems
May 18, 2026

Table of contents

Table of contents	i
1 Get Started	3
1.1 Introduction	3
1.2 What You Need	3
1.2.1 Hardware	3
1.2.2 Software	3
1.3 Installation	4
1.3.1 IDE	4
1.3.2 Manual Installation	4
1.4 Build Your First Project	36
1.5 Uninstall ESP-IDF	36
2 API Reference	37
2.1 API Conventions	37
2.1.1 Error Handling	37
2.1.2 Configuration Structures	37
2.1.3 Private APIs	39
2.1.4 Components in Example Projects	39
2.1.5 API Stability	39
2.2 Application Protocols	40
2.2.1 ASIO Port	40
2.2.2 ESP-Modbus	41
2.2.3 ESP-MQTT	41
2.2.4 ESP-TLS	60
2.2.5 ESP HTTP Client	77
2.2.6 ESP Local Control	95
2.2.7 ESP Serial Slave Link	104
2.2.8 ESP x509 Certificate Bundle	119
2.2.9 HTTP Server	122
2.2.10 HTTPS Server	150
2.2.11 ICMP Echo	154
2.2.12 mDNS Service	159
2.2.13 Mbed TLS	160
2.2.14 IP Network Layer	161
2.3 Error Codes Reference	162
2.4 Networking APIs	168
2.4.1 Ethernet	168
2.4.2 Thread	205
2.4.3 ESP-NETIF	216
2.4.4 IP Network Layer	252
2.4.5 Application Layer	254
2.5 Peripherals API	255
2.5.1 Analog Comparator	255
2.5.2 Clock Tree	265
2.5.3 Elliptic Curve Digital Signature Algorithm (ECDSA)	278
2.5.4 Event Task Matrix (ETM)	281

2.5.5	GPIO & RTC GPIO	290
2.5.6	General Purpose Timer (GPTimer)	311
2.5.7	Hash-Based Message Authentication Code (HMAC)	327
2.5.8	Digital Signature (DS)	331
2.5.9	Inter-Integrated Circuit (I2C)	337
2.5.10	Inter-IC Sound (I2S)	365
2.5.11	LCD	414
2.5.12	LED Control (LEDC)	430
2.5.13	Motor Control Pulse Width Modulator (MCPWM)	454
2.5.14	Parallel IO	510
2.5.15	Pulse Counter (PCNT)	517
2.5.16	Remote Control Transceiver (RMT)	533
2.5.17	SD Pull-up Requirements	561
2.5.18	SDMMC Host Driver	562
2.5.19	SD SPI Host Driver	568
2.5.20	SPI Flash API	574
2.5.21	SPI Master Driver	605
2.5.22	SPI Slave Driver	629
2.5.23	Universal Asynchronous Receiver/Transmitter (UART)	635
2.6	Project Configuration	661
2.6.1	Introduction	661
2.6.2	Project Configuration Menu	661
2.6.3	Using <code>sdkconfig.defaults</code>	662
2.6.4	Kconfig Format Rules	662
2.6.5	Backward Compatibility of Kconfig Options	662
2.6.6	Configuration Options Reference	663
2.7	Provisioning API	1080
2.7.1	Protocol Communication	1080
2.8	Storage API	1096
2.8.1	FAT Filesystem Support	1097
2.8.2	Generating and Parsing FATFS on Host	1106
2.8.3	Manufacturing Utility	1113
2.8.4	Non-Volatile Storage Library	1118
2.8.5	NVS Encryption	1142
2.8.6	NVS Partition Generator Utility	1148
2.8.7	NVS Partition Parser Utility	1154
2.8.8	SD/SDIO/MMC Driver	1155
2.8.9	Partitions API	1170
2.8.10	SPIFFS Filesystem	1180
2.8.11	Virtual Filesystem Component	1184
2.8.12	Wear Levelling API	1200
2.9	System API	1203
2.9.1	App Image Format	1203
2.9.2	Bootloader Image Format	1210
2.9.3	Application Level Tracing	1212
2.9.4	Call Function with External Stack	1217
2.9.5	Chip Revision	1219
2.9.6	Console	1223
2.9.7	eFuse Manager	1232
2.9.8	Error Code and Helper Functions	1262
2.9.9	ESP HTTPS OTA	1265
2.9.10	Event Loop Library	1272
2.9.11	FreeRTOS Overview	1285
2.9.12	FreeRTOS (IDF)	1287
2.9.13	FreeRTOS (Supplemental Features)	1403
2.9.14	Heap Memory Allocation	1431
2.9.15	Memory Management for MMU Supported Memory	1445
2.9.16	Memory Synchronization	1450

2.9.17	Heap Memory Debugging	1456
2.9.18	High Resolution Timer (ESP Timer)	1469
2.9.19	Internal and Unstable APIs	1476
2.9.20	Inter-Processor Call (IPC)	1477
2.9.21	Interrupt Allocation	1482
2.9.22	Logging library	1490
2.9.23	Miscellaneous System APIs	1498
2.9.24	Over The Air Updates (OTA)	1515
2.9.25	Power Management	1529
2.9.26	POSIX Threads Support	1534
2.9.27	Random Number Generation	1539
2.9.28	Sleep Modes	1542
2.9.29	SoC Capabilities	1554
2.9.30	System Time	1570
2.9.31	Asynchronous Memory Copy	1576
2.9.32	Watchdogs	1580
3	Hardware Reference	1587
4	API Guides	1589
4.1	Application Level Tracing Library	1589
4.1.1	Overview	1589
4.1.2	Modes of Operation	1589
4.1.3	Configuration Options and Dependencies	1590
4.1.4	How to Use This Library	1591
4.2	Application Startup Flow	1600
4.2.1	First Stage Bootloader	1600
4.2.2	Second Stage Bootloader	1600
4.2.3	Application Startup	1601
4.3	Bootloader	1603
4.3.1	Bootloader Compatibility	1603
4.3.2	Log Level	1603
4.3.3	Factory Reset	1603
4.3.4	Boot from Test Firmware	1604
4.3.5	Rollback	1605
4.3.6	Watchdog	1605
4.3.7	Bootloader Size	1605
4.3.8	Fast Boot from Deep-Sleep	1605
4.3.9	Custom Bootloader	1605
4.4	Build System	1606
4.4.1	Overview	1606
4.4.2	Using the Build System	1607
4.4.3	Example Project	1608
4.4.4	Project CMakeLists File	1609
4.4.5	Component CMakeLists Files	1611
4.4.6	Component Configuration	1613
4.4.7	Preprocessor Definitions	1613
4.4.8	Component Requirements	1613
4.4.9	Overriding Parts of the Project	1618
4.4.10	Configuration-Only Components	1619
4.4.11	Debugging CMake	1619
4.4.12	Example Component CMakeLists	1620
4.4.13	Custom Sdkconfig Defaults	1624
4.4.14	Flash Arguments	1624
4.4.15	Building the Bootloader	1625
4.4.16	Writing Pure CMake Components	1625
4.4.17	Using Third-Party CMake Projects with Components	1625
4.4.18	Using Prebuilt Libraries with Components	1626

4.4.19	Using ESP-IDF in Custom CMake Projects	1627
4.4.20	ESP-IDF CMake Build System API	1627
4.4.21	File Globbing & Incremental Builds	1631
4.4.22	Build System Metadata	1632
4.4.23	Build System Internals	1633
4.4.24	Migrating from ESP-IDF GNU Make System	1634
4.5	Core Dump	1636
4.5.1	Overview	1636
4.5.2	Configurations	1636
4.5.3	Core Dump to Flash	1637
4.5.4	Core Dump to UART	1637
4.5.5	Core Dump Commands	1639
4.5.6	ROM Functions in Backtraces	1639
4.5.7	Dumping Variables on Demand	1639
4.5.8	Running <code>idf.py coredump-info</code> and <code>idf.py coredump-debug</code>	1640
4.6	C++ Support	1640
4.6.1	<code>esp-idf-cxx</code> Component	1642
4.6.2	C++ Language Standard	1642
4.6.3	Multithreading	1643
4.6.4	Exception Handling	1643
4.6.5	Runtime Type Information (RTTI)	1643
4.6.6	Developing in C++	1644
4.6.7	Limitations	1645
4.6.8	What to Avoid	1645
4.7	Current Consumption Measurement of Modules	1645
4.7.1	Notes to Measurement	1645
4.7.2	Hardware Connection	1646
4.7.3	Measurement Steps	1646
4.8	Deep Sleep Wake Stubs	1650
4.8.1	Rules for Wake Stubs	1650
4.8.2	Implementing A Stub	1650
4.8.3	Loading Code Into RTC Memory	1650
4.8.4	Loading Data Into RTC Memory	1651
4.8.5	CRC Check For Wake Stubs	1651
4.8.6	Example	1651
4.9	Error Handling	1652
4.9.1	Overview	1652
4.9.2	Error Codes	1652
4.9.3	Converting Error Codes to Error Messages	1652
4.9.4	<code>ESP_ERROR_CHECK</code> Macro	1653
4.9.5	<code>ESP_ERROR_CHECK_WITHOUT_ABORT</code> Macro	1653
4.9.6	<code>ESP_RETURN_ON_ERROR</code> Macro	1653
4.9.7	<code>ESP_GOTO_ON_ERROR</code> Macro	1653
4.9.8	<code>ESP_RETURN_ON_FALSE</code> Macro	1653
4.9.9	<code>ESP_GOTO_ON_FALSE</code> Macro	1654
4.9.10	<code>CHECK_MACROS</code> Examples	1654
4.9.11	Error Handling Patterns	1654
4.9.12	C++ Exceptions	1655
4.10	Support for External RAM	1655
4.10.1	Introduction	1655
4.10.2	Hardware	1655
4.10.3	Configuring External RAM	1656
4.10.4	Restrictions	1657
4.10.5	Failure to Initialize	1658
4.10.6	Encryption	1658
4.11	Fatal Errors	1658
4.11.1	Overview	1658
4.11.2	Panic Handler	1658

4.11.3	Register Dump and Backtrace	1659
4.11.4	GDB Stub	1662
4.11.5	RTC Watchdog Timeout	1662
4.11.6	Guru Meditation Errors	1663
4.11.7	Other Fatal Errors	1664
4.12	Hardware Abstraction	1666
4.12.1	Architecture	1667
4.12.2	LL (Low Level) Layer	1668
4.12.3	HAL (Hardware Abstraction Layer)	1669
4.13	JTAG Debugging	1670
4.13.1	Introduction	1670
4.13.2	How it Works?	1671
4.13.3	Selecting JTAG Adapter	1671
4.13.4	Setup of OpenOCD	1672
4.13.5	Configuring ESP32-P4 Target	1672
4.13.6	Launching Debugger	1675
4.13.7	Debugging Examples	1675
4.13.8	Building OpenOCD from Sources	1676
4.13.9	Tips and Quirks	1680
4.13.10	Related Documents	1684
4.14	Linker Script Generation	1710
4.14.1	Overview	1710
4.14.2	Quick Start	1711
4.14.3	Linker Script Generation Internals	1714
4.15	lwIP	1720
4.15.1	Supported APIs	1720
4.15.2	BSD Sockets API	1721
4.15.3	Netconn API	1724
4.15.4	lwIP FreeRTOS Task	1725
4.15.5	IPv6 Support	1725
4.15.6	ESP-lwIP Custom Modifications	1726
4.15.7	Performance Optimization	1727
4.16	Memory Types	1729
4.16.1	DRAM (Data RAM)	1729
4.16.2	IRAM (Instruction RAM)	1729
4.16.3	IROM (Code Executed from flash)	1730
4.16.4	DROM (Data Stored in flash)	1731
4.16.5	RTC FAST Memory	1731
4.16.6	TCM (Tightly-Coupled Memory)	1731
4.16.7	DMA-Capable Requirement	1731
4.16.8	DMA Buffer in the Stack	1732
4.17	OpenThread	1732
4.17.1	Modes of the OpenThread Stack	1732
4.17.2	How to Write an OpenThread Application	1733
4.17.3	OpenThread Macro Definitions	1734
4.17.4	The OpenThread Border Router	1734
4.18	Partition Tables	1735
4.18.1	Overview	1735
4.18.2	Built-in Partition Tables	1735
4.18.3	Creating Custom Tables	1736
4.18.4	Generating Binary Partition Table	1738
4.18.5	Partition Size Checks	1739
4.18.6	Flashing the Partition Table	1739
4.18.7	Partition Tool (parttool.py)	1739
4.19	Performance	1741
4.19.1	How to Optimize Performance	1741
4.19.2	Guides	1741
4.20	Reproducible Builds	1759



4.20.1	Introduction	1759
4.20.2	Reasons for Non-Reproducible Builds	1759
4.20.3	Enabling Reproducible Builds in ESP-IDF	1759
4.20.4	How Reproducible Builds Are Achieved	1759
4.20.5	Reproducible Builds and Debugging	1760
4.20.6	Factors Which Still Affect Reproducible Builds	1760
4.21	Thread Local Storage	1760
4.21.1	Overview	1760
4.21.2	FreeRTOS Native APIs	1760
4.21.3	Pthread APIs	1761
4.21.4	C11 Standard	1761
4.22	Tools	1761
4.22.1	IDF Frontend - <code>idf.py</code>	1761
4.22.2	IDF Docker Image	1765
4.22.3	IDF Windows Installer	1768
4.22.4	IDF Component Manager	1769
4.22.5	IDF Clang-Tidy	1772
4.22.6	Downloadable IDF Tools	1773
4.23	Unit Testing in ESP32-P4	1785
4.23.1	Normal Test Cases	1785
4.23.2	Multi-device Test Cases	1786
4.23.3	Multi-stage Test Cases	1787
4.23.4	Tests For Different Targets	1787
4.23.5	Building Unit Test App	1788
4.23.6	Running Unit Tests	1789
4.23.7	Timing Code with Cache Compensated Timer	1790
4.23.8	Mocks	1790
4.24	Running ESP-IDF Applications on Host	1793
4.24.1	Introduction	1793
4.24.2	Requirements for Using Mocks	1794
4.24.3	Build and Run	1794
4.24.4	Component Linux/Mock Support Overview	1794
4.25	Low Power Mode User Guide	1795
5	Security Guides	1797
5.1	Overview	1797
5.1.1	Security	1797
5.2	Features	1800
5.2.1	Flash Encryption	1801
5.2.2	Secure Boot V2	1811
5.3	Workflows	1821
5.3.1	Host-Based Security Workflows	1821
6	Migration Guides	1829
6.1	ESP-IDF 5.x Migration Guide	1829
6.1.1	Migration from 4.4 to 5.0	1829
6.1.2	Migration from 5.0 to 5.1	1858
6.1.3	Migration from 5.1 to 5.2	1861
7	Libraries and Frameworks	1865
7.1	Cloud Frameworks	1865
7.1.1	ESP RainMaker	1865
7.1.2	AWS IoT	1865
7.1.3	Azure IoT	1865
7.1.4	Google IoT Core	1865
7.1.5	Aliyun IoT	1865
7.1.6	Joylink IoT	1865
7.1.7	Tencent IoT	1866
7.1.8	Tencentyun IoT	1866

7.1.9	Baidu IoT	1866
7.2	Espressif's Frameworks	1866
7.2.1	Espressif Audio Development Framework	1866
7.2.2	ESP-CSI	1866
7.2.3	Espressif DSP Library	1866
7.2.4	ESP-WIFI-MESH Development Framework	1867
7.2.5	ESP-WHO	1867
7.2.6	ESP RainMaker	1867
7.2.7	ESP-IoT-Solution	1867
7.2.8	ESP-Protocols	1867
7.2.9	ESP-BSP	1868
7.2.10	ESP-IDF-CXX	1868
8	Contributions Guide	1869
8.1	How to Contribute	1869
8.2	Before Contributing	1869
8.3	Pull Request Process	1869
8.4	Legal Part	1870
8.5	Related Documents	1870
8.5.1	Espressif IoT Development Framework Style Guide	1870
8.5.2	Install Pre-commit Hook for ESP-IDF Project	1878
8.5.3	Documenting Code	1879
8.5.4	Creating Examples	1884
8.5.5	API Documentation Template	1885
8.5.6	Contributor Agreement	1887
8.5.7	Copyright Header Guide	1889
8.5.8	pytest in ESP-IDF	1891
9	ESP-IDF Versions	1903
9.1	Releases	1903
9.2	Which Version Should I Start With?	1903
9.3	Versioning Scheme	1904
9.4	Support Periods	1904
9.5	Checking the Current Version	1905
9.6	Git Workflow	1906
9.7	Updating ESP-IDF	1906
9.7.1	Updating to Stable Release	1907
9.7.2	Updating to a Pre-Release Version	1907
9.7.3	Updating to Master Branch	1907
9.7.4	Updating to a Release Branch	1908
10	Resources	1909
10.1	PlatformIO	1909
10.1.1	What Is PlatformIO?	1909
10.1.2	Installation	1909
10.1.3	Configuration	1910
10.1.4	Tutorials	1910
10.1.5	Project Examples	1910
10.1.6	Next Steps	1910
10.2	CLion	1910
10.2.1	What Is CLion?	1910
10.2.2	Installation	1910
10.2.3	Configuration	1910
10.2.4	Resources	1910
10.3	VisualGDB	1910
10.3.1	What Is VisualGDB?	1910
10.3.2	Installation	1911
10.3.3	Configuration	1911
10.3.4	Resources	1911

10.4 Useful Links	1911
11 Copyrights and Licenses	1913
11.1 Software Copyrights	1913
11.1.1 Firmware Components	1913
11.1.2 Documentation	1914
11.2 ROM Source Code Copyrights	1914
11.3 Xtensa libhal MIT License	1915
11.4 TinyBasic Plus MIT License	1915
11.5 TJpgDec License	1915
12 About	1917
13 Switch Between Languages	1919
Index	1921
Index	1921

This is the documentation for Espressif IoT Development Framework ([esp-idf](#)). ESP-IDF is the official development framework for the [ESP32](#), [ESP32-S](#), [ESP32-C](#), [ESP32-H](#) and [ESP32-P](#) Series SoCs.

This document describes using ESP-IDF with the ESP32-P4 SoC.

		
Get Started	API Reference	API Guides

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Chapter 1

Get Started

This document is intended to help you set up the software development environment for the hardware based on the ESP32-P4 chip by Espressif. After that, a simple example will show you how to use ESP-IDF (Espressif IoT Development Framework) for menu configuration, then for building and flashing firmware onto an ESP32-P4 board.

Note: This is documentation for stable version v5.2.7 of ESP-IDF. Other [ESP-IDF Versions](#) are also available.

1.1 Introduction

ESP32-P4 is a system on a chip that integrates the following features:

Powered by 40 nm technology, ESP32-P4 provides a robust, highly integrated platform, which helps meet the continuous demands for efficient power usage, compact design, security, high performance, and reliability.

Espressif provides basic hardware and software resources to help application developers realize their ideas using the ESP32-P4 series hardware. The software development framework by Espressif is intended for development of Internet-of-Things (IoT) applications with Wi-Fi, Bluetooth, power management and several other system features.

1.2 What You Need

1.2.1 Hardware

- An **ESP32-P4** board.
- **USB cable** - USB A / micro USB B.
- **Computer** running Windows, Linux, or macOS.


Note: Currently, some of the development boards are using USB Type C connectors. Be sure you have the correct cable to connect your board!

If you have one of ESP32-P4 official development boards listed below, you can click on the link to learn more about the hardware.

1.2.2 Software

To start using ESP-IDF on **ESP32-P4**, install the following software:

- **Toolchain** to compile code for ESP32-P4
- **Build tools** - CMake and Ninja to build a full **Application** for ESP32-P4
- **ESP-IDF** that essentially contains API (software libraries and source code) for ESP32-P4 and scripts to operate the **Toolchain**



1.3 Installation

To install all the required software, we offer some different ways to facilitate this task. Choose from one of the available options.

1.3.1 IDE

Note: We highly recommend installing the ESP-IDF through your favorite IDE.

- [Eclipse Plugin](#)
- [VSCode Extension](#)

1.3.2 Manual Installation

For the manual procedure, please select according to your operating system.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Standard Setup of Toolchain for Windows

Introduction ESP-IDF requires some prerequisite tools to be installed so you can build firmware for supported chips. The prerequisite tools include Python, Git, cross-compilers, CMake and Ninja build tools.

For this Getting Started we are going to use the Command Prompt, but after ESP-IDF is installed you can use [Eclipse Plugin](#) or another graphical IDE with CMake support instead.

Note: Limitations: - The installation path of ESP-IDF and ESP-IDF Tools must not be longer than 90 characters. Too long installation paths might result in a failed build. - The installation path of Python or ESP-IDF must not contain white spaces or parentheses. - The installation path of Python or ESP-IDF should not contain special characters (non-ASCII) unless the operating system is configured with "Unicode UTF-8" support.

System Administrator can enable the support via Control Panel > Change date, time, or number formats > Administrative tab > Change system locale > check the option Beta: Use Unicode UTF-8 for worldwide language support > Ok > reboot the computer.

ESP-IDF Tools Installer The easiest way to install ESP-IDF's prerequisites is to download one of ESP-IDF Tools Installers.



What Is the Usecase for Online and Offline Installer Online Installer is very small and allows the installation of all available releases of ESP-IDF. The installer downloads only necessary dependencies including [Git For Windows](#) during the installation process. The installer stores downloaded files in the cache directory `%userprofile%\espressif`

Offline Installer does not require any network connection. The installer contains all required dependencies including [Git For Windows](#).

Components of the Installation The installer deploys the following components:

- Embedded Python
- Cross-compilers
- OpenOCD
- [CMake](#) and [Ninja](#) build tools
- ESP-IDF

The installer also allows reusing the existing directory with ESP-IDF. The recommended directory is `%userprofile%\Desktop\esp-idf` where `%userprofile%` is your home directory.

Launching ESP-IDF Environment At the end of the installation process you can check out option Run ESP-IDF PowerShell Environment or Run ESP-IDF Command Prompt (`cmd.exe`). The installer launches ESP-IDF environment in selected prompt.

Run ESP-IDF PowerShell Environment:




Fig. 1: Completing the ESP-IDF Tools Setup Wizard with Run ESP-IDF PowerShell Environment

Run ESP-IDF Command Prompt (`cmd.exe`):

Using the Command Prompt For the remaining Getting Started steps, we are going to use the Windows Command Prompt.

ESP-IDF Tools Installer also creates a shortcut in the Start menu to launch the ESP-IDF Command Prompt. This shortcut launches the Command Prompt (`cmd.exe`) and runs `export.bat` script to set up the environment variables (`PATH`, `IDF_PATH` and others). Inside this command prompt, all the installed tools are available.

Note that this shortcut is specific to the ESP-IDF directory selected in the ESP-IDF Tools Installer. If you have multiple ESP-IDF directories on the computer (for example, to work with different versions of ESP-IDF), you have two options to use them:

1. Create a copy of the shortcut created by the ESP-IDF Tools Installer, and change the working directory of the new shortcut to the ESP-IDF directory you wish to use.
2. Alternatively, run `cmd.exe`, then change to the ESP-IDF directory you wish to use, and run `export.bat`. Note that unlike the previous option, this way requires Python and Git to be present in `PATH`. If you get errors related to Python or Git not being found, use the first option.

```
ESP-IDF PowerShell

Using Python in C:/Users/developer/.espressif/python_env/idf4.1_py3.8_env/Scripts
Python 3.8.7
Using Git in C:/Program Files/Git/cmd/
git version 2.29.2.windows.1
Setting IDF_PATH: C:/Users/developer/Desktop/esp-idf
Adding ESP-IDF tools to PATH...
C:/Users/developer/.espressif/tools/xtensa-esp32-elf/esp-2020r3-8.4.0/xtensa-esp32-elf/bin
C:/Users/developer/.espressif/tools/xtensa-esp32s2-elf/esp-2020r3-8.4.0/xtensa-esp32s2-elf/bin
C:/Users/developer/.espressif/tools/esp32ulp-elf/2.28.51-esp-20191205/esp32ulp-elf-binutils/bin
C:/Users/developer/.espressif/tools/esp32s2ulp-elf/2.28.51-esp-20191205/esp32s2ulp-elf-binutils/bin
C:/Users/developer/.espressif/tools/cmake/3.13.4/bin
C:/Users/developer/.espressif/tools/openocd-esp32/v0.10.0-esp32-20200709/openocd-esp32/bin
C:/Users/developer/.espressif/tools/ninja/1.9.0/
C:/Users/developer/.espressif/tools/idf-exe/1.0.1/
C:/Users/developer/.espressif/tools/ccache/3.7/
C:/Users/developer/Desktop/esp-idf/tools
Checking if Python packages are up to date...
Python requirements from C:/Users/developer/Desktop/esp-idf/requirements.txt are satisfied.
Done! You can now compile ESP-IDF projects.
Go to the project directory and run:
    idf.py build

PS C:/Users/developer/Desktop/esp-idf>
```

Fig. 2: ESP-IDF PowerShell





Fig. 3: Completing the ESP-IDF Tools Setup Wizard with Run ESP-IDF Command Prompt (cmd.exe)

A screenshot of a Windows Command Prompt window titled "ESP-IDF Command Prompt (cmd.exe)". The window has a blue title bar with standard Windows window controls (minimize, maximize, close). The background is black with white text. The text shows the execution of a script to set up the ESP-IDF environment. It starts with checking Python and Git versions, then sets the IDF_PATH. A large block of paths is added to the PATH environment variable. It then checks if Python packages are up to date and finally prompts the user to run 'idf.py build' in the project directory.

```
ESP-IDF Command Prompt (cmd.exe)

Using Python in C:\Users\test\AppData\Local\Programs\Python\Python37\
Python 3.7.8
Using Git in C:\Users\test\Git\cmd\
git version 2.30.0.windows.1
Setting IDF_PATH: C:\Users\test\esp\esp-idf

Adding ESP-IDF tools to PATH...
    C:\Users\test\.espressif\tools\xtensa-esp32-elf\esp-2020r3-8.4.0\xtensa-esp32-elf\bin
    C:\Users\test\.espressif\tools\xtensa-esp32s2-elf\esp-2020r3-8.4.0\xtensa-esp32s2-elf\bin
    C:\Users\test\.espressif\tools\xtensa-esp32s3-elf\esp-2020r3-8.4.0\xtensa-esp32s3-elf\bin
    C:\Users\test\.espressif\tools\riscv32-esp-elf\1.24.0.123_64eb9ff-8.4.0\riscv32-esp-elf\bin
    C:\Users\test\.espressif\tools\esp32ulp-elf\2.28.51-esp-20191205\esp32ulp-elf-binutils\bin
    C:\Users\test\.espressif\tools\esp32s2ulp-elf\2.28.51-esp-20191205\esp32s2ulp-elf-binutils\bin
    C:\Users\test\.espressif\tools\cmake\3.16.4\bin
    C:\Users\test\.espressif\tools\openocd-esp32\v0.10.0-esp32-20200709\openocd-esp32\bin
    C:\Users\test\.espressif\tools\ninja\1.10.0\
    C:\Users\test\.espressif\tools\idf-exe\1.0.1\
    C:\Users\test\.espressif\tools\ccache\3.7\
    C:\Users\test\.espressif\tools\dfu-util\0.9\dfu-util-0.9-win64
    C:\Users\test\.espressif\python_env\idf4.3_py3.7_env\Scripts
    C:\Users\test\esp\esp-idf\tools

Checking if Python packages are up to date...
Python requirements from C:\Users\test\esp\esp-idf\requirements.txt are satisfied.

Done! You can now compile ESP-IDF projects.
Go to the project directory and run:

    idf.py build

C:\Users\test\esp\esp-idf>
```

Fig. 4: ESP-IDF Command Prompt

First Steps on ESP-IDF Now since all requirements are met, the next topic guides you on how to start your first project.

This guide helps you on the first steps using ESP-IDF. Follow this guide to start a new project on the ESP32-P4 and build, flash, and monitor the device output.

Note: If you have not yet installed ESP-IDF, please go to [Installation](#) and follow the instruction in order to get all the software needed to use this guide.

Start a Project Now you are ready to prepare your application for ESP32-P4. You can start with [get-started/hello_world](#) project from [examples](#) directory in ESP-IDF.

Important: The ESP-IDF build system does not support spaces in the paths to either ESP-IDF or to projects.

Copy the project [get-started/hello_world](#) to `~/esp` directory:

```
cd %userprofile%\esp
xcopy /e /i %IDF_PATH%\examples\get-started\hello_world hello_world
```

Note: There is a range of example projects in the [examples](#) directory in ESP-IDF. You can copy any project in the same way as presented above and run it. It is also possible to build examples in-place without copying them first.

Connect Your Device Now connect your ESP32-P4 board to the computer and check under which serial port the board is visible.

Serial port names start with COM in Windows.

If you are not sure how to check the serial port name, please refer to [Establish Serial Connection with ESP32-P4](#) for full details.

Note: Keep the port name handy as it is needed in the next steps.

Configure Your Project Navigate to your `hello_world` directory, set ESP32-P4 as the target, and run the project configuration utility `menuconfig`.

Windows

```
cd %userprofile%\esp\hello_world
idf.py set-target esp32p4
idf.py menuconfig
```

After opening a new project, you should first set the target with `idf.py set-target esp32p4`. Note that existing builds and configurations in the project, if any, are cleared and initialized in this process. The target may be saved in the environment variable to skip this step at all. See [Select the Target Chip: set-target](#) for additional information.

If the previous steps have been done correctly, the following menu appears:

You are using this menu to set up project specific variables, e.g., Wi-Fi network name and password, the processor speed, etc. Setting up the project with `menuconfig` may be skipped for "hello_word", since this example runs with default configuration.




Fig. 5: Project configuration - Home window

Note: The colors of the menu could be different in your terminal. You can change the appearance with the option `--style`. Please run `idf.py menuconfig --help` for further information.

Build the Project Build the project by running:

```
idf.py build
```

This command compiles the application and all ESP-IDF components, then it generates the bootloader, partition table, and application binaries.

```
$ idf.py build
Running cmake in directory /path/to/hello_world/build
Executing "cmake -G Ninja --warn-uninitialized /path/to/hello_world"...
Warn about uninitialized values.
-- Found Git: /usr/bin/git (found version "2.17.0")
-- Building empty aws_iot component due to configuration
-- Component names: ...
-- Component paths: ...

... (more lines of build system output)

[527/527] Generating hello_world.bin
esptool.py v2.3.1

Project build complete. To flash, run this command:
../../components/esptool_py/esptool/esptool.py -p (PORT) -b 921600 write_flash -
↪-flash_mode dio --flash_size detect --flash_freq 40m 0x10000 build/hello_world.
↪bin build 0x1000 build/bootloader/bootloader.bin 0x8000 build/partition-table/
↪partition-table.bin
or run 'idf.py -p PORT flash'
```

If there are no errors, the build finishes by generating the firmware binary `.bin` files.

Flash onto the Device To flash the binaries that you just built for the ESP32-P4 in the previous step, you need to run the following command:

```
idf.py -p PORT flash
```

Replace `PORT` with your ESP32-P4 board's USB port name. If the `PORT` is not defined, the *idf.py* will try to connect automatically using the available USB ports.

For more information on `idf.py` arguments, see *idf.py*.

Note: The option `flash` automatically builds and flashes the project, so running `idf.py build` is not necessary.

Encountered Issues While Flashing? See the "Additional Tips" below. You can also refer to *Flashing Troubleshooting* page or *Establish Serial Connection with ESP32-P4* for more detailed information.

Normal Operation When flashing, you will see the output log similar to the following:

If there are no issues by the end of the flash process, the board will reboot and start up the "hello_world" application.

If you would like to use the Eclipse or VS Code IDE instead of running `idf.py`, check out *Eclipse Plugin*, *VSCode Extension*.

Monitor the Output To check if "hello_world" is indeed running, type `idf.py -p PORT monitor` (Do not forget to replace `PORT` with your serial port name).

This command launches the *IDF Monitor* application:

```
$ idf.py -p <PORT> monitor
Running idf_monitor in directory [...]/esp/hello_world/build
Executing "python [...]/esp-idf/tools/idf_monitor.py -b 115200 [...]/esp/hello_
↪world/build/hello_world.elf"...
--- idf_monitor on <PORT> 115200 ---
--- Quit: Ctrl+] | Menu: Ctrl+T | Help: Ctrl+T followed by Ctrl+H ---
ets Jun  8 2016 00:22:57

rst:0x1 (POWERON_RESET),boot:0x13 (SPI_FAST_FLASH_BOOT)
ets Jun  8 2016 00:22:57
...
```

After startup and diagnostic logs scroll up, you should see "Hello world!" printed out by the application.

```
...
Hello world!
Restarting in 10 seconds...
This is esp32p4 chip with 2 CPU core(s), [NEEDS TO BE UPDATED]
Minimum free heap size: [NEEDS TO BE UPDATED] bytes
Restarting in 9 seconds...
Restarting in 8 seconds...
Restarting in 7 seconds...
```

To exit IDF monitor use the shortcut `Ctrl+]`.

Note: You can combine building, flashing and monitoring into one step by running:

```
idf.py -p PORT flash monitor
```

See also:

- *IDF Monitor* for handy shortcuts and more details on using IDF monitor.
- *idf.py* for a full reference of `idf.py` commands and options.

That is all that you need to get started with ESP32-P4!

Now you are ready to try some other [examples](#), or go straight to developing your own applications.

Important: Some of examples do not support ESP32-P4 because required hardware is not included in ESP32-P4 so it cannot be supported.

If building an example, please check the README file for the `Supported Targets` table. If this is present including ESP32-P4 target, or the table does not exist at all, the example will work on ESP32-P4.

Additional Tips

Permission Denied Issue With some Linux distributions, you may get the error message similar to `Could not open port <PORT>: Permission denied: '<PORT>'` when flashing the ESP32-P4. *This can be solved by adding the current user to the specific group*, such as `dialout` or `uucp` group.

Python Compatibility ESP-IDF supports Python 3.8 or newer. It is recommended to upgrade your operating system to a recent version satisfying this requirement. Other options include the installation of Python from [sources](#) or the use of a Python version management system such as [pyenv](#).

Flash Erase Erasing the flash is also possible. To erase the entire flash memory you can run the following command:

```
idf.py -p PORT erase-flash
```

For erasing the OTA data, if present, you can run this command:

```
idf.py -p PORT erase-otadata
```

The flash erase command can take a while to be done. Do not disconnect your device while the flash erasing is in progress.

Related Documents For advanced users who want to customize the install process:

- [Updating ESP-IDF Tools on Windows](#)
- [Establish Serial Connection with ESP32-P4](#)
- [Eclipse Plugin](#)
- [VSCode Extension](#)
- [IDF Monitor](#)

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Updating ESP-IDF Tools on Windows

Install ESP-IDF Tools Using a Script From the Windows Command Prompt, change to the directory where ESP-IDF is installed. Then run:

```
install.bat
```

For Powershell, change to the directory where ESP-IDF is installed. Then run:

```
install.ps1
```

This downloads and installs the tools necessary to use ESP-IDF. If the specific version of the tool is already installed, no action will be taken. The tools are downloaded and installed into a directory specified during ESP-IDF Tools Installer process. By default, this is `C:\Users\username\.espressif`.

Add ESP-IDF Tools to PATH Using an Export Script ESP-IDF tools installer creates a Start menu shortcut for "ESP-IDF Command Prompt". This shortcut opens a Command Prompt window where all the tools are already available.

In some cases, you may want to work with ESP-IDF in a Command Prompt window which was not started using that shortcut. If this is the case, follow the instructions below to add ESP-IDF tools to PATH.

In the command prompt where you need to use ESP-IDF, change to the directory where ESP-IDF is installed, then execute `export.bat`:

```
cd %userprofile%\esp\esp-idf
export.bat
```

Alternatively in the Powershell where you need to use ESP-IDF, change to the directory where ESP-IDF is installed, then execute `export.ps1`:

```
cd ~/esp/esp-idf
export.ps1
```

When this is done, the tools will be available in this command prompt.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Establish Serial Connection with ESP32-P4

Establishing a serial connection with the ESP32-P4 target device could be done using a USB-to-UART bridge.

Some development boards have the USB-to-UART bridge installed. If a board does not have a bridge then an external bridge may be used.

USB-to-UART Bridge on Development Board For boards with an installed USB-to-UART bridge, the connection between the personal computer and the bridge is USB and between the bridge and ESP32-P4 is UART.




Fig. 6: Development Board with USB-to-UART Bridge

External USB-to-UART Bridge Sometimes the USB-to-UART bridge is external. This is often used in small development boards or finished products when space and costs are crucial.




Fig. 7: External USB-to-UART Bridge

Flash Using UART This section provides guidance on how to establish a serial connection between ESP32-P4 and PC using USB-to-UART Bridge, either installed on the development board or external.

Connect ESP32-P4 to PC Connect the ESP32-P4 board to the PC using the USB cable. If device driver does not install automatically, identify USB-to-UART bridge on your ESP32-P4 board (or external converter dongle), search for drivers in internet and install them.

Below is the list of USB to serial converter chips installed on most of the ESP32-P4 boards produced by Espressif together with links to the drivers:

- CP210x: [CP210x USB to UART Bridge VCP Drivers](#)
- FTDI: [FTDI Virtual COM Port Drivers](#)

Please check the board user guide for specific USB-to-UART bridge chip used. The drivers above are primarily for reference. Under normal circumstances, the drivers should be bundled with an operating system and automatically installed upon connecting the board to the PC.

For devices downloaded using a USB-to-UART bridge, you can run the following command including the optional argument to define the baud rate.

```
idf.py -p PORT [-b BAUD] flash
```

You can change the flasher baud rate by replacing BAUD with the baud rate you need. The default baud rate is 460800.

Note: If the device does not support the auto download mode, you need to get into the download mode manually. To do so, press and hold the `BOOT` button and then press the `RESET` button once. After that release the `BOOT` button.

Check Port on Windows Check the list of identified COM ports in the Windows Device Manager. Disconnect ESP32-P4 and connect it back, to verify which port disappears from the list and then shows back again.

Figures below show serial port for ESP32 DevKitC and ESP32 WROVER KIT

Check Port on Linux and macOS To check the device name for the serial port of your ESP32-P4 board (or external converter dongle), run this command two times, first with the board/dongle unplugged, then with plugged in. The port which appears the second time is the one you need:

Linux




Fig. 8: USB to UART bridge of ESP32-DevKitC in Windows Device Manager




Fig. 9: Two USB Serial Ports of ESP-WROVER-KIT in Windows Device Manager

```
ls /dev/tty*
```

macOS

```
ls /dev/cu.*
```

Note: macOS users: if you do not see the serial port then check you have the USB/serial drivers installed. See Section [Connect ESP32-P4 to PC](#) for links to drivers. For macOS High Sierra (10.13), you may also have to explicitly allow the drivers to load. Open System Preferences -> Security & Privacy -> General and check if there is a message shown here about "System Software from developer ..." where the developer name is Silicon Labs or FTDI.

Adding User to dialout or uucp on Linux The currently logged user should have read and write access the serial port over USB. On most Linux distributions, this is done by adding the user to `dialout` group with the following command:

```
sudo usermod -a -G dialout $USER
```

on Arch Linux this is done by adding the user to `uucp` group with the following command:

```
sudo usermod -a -G uucp $USER
```

Make sure you re-login to enable read and write permissions for the serial port.

Verify Serial Connection Now verify that the serial connection is operational. You can do this using a serial terminal program by checking if you get any output on the terminal after resetting ESP32-P4.

The default console baud rate on ESP32-P4 is 115200.

Windows and Linux In this example, we use [PuTTY SSH Client](#) that is available for both Windows and Linux. You can use other serial programs and set communication parameters like below.

Run terminal and set identified serial port. Baud rate = 115200 (if needed, change this to the default baud rate of the chip in use), data bits = 8, stop bits = 1, and parity = N. Below are example screenshots of setting the port and such transmission parameters (in short described as 115200-8-1-N) on Windows and Linux. Remember to select exactly the same serial port you have identified in steps above.

Then open serial port in terminal and check, if you see any log printed out by ESP32-P4. The log contents depend on application loaded to ESP32-P4, see [Example Output](#). Reset the board if no log has been printed out.

Note: Close the serial terminal after verification that communication is working. If you keep the terminal session open, the serial port will be inaccessible for uploading firmware later.

Note: If there is no log output, check

- if the required power is supplied to ESP32-P4
- if the board was reset after starting the terminal program
- if the selected serial port is the correct one by using the method stated in [Check Port on Windows](#) and [Check Port on Linux and macOS](#)
- if the serial port is not being used by another program
- if the identified port has been selected in serial terminal programs you are using, as stated in [Windows and Linux](#)
- if settings of the serial port in serial terminal programs are applicable to corresponding applications
- if the correct USB connector (UART) is used on the development board
- if your application is expected to output some log




Fig. 10: Setting Serial Communication in PuTTY on Windows




Fig. 11: Setting Serial Communication in PuTTY on Linux

- if the log output has not been disabled (use [hello world application](#) to test)
-

macOS To spare you the trouble of installing a serial terminal program, macOS offers the **screen** command.

- As discussed in [Check port on Linux and macOS](#), run:

```
ls /dev/cu.*
```

- You should see similar output:

```
/dev/cu.Bluetooth-Incoming-Port /dev/cu.SLAB_USBtoUART /dev/cu.SLAB_
↪USBtoUART7
```

- The output varies depending on the type and the number of boards connected to your PC. Then pick the device name of your board and run (if needed, change "115200" to the default baud rate of the chip in use):

```
screen /dev/cu.device_name 115200
```

Replace `device_name` with the name found running `ls /dev/cu.*`.

- What you are looking for is some log displayed by the **screen**. The log contents depend on application loaded to ESP32-P4, see [Example Output](#). To exit the current **screen** session, type `Ctrl-A + K`.

Note: Do not forget to **exit the current screen session** after verifying that the communication is working. If you fail to do it and just close the terminal window, the serial port will be inaccessible for uploading firmware later.

Example Output An example log is shown below. Reset the board if you do not see anything.

```
ets Jun  8 2016 00:22:57

rst:0x5 (DEEPSLEEP_RESET),boot:0x13 (SPI_FAST_FLASH_BOOT)
ets Jun  8 2016 00:22:57

rst:0x7 (TGWDT_SYS_RESET),boot:0x13 (SPI_FAST_FLASH_BOOT)
configsip: 0, SPIWP:0x00
clk_drv:0x00,q_drv:0x00,d_drv:0x00,cs0_drv:0x00,hd_drv:0x00,wp_drv:0x00
mode:DIO, clock div:2
load:0x3fff0008,len:8
load:0x3fff0010,len:3464
load:0x40078000,len:7828
load:0x40080000,len:252
entry 0x40080034
I (44) boot: ESP-IDF v2.0-rc1-401-gf9fba35 2nd stage bootloader
I (45) boot: compile time 18:48:10
...
```

If you can see readable log output, it means serial connection is working and you are ready to proceed with installation and finally upload an application to ESP32-P4.

Note: For some serial port wiring configurations, the serial RTS & DTR pins need to be disabled in the terminal program before the ESP32-P4 booting and producing serial output. This depends on the hardware itself, most development boards (including all Espressif boards) *do not* have this issue. The issue is present if RTS & DTR are wired directly to the EN & GPIO0 pins. See the [esptool documentation](#) for more details.

If you got here from [Step 5. First Steps on ESP-IDF](#) when installing s/w for ESP32-P4 development, then you can continue with [Step 5. First Steps on ESP-IDF](#).

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Flashing Troubleshooting

Failed to Connect If you run the given command and see errors such as "Failed to connect", there might be several reasons for this. One of the reasons might be issues encountered by `esptool.py`, the utility that is called by the build system to reset the chip, interact with the ROM bootloader, and flash firmware. One simple solution to try is to manually reset as described below. If it does not help, you can find more details about possible issues in the [esptool troubleshooting](#) page.

`esptool.py` resets ESP32-P4 automatically by asserting DTR and RTS control lines of the USB-to-UART bridge, i.e., FTDI or CP210x (for more information, see [Establish Serial Connection with ESP32-P4](#)). The DTR and RTS control lines are in turn connected to `[NEEDS TO BE UPDATED]` and `CHIP_PU (EN)` pins of ESP32-P4, thus changes in the voltage levels of DTR and RTS will boot ESP32-P4 into Firmware Download mode. As an example, check the [schematic](#) for the ESP32 DevKitC development board.

In general, you should have no problems with the [official esp-idf development boards](#). However, `esptool.py` is not able to reset your hardware automatically in the following cases:

- Your hardware does not have the DTR and RTS lines connected to `[NEEDS TO BE UPDATED]` and `CHIP_PU`.
- The DTR and RTS lines are configured differently.
- There are no such serial control lines at all.

Depending on the kind of hardware you have, it may also be possible to manually put your ESP32-P4 board into Firmware Download mode (reset).

- For development boards produced by Espressif, this information can be found in the respective getting started guides or user guides. For example, to manually reset an ESP-IDF development board, hold down the `Boot` button (`[NEEDS TO BE UPDATED]`) and press the `EN` button (`CHIP_PU`).
- For other types of hardware, try pulling `[NEEDS TO BE UPDATED]` down.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

IDF Monitor

IDF Monitor uses the [esp-idf-monitor](#) package as a serial terminal program which relays serial data to and from the target device's serial port. It also provides some ESP-IDF-specific features.

IDF Monitor can be launched from an ESP-IDF project by running `idf.py monitor`.

Keyboard Shortcuts For easy interaction with IDF Monitor, use the keyboard shortcuts given in the table. These keyboard shortcuts can be customized, for more details see [Configuration File](#) section.

Keyboard Shortcut	Action	Description
Ctrl +]	Exit the program	
Ctrl + T	Menu escape key	Press and follow it by one of the keys given below.
• Ctrl + T	Send the menu character itself to remote	
• Ctrl +]	Send the exit character itself to remote	
• Ctrl + P	Reset target into bootloader to pause app via RTS line	Resets the target, into bootloader via the RTS line (if connected), so that the board runs nothing. Useful when you need to wait for another device to startup.
• Ctrl + R	Reset target board via RTS	Resets the target board and re-starts the application via the RTS line (if connected).
• Ctrl + F	Build and flash the project	Pauses idf_monitor to run the project flash target, then resumes idf_monitor. Any changed source files are recompiled and then re-flashed. Target encrypted-flash is run if idf_monitor was started with argument -E.
• Ctrl + A (or A)	Build and flash the app only	Pauses idf_monitor to run the app-flash target, then resumes idf_monitor. Similar to the flash target, but only the main app is built and re-flashed. Target encrypted-app-flash is run if idf_monitor was started with argument -E.
• Ctrl + Y	Stop/resume log output printing on screen	Discards all incoming serial data while activated. Allows to quickly pause and examine log output without quitting the monitor.
• Ctrl + L	Stop/resume log output saved to file	Creates a file in the project directory and the output is written to that file until this is disabled with the same keyboard shortcut (or IDF Monitor exits).
• Ctrl + I (or I)	Stop/resume printing timestamps	IDF Monitor can print a timestamp in the beginning of each line. The timestamp format can be changed by the --timestamp-format command line argument.
• Ctrl + H (or H)	Display all keyboard shortcuts	
• Ctrl + X (or X)	Exit the program	
Ctrl + C	Interrupt running application	Pauses IDF Monitor and runs GDB project debugger to debug the application at runtime. This requires CONFIG_ESP_SYSTEM_GDBSTUB_RUNTIME option to be enabled.

Any keys pressed, other than Ctrl-] and Ctrl-T, will be sent through the serial port.

ESP-IDF-specific Features

Automatic Address Decoding Whenever the chip outputs a hexadecimal address that points to executable code, IDF monitor looks up the location in the source code (file name and line number) and prints the location on the next line in yellow.

If an ESP-IDF app crashes and panics, a register dump and backtrace are produced, such as the following:

```
abort() was called at PC 0x42067cd5 on core 0
```

(continues on next page)

(continued from previous page)

```

Stack dump detected
Core 0 register dump:
MEPC      : 0x40386488  RA      : 0x40386b02  SP      : 0x3fc9a350  GP      : 0x3fc923c0
TP        : 0xa5a5a5a5  T0      : 0x37363534  T1      : 0x7271706f  T2      : 0x33323130
S0/FP     : 0x00000004  S1      : 0x3fc9a3b4  A0      : 0x3fc9a37c  A1      : 0x3fc9a3b2
A2        : 0x00000000  A3      : 0x3fc9a3a9  A4      : 0x00000001  A5      : 0x3fc99000
A6        : 0x7a797877  A7      : 0x76757473  S2      : 0xa5a5a5a5  S3      : 0xa5a5a5a5
S4        : 0xa5a5a5a5  S5      : 0xa5a5a5a5  S6      : 0xa5a5a5a5  S7      : 0xa5a5a5a5
S8        : 0xa5a5a5a5  S9      : 0xa5a5a5a5  S10     : 0xa5a5a5a5  S11     : 0xa5a5a5a5
T3        : 0x6e6d6c6b  T4      : 0x6a696867  T5      : 0x66656463  T6      : 0x62613938
MSTATUS   : 0x00001881  MTVEC   : 0x40380001  MCAUSE  : 0x00000007  MTVAL   : 0x00000000

MHARTID   : 0x00000000

Stack memory:
3fc9a350: 0xa5a5a5a5 0xa5a5a5a5 0x3fc9a3b0 0x403906cc 0xa5a5a5a5 0xa5a5a5a5
3fc9a370: 0x3fc9a3b4 0x3fc9423c 0x3fc9a3b0 0x726f6261 0x20292874 0x20736177
3fc9a390: 0x43502074 0x34783020 0x37363032 0x20356463 0x63206e6f 0x2065726f
3fc9a3b0: 0x00000030 0x36303234 0x35646337 0x3c093700 0x0000002a 0xa5a5a5a5
3fc9a3d0: 0x00000001 0x3c0917f8 0x3c0937d4 0x0000002a 0xa5a5a5a5 0xa5a5a5a5
3fc9a3f0: 0x0001f24c 0x000000c8 0x00000000 0x0001c200 0xffffffff 0xffffffff
3fc9a410: 0x00001000 0x00000002 0x3c093818 0x3fccb470 0xa5a5a5a5 0xa5a5a5a5
.....

```

IDF Monitor adds more details to the dump by analyzing the stack dump:

```

abort() was called at PC 0x42067cd5 on core 0
0x42067cd5: __assert_func at /builds/idf/crosstool-NG/.build/riscv32-esp-elf/src/
newlib/newlib/libc/stdlib/assert.c:62 (discriminator 8)

Stack dump detected
Core 0 register dump:
MEPC      : 0x40386488  RA      : 0x40386b02  SP      : 0x3fc9a350  GP      : 0x3fc923c0
0x40386488: panic_abort at /home/marius/esp-idf_2/components/esp_system/panic.c:367

0x40386b02: rtos_int_enter at /home/marius/esp-idf_2/components/freertos/port/
riscv/portasm.S:35

TP        : 0xa5a5a5a5  T0      : 0x37363534  T1      : 0x7271706f  T2      : 0x33323130
S0/FP     : 0x00000004  S1      : 0x3fc9a3b4  A0      : 0x3fc9a37c  A1      : 0x3fc9a3b2
A2        : 0x00000000  A3      : 0x3fc9a3a9  A4      : 0x00000001  A5      : 0x3fc99000

```

(continues on next page)

(continued from previous page)

```

A6      : 0x7a797877  A7      : 0x76757473  S2      : 0xa5a5a5a5  S3      : _
↪0xa5a5a5a5
S4      : 0xa5a5a5a5  S5      : 0xa5a5a5a5  S6      : 0xa5a5a5a5  S7      : _
↪0xa5a5a5a5
S8      : 0xa5a5a5a5  S9      : 0xa5a5a5a5  S10     : 0xa5a5a5a5  S11     : _
↪0xa5a5a5a5
T3      : 0x6e6d6c6b  T4      : 0x6a696867  T5      : 0x66656463  T6      : _
↪0x62613938
MSTATUS : 0x00001881  MTVEC   : 0x40380001  MCAUSE  : 0x00000007  MTVAL   : _
↪0x00000000

MHARTID : 0x00000000

Backtrace:
panic_abort (details=details@entry=0x3fc9a37c "abort() was called at PC 0x42067cd5_
↪on core 0") at /home/marius/esp-idf_2/components/esp_system/panic.c:367
367      *((int *) 0) = 0; // NOLINT(clang-analyzer-core.NullDereference) should be_
↪an invalid operation on targets
#0  panic_abort (details=details@entry=0x3fc9a37c "abort() was called at PC_
↪0x42067cd5 on core 0") at /home/marius/esp-idf_2/components/esp_system/panic.
↪c:367
#1  0x40386b02 in esp_system_abort (details=details@entry=0x3fc9a37c "abort() was_
↪called at PC 0x42067cd5 on core 0") at /home/marius/esp-idf_2/components/esp_
↪system/system_api.c:108
#2  0x403906cc in abort () at /home/marius/esp-idf_2/components/newlib/abort.c:46
#3  0x42067cd8 in __assert_func (file=file@entry=0x3c0937f4 "", line=line@entry=42,
↪ func=func@entry=0x3c0937d4 <__func__.8540> "", _
↪failedexpr=failedexpr@entry=0x3c0917f8 "") at /builds/idf/crosstool-NG/.build/
↪riscv32-esp-elf/src/newlib/newlib/libc/stdlib/assert.c:62
#4  0x4200729e in app_main () at ../main/iperf_example_main.c:42
#5  0x42086cd6 in main_task (args=<optimized out>) at /home/marius/esp-idf_2/
↪components/freertos/port/port_common.c:133
#6  0x40389f3a in vPortEnterCritical () at /home/marius/esp-idf_2/components/
↪freertos/port/riscv/port.c:129

```

To decode each address, IDF Monitor runs the following command in the background:

```
riscv32-esp-elf-addr2line -pfiaC -e build/PROJECT.elf ADDRESS
```

If an address is not matched in the app source code, IDF monitor also checks the ROM code. Instead of printing the source file name and line number, only the function name followed by in ROM is displayed:

```

abort() was called at PC 0x400481c1 on core 0
0x400481c1: ets_rsa_pss_verify in ROM

Stack dump detected
Core 0 register dump:
MEPC    : 0x4038051c  RA      : 0x40383840  SP      : 0x3fc8f6b0  GP      : _
↪0x3fc8b000
0x4038051c: panic_abort at /Users/espressif/esp-idf/components/esp_system/panic.
↪c:452
0x40383840: __ubsan_include at /Users/espressif/esp-idf/components/esp_system/
↪ubsan.c:313

TP      : 0x3fc8721c  T0      : 0x37363534  T1      : 0x7271706f  T2      : _
↪0x33323130
S0/FP   : 0x00000004  S1      : 0x3fc8f714  A0      : 0x3fc8f6dc  A1      : _
↪0x3fc8f712
A2      : 0x00000000  A3      : 0x3fc8f709  A4      : 0x00000001  A5      : _
↪0x3fc8c000
A6      : 0x7a797877  A7      : 0x76757473  S2      : 0x00000000  S3      : _
↪0x3fc8f750

```

(continues on next page)

(continued from previous page)

```

S4      : 0x3fc8f7e4  S5      : 0x00000000  S6      : 0x400481b0  S7      : ↵
↵0x3c025841
0x400481b0: ets_rsa_pss_verify in ROM
.....

```

The ROM ELF file is automatically loaded from a location based on the `IDF_PATH` and `ESP_ROM_ELF_DIR` environment variables. This can be overridden by calling `esp_idf_monitor` and providing a path to a specific ROM ELF file: `python -m esp_idf_monitor --rom-elf-file [path to ROM ELF file]`.

Note: Set environment variable `ESP_MONITOR_DECODE` to 0 or call `esp_idf_monitor` with specific command line option: `python -m esp_idf_monitor --disable-address-decoding` to disable address decoding.

Target Reset on Connection By default, IDF Monitor will reset the target when connecting to it. The reset of the target chip is performed using the DTR and RTS serial lines. To prevent IDF Monitor from automatically resetting the target on connection, call IDF Monitor with the `--no-reset` option (e.g., `idf.py monitor --no-reset`).

Note: The `--no-reset` option applies the same behavior even when connecting IDF Monitor to a particular port (e.g., `idf.py monitor --no-reset -p [PORT]`).

Launching GDB with GDBStub GDBStub is a useful runtime debugging feature that runs on the target and connects to the host over the serial port to receive debugging commands. GDBStub supports commands such as reading memory and variables, examining call stack frames etc. Although GDBStub is less versatile than JTAG debugging, it does not require any special hardware (such as a JTAG to USB bridge) as communication is done entirely over the serial port.

A target can be configured to run GDBStub in the background by setting the [CONFIG_ESP_SYSTEM_GDBSTUB_RUNTIME](#). GDBStub will run in the background until a `Ctrl+C` message is sent over the serial port and causes the GDBStub to break (i.e., stop the execution of) the program, thus allowing GDBStub to handle debugging commands.

Furthermore, the panic handler can be configured to run GDBStub on a crash by setting the [CONFIG_ESP_SYSTEM_PANIC](#) to GDBStub on panic. When a crash occurs, GDBStub will output a special string pattern over the serial port to indicate that it is running.

In both cases (i.e., sending the `Ctrl+C` message, or receiving the special string pattern), IDF Monitor will automatically launch GDB in order to allow the user to send debugging commands. After GDB exits, the target is reset via the RTS serial line. If this line is not connected, users can reset their target (by pressing the board's Reset button).

Note: In the background, IDF Monitor runs the following command to launch GDB:

```

riscv32-esp-elf-gdb -ex "set serial baud BAUD" -ex "target remote PORT" -ex ↵
↵interrupt build/PROJECT.elf :idf_target:`Hello NAME chip`

```

Output Filtering IDF monitor can be invoked as `idf.py monitor --print-filter="xyz"`, where `--print-filter` is the parameter for output filtering. The default value is an empty string, which means that everything is printed. Filtering can also be configured using the `ESP_IDF_MONITOR_PRINT_FILTER` environment variable.

Note: When using both the environment variable `ESP_IDF_MONITOR_PRINT_FILTER` and the argument `--print-filter`, the setting from the CLI argument will take precedence.

Restrictions on what to print can be specified as a series of `<tag>:<log_level>` items where `<tag>` is the tag string and `<log_level>` is a character from the set {N, E, W, I, D, V, *} referring to a level for [logging](#).

For example, `--print_filter="tag1:W"` matches and prints only the outputs written with `ESP_LOGW("tag1", ...)` or at lower verbosity level, i.e., `ESP_LOGE("tag1", ...)`. Not specifying a `<log_level>` or using `*` defaults to a Verbose level.

Note: Use primary logging to disable at compilation the outputs you do not need through the [logging library](#). Output filtering with the IDF monitor is a secondary solution that can be useful for adjusting the filtering options without recompiling the application.

Your app tags must not contain spaces, asterisks `*`, or colons `:` to be compatible with the output filtering feature.

If the last line of the output in your app is not followed by a carriage return, the output filtering might get confused, i.e., the monitor starts to print the line and later finds out that the line should not have been written. This is a known issue and can be avoided by always adding a carriage return (especially when no output follows immediately afterwards).

Examples of Filtering Rules:

- `*` can be used to match any tags. However, the string `--print_filter="*:I tag1:E"` with regards to `tag1` prints errors only, because the rule for `tag1` has a higher priority over the rule for `*`.
- The default (empty) rule is equivalent to `*:V` because matching every tag at the Verbose level or lower means matching everything.
- `"*:N"` suppresses not only the outputs from logging functions, but also the prints made by `printf`, etc. To avoid this, use `*:E` or a higher verbosity level.
- Rules `"tag1:V"`, `"tag1:v"`, `"tag1:"`, `"tag1:*"`, and `"tag1"` are equivalent.
- Rule `"tag1:W tag1:E"` is equivalent to `"tag1:E"` because any consequent occurrence of the same tag name overwrites the previous one.
- Rule `"tag1:I tag2:W"` only prints `tag1` at the Info verbosity level or lower and `tag2` at the Warning verbosity level or lower.
- Rule `"tag1:I tag2:W tag3:N"` is essentially equivalent to the previous one because `tag3:N` specifies that `tag3` should not be printed.
- `tag3:N` in the rule `"tag1:I tag2:W tag3:N *:V"` is more meaningful because without `tag3:N` the `tag3` messages could have been printed; the errors for `tag1` and `tag2` will be printed at the specified (or lower) verbosity level and everything else will be printed by default.

A More Complex Filtering Example The following log snippet was acquired without any filtering options:

```
load:0x40078000,len:13564
entry 0x40078d4c
E (31) esp_image: image at 0x30000 has invalid magic byte
W (31) esp_image: image at 0x30000 has invalid SPI mode 255
E (39) boot: Factory app partition is not bootable
I (568) cpu_start: Pro cpu up.
I (569) heap_init: Initializing. RAM available for dynamic allocation:
I (603) cpu_start: Pro cpu start user code
D (309) light_driver: [light_init, 74]:status: 1, mode: 2
D (318) vfs: esp_vfs_register_fd_range is successful for range <54; 64) and VFS ID_
↪1
I (328) wifi: wifi driver task: 3ffdbf84, prio:23, stack:4096, core=0
```

The captured output for the filtering options `--print_filter="wifi esp_image:E light_driver:I"` is given below:

```
E (31) esp_image: image at 0x30000 has invalid magic byte
I (328) wifi: wifi driver task: 3ffdbf84, prio:23, stack:4096, core=0
```

The options `--print_filter="light_driver:D esp_image:N boot:N cpu_start:N vfs:N wifi:N *:V"` show the following output:

```
load:0x40078000,len:13564
entry 0x40078d4c
I (569) heap_init: Initializing. RAM available for dynamic allocation:
D (309) light_driver: [light_init, 74]:status: 1, mode: 2
```

Configuration File `esp-idf-monitor` is using [C0 control codes](#) to interact with the console. Characters from the config file are converted to their C0 control codes. Available characters include the English alphabet (A-Z) and special symbols: `[, ], \, ^, _`.

Warning: Please note that some characters may not work on all platforms or can be already reserved as a shortcut for something else. Use this feature with caution!

File Location The default name for a configuration file is `esp-idf-monitor.cfg`. First, the same directory `esp-idf-monitor` is being run if is inspected.

If a configuration file is not found here, the current user's OS configuration directory is inspected next:

- Linux: `/home/<user>/ .config/esp-idf-monitor/`
- MacOS `/Users/<user>/ .config/esp-idf-monitor/`
- Windows: `c:\Users\<user>\AppData\Local\esp-idf-monitor\`

If a configuration file is still not found, the last inspected location is the home directory:

- Linux: `/home/<user>/`
- MacOS `/Users/<user>/`
- Windows: `c:\Users\<user>\`

On Windows, the home directory can be set with the `HOME` or `USERPROFILE` environment variables. Therefore, the Windows configuration directory location also depends on these.

A different location for the configuration file can be specified with the `ESP_IDF_MONITOR_CFGFILE` environment variable, e.g., `ESP_IDF_MONITOR_CFGFILE = ~/custom_config.cfg`. This overrides the search priorities described above.

`esp-idf-monitor` will read settings from other usual configuration files if no other configuration file is used. It automatically reads from `setup.cfg` or `tox.ini` if they exist.

Configuration Options Below is a table listing the available configuration options:

Option Name	Description	Default Value
<code>menu_key</code>	Key to access the main menu.	T
<code>exit_key</code>	Key to exit the monitor.	]
<code>chip_reset_key</code>	Key to initiate a chip reset.	R
<code>recompile_upload_key</code>	Key to recompile and upload.	F
<code>recompile_upload_app_key</code>	Key to recompile and upload just the application.	A
<code>toggle_output_key</code>	Key to toggle the output display.	Y
<code>toggle_log_key</code>	Key to toggle the logging feature.	L
<code>toggle_timestamp_key</code>	Key to toggle timestamp display.	I
<code>chip_reset_bootloader_key</code>	Key to reset the chip to bootloader mode.	P
<code>exit_menu_key</code>	Key to exit the monitor from the menu.	X
<code>skip_menu_key</code>	Pressing the menu key can be skipped for menu commands.	False

Syntax The configuration file is in .ini file format: it must be introduced by an `[esp-idf-monitor]` header to be recognized as valid. This section then contains `name = value` entries. Lines beginning with `#` or `;` are ignored as comments.


```
# esp-idf-monitor.cfg file to configure internal settings of esp-idf-monitor
[esp-idf-monitor]
menu_key = T
exit_key = ]
chip_reset_key = R
recompile_upload_key = F
recompile_upload_app_key = A
toggle_output_key = Y
toggle_log_key = L
toggle_timestamp_key = I
chip_reset_bootloader_key = P
exit_menu_key = X
skip_menu_key = False
```

Known Issues with IDF Monitor

Issues Observed on Windows

- Arrow keys, as well as some other keys, do not work in GDB due to Windows Console limitations.
- Occasionally, when "idf.py" exits, it might stall for up to 30 seconds before IDF Monitor resumes.
- When "gdb" is run, it might stall for a short time before it begins communicating with the GDBStub.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Standard Toolchain Setup for Linux and macOS

Installation Step by Step This is a detailed roadmap to walk you through the installation process.

Setting up Development Environment These are the steps for setting up the ESP-IDF for your ESP32-P4.

- [Step 1. Install Prerequisites](#)
- [Step 2. Get ESP-IDF](#)
- [Step 3. Set up the Tools](#)
- [Step 4. Set up the Environment Variables](#)
- [Step 5. First Steps on ESP-IDF](#)

Step 1. Install Prerequisites In order to use ESP-IDF with the ESP32-P4, you need to install some software packages based on your Operating System. This setup guide helps you on getting everything installed on Linux and macOS based systems.

For Linux Users To compile using ESP-IDF, you need to get the following packages. The command to run depends on which distribution of Linux you are using:

- Ubuntu and Debian:

```
sudo apt-get install git wget flex bison gperf python3 python3-pip python3-  
➔venv cmake ninja-build ccache libffi-dev libssl-dev dfu-util libusb-1.0-0
```

- CentOS 7 & 8:

```
sudo yum -y update && sudo yum install git wget flex bison gperf python3 cmake_
↪ninja-build ccache dfu-util libusb
```

CentOS 7 is still supported but CentOS version 8 is recommended for a better user experience.

- Arch:

```
sudo pacman -S --needed gcc git make flex bison gperf python cmake ninja_
↪ccache dfu-util libusb
```

Note:

- CMake version 3.16 or newer is required for use with ESP-IDF. Run "tools/idf_tools.py install cmake" to install a suitable version if your OS versions does not have one.
 - If you do not see your Linux distribution in the above list then please check its documentation to find out which command to use for package installation.
-

For macOS Users ESP-IDF uses the version of Python installed by default on macOS.

- Install CMake & Ninja build:
 - If you have [HomeBrew](#), you can run:

```
brew install cmake ninja dfu-util
```

- If you have [MacPorts](#), you can run:

```
sudo port install cmake ninja dfu-util
```

- Otherwise, consult the [CMake](#) and [Ninja](#) home pages for macOS installation downloads.
 - It is strongly recommended to also install [ccache](#) for faster builds. If you have [HomeBrew](#), this can be done via `brew install ccache` or `sudo port install ccache` on [MacPorts](#).
-

Note: If an error like this is shown during any step:

```
xcrun: error: invalid active developer path (/Library/Developer/CommandLineTools), _
↪missing xcrun at: /Library/Developer/CommandLineTools/usr/bin/xcrun
```

Then you need to install the XCode command line tools to continue. You can install these by running `xcode-select --install`.

Apple M1 Users If you use Apple M1 platform and see an error like this:

```
WARNING: directory for tool xtensa-esp32-elf version esp-2021r2-patch3-8.4.0 is_
↪present, but tool was not found
ERROR: tool xtensa-esp32-elf has no installed versions. Please run 'install.sh' to_
↪install it.
```

or:

```
zsh: bad CPU type in executable: ~/.espressif/tools/xtensa-esp32-elf/esp-2021r2-_
↪patch3-8.4.0/xtensa-esp32-elf/bin/xtensa-esp32-elf-gcc
```

Then you need to install Apple Rosetta 2 by running

```
/usr/sbin/softwareupdate --install-rosetta --agree-to-license
```

Installing Python 3 Based on macOS [Catalina 10.15 release notes](#), use of Python 2.7 is not recommended and Python 2.7 is not included by default in future versions of macOS. Check what Python you currently have:

```
python --version
```

If the output is like `Python 2.7.17`, your default interpreter is Python 2.7. If so, also check if Python 3 is not already installed on your computer:

```
python3 --version
```

If the above command returns an error, it means Python 3 is not installed.

Below is an overview of the steps to install Python 3.

- Installing with [HomeBrew](#) can be done as follows:

```
brew install python3
```

- If you have [MacPorts](#), you can run:

```
sudo port install python38
```

Step 2. Get ESP-IDF To build applications for the ESP32-P4, you need the software libraries provided by Espressif in [ESP-IDF repository](#).

To get ESP-IDF, navigate to your installation directory and clone the repository with `git clone`, following instructions below specific to your operating system.

Open Terminal, and run the following commands:

```
mkdir -p ~/esp
cd ~/esp
git clone -b v5.2.7 --recursive https://github.com/espressif/esp-idf.git
```

ESP-IDF is downloaded into `~/esp/esp-idf`.

Consult [ESP-IDF Versions](#) for information about which ESP-IDF version to use in a given situation.

Step 3. Set up the Tools Aside from the ESP-IDF, you also need to install the tools used by ESP-IDF, such as the compiler, debugger, Python packages, etc, for projects supporting ESP32-P4.

```
cd ~/esp/esp-idf
./install.sh esp32p4
```

or with Fish shell

```
cd ~/esp/esp-idf
./install.fish esp32p4
```

The above commands install tools for ESP32-P4 only. If you intend to develop projects for more chip targets then you should list all of them and run for example:

```
cd ~/esp/esp-idf
./install.sh esp32,esp32s2
```

or with Fish shell

```
cd ~/esp/esp-idf
./install.fish esp32,esp32s2
```

In order to install tools for all supported targets please run the following command:

```
cd ~/esp/esp-idf
./install.sh all
```

or with Fish shell

```
cd ~/esp/esp-idf
./install.fish all
```

Note: For macOS users, if an error like this is shown during any step:

```
<urlopen error [SSL: CERTIFICATE_VERIFY_FAILED] certificate verify failed: unable
to get local issuer certificate (_ssl.c:xxx)
```

You may run `Install Certificates.command` in the Python folder of your computer to install certificates. For details, see [Download Error While Installing ESP-IDF Tools](#).

Alternative File Downloads The tools installer downloads a number of files attached to GitHub Releases. If accessing GitHub is slow then it is possible to set an environment variable to prefer Espressif's download server for GitHub asset downloads.

Note: This setting only controls individual tools downloaded from GitHub releases, it does not change the URLs used to access any Git repositories.

To prefer the Espressif download server when installing tools, use the following sequence of commands when running `install.sh`:

```
cd ~/esp/esp-idf
export IDF_GITHUB_ASSETS="dl.espressif.com/github_assets"
./install.sh
```

Note: For users in China, we recommend using our download server located in China for faster download speed. Alternatively, you can set as well the environment variable `PIP_INDEX_URL` to the China mirror of PyPI.

```
cd ~/esp/esp-idf
export IDF_GITHUB_ASSETS="dl.espressif.cn/github_assets"
export PIP_INDEX_URL="https://mirrors.tuna.tsinghua.edu.cn/pypi/web/simple"
./install.sh
```

Customizing the Tools Installation Path The scripts introduced in this step install compilation tools required by ESP-IDF inside the user home directory: `$HOME/.espressif` on Linux. If you wish to install the tools into a different directory, set the environment variable `IDF_TOOLS_PATH` before running the installation scripts. Make sure that your user account has sufficient permissions to read and write this path.

If changing the `IDF_TOOLS_PATH`, make sure it is set to the same value every time the Install script (`install.bat`, `install.ps1` or `install.sh`) and an Export script (`export.bat`, `export.ps1` or `export.sh`) are executed.

Step 4. Set up the Environment Variables The installed tools are not yet added to the `PATH` environment variable. To make the tools usable from the command line, some environment variables must be set. ESP-IDF provides another script which does that.

In the terminal where you are going to use ESP-IDF, run:

```
. $HOME/esp/esp-idf/export.sh
```

or for fish (supported only since fish version 3.0.0):

```
. $HOME/esp/esp-idf/export.fish
```

Note the space between the leading dot and the path!

If you plan to use esp-idf frequently, you can create an alias for executing `export .sh`:

1. Copy and paste the following command to your shell's profile (`.profile`, `.bashrc`, `.zprofile`, etc.)

```
alias get_idf='. $HOME/esp/esp-idf/export.sh'
```

2. Refresh the configuration by restarting the terminal session or by running `source [path to profile]`, for example, `source ~/.bashrc`.

Now you can run `get_idf` to set up or refresh the esp-idf environment in any terminal session.

Technically, you can add `export.sh` to your shell's profile directly; however, it is not recommended. Doing so activates IDF virtual environment in every terminal session (including those where IDF is not needed), defeating the purpose of the virtual environment and likely affecting other software.

Step 5. First Steps on ESP-IDF Now since all requirements are met, the next topic will guide you on how to start your first project.

This guide helps you on the first steps using ESP-IDF. Follow this guide to start a new project on the ESP32-P4 and build, flash, and monitor the device output.

Note: If you have not yet installed ESP-IDF, please go to [Installation](#) and follow the instruction in order to get all the software needed to use this guide.

Start a Project Now you are ready to prepare your application for ESP32-P4. You can start with [get-started/hello_world](#) project from [examples](#) directory in ESP-IDF.

Important: The ESP-IDF build system does not support spaces in the paths to either ESP-IDF or to projects.

Copy the project [get-started/hello_world](#) to `~/esp` directory:

```
cd ~/esp
cp -r $IDF_PATH/examples/get-started/hello_world .
```

Note: There is a range of example projects in the [examples](#) directory in ESP-IDF. You can copy any project in the same way as presented above and run it. It is also possible to build examples in-place without copying them first.

Connect Your Device Now connect your ESP32-P4 board to the computer and check under which serial port the board is visible.

Serial ports have the following naming patterns:

- **Linux:** starting with `/dev/tty`
- **macOS:** starting with `/dev/cu`.

If you are not sure how to check the serial port name, please refer to [Establish Serial Connection with ESP32-P4](#) for full details.

Note: Keep the port name handy as it is needed in the next steps.

Configure Your Project Navigate to your `hello_world` directory, set ESP32-P4 as the target, and run the project configuration utility `menuconfig`.

```
cd ~/esp/hello_world
idf.py set-target esp32p4
idf.py menuconfig
```

After opening a new project, you should first set the target with `idf.py set-target esp32p4`. Note that existing builds and configurations in the project, if any, are cleared and initialized in this process. The target may be saved in the environment variable to skip this step at all. See [Select the Target Chip: set-target](#) for additional information.

If the previous steps have been done correctly, the following menu appears:




Fig. 12: Project configuration - Home window

You are using this menu to set up project specific variables, e.g., Wi-Fi network name and password, the processor speed, etc. Setting up the project with `menuconfig` may be skipped for "hello_world", since this example runs with default configuration.

Note: The colors of the menu could be different in your terminal. You can change the appearance with the option `--style`. Please run `idf.py menuconfig --help` for further information.

Build the Project Build the project by running:

```
idf.py build
```

This command compiles the application and all ESP-IDF components, then it generates the bootloader, partition table, and application binaries.

```
$ idf.py build
Running cmake in directory /path/to/hello_world/build
Executing "cmake -G Ninja --warn-uninitialized /path/to/hello_world"...
```

(continues on next page)

(continued from previous page)

```
Warn about uninitialized values.
-- Found Git: /usr/bin/git (found version "2.17.0")
-- Building empty aws_iot component due to configuration
-- Component names: ...
-- Component paths: ...

... (more lines of build system output)

[527/527] Generating hello_world.bin
esptool.py v2.3.1

Project build complete. To flash, run this command:
../../components/esptool_py/esptool/esptool.py -p (PORT) -b 921600 write_flash -
↪-flash_mode dio --flash_size detect --flash_freq 40m 0x10000 build/hello_world.
↪bin build 0x1000 build/bootloader/bootloader.bin 0x8000 build/partition_table/
↪partition-table.bin
or run 'idf.py -p PORT flash'
```

If there are no errors, the build finishes by generating the firmware binary .bin files.

Flash onto the Device To flash the binaries that you just built for the ESP32-P4 in the previous step, you need to run the following command:

```
idf.py -p PORT flash
```

Replace PORT with your ESP32-P4 board's USB port name. If the PORT is not defined, the *idf.py* will try to connect automatically using the available USB ports.

For more information on *idf.py* arguments, see *idf.py*.

Note: The option *flash* automatically builds and flashes the project, so running *idf.py build* is not necessary.

Encountered Issues While Flashing? See the "Additional Tips" below. You can also refer to *Flashing Troubleshooting* page or *Establish Serial Connection with ESP32-P4* for more detailed information.

Normal Operation When flashing, you will see the output log similar to the following:

If there are no issues by the end of the flash process, the board will reboot and start up the "hello_world" application.

If you would like to use the Eclipse or VS Code IDE instead of running *idf.py*, check out *Eclipse Plugin*, *VSCode Extension*.

Monitor the Output To check if "hello_world" is indeed running, type *idf.py -p PORT monitor* (Do not forget to replace PORT with your serial port name).

This command launches the *IDF Monitor* application:

```
$ idf.py -p <PORT> monitor
Running idf_monitor in directory [...]/esp/hello_world/build
Executing "python [...]/esp-idf/tools/idf_monitor.py -b 115200 [...]/esp/hello_
↪world/build/hello_world.elf"...
--- idf_monitor on <PORT> 115200 ---
--- Quit: Ctrl+] | Menu: Ctrl+T | Help: Ctrl+T followed by Ctrl+H ---
ets Jun  8 2016 00:22:57

rst:0x1 (POWERON_RESET),boot:0x13 (SPI_FAST_FLASH_BOOT)
ets Jun  8 2016 00:22:57
...
```

After startup and diagnostic logs scroll up, you should see "Hello world!" printed out by the application.

```
...
Hello world!
Restarting in 10 seconds...
This is esp32p4 chip with 2 CPU core(s), [NEEDS TO BE UPDATED]
Minimum free heap size: [NEEDS TO BE UPDATED] bytes
Restarting in 9 seconds...
Restarting in 8 seconds...
Restarting in 7 seconds...
```

To exit IDF monitor use the shortcut `Ctrl+J`.

Note: You can combine building, flashing and monitoring into one step by running:

```
idf.py -p PORT flash monitor
```

See also:

- [IDF Monitor](#) for handy shortcuts and more details on using IDF monitor.
- [idf.py](#) for a full reference of `idf.py` commands and options.

That is all that you need to get started with ESP32-P4!

Now you are ready to try some other [examples](#), or go straight to developing your own applications.

Important: Some of examples do not support ESP32-P4 because required hardware is not included in ESP32-P4 so it cannot be supported.

If building an example, please check the README file for the Supported Targets table. If this is present including ESP32-P4 target, or the table does not exist at all, the example will work on ESP32-P4.

Additional Tips

Permission Denied Issue With some Linux distributions, you may get the error message similar to `Could not open port <PORT>: Permission denied: '<PORT>'` when flashing the ESP32-P4. *This can be solved by adding the current user to the specific group*, such as `dialout` or `uucp` group.

Python Compatibility ESP-IDF supports Python 3.8 or newer. It is recommended to upgrade your operating system to a recent version satisfying this requirement. Other options include the installation of Python from [sources](#) or the use of a Python version management system such as [pyenv](#).

Flash Erase Erasing the flash is also possible. To erase the entire flash memory you can run the following command:

```
idf.py -p PORT erase-flash
```

For erasing the OTA data, if present, you can run this command:

```
idf.py -p PORT erase-otadata
```

The flash erase command can take a while to be done. Do not disconnect your device while the flash erasing is in progress.

Tip: Updating ESP-IDF It is recommended to update ESP-IDF from time to time, as newer versions fix bugs and/or provide new features. Please note that each ESP-IDF major and minor release version has an associated support period, and when one release branch is approaching end of life (EOL), all users are encouraged to upgrade their projects to more recent ESP-IDF releases, to find out more about support periods, see [ESP-IDF Versions](#).

The simplest way to do the update is to delete the existing `esp-idf` folder and clone it again, as if performing the initial installation described in [Step 2. Get ESP-IDF](#).

Another solution is to update only what has changed. *The update procedure depends on the version of ESP-IDF you are using.*

After updating ESP-IDF, execute the Install script again, in case the new ESP-IDF version requires different versions of tools. See instructions at [Step 3. Set up the Tools](#).

Once the new tools are installed, update the environment using the Export script. See instructions at [Step 4. Set up the Environment Variables](#).

Related Documents

- [Establish Serial Connection with ESP32-P4](#)
- [Eclipse Plugin](#)
- [VSCode Extension](#)
- [IDF Monitor](#)

1.4 Build Your First Project

If you already have the ESP-IDF installed and not using IDE, you can build your first project from the command line following the [Start a Project on Windows](#) or [Start a Project on Linux and macOS](#).

1.5 Uninstall ESP-IDF

If you want to remove ESP-IDF, please follow [Uninstall ESP-IDF](#).

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Chapter 2

API Reference

2.1 API Conventions

This document describes conventions and assumptions common to ESP-IDF Application Programming Interfaces (APIs).

ESP-IDF provides several kinds of programming interfaces:

- C functions, structures, enums, type definitions, and preprocessor macros declared in public header files of ESP-IDF components. Various pages in the API Reference section of the programming guide contain descriptions of these functions, structures, and types.
- Build system functions, predefined variables, and options. These are documented in the [ESP-IDF CMake Build System API](#).
- *Kconfig* options can be used in code and in the build system (`CMakeLists.txt`) files.
- *Host tools* and their command line parameters are also part of the ESP-IDF interfaces.

ESP-IDF is made up of multiple components where these components either contain code specifically written for ESP chips, or contain a third-party library (i.e., a third-party component). In some cases, third-party components contain an "ESP-IDF specific" wrapper in order to provide an interface that is either simpler or better integrated with the rest of ESP-IDF's features. In other cases, third-party components present the original API of the underlying library directly.

The following sections explain some of the aspects of ESP-IDF APIs and their usage.

2.1.1 Error Handling

Most ESP-IDF APIs return error codes defined with the `esp_err_t` type. See [Error Handling](#) section for more information about error handling approaches. [Error Codes Reference](#) contains the list of error codes returned by ESP-IDF components.

2.1.2 Configuration Structures

Important: Correct initialization of configuration structures is an important part of making the application compatible with future versions of ESP-IDF.

Most initialization, configuration, and installation functions in ESP-IDF (typically named `..._init()`, `..._config()`, and `..._install()`) take a configuration structure pointer as an argument. For example:

```
const esp_timer_create_args_t my_timer_args = {
    .callback = &my_timer_callback,
    .arg = callback_arg,
    .name = "my_timer"
};
esp_timer_handle_t my_timer;
esp_err_t err = esp_timer_create(&my_timer_args, &my_timer);
```

These functions never store the pointer to the configuration structure, so it is safe to allocate the structure on the stack.

The application must initialize all fields of the structure. The following is incorrect:

```
esp_timer_create_args_t my_timer_args;
my_timer_args.callback = &my_timer_callback;
/* Incorrect! Fields .arg and .name are not initialized */
esp_timer_create(&my_timer_args, &my_timer);
```

Most ESP-IDF examples use C99 [designated initializers](#) for structure initialization since they provide a concise way of setting a subset of fields, and zero-initializing the remaining fields:

```
const esp_timer_create_args_t my_timer_args = {
    .callback = &my_timer_callback,
    /* Correct, fields .arg and .name are zero-initialized */
};
```

The C++ language supports designated initializer syntax, too, but the initializers must be in the order of declaration. When using ESP-IDF APIs in C++ code, you may consider using the following pattern:

```
/* Correct, fields .dispatch_method, .name and .skip_unhandled_events are zero-
↳ initialized */
const esp_timer_create_args_t my_timer_args = {
    .callback = &my_timer_callback,
    .arg = &my_arg,
};

/* Incorrect, .arg is declared after .callback in esp_timer_create_args_t */
//const esp_timer_create_args_t my_timer_args = {
//    .arg = &my_arg,
//    .callback = &my_timer_callback,
//};
```

For more information on designated initializers, see [Designated Initializers](#). Note that C++ language versions older than C++20, which are not the default in the current version of ESP-IDF, do not support designated initializers. If you have to compile code with an older C++ standard than C++20, you may use GCC extensions to produce the following pattern:

```
esp_timer_create_args_t my_timer_args = {};
/* All the fields are zero-initialized */
my_timer_args.callback = &my_timer_callback;
```

Default Initializers

For some configuration structures, ESP-IDF provides macros for setting default values of fields:

```
httpd_config_t config = HTTPD_DEFAULT_CONFIG();
/* HTTPD_DEFAULT_CONFIG expands to a designated initializer. Now all fields are
↳ set to the default values, and any field can still be modified: */
config.server_port = 8081;
```

(continues on next page)

(continued from previous page)

```
httpd_handle_t server;  
esp_err_t err = httpd_start(&server, &config);
```

It is recommended to use default initializer macros whenever they are provided for a particular configuration structure.

2.1.3 Private APIs

Certain header files in ESP-IDF contain APIs intended to be used only in ESP-IDF source code rather than by the applications. Such header files often contain `private` or `esp_private` in their name or path. Certain components, such as [hal](#) only contain private APIs.

Private APIs may be removed or changed in an incompatible way between minor or patch releases.

2.1.4 Components in Example Projects

ESP-IDF examples contain a variety of projects demonstrating the usage of ESP-IDF APIs. In order to reduce code duplication in the examples, a few common helpers are defined inside components that are used by multiple examples. This includes components located in [common_components](#) directory, as well as some of the components located in the examples themselves. These components are not considered to be part of the ESP-IDF API.

It is not recommended to reference these components directly in custom projects (via `EXTRA_COMPONENT_DIRS` build system variable), as they may change significantly between ESP-IDF versions. When starting a new project based on an ESP-IDF example, copy both the project and the common components it depends on out of ESP-IDF, and treat the common components as part of the project. Note that the common components are written with examples in mind, and might not include all the error handling required for production applications. Before using, take time to read the code and understand if it is applicable to your use case.

2.1.5 API Stability

ESP-IDF uses [Semantic Versioning](#) as explained in the [Versioning Scheme](#).

Minor and bugfix releases of ESP-IDF guarantee compatibility with previous releases. The sections below explain different aspects and limitations to compatibility.

Source-level Compatibility

ESP-IDF guarantees source-level compatibility of C functions, structures, enums, type definitions, and preprocessor macros declared in public header files of ESP-IDF components. Source-level compatibility implies that the application source code can be recompiled with the newer version of ESP-IDF without changes.

The following changes are allowed between minor versions and do not break source-level compatibility:

- Deprecating functions (using the `deprecated` attribute) and header files (using a preprocessor `#warning`). Deprecations are listed in ESP-IDF release notes. It is recommended to update the source code to use the newer functions or files that replace the deprecated ones, however, this is not mandatory. Deprecated functions and files can be removed from major versions of ESP-IDF.
- Renaming components, moving source and header files between components —provided that the build system ensures that correct files are still found.
- Renaming Kconfig options. Kconfig system's [backward compatibility](#) ensures that the original Kconfig option names can still be used by the application in `sdkconfig` file, CMake files, and source code.

Lack of Binary Compatibility

ESP-IDF does not guarantee binary compatibility between releases. This means that if a precompiled library is built with one ESP-IDF version, it is not guaranteed to work the same way with the next minor or bugfix release. The following are the possible changes that keep source-level compatibility but not binary compatibility:

- Changing numerical values for C enum members.
- Adding new structure members or changing the order of members. See [Configuration Structures](#) for tips that help ensure compatibility.
- Replacing an `extern` function with a `static inline` one with the same signature, or vice versa.
- Replacing a function-like macro with a compatible C function.

Other Exceptions from Compatibility

While we try to make upgrading to a new ESP-IDF version easy, there are parts of ESP-IDF that may change between minor versions in an incompatible way. We appreciate issuing reports about any unintended breaking changes that do not fall into the categories below.

- [Private APIs](#).
- [Components in Example Projects](#).
- Features clearly marked as "beta", "preview", or "experimental".
- Changes made to mitigate security issues or to replace insecure default behaviors with secure ones.
- Features that were never functional. For example, if it was never possible to use a certain function or an enumeration value, it may get renamed (as part of fixing it) or removed. This includes software features that depend on non-functional chip hardware features.
- Unexpected or undefined behavior that is not documented explicitly may be fixed/changed, such as due to missing validation of argument ranges.
- Location of [Kconfig](#) options in `menuconfig`.
- Location and names of example projects.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.2 Application Protocols

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.2.1 ASIO Port

ASIO is a cross-platform C++ library, see <https://think-async.com/Asio/>. It provides a consistent asynchronous model using a modern C++ approach.

The ESP-IDF component `ASIO` has been moved from ESP-IDF since version v5.0 to a separate repository:

- [ASIO component on GitHub](#)

To add ASIO component in your project, please run `idf.py add-dependency espressif/asio`.

Hosted Documentation

The documentation can be found on the link below:

- [ASIO documentation \(English\)](#)

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.2.2 ESP-Modbus

The Espressif ESP-Modbus Library (`esp-modbus`) supports Modbus communication in the networks based on RS485, Wi-Fi, and Ethernet interfaces. Since ESP-IDF version v5.0, the component `freemodbus` has been moved from ESP-IDF to a separate repository:

- [ESP-Modbus component on GitHub](#)

Hosted Documentation

The documentation can be found through the link below:

- [ESP-Modbus documentation \(English\)](#)

Application Example

The examples below demonstrate the ESP-Modbus library of serial and TCP ports for both slave and master implementations respectively.

- [protocols/modbus/serial/mb_slave](#)
- [protocols/modbus/serial/mb_master](#)
- [protocols/modbus/tcp/mb_tcp_slave](#)
- [protocols/modbus/tcp/mb_tcp_master](#)

Please refer to the `README.md` documents of each specific example for details.

Protocol References

- For the detailed protocol specifications, see [The Modbus Organization](#).

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.2.3 ESP-MQTT

Overview

ESP-MQTT is an implementation of [MQTT](#) protocol client, which is a lightweight publish/subscribe messaging protocol. Now ESP-MQTT supports [MQTT v5.0](#).

Features

- Support MQTT over TCP, SSL with Mbed TLS, MQTT over WebSocket, and MQTT over WebSocket Secure
- Easy to setup with URI
- Multiple instances (multiple clients in one application)
- Support subscribing, publishing, authentication, last will messages, keep alive pings, and all 3 Quality of Service (QoS) levels (it should be a fully functional client)

Application Examples

- [protocols/mqtt/tcp](#): MQTT over TCP, default port 1883
- [protocols/mqtt/ssl](#): MQTT over TLS, default port 8883
- [protocols/mqtt/ssl_ds](#): MQTT over TLS using digital signature peripheral for authentication, default port 8883
- [protocols/mqtt/ssl_mutual_auth](#): MQTT over TLS using certificates for authentication, default port 8883
- [protocols/mqtt/ssl_psk](#): MQTT over TLS using pre-shared keys for authentication, default port 8883
- [protocols/mqtt/ws](#): MQTT over WebSocket, default port 80
- [protocols/mqtt/wss](#): MQTT over WebSocket Secure, default port 443
- [protocols/mqtt5](#): Uses ESP-MQTT library to connect to broker with MQTT v5.0

MQTT Message Retransmission

A new MQTT message is created by calling [esp_mqtt_client_publish](#) or its non blocking counterpart [esp_mqtt_client_enqueue](#).

Messages with QoS 0 is sent only once. QoS 1 and 2 have different behaviors since the protocol requires extra steps to complete the process.

The ESP-MQTT library opts to always retransmit unacknowledged QoS 1 and 2 publish messages to avoid losses in faulty connections, even though the MQTT specification requires the re-transmission only on reconnect with Clean Session flag been set to 0 (set [disable_clean_session](#) to true for this behavior).

QoS 1 and 2 messages that may need retransmission are always enqueued, but first transmission try occurs immediately if [esp_mqtt_client_publish](#) is used. A transmission retry for unacknowledged messages will occur after [message_retransmit_timeout](#). After [CONFIG_MQTT_OUTBOX_EXPIRED_TIMEOUT_MS](#) messages will expire and be deleted. If [CONFIG_MQTT_REPORT_DELETED_MESSAGES](#) is set, an event will be sent to notify the user.

Configuration

The configuration is made by setting fields in [esp_mqtt_client_config_t](#) struct. The configuration struct has the following sub structs to configure different aspects of the client operation.

- [esp_mqtt_client_config_t::broker_t](#) - Allow to set address and security verification.
- [esp_mqtt_client_config_t::credentials_t](#) - Client credentials for authentication.
- [esp_mqtt_client_config_t::session_t](#) - Configuration for MQTT session aspects.
- [esp_mqtt_client_config_t::network_t](#) - Networking related configuration.
- [esp_mqtt_client_config_t::task_t](#) - Allow to configure FreeRTOS task.
- [esp_mqtt_client_config_t::buffer_t](#) - Buffer size for input and output.

In the following sections, the most common aspects are detailed.

Broker

Address Broker address can be set by usage of [address](#) struct. The configuration can be made by usage of [uri](#) field or the combination of [hostname](#), [transport](#) and [port](#). Optionally, [path](#) could be set, this field is useful in WebSocket connections.

The [uri](#) field is used in the format `scheme://hostname:port/path`.

- Currently support mqtt, mqtts, ws, wss schemes
- MQTT over TCP samples:
 - `mqtt://mqtt.eclipseprojects.io`: MQTT over TCP, default port 1883
 - `mqtt://mqtt.eclipseprojects.io:1884`: MQTT over TCP, port 1884
 - `mqtt://username:password@mqtt.eclipseprojects.io:1884`: MQTT over TCP, port 1884, with username and password
- MQTT over SSL samples:
 - `mqtt://mqtt.eclipseprojects.io`: MQTT over SSL, port 8883
 - `mqtt://mqtt.eclipseprojects.io:8884`: MQTT over SSL, port 8884
- MQTT over WebSocket samples:
 - `ws://mqtt.eclipseprojects.io:80/mqtt`
- MQTT over WebSocket Secure samples:
 - `wss://mqtt.eclipseprojects.io:443/mqtt`
- Minimal configurations:

```
const esp_mqtt_client_config_t mqtt_cfg = {
    .broker.address.uri = "mqtt://mqtt.eclipseprojects.io",
};
esp_mqtt_client_handle_t client = esp_mqtt_client_init(&mqtt_cfg);
esp_mqtt_client_register_event(client, ESP_EVENT_ANY_ID, mqtt_event_handler,
    ↪client);
esp_mqtt_client_start(client);
```

Note: By default MQTT client uses event loop library to post related MQTT events (connected, subscribed, published, etc.).

Verification For secure connections with TLS used, and to guarantee Broker's identity, the `verification` struct must be set. The broker certificate may be set in PEM or DER format. To select DER, the equivalent `certificate_len` field must be set. Otherwise, a null-terminated string in PEM format should be provided to `certificate` field.

- Get certificate from server, example: `mqtt.eclipseprojects.io`

```
openssl s_client -showcerts -connect mqtt.eclipseprojects.io:8883 < /dev/
    ↪null \
2> /dev/null | openssl x509 -outform PEM > mqtt_eclipse_org.pem
```

- Check the sample application: [protocols/mqtt/ssl](#)
- Configuration:

```
const esp_mqtt_client_config_t mqtt_cfg = {
    .broker = {
        .address.uri = "mqtt://mqtt.eclipseprojects.io:8883",
        .verification.certificate = (const char *)mqtt_eclipse_org_pem_start,
    },
};
```

For details about other fields, please check the [API Reference](#) and [TLS Server Verification](#).

Client Credentials All client related credentials are under the `credentials` field.

- `username`: pointer to the username used for connecting to the broker, can also be set by URI
- `client_id`: pointer to the client ID, defaults to `ESP32_%CHIPID%` where `%CHIPID%` are the last 3 bytes of MAC address in hex format

Authentication It is possible to set authentication parameters through the `authentication` field. The client supports the following authentication methods:

- `password`: use a password by setting
- `certificate` and `key`: mutual authentication with TLS, and both can be provided in PEM or DER format
- `use_secure_element`: use secure element available in ESP32-WROOM-32SE
- `ds_data`: use Digital Signature Peripheral available in some Espressif devices

Session For MQTT session related configurations, `session` fields should be used.

Last Will and Testament MQTT allows for a last will and testament (LWT) message to notify other clients when a client ungracefully disconnects. This is configured by the following fields in the `last_will` struct.

- `topic`: pointer to the LWT message topic
- `msg`: pointer to the LWT message
- `msg_len`: length of the LWT message, required if `msg` is not null-terminated
- `qos`: quality of service for the LWT message
- `retain`: specifies the retain flag of the LWT message

Change Settings in Project Configuration Menu The settings for MQTT can be found using `idf.py menu-config`, under Component `config>ESP-MQTT Configuration`.

The following settings are available:

- `CONFIG_MQTT_PROTOCOL_311`: enable 3.1.1 version of MQTT protocol
- `CONFIG_MQTT_TRANSPORT_SSL` and `CONFIG_MQTT_TRANSPORT_WEBSOCKET`: enable specific MQTT transport layer, such as SSL, WEBSOCKET, and WEBSOCKET_SECURE
- `CONFIG_MQTT_CUSTOM_OUTBOX`: disable default implementation of `mqtt_outbox`, so a specific implementation can be supplied

Events

The following events may be posted by the MQTT client:

- `MQTT_EVENT_BEFORE_CONNECT`: The client is initialized and about to start connecting to the broker.
- `MQTT_EVENT_CONNECTED`: The client has successfully established a connection to the broker. The client is now ready to send and receive data.
- `MQTT_EVENT_DISCONNECTED`: The client has aborted the connection due to being unable to read or write data, e.g., because the server is unavailable.
- `MQTT_EVENT_SUBSCRIBED`: The broker has acknowledged the client's subscribe request. The event data contains the message ID of the subscribe message.
- `MQTT_EVENT_UNSUBSCRIBED`: The broker has acknowledged the client's unsubscribe request. The event data contains the message ID of the unsubscribe message.
- `MQTT_EVENT_PUBLISHED`: The broker has acknowledged the client's publish message. This is only posted for QoS level 1 and 2, as level 0 does not use acknowledgements. The event data contains the message ID of the publish message.
- `MQTT_EVENT_DATA`: The client has received a publish message. The event data contains: message ID, name of the topic it was published to, received data and its length. For data that exceeds the internal buffer, multiple `MQTT_EVENT_DATA` events are posted and `current_data_offset` and `total_data_len` from event data updated to keep track of the fragmented message.
- `MQTT_EVENT_ERROR`: The client has encountered an error. The field `error_handle` in the event data contains `error_type` that can be used to identify the error. The type of error determines which parts of the `error_handle` struct is filled.

API Reference

Header File

- `components/mqtt/esp-mqtt/include/mqtt_client.h`
- This header file can be included with:

```
#include "mqtt_client.h"
```

- This header file is a part of the API provided by the `mqtt` component. To declare that your component depends on `mqtt`, add the following to your `CMakeLists.txt`:

```
REQUIRES mqtt
```

or

```
PRIV_REQUIRES mqtt
```

Functions

esp_mqtt_client_handle_t **esp_mqtt_client_init** (const *esp_mqtt_client_config_t* *config)

Creates *MQTT* client handle based on the configuration.

Parameters `config` -- *MQTT* configuration structure

Returns `mqtt_client_handle` if successfully created, `NULL` on error

esp_err_t **esp_mqtt_client_set_uri** (*esp_mqtt_client_handle_t* client, const char *uri)

Sets *MQTT* connection URI. This API is usually used to overrides the URI configured in `esp_mqtt_client_init`.

Parameters

- `client` -- *MQTT* client handle
- `uri` --

Returns `ESP_FAIL` if URI parse error, `ESP_OK` on success

esp_err_t **esp_mqtt_client_start** (*esp_mqtt_client_handle_t* client)

Starts *MQTT* client with already created client handle.

Parameters `client` -- *MQTT* client handle

Returns `ESP_OK` on success `ESP_ERR_INVALID_ARG` on wrong initialization `ESP_FAIL` on other error

esp_err_t **esp_mqtt_client_reconnect** (*esp_mqtt_client_handle_t* client)

This api is typically used to force reconnection upon a specific event.

Parameters `client` -- *MQTT* client handle

Returns `ESP_OK` on success `ESP_ERR_INVALID_ARG` on wrong initialization `ESP_FAIL` if client is in invalid state

esp_err_t **esp_mqtt_client_disconnect** (*esp_mqtt_client_handle_t* client)

This api is typically used to force disconnection from the broker.

Parameters `client` -- *MQTT* client handle

Returns `ESP_OK` on success `ESP_ERR_INVALID_ARG` on wrong initialization

esp_err_t **esp_mqtt_client_stop** (*esp_mqtt_client_handle_t* client)

Stops *MQTT* client tasks.

- Notes:
- Cannot be called from the *MQTT* event handler

Parameters `client` -- *MQTT* client handle

Returns `ESP_OK` on success `ESP_ERR_INVALID_ARG` on wrong initialization `ESP_FAIL` if client is in invalid state

int **esp_mqtt_client_subscribe_single** (*esp_mqtt_client_handle_t* client, const char *topic, int qos)

Subscribe the client to defined topic with defined qos.

Notes:

- Client must be connected to send subscribe message

- This API could be executed from a user task or from a *MQTT* event callback i.e. internal *MQTT* task (API is protected by internal mutex, so it might block if a longer data receive operation is in progress).
- `esp_mqtt_client_subscribe` could be used to call this function.

Parameters

- **client** -- *MQTT* client handle
- **topic** -- topic filter to subscribe
- **qos** -- Max qos level of the subscription

Returns message_id of the subscribe message on success -1 on failure -2 in case of full outbox.

int **esp_mqtt_client_subscribe_multiple** (*esp_mqtt_client_handle_t* client, const *esp_mqtt_topic_t* *topic_list, int size)

Subscribe the client to a list of defined topics with defined qos.

Notes:

- Client must be connected to send subscribe message
- This API could be executed from a user task or from a *MQTT* event callback i.e. internal *MQTT* task (API is protected by internal mutex, so it might block if a longer data receive operation is in progress).
- `esp_mqtt_client_subscribe` could be used to call this function.

Parameters

- **client** -- *MQTT* client handle
- **topic_list** -- List of topics to subscribe
- **size** -- size of topic_list

Returns message_id of the subscribe message on success -1 on failure -2 in case of full outbox.

int **esp_mqtt_client_unsubscribe** (*esp_mqtt_client_handle_t* client, const char *topic)

Unsubscribe the client from defined topic.

Notes:

- Client must be connected to send unsubscribe message
- It is thread safe, please refer to `esp_mqtt_client_subscribe_single` for details

Parameters

- **client** -- *MQTT* client handle
- **topic** --

Returns message_id of the subscribe message on success -1 on failure

int **esp_mqtt_client_publish** (*esp_mqtt_client_handle_t* client, const char *topic, const char *data, int len, int qos, int retain)

Client to send a publish message to the broker.

Notes:

- This API might block for several seconds, either due to network timeout (10s) or if publishing payloads longer than internal buffer (due to message fragmentation)
- Client doesn't have to be connected for this API to work, enqueueing the messages with qos>1 (returning -1 for all the qos=0 messages if disconnected). If `MQTT_SKIP_PUBLISH_IF_DISCONNECTED` is enabled, this API will not attempt to publish when the client is not connected and will always return -1.
- It is thread safe, please refer to `esp_mqtt_client_subscribe` for details

Parameters

- **client** -- *MQTT* client handle
- **topic** -- topic string
- **data** -- payload string (set to NULL, sending empty payload message)
- **len** -- data length, if set to 0, length is calculated from payload string
- **qos** -- QoS of publish message
- **retain** -- retain flag

Returns message_id of the publish message (for QoS 0 message_id will always be zero) on success.
-1 on failure, -2 in case of full outbox.

int **esp_mqtt_client_enqueue** (*esp_mqtt_client_handle_t* client, const char *topic, const char *data, int len, int qos, int retain, bool store)

Enqueue a message to the outbox, to be sent later. Typically used for messages with qos>0, but could be also used for qos=0 messages if store=true.

This API generates and stores the publish message into the internal outbox and the actual sending to the network is performed in the mqtt-task context (in contrast to the esp_mqtt_client_publish() which sends the publish message immediately in the user task's context). Thus, it could be used as a non blocking version of esp_mqtt_client_publish().

Parameters

- **client** -- *MQTT* client handle
- **topic** -- topic string
- **data** -- payload string (set to NULL, sending empty payload message)
- **len** -- data length, if set to 0, length is calculated from payload string
- **qos** -- QoS of publish message
- **retain** -- retain flag
- **store** -- if true, all messages are enqueued; otherwise only QoS 1 and QoS 2 are enqueued

Returns message_id if queued successfully, -1 on failure, -2 in case of full outbox.

esp_err_t **esp_mqtt_client_destroy** (*esp_mqtt_client_handle_t* client)

Destroys the client handle.

Notes:

- Cannot be called from the *MQTT* event handler

Parameters **client** -- *MQTT* client handle

Returns ESP_OK ESP_ERR_INVALID_ARG on wrong initialization

esp_err_t **esp_mqtt_set_config** (*esp_mqtt_client_handle_t* client, const *esp_mqtt_client_config_t* *config)

Set configuration structure, typically used when updating the config (i.e. on "before_connect" event).

Notes:

- When calling this function make sure to have all the intendend configurations set, otherwise default values are set.

Parameters

- **client** -- *MQTT* client handle
- **config** -- *MQTT* configuration structure

Returns ESP_ERR_NO_MEM if failed to allocate ESP_ERR_INVALID_ARG if conflicts on transport configuration. ESP_OK on success

esp_err_t **esp_mqtt_client_register_event** (*esp_mqtt_client_handle_t* client, *esp_mqtt_event_id_t* event, *esp_event_handler_t* event_handler, void *event_handler_arg)

Registers *MQTT* event.

Parameters

- **client** -- *MQTT* client handle
- **event** -- event type
- **event_handler** -- handler callback
- **event_handler_arg** -- handlers context

Returns ESP_ERR_NO_MEM if failed to allocate ESP_ERR_INVALID_ARG on wrong initialization ESP_OK on success

esp_err_t **esp_mqtt_client_unregister_event** (*esp_mqtt_client_handle_t* client, *esp_mqtt_event_id_t* event, *esp_event_handler_t* event_handler)

Unregisters mqtt event.

Parameters

- **client** -- mqtt client handle
- **event** -- event ID
- **event_handler** -- handler to unregister

Returns ESP_ERR_NO_MEM if failed to allocate ESP_ERR_INVALID_ARG on invalid event ID ESP_OK on success

int **esp_mqtt_client_get_outbox_size** (*esp_mqtt_client_handle_t* client)

Get outbox size.

Parameters **client** -- *MQTT* client handle

Returns outbox size 0 on wrong initialization

esp_err_t **esp_mqtt_dispatch_custom_event** (*esp_mqtt_client_handle_t* client, *esp_mqtt_event_t* *event)

Dispatch user event to the mqtt internal event loop.

Parameters

- **client** -- *MQTT* client handle
- **event** -- *MQTT* event handle structure

Returns ESP_OK on success ESP_ERR_TIMEOUT if the event couldn't be queued (ref also CONFIG_MQTT_EVENT_QUEUE_SIZE)

Structures

struct **esp_mqtt_error_codes**

MQTT error code structure to be passed as a contextual information into ERROR event

Important: This structure extends *esp_tls_last_error* error structure and is backward compatible with it (so might be down-casted and treated as *esp_tls_last_error* error, but recommended to update applications if used this way previously)

Use this structure directly checking error_type first and then appropriate error code depending on the source of the error:

error_type	related member variables	note
MQTT_ERROR_TYPE_TCP_TRANSPORT	esp_tls_last_esp_err, esp_tls_stack_err, esp_tls_cert_verify_flags, sock_errno	Error reported from tcp_transport/esp-tls
MQTT_ERROR_TYPE_CONNECTION_REFUSED	connect_return_code	Internal error reported from <i>MQTT</i> broker on connection

Public Members

esp_err_t **esp_tls_last_esp_err**

last esp_err code reported from esp-tls component

int **esp_tls_stack_err**

tls specific error code reported from underlying tls stack

int **esp_tls_cert_verify_flags**

tls flags reported from underlying tls stack during certificate verification

esp_mqtt_error_type_t **error_type**

error type referring to the source of the error

esp_mqtt_connect_return_code_t **connect_return_code**

connection refused error code reported from MQTT* broker on connection

int **esp_transport_sock_errno**

errno from the underlying socket

struct **esp_mqtt_event_t**

MQTT event configuration structure

Public Members

esp_mqtt_event_id_t **event_id**

MQTT event type

esp_mqtt_client_handle_t **client**

MQTT client handle for this event

char ***data**

Data associated with this event

int **data_len**

Length of the data for this event

int **total_data_len**

Total length of the data (longer data are supplied with multiple events)

int **current_data_offset**

Actual offset for the data associated with this event

char ***topic**

Topic associated with this event

int **topic_len**

Length of the topic for this event associated with this event

int **msg_id**

MQTT messaged id of message

int **session_present**

MQTT session_present flag for connection event

esp_mqtt_error_codes_t ***error_handle**

esp-mqtt error handle including esp-tls errors as well as internal *MQTT* errors

bool **retain**

Retained flag of the message associated with this event

int **qos**

QoS of the messages associated with this event

bool **dup**

dup flag of the message associated with this event

esp_mqtt_protocol_ver_t **protocol_ver**

MQTT protocol version used for connection, defaults to value from menuconfig

struct **esp_mqtt_client_config_t**

MQTT client configuration structure

- Default values can be set via menuconfig
- All certificates and key data could be passed in PEM or DER format. PEM format must have a terminating NULL character and the related len field set to 0. DER format requires a related len field set to the correct length.

Public Members

struct *esp_mqtt_client_config_t::broker_t* **broker**

Broker address and security verification

struct *esp_mqtt_client_config_t::credentials_t* **credentials**

User credentials for broker

struct *esp_mqtt_client_config_t::session_t* **session**

MQTT session configuration.

struct *esp_mqtt_client_config_t::network_t* **network**

Network configuration

struct *esp_mqtt_client_config_t::task_t* **task**

FreeRTOS task configuration.

struct *esp_mqtt_client_config_t::buffer_t* **buffer**

Buffer size configuration.

struct *esp_mqtt_client_config_t::outbox_config_t* **outbox**

Outbox configuration.

struct **broker_t**

Broker related configuration

Public Members

struct *esp_mqtt_client_config_t::broker_t::address_t* **address**

Broker address configuration

```
struct esp_mqtt_client_config_t::broker_t::verification_t verification
```

Security verification of the broker

```
struct address_t
```

Broker address

- uri have precedence over other fields
- If uri isn't set at least hostname, transport and port should.

Public Members

```
const char *uri
```

Complete *MQTT* broker URI

```
const char *hostname
```

Hostname, to set ipv4 pass it as string)

```
esp_mqtt_transport_t transport
```

Selects transport

```
const char *path
```

Path in the URI

```
uint32_t port
```

MQTT server port

```
struct verification_t
```

Broker identity verification

If fields are not set broker's identity isn't verified. it's recommended to set the options in this struct for security reasons.

Public Members

```
bool use_global_ca_store
```

Use a global ca_store, look esp-tls documentation for details.

```
esp_err_t (*crt_bundle_attach)(void *conf)
```

Pointer to ESP x509 Certificate Bundle attach function for the usage of certificate bundles. Client only attach the bundle, the clean up must be done by the user.

```
const char *certificate
```

Certificate data, default is NULL. It's not copied nor freed by the client, user needs to clean up.

```
size_t certificate_len
```

Length of the buffer pointed to by certificate.

const struct *psk_key_hint* ***psk_hint_key**

Pointer to PSK struct defined in `esp_tls.h` to enable PSK authentication (as alternative to certificate verification). PSK is enabled only if there are no other ways to verify broker. It's not copied nor freed by the client, user needs to clean up.

bool **skip_cert_common_name_check**

Skip any validation of server certificate CN field, this reduces the security of TLS and makes the *MQTT* client susceptible to MITM attacks

const char ****alpn_protos**

NULL-terminated list of supported application protocols to be used for ALPN.

const char ***common_name**

Pointer to the string containing server certificate common name. If non-NULL, server certificate CN must match this name, If NULL, server certificate CN must match hostname. This is ignored if `skip_cert_common_name_check=true`. It's not copied nor freed by the client, user needs to clean up.

struct **buffer_t**

Client buffer size configuration

Client have two buffers for input and output respectively.

Public Members

int **size**

size of *MQTT* send/receive buffer

int **out_size**

size of *MQTT* output buffer. If not defined, defaults to the size defined by `buffer_size`

struct **credentials_t**

Client related credentials for authentication.

Public Members

const char ***username**

MQTT username

const char ***client_id**

Set *MQTT* client identifier. Ignored if `set_null_client_id == true` If NULL set the default client id. Default client id is `ESP32_CHIPID%` where `CHIPID%` are last 3 bytes of MAC address in hex format

bool **set_null_client_id**

Selects a NULL client id

struct *esp_mqtt_client_config_t::credentials_t::authentication_t* **authentication**

Client authentication

struct **authentication_t**

Client authentication

Fields related to client authentication by broker

For mutual authentication using TLS, user could select certificate and key, secure element or digital signature peripheral if available.

Public Members

const char ***password**

MQTT password

const char ***certificate**

Certificate for ssl mutual authentication, not required if mutual authentication is not needed. Must be provided with `key`. It's not copied nor freed by the client, user needs to clean up.

size_t **certificate_len**

Length of the buffer pointed to by certificate.

const char ***key**

Private key for SSL mutual authentication, not required if mutual authentication is not needed. If it is not NULL, also `certificate` has to be provided. It's not copied nor freed by the client, user needs to clean up.

size_t **key_len**

Length of the buffer pointed to by key.

const char ***key_password**

Client key decryption password, not PEM nor DER, if provided `key_password_len` must be correctly set.

int **key_password_len**

Length of the password pointed to by `key_password`

bool **use_secure_element**

Enable secure element, available in ESP32-ROOM-32SE, for SSL connection

void ***ds_data**

Carrier of handle for digital signature parameters, digital signature peripheral is available in some Espressif devices. It's not copied nor freed by the client, user needs to clean up.

struct **network_t**

Network related configuration

Public Members

int **reconnect_timeout_ms**

Reconnect to the broker after this value in milliseconds if auto reconnect is not disabled (defaults to 10s)

int **timeout_ms**

Abort network operation if it is not completed after this value, in milliseconds (defaults to 10s).

int **refresh_connection_after_ms**

Refresh connection after this value (in milliseconds)

bool **disable_auto_reconnect**

Client will reconnect to server (when errors/disconnect). Set `disable_auto_reconnect=true` to disable

esp_transport_handle_t **transport**

Custom transport handle to use. Warning: The transport should be valid during the client lifetime and is destroyed when `esp_mqtt_client_destroy` is called.

struct ifreq ***if_name**

The name of interface for data to go through. Use the default interface without setting

struct **outbox_config_t**

Client outbox configuration options.

Public Members

uint64_t **limit**

Size limit for the outbox in bytes.

struct **session_t**

MQTT Session related configuration

Public Members

struct *esp_mqtt_client_config_t::session_t::last_will_t* **last_will**

Last will configuration

bool **disable_clean_session**

MQTT clean session, default `clean_session` is true

int **keepalive**

MQTT keepalive, default is 120 seconds When configuring this value, keep in mind that the client attempts to communicate with the broker at half the interval that is actually set. This conservative approach allows for more attempts before the broker's timeout occurs

bool **disable_keepalive**

Set `disable_keepalive=true` to turn off keep-alive mechanism, keepalive is active by default. Note: setting the config value `keepalive` to 0 doesn't disable keepalive feature, but uses a default keepalive period

esp_mqtt_protocol_ver_t **protocol_ver**

MQTT protocol version used for connection.

int **message_retransmit_timeout**
timeout for retransmitting of failed packet

struct **last_will_t**
Last Will and Testament message configuration.

Public Members

const char ***topic**
LWT (Last Will and Testament) message topic

const char ***msg**
LWT message, may be NULL terminated

int **msg_len**
LWT message length, if msg isn't NULL terminated must have the correct length

int **qos**
LWT message QoS

int **retain**
LWT retained message flag

struct **task_t**
Client task configuration

Public Members

int **priority**
MQTT task priority

int **stack_size**
MQTT task stack size

struct **topic_t**
Topic definition struct

Public Members

const char ***filter**
Topic filter to subscribe

int **qos**
Max QoS level of the subscription

Macros

MQTT_ERROR_TYPE_ESP_TLS

MQTT_ERROR_TYPE_TCP_TRANSPORT error type hold all sorts of transport layer errors, including ESP-TLS error, but in the past only the errors from MQTT_ERROR_TYPE_ESP_TLS layer were reported, so the ESP-TLS error type is re-defined here for backward compatibility

esp_mqtt_client_subscribe (client_handle, topic_type, qos_or_size)

Convenience macro to select subscribe function to use.

Notes:

- Usage of `esp_mqtt_client_subscribe_single` is the same as previous `esp_mqtt_client_subscribe`, refer to it for details.

Parameters

- **client_handle** -- MQTT client handle
- **topic_type** -- Needs to be char* for single subscription or `esp_mqtt_topic_t` for multiple topics
- **qos_or_size** -- It's either a qos when subscribing to a single topic or the size of the subscription array when subscribing to multiple topics.

Returns message_id of the subscribe message on success -1 on failure -2 in case of full outbox.

Type Definitions

```
typedef struct esp_mqtt_client *esp_mqtt_client_handle_t
```

```
typedef enum esp_mqtt_event_id_t esp_mqtt_event_id_t
```

MQTT event types.

User event handler receives context data in `esp_mqtt_event_t` structure with

- client - MQTT client handle
- various other data depending on event type

```
typedef enum esp_mqtt_connect_return_code_t esp_mqtt_connect_return_code_t
```

MQTT connection error codes propagated via ERROR event

```
typedef enum esp_mqtt_error_type_t esp_mqtt_error_type_t
```

MQTT connection error codes propagated via ERROR event

```
typedef enum esp_mqtt_transport_t esp_mqtt_transport_t
```

```
typedef enum esp_mqtt_protocol_ver_t esp_mqtt_protocol_ver_t
```

MQTT protocol version used for connection

```
typedef struct esp_mqtt_error_codes esp_mqtt_error_codes_t
```

MQTT error code structure to be passed as a contextual information into ERROR event

Important: This structure extends `esp_tls_last_error` error structure and is backward compatible with it (so might be down-casted and treated as `esp_tls_last_error` error, but recommended to update applications if used this way previously)

Use this structure directly checking error_type first and then appropriate error code depending on the source of the error:

error_type	related member variables	note
MQTT_ERROR_TYPE_TCP_TRANSPORT	esp_tls_last_esp_err, esp_tls_stack_err, esp_tls_cert_verify_flags, sock_errno	Error reported from

tcp_transport/esp-tls || MQTT_ERROR_TYPE_CONNECTION_REFUSED | connect_return_code | Internal error reported from *MQTT* broker on connection |

typedef struct *esp_mqtt_event_t* **esp_mqtt_event_t**

MQTT event configuration structure

typedef *esp_mqtt_event_t* ***esp_mqtt_event_handle_t**

typedef struct *esp_mqtt_client_config_t* **esp_mqtt_client_config_t**

MQTT client configuration structure

- Default values can be set via menuconfig
- All certificates and key data could be passed in PEM or DER format. PEM format must have a terminating NULL character and the related len field set to 0. DER format requires a related len field set to the correct length.

typedef struct *topic_t* **esp_mqtt_topic_t**

Topic definition struct

Enumerations

enum **esp_mqtt_event_id_t**

MQTT event types.

User event handler receives context data in *esp_mqtt_event_t* structure with

- client - *MQTT* client handle
- various other data depending on event type

Values:

enumerator **MQTT_EVENT_ANY**

enumerator **MQTT_EVENT_ERROR**

on error event, additional context: connection return code, error handle from esp_tls (if supported)

enumerator **MQTT_EVENT_CONNECTED**

connected event, additional context: session_present flag

enumerator **MQTT_EVENT_DISCONNECTED**

disconnected event

enumerator **MQTT_EVENT_SUBSCRIBED**

subscribed event, additional context:

- msg_id message id
- error_handle *error_type* in case subscribing failed
- data pointer to broker response, check for errors.
- data_len length of the data for this event

enumerator **MQTT_EVENT_UNSUBSCRIBED**

unsubscribed event, additional context: msg_id

enumerator **MQTT_EVENT_PUBLISHED**

published event, additional context: msg_id

enumerator **MQTT_EVENT_DATA**

data event, additional context:

- msg_id message id
- topic pointer to the received topic
- topic_len length of the topic
- data pointer to the received data
- data_len length of the data for this event
- current_data_offset offset of the current data for this event
- total_data_len total length of the data received
- retain retain flag of the message
- qos QoS level of the message
- dup dup flag of the message Note: Multiple MQTT_EVENT_DATA could be fired for one message, if it is longer than internal buffer. In that case only first event contains topic pointer and length, other contain data only with current data length and current data offset updating.

enumerator **MQTT_EVENT_BEFORE_CONNECT**

The event occurs before connecting

enumerator **MQTT_EVENT_DELETED**

Notification on delete of one message from the internal outbox, if the message couldn't have been sent and acknowledged before expiring defined in OUTBOX_EXPIRED_TIMEOUT_MS. (events are not posted upon deletion of successfully acknowledged messages)

- This event id is posted only if MQTT_REPORT_DELETED_MESSAGES==1
- Additional context: msg_id (id of the deleted message).

enumerator **MQTT_USER_EVENT**

Custom event used to queue tasks into mqtt event handler All fields from the *esp_mqtt_event_t* type could be used to pass an additional context data to the handler.

enum **esp_mqtt_connect_return_code_t**

MQTT connection error codes propagated via ERROR event

Values:

enumerator **MQTT_CONNECTION_ACCEPTED**

Connection accepted

enumerator **MQTT_CONNECTION_REFUSE_PROTOCOL**

MQTT connection refused reason: Wrong protocol

enumerator **MQTT_CONNECTION_REFUSE_ID_REJECTED**

MQTT connection refused reason: ID rejected

enumerator **MQTT_CONNECTION_REFUSE_SERVER_UNAVAILABLE**

MQTT connection refused reason: Server unavailable

enumerator **MQTT_CONNECTION_REFUSE_BAD_USERNAME**

MQTT connection refused reason: Wrong user

enumerator **MQTT_CONNECTION_REFUSE_NOT_AUTHORIZED**

MQTT connection refused reason: Wrong username or password

enum **esp_mqtt_error_type_t**

MQTT connection error codes propagated via ERROR event

Values:

enumerator **MQTT_ERROR_TYPE_NONE**

enumerator **MQTT_ERROR_TYPE_TCP_TRANSPORT**

enumerator **MQTT_ERROR_TYPE_CONNECTION_REFUSED**

enumerator **MQTT_ERROR_TYPE_SUBSCRIBE_FAILED**

enum **esp_mqtt_transport_t**

Values:

enumerator **MQTT_TRANSPORT_UNKNOWN**

enumerator **MQTT_TRANSPORT_OVER_TCP**

MQTT over TCP, using scheme: **MQTT**

enumerator **MQTT_TRANSPORT_OVER_SSL**

MQTT over SSL, using scheme: **MQTTS**

enumerator **MQTT_TRANSPORT_OVER_WS**

MQTT over Websocket, using scheme:: **ws**

enumerator **MQTT_TRANSPORT_OVER_WSS**

MQTT over Websocket Secure, using scheme: **wss**

enum **esp_mqtt_protocol_ver_t**

MQTT protocol version used for connection

Values:

enumerator **MQTT_PROTOCOL_UNDEFINED**

enumerator **MQTT_PROTOCOL_V_3_1**

enumerator **MQTT_PROTOCOL_V_3_1_1**

enumerator **MQTT_PROTOCOL_V_5**

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.2.4 ESP-TLS

Overview

The ESP-TLS component provides a simplified API interface for accessing the commonly used TLS functions. It supports common scenarios like CA certification validation, SNI, ALPN negotiation, and non-blocking connection among others. All the configurations can be specified in the `esp_tls_cfg_t` data structure. Once done, TLS communication can be conducted using the following APIs:

- `esp_tls_init()`: for initializing the TLS connection handle.
- `esp_tls_conn_new_sync()`: for opening a new blocking TLS connection.
- `esp_tls_conn_new_async()`: for opening a new non-blocking TLS connection.
- `esp_tls_conn_read()`: for reading from the connection.
- `esp_tls_conn_write()`: for writing into the connection.
- `esp_tls_conn_destroy()`: for freeing up the connection.

Any application layer protocol like HTTP1, HTTP2, etc can be executed on top of this layer.

Application Example

Simple HTTPS example that uses ESP-TLS to establish a secure socket connection: [protocols/https_request](#).

Tree Structure for ESP-TLS Component

```

├─ esp_tls.c
├─ esp_tls.h
├─ esp_tls_mbedtls.c
├─ esp_tls_wolfssl.c
├─ private_include
│   └─ esp_tls_mbedtls.h
│   └─ esp_tls_wolfssl.h

```

The ESP-TLS component has a file `esp-tls/esp_tls.h` which contains the public API headers for the component. Internally, the ESP-TLS component operates using either MbedTLS or WolfSSL, which are SSL/TLS libraries. APIs specific to MbedTLS are present in `esp-tls/private_include/esp_tls_mbedtls.h` and APIs specific to WolfSSL are present in `esp-tls/private_include/esp_tls_wolfssl.h`.

TLS Server Verification

ESP-TLS provides multiple options for TLS server verification on the client side. The ESP-TLS client can verify the server by validating the peer's server certificate or with the help of pre-shared keys. The user should select only one of the following options in the `esp_tls_cfg_t` structure for TLS server verification. If no option is selected, the client will return a fatal error by default during the TLS connection setup.

- **cacert_buf** and **cacert_bytes**: The CA certificate can be provided in a buffer to the `esp_tls_cfg_t` structure. The ESP-TLS uses the CA certificate present in the buffer to verify the server. The following variables in the `esp_tls_cfg_t` structure must be set.
 - `cacert_buf` - pointer to the buffer which contains the CA certification.
 - `cacert_bytes` - the size of the CA certificate in bytes.
- **use_global_ca_store**: The `global_ca_store` can be initialized and set at once. Then it can be used to verify the server for all the ESP-TLS connections which have set `use_global_ca_store = true` in their respective `esp_tls_cfg_t` structure. See the API Reference section below for information regarding different APIs used for initializing and setting up the `global_ca_store`.

- **crt_bundle_attach:** The ESP x509 Certificate Bundle API provides an easy way to include a bundle of custom x509 root certificates for TLS server verification. More details can be found at [ESP x509 Certificate Bundle](#).
- **psk_hint_key:** To use pre-shared keys for server verification, [CONFIG_ESP_TLS_PSK_VERIFICATION](#) should be enabled in the ESP-TLS menuconfig. Then the pointer to the PSK hint and key should be provided to the `esp_tls_cfg_t` structure. The ESP-TLS will use the PSK for server verification only when no other option regarding server verification is selected.
- **skip server verification:** This is an insecure option provided in the ESP-TLS for testing purposes. The option can be set by enabling [CONFIG_ESP_TLS_INSECURE](#) and [CONFIG_ESP_TLS_SKIP_SERVER_CERT_VERIFY](#) in the ESP-TLS menuconfig. When this option is enabled the ESP-TLS will skip server verification by default when no other options for server verification are selected in the `esp_tls_cfg_t` structure.

Warning: Enabling this option comes with a potential risk of establishing a TLS connection with a server that has a fake identity, provided that the server certificate is not provided either through API or other mechanisms like `ca_store` etc.

ESP-TLS Server Cert Selection Hook

The ESP-TLS component provides an option to set the server certification selection hook when using the MbedTLS stack. This provides an ability to configure and use a certificate selection callback during server handshake. The callback helps to select a certificate to present to the client based on the TLS extensions supplied in the client hello message, such as ALPN and SNI. To enable this feature, please enable [CONFIG_ESP_TLS_SERVER_CERT_SELECT_HOOK](#) in the ESP-TLS menuconfig.

The certificate selection callback can be configured in the `esp_tls_cfg_t` structure as follows:

```
int cert_selection_callback(mbedtls_ssl_context *ssl)
{
    /* Code that the callback should execute */
    return 0;
}

esp_tls_cfg_t cfg = {
    cert_select_cb = cert_section_callback,
};
```

Underlying SSL/TLS Library Options

The ESP-TLS component offers the option to use MbedTLS or WolfSSL as its underlying SSL/TLS library. By default, only MbedTLS is available and used, WolfSSL SSL/TLS library is also available publicly at <https://github.com/espressif/esp-wolfssl>. The repository provides the WolfSSL component in binary format, and it also provides a few examples that are useful for understanding the API. Please refer to the repository `README.md` for information on licensing and other options. Please see the below section for instructions on how to use WolfSSL in your project.

Note: As the library options are internal to ESP-TLS, switching the libraries will not change ESP-TLS specific code for a project.

How to Use WolfSSL with ESP-IDF

There are two ways to use WolfSSL in your project:

- 1) Directly add WolfSSL as a component in your project with the following three commands:

```
(First, change the directory (cd) to your project directory)
mkdir components
cd components
git clone --recursive https://github.com/espressif/esp-wolfssl.git
```

2) Add WolfSSL as an extra component in your project.

- Download WolfSSL with:

```
git clone --recursive https://github.com/espressif/esp-wolfssl.git
```

- Include ESP-WolfSSL in ESP-IDF with setting `EXTRA_COMPONENT_DIRS` in `CMakeLists.txt` of your project as done in [wolfssl/examples](#). For reference see [Optional Project Variables](#) in [build-system](#).

After the above steps, you will have the option to choose WolfSSL as the underlying SSL/TLS library in the configuration menu of your project as follows:

```
idf.py menuconfig > ESP-TLS > SSL/TLS Library > Mbedtls/Wolfssl
```

Comparison Between MbedTLS and WolfSSL

The following table shows a typical comparison between WolfSSL and MbedTLS when the [protocols/https_request](#) example (which includes server authentication) is running with both SSL/TLS libraries and with all respective configurations set to default. For MbedTLS, the `IN_CONTENT` length and `OUT_CONTENT` length are set to 16384 bytes and 4096 bytes respectively.

Property	WolfSSL	MbedTLS
Total Heap Consumed	~ 19 KB	~ 37 KB
Task Stack Used	~ 2.2 KB	~ 3.6 KB
Bin size	~ 858 KB	~ 736 KB

Note: These values can vary based on configuration options and version of respective libraries.

Digital Signature with ESP-TLS

ESP-TLS provides support for using the Digital Signature (DS) with ESP32-P4. Use of the DS for TLS is supported only when ESP-TLS is used with MbedTLS (default stack) as its underlying SSL/TLS stack. For more details on Digital Signature, please refer to the [Digital Signature \(DS\)](#). The technical details of Digital Signature such as how to calculate private key parameters can be found in **ESP32-P4 Technical Reference Manual > Digital Signature (DS)** [\[PDF\]](#). The DS peripheral must be configured before it can be used to perform Digital Signature, see [Configure the DS peripheral for a TLS connection](#).

The DS peripheral must be initialized with the required encrypted private key parameters, which are obtained when the DS peripheral is configured. ESP-TLS internally initializes the DS peripheral when provided with the required DS context, i.e., DS parameters. Please see the below code snippet for passing the DS context to the ESP-TLS context. The DS context passed to the ESP-TLS context should not be freed till the TLS connection is deleted.

```
#include "esp_tls.h"
esp_ds_data_ctx_t *ds_ctx;
/* initialize ds_ctx with encrypted private key parameters, which can be read from
↳ the nvs or provided through the application code */
esp_tls_cfg_t cfg = {
    .clientcert_buf = /* the client certification */,
    .clientcert_bytes = /* length of the client certification */,
    /* other configurations options */
    .ds_data = (void *)ds_ctx,
};
```

Note: When using Digital Signature for the TLS connection, along with the other required params, only the client certification (*clientcert_buf*) and the DS params (*ds_data*) are required and the client key (*clientkey_buf*) can be set to NULL.

- An example of mutual authentication with the DS peripheral can be found at [ssl mutual auth](#) which internally uses (ESP-TLS) for the TLS connection.

ECDSA Peripheral with ESP-TLS

ESP-TLS provides support for using the ECDSA peripheral with ESP32-P4. The use of ECDSA peripheral is supported only when ESP-TLS is used with MbedTLS as its underlying SSL/TLS stack. The ECDSA private key should be present in the eFuse for using the ECDSA peripheral. Please refer to [espefuse.py](#) documentation for programming the ECDSA key in the efuse. To use ECDSA peripheral with ESP-TLS, set `esp_tls_cfg_t::use_ecdsa_peripheral` to `true`, and set `esp_tls_cfg_t::ecdsa_key_efuse_blk` to the eFuse block ID in which ECDSA private key is stored. This will enable the use of ECDSA peripheral for private key operations. As the client private key is already present in the eFuse, it needs not be supplied to the `esp_tls_cfg_t` structure.

```
#include "esp_tls.h"
esp_tls_cfg_t cfg = {
    .use_ecdsa_peripheral = true,
    .ecdsa_key_efuse_blk = /* efuse block with ecdsa private key */,
};
```

Note: When using ECDSA peripheral with TLS, only `MBEDTLS_TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256` ciphersuite is supported. If using TLS v1.3, `MBEDTLS_TLS1_3_AES_128_GCM_SHA256` ciphersuite is supported.

TLS Ciphersuites

ESP-TLS provides the ability to set a ciphersuites list in client mode. The TLS ciphersuites list informs the server about the supported ciphersuites for the specific TLS connection regardless of the TLS stack configuration. If the server supports any ciphersuite from this list, then the TLS connection will succeed; otherwise, it will fail.

You can set `ciphersuites_list` in the `esp_tls_cfg_t` structure during client connection as follows:

```
/* ciphersuites_list must end with 0 and must be available in the memory scope_
↳ active during the entire TLS connection */
static const int ciphersuites_list[] = {MBEDTLS_TLS_ECDHE_ECDSA_WITH_AES_256_GCM_
↳ SHA384, MBEDTLS_TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, 0};
esp_tls_cfg_t cfg = {
    .ciphersuites_list = ciphersuites_list,
};
```

ESP-TLS will not check the validity of `ciphersuites_list` that was set, you should call `esp_tls_get_ciphersuites_list()` to get ciphersuites list supported in the TLS stack and cross-check it against the supplied list.

Note: This feature is supported only in the MbedTLS stack.

API Reference

Header File

- [components/esp-tls/esp_tls.h](#)
- This header file can be included with:

```
#include "esp_tls.h"
```

- This header file is a part of the API provided by the `esp-tls` component. To declare that your component depends on `esp-tls`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp-tls
```

or

```
PRIV_REQUIRES esp-tls
```

Functions

esp_tls_t ***esp_tls_init** (void)

Create TLS connection.

This function allocates and initializes `esp-tls` structure handle.

Returns `tls` Pointer to `esp-tls` as `esp-tls` handle if successfully initialized, NULL if allocation error

esp_tls_t ***esp_tls_conn_http_new** (const char *url, const *esp_tls_cfg_t* *cfg)

Create a new blocking TLS/SSL connection with a given "HTTP" url.

Note: This API is present for backward compatibility reasons. Alternative function with the same functionality is `esp_tls_conn_http_new_sync` (and its asynchronous version `esp_tls_conn_http_new_async`)

Parameters

- **url** -- [in] url of host.
- **cfg** -- [in] TLS configuration as `esp_tls_cfg_t`. If you wish to open non-TLS connection, keep this NULL. For TLS connection, a pass pointer to '`esp_tls_cfg_t`'. At a minimum, this structure should be zero-initialized.

Returns pointer to `esp_tls_t`, or NULL if connection couldn't be opened.

int **esp_tls_conn_new_sync** (const char *hostname, int hostlen, int port, const *esp_tls_cfg_t* *cfg, *esp_tls_t* *tls)

Create a new blocking TLS/SSL connection.

This function establishes a TLS/SSL connection with the specified host in blocking manner.

Parameters

- **hostname** -- [in] Hostname of the host.
- **hostlen** -- [in] Length of hostname.
- **port** -- [in] Port number of the host.
- **cfg** -- [in] TLS configuration as `esp_tls_cfg_t`. If you wish to open non-TLS connection, keep this NULL. For TLS connection, a pass pointer to `esp_tls_cfg_t`. At a minimum, this structure should be zero-initialized.
- **tls** -- [in] Pointer to `esp-tls` as `esp-tls` handle.

Returns

- -1 If connection establishment fails.
- 1 If connection establishment is successful.
- 0 If connection state is in progress.

int **esp_tls_conn_http_new_sync** (const char *url, const *esp_tls_cfg_t* *cfg, *esp_tls_t* *tls)

Create a new blocking TLS/SSL connection with a given "HTTP" url.

The behaviour is same as `esp_tls_conn_new_sync()` API. However this API accepts host's url.

Parameters

- **url** -- [in] url of host.

- **cfg** -- [in] TLS configuration as `esp_tls_cfg_t`. If you wish to open non-TLS connection, keep this NULL. For TLS connection, a pass pointer to 'esp_tls_cfg_t'. At a minimum, this structure should be zero-initialized.
- **tls** -- [in] Pointer to esp-tls as esp-tls handle.

Returns

- -1 If connection establishment fails.
- 1 If connection establishment is successful.
- 0 If connection state is in progress.

int **esp_tls_conn_new_async** (const char *hostname, int hostlen, int port, const *esp_tls_cfg_t* *cfg, *esp_tls_t* *tls)

Create a new non-blocking TLS/SSL connection.

This function initiates a non-blocking TLS/SSL connection with the specified host, but due to its non-blocking nature, it doesn't wait for the connection to get established.

Parameters

- **hostname** -- [in] Hostname of the host.
- **hostlen** -- [in] Length of hostname.
- **port** -- [in] Port number of the host.
- **cfg** -- [in] TLS configuration as `esp_tls_cfg_t`. `non_block` member of this structure should be set to be true.
- **tls** -- [in] pointer to esp-tls as esp-tls handle.

Returns

- -1 If connection establishment fails.
- 0 If connection establishment is in progress.
- 1 If connection establishment is successful.

int **esp_tls_conn_http_new_async** (const char *url, const *esp_tls_cfg_t* *cfg, *esp_tls_t* *tls)

Create a new non-blocking TLS/SSL connection with a given "HTTP" url.

The behaviour is same as `esp_tls_conn_new_async()` API. However this API accepts host's url.

Parameters

- **url** -- [in] url of host.
- **cfg** -- [in] TLS configuration as `esp_tls_cfg_t`.
- **tls** -- [in] pointer to esp-tls as esp-tls handle.

Returns

- -1 If connection establishment fails.
- 0 If connection establishment is in progress.
- 1 If connection establishment is successful.

ssize_t **esp_tls_conn_write** (*esp_tls_t* *tls, const void *data, size_t datalen)

Write from buffer 'data' into specified tls connection.

Parameters

- **tls** -- [in] pointer to esp-tls as esp-tls handle.
- **data** -- [in] Buffer from which data will be written.
- **datalen** -- [in] Length of data buffer.

Returns

- ≥ 0 if write operation was successful, the return value is the number of bytes actually written to the TLS/SSL connection.
- < 0 if write operation was not successful, because either an error occurred or an action must be taken by the calling process.
- `ESP_TLS_ERR_SSL_WANT_READ/ ESP_TLS_ERR_SSL_WANT_WRITE`. if the handshake is incomplete and waiting for data to be available for reading. In this case this functions needs to be called again when the underlying transport is ready for operation.

ssize_t **esp_tls_conn_read** (*esp_tls_t* *tls, void *data, size_t datalen)

Read from specified tls connection into the buffer 'data'.

Parameters

- **tls** -- **[in]** pointer to esp-tls as esp-tls handle.
- **data** -- **[in]** Buffer to hold read data.
- **datalen** -- **[in]** Length of data buffer.

Returns

- >0 if read operation was successful, the return value is the number of bytes actually read from the TLS/SSL connection.
- 0 if read operation was not successful. The underlying connection was closed.
- <0 if read operation was not successful, because either an error occurred or an action must be taken by the calling process.

int **esp_tls_conn_destroy** (*esp_tls_t* *tls)

Close the TLS/SSL connection and free any allocated resources.

This function should be called to close each tls connection opened with `esp_tls_conn_new_sync()` (or `esp_tls_conn_http_new_sync()`) and `esp_tls_conn_new_async()` (or `esp_tls_conn_http_new_async()`) APIs.

Parameters **tls** -- **[in]** pointer to esp-tls as esp-tls handle.

Returns - 0 on success

- -1 if socket error or an invalid argument

ssize_t **esp_tls_get_bytes_avail** (*esp_tls_t* *tls)

Return the number of application data bytes remaining to be read from the current record.

This API is a wrapper over mbedtls's `mbedtls_ssl_get_bytes_avail()` API.

Parameters **tls** -- **[in]** pointer to esp-tls as esp-tls handle.

Returns

- -1 in case of invalid arg
- bytes available in the application data record read buffer

esp_err_t **esp_tls_get_conn_sockfd** (*esp_tls_t* *tls, int *sockfd)

Returns the connection socket file descriptor from esp_tls session.

Parameters

- **tls** -- **[in]** handle to esp_tls context
- **sockfd** -- **[out]** int pointer to sockfd value.

Returns - ESP_OK on success and value of sockfd will be updated with socket file descriptor for connection

- ESP_ERR_INVALID_ARG if (tls == NULL || sockfd == NULL)

esp_err_t **esp_tls_set_conn_sockfd** (*esp_tls_t* *tls, int sockfd)

Sets the connection socket file descriptor for the esp_tls session.

Parameters

- **tls** -- **[in]** handle to esp_tls context
- **sockfd** -- **[in]** sockfd value to set.

Returns - ESP_OK on success and value of sockfd for the tls connection shall updated with the provided value

- ESP_ERR_INVALID_ARG if (tls == NULL || sockfd < 0)

esp_err_t **esp_tls_get_conn_state** (*esp_tls_t* *tls, *esp_tls_conn_state_t* *conn_state)

Gets the connection state for the esp_tls session.

Parameters

- **tls** -- **[in]** handle to esp_tls context
- **conn_state** -- **[out]** pointer to the connection state value.

Returns - ESP_OK on success and value of sockfd for the tls connection shall updated with the provided value

- ESP_ERR_INVALID_ARG (Invalid arguments)

esp_err_t **esp_tls_set_conn_state** (*esp_tls_t* *tls, *esp_tls_conn_state_t* conn_state)

Sets the connection state for the esp_tls session.

Parameters

- **tls** -- [in] handle to esp_tls context
- **conn_state** -- [in] connection state value to set.

Returns - ESP_OK on success and value of sockfd for the tls connection shall updated with the provided value

- ESP_ERR_INVALID_ARG (Invalid arguments)

void **esp_tls_get_ssl_context** (*esp_tls_t* *tls)

Returns the ssl context.

Parameters **tls** -- [in] handle to esp_tls context

Returns - ssl_ctx pointer to ssl context of underlying TLS layer on success

- NULL in case of error

esp_err_t **esp_tls_init_global_ca_store** (void)

Create a global CA store, initially empty.

This function should be called if the application wants to use the same CA store for multiple connections. This function initialises the global CA store which can be then set by calling `esp_tls_set_global_ca_store()`. To be effective, this function must be called before any call to `esp_tls_set_global_ca_store()`.

Returns

- ESP_OK if creating global CA store was successful.
- ESP_ERR_NO_MEM if an error occurred when allocating the mbedTLS resources.

esp_err_t **esp_tls_set_global_ca_store** (const unsigned char *cacert_pem_buf, const unsigned int cacert_pem_bytes)

Set the global CA store with the buffer provided in pem format.

This function should be called if the application wants to set the global CA store for multiple connections i.e. to add the certificates in the provided buffer to the certificate chain. This function implicitly calls `esp_tls_init_global_ca_store()` if it has not already been called. The application must call this function before calling `esp_tls_conn_new()`.

Parameters

- **cacert_pem_buf** -- [in] Buffer which has certificates in pem format. This buffer is used for creating a global CA store, which can be used by other tls connections.
- **cacert_pem_bytes** -- [in] Length of the buffer.

Returns

- ESP_OK if adding certificates was successful.
- Other if an error occurred or an action must be taken by the calling process.

void **esp_tls_free_global_ca_store** (void)

Free the global CA store currently being used.

The memory being used by the global CA store to store all the parsed certificates is freed up. The application can call this API if it no longer needs the global CA store.

esp_err_t **esp_tls_get_and_clear_last_error** (*esp_tls_error_handle_t* h, int *esp_tls_code, int *esp_tls_flags)

Returns last error in esp_tls with detailed mbedtls related error codes. The error information is cleared internally upon return.

Parameters

- **h** -- [in] esp-tls error handle.
- **esp_tls_code** -- [out] last error code returned from mbedtls api (set to zero if none) This pointer could be NULL if caller does not care about esp_tls_code
- **esp_tls_flags** -- [out] last certification verification flags (set to zero if none) This pointer could be NULL if caller does not care about esp_tls_code

Returns

- ESP_ERR_INVALID_STATE if invalid parameters
- ESP_OK (0) if no error occurred
- specific error code (based on ESP_ERR_ESP_TLS_BASE) otherwise

esp_err_t **esp_tls_get_and_clear_error_type** (*esp_tls_error_handle_t* h, *esp_tls_error_type_t* err_type, int *error_code)

Returns the last error captured in esp_tls of a specific type. The error information is cleared internally upon return.

Parameters

- **h** -- [in] esp-tls error handle.
- **err_type** -- [in] specific error type
- **error_code** -- [out] last error code returned from mbedtls api (set to zero if none). This pointer could be NULL if caller does not care about esp_tls_code

Returns

- ESP_ERR_INVALID_STATE if invalid parameters
- ESP_OK if a valid error returned and was cleared

esp_err_t **esp_tls_get_error_handle** (*esp_tls_t* *tls, *esp_tls_error_handle_t* *error_handle)

Returns the ESP-TLS error_handle.

Parameters

- **tls** -- [in] handle to esp_tls context
- **error_handle** -- [out] pointer to the error handle.

Returns

- ESP_OK on success and error_handle will be updated with the ESP-TLS error handle.
- ESP_ERR_INVALID_ARG if (tls == NULL || error_handle == NULL)

mbedtls_x509_crt ***esp_tls_get_global_ca_store** (void)

Get the pointer to the global CA store currently being used.

The application must first call esp_tls_set_global_ca_store(). Then the same CA store could be used by the application for APIs other than esp_tls.

Note: Modifying the pointer might cause a failure in verifying the certificates.

Returns

- Pointer to the global CA store currently being used if successful.
- NULL if there is no global CA store set.

const int ***esp_tls_get_ciphersuites_list** (void)

Get supported TLS ciphersuites list.

See <https://www.iana.org/assignments/tls-parameters/tls-parameters.xhtml#tls-parameters-4> for the list of ciphersuites

Returns Pointer to a zero-terminated array of IANA identifiers of TLS ciphersuites.

esp_err_t **esp_tls_plain_tcp_connect** (const char *host, int hostlen, int port, const *esp_tls_cfg_t* *cfg, *esp_tls_error_handle_t* error_handle, int *sockfd)

Creates a plain TCP connection, returning a valid socket fd on success or an error handle.

Parameters

- **host** -- [in] Hostname of the host.
- **hostlen** -- [in] Length of hostname.
- **port** -- [in] Port number of the host.
- **cfg** -- [in] ESP-TLS configuration as esp_tls_cfg_t.
- **error_handle** -- [out] ESP-TLS error handle holding potential errors occurred during connection
- **sockfd** -- [out] Socket descriptor if successfully connected on TCP layer

Returns ESP_OK on success, ESP_ERR_INVALID_ARG if invalid output parameters, ESP-TLS based error codes on failure

Structures

struct **psk_key_hint**

ESP-TLS preshared key and hint structure.

Public Members

const uint8_t ***key**

key in PSK authentication mode in binary format

const size_t **key_size**

length of the key

const char ***hint**

hint in PSK authentication mode in string format

struct **tls_keep_alive_cfg**

esp-tls client session ticket ctx

Keep alive parameters structure

Public Members

bool **keep_alive_enable**

Enable keep-alive timeout

int **keep_alive_idle**

Keep-alive idle time (second)

int **keep_alive_interval**

Keep-alive interval time (second)

int **keep_alive_count**

Keep-alive packet retry send count

struct **esp_tls_cfg**

ESP-TLS configuration parameters.

Note: Note about format of certificates:

- This structure includes certificates of a Certificate Authority, of client or server as well as private keys, which may be of PEM or DER format. In case of PEM format, the buffer must be NULL terminated (with NULL character included in certificate size).
 - Certificate Authority's certificate may be a chain of certificates in case of PEM format, but could be only one certificate in case of DER format
 - Variables names of certificates and private key buffers and sizes are defined as unions providing backward compatibility for legacy *_pem_buf and *_pem_bytes names which suggested only PEM format was supported. It is encouraged to use generic names such as cacert_buf and cacert_bytes.
-

Public Members

const char ****alpn_protos**

Application protocols required for HTTP2. If HTTP2/ALPN support is required, a list of protocols that should be negotiated. The format is length followed by protocol name. For the most common cases the following is ok: const char ****alpn_protos** = { "h2", NULL };

- where 'h2' is the protocol name

const unsigned char ***cacert_buf**

Certificate Authority's certificate in a buffer. Format may be PEM or DER, depending on mbedtls-support This buffer should be NULL terminated in case of PEM

const unsigned char ***cacert_pem_buf**

CA certificate buffer legacy name

unsigned int **cacert_bytes**

Size of Certificate Authority certificate pointed to by cacert_buf (including NULL-terminator in case of PEM format)

unsigned int **cacert_pem_bytes**

Size of Certificate Authority certificate legacy name

const unsigned char ***clientcert_buf**

Client certificate in a buffer Format may be PEM or DER, depending on mbedtls-support This buffer should be NULL terminated in case of PEM

const unsigned char ***clientcert_pem_buf**

Client certificate legacy name

unsigned int **clientcert_bytes**

Size of client certificate pointed to by clientcert_pem_buf (including NULL-terminator in case of PEM format)

unsigned int **clientcert_pem_bytes**

Size of client certificate legacy name

const unsigned char ***clientkey_buf**

Client key in a buffer Format may be PEM or DER, depending on mbedtls-support This buffer should be NULL terminated in case of PEM

const unsigned char ***clientkey_pem_buf**

Client key legacy name

unsigned int **clientkey_bytes**

Size of client key pointed to by clientkey_pem_buf (including NULL-terminator in case of PEM format)

unsigned int **clientkey_pem_bytes**

Size of client key legacy name

const unsigned char ***clientkey_password**

Client key decryption password string

unsigned int **clientkey_password_len**

String length of the password pointed to by clientkey_password

bool **use_ecdsa_peripheral**

Use the ECDSA peripheral for the private key operations

uint8_t **ecdsa_key_efuse_blk**

The efuse block where the ECDSA key is stored

bool **non_block**

Configure non-blocking mode. If set to true the underneath socket will be configured in non blocking mode after tls session is established

bool **use_secure_element**

Enable this option to use secure element or atec608a chip (Integrated with ESP32-WROOM-32SE)

int **timeout_ms**

Network timeout in milliseconds. Note: If this value is not set, by default the timeout is set to 10 seconds. If you wish that the session should wait indefinitely then please use a larger value e.g., INT32_MAX

bool **use_global_ca_store**

Use a global ca_store for all the connections in which this bool is set.

const char ***common_name**

If non-NULL, server certificate CN must match this name. If NULL, server certificate CN must match hostname.

bool **skip_common_name**

Skip any validation of server certificate CN field

tls_keep_alive_cfg_t ***keep_alive_cfg**

Enable TCP keep-alive timeout for SSL connection

const *psk_hint_key_t* ***psk_hint_key**

Pointer to PSK hint and key. if not NULL (and certificates are NULL) then PSK authentication is enabled with configured setup. Important note: the pointer must be valid for connection

esp_err_t (***crt_bundle_attach**)(void *conf)

Function pointer to esp_crt_bundle_attach. Enables the use of certification bundle for server verification, must be enabled in menuconfig

void ***ds_data**

Pointer for digital signature peripheral context

bool **is_plain_tcp**

Use non-TLS connection: When set to true, the esp-tls uses plain TCP transport rather than TLS/SSL connection. Note, that it is possible to connect using a plain tcp transport directly with esp_tls_plain_tcp_connect() API

struct ifreq ***if_name**

The name of interface for data to go through. Use the default interface without setting

esp_tls_addr_family_t **addr_family**

The address family to use when connecting to a host.

const int ***ciphersuites_list**

Pointer to a zero-terminated array of IANA identifiers of TLS ciphersuites. Please check the list validity by `esp_tls_get_ciphersuites_list()` API

esp_tls_proto_ver_t **tls_version**

TLS protocol version of the connection, e.g., TLS 1.2, TLS 1.3 (default - no preference)

Type Definitions

typedef enum *esp_tls_conn_state* **esp_tls_conn_state_t**

ESP-TLS Connection State.

typedef enum *esp_tls_role* **esp_tls_role_t**

typedef struct *psk_key_hint* **psk_hint_key_t**

ESP-TLS preshared key and hint structure.

typedef struct *tls_keep_alive_cfg* **tls_keep_alive_cfg_t**

esp-tls client session ticket ctx

Keep alive parameters structure

typedef enum *esp_tls_addr_family* **esp_tls_addr_family_t**

typedef struct *esp_tls_cfg* **esp_tls_cfg_t**

ESP-TLS configuration parameters.

Note: Note about format of certificates:

- This structure includes certificates of a Certificate Authority, of client or server as well as private keys, which may be of PEM or DER format. In case of PEM format, the buffer must be NULL terminated (with NULL character included in certificate size).
 - Certificate Authority's certificate may be a chain of certificates in case of PEM format, but could be only one certificate in case of DER format
 - Variables names of certificates and private key buffers and sizes are defined as unions providing backward compatibility for legacy `*_pem_buf` and `*_pem_bytes` names which suggested only PEM format was supported. It is encouraged to use generic names such as `cacert_buf` and `cacert_bytes`.
-

typedef struct esp_tls **esp_tls_t**

Enumerations

enum **esp_tls_conn_state**

ESP-TLS Connection State.

Values:

enumerator **ESP_TLS_INIT**

enumerator **ESP_TLS_CONNECTING**

enumerator **ESP_TLS_HANDSHAKE**

enumerator **ESP_TLS_FAIL**

enumerator **ESP_TLS_DONE**

enum **esp_tls_role**

Values:

enumerator **ESP_TLS_CLIENT**

enumerator **ESP_TLS_SERVER**

enum **esp_tls_addr_family**

Values:

enumerator **ESP_TLS_AF_UNSPEC**

Unspecified address family.

enumerator **ESP_TLS_AF_INET**

IPv4 address family.

enumerator **ESP_TLS_AF_INET6**

IPv6 address family.

enum **esp_tls_proto_ver_t**

Values:

enumerator **ESP_TLS_VER_ANY**

enumerator **ESP_TLS_VER_TLS_1_2**

enumerator **ESP_TLS_VER_TLS_1_3**

enumerator **ESP_TLS_VER_TLS_MAX**

Header File

- [components/esp-tls/esp_tls_errors.h](#)
- This header file can be included with:

```
#include "esp_tls_errors.h"
```

- This header file is a part of the API provided by the `esp-tls` component. To declare that your component depends on `esp-tls`, add the following to your `CMakeLists.txt`:

REQUIRES esp-tls

or

PRIV_REQUIRES esp-tls

Structures

struct **esp_tls_last_error**

Error structure containing relevant errors in case tls error occurred.

Public Members

esp_err_t **last_error**

error code (based on ESP_ERR_ESP_TLS_BASE) of the last occurred error

int **esp_tls_error_code**

esp_tls error code from last esp_tls failed api

int **esp_tls_flags**

last certification verification flags

Macros

ESP_ERR_ESP_TLS_BASE

Starting number of ESP-TLS error codes

ESP_ERR_ESP_TLS_CANNOT_RESOLVE_HOSTNAME

Error if hostname couldn't be resolved upon tls connection

ESP_ERR_ESP_TLS_CANNOT_CREATE_SOCKET

Failed to create socket

ESP_ERR_ESP_TLS_UNSUPPORTED_PROTOCOL_FAMILY

Unsupported protocol family

ESP_ERR_ESP_TLS_FAILED_CONNECT_TO_HOST

Failed to connect to host

ESP_ERR_ESP_TLS_SOCKET_SETOPT_FAILED

failed to set/get socket option

ESP_ERR_ESP_TLS_CONNECTION_TIMEOUT

new connection in esp_tls_low_level_conn connection timeouted

ESP_ERR_ESP_TLS_SE_FAILED

ESP_ERR_ESP_TLS_TCP_CLOSED_FIN

ESP_ERR_MBEDTLS_CERT_PARTLY_OK

mbedtls parse certificates was partly successful

ESP_ERR_MBEDTLS_CTR_DRBG_SEED_FAILED

mbedtls api returned error

ESP_ERR_MBEDTLS_SSL_SET_HOSTNAME_FAILED

mbedtls api returned error

ESP_ERR_MBEDTLS_SSL_CONFIG_DEFAULTS_FAILED

mbedtls api returned error

ESP_ERR_MBEDTLS_SSL_CONF_ALPN_PROTOCOLS_FAILED

mbedtls api returned error

ESP_ERR_MBEDTLS_X509_CRT_PARSE_FAILED

mbedtls api returned error

ESP_ERR_MBEDTLS_SSL_CONF_OWN_CERT_FAILED

mbedtls api returned error

ESP_ERR_MBEDTLS_SSL_SETUP_FAILED

mbedtls api returned error

ESP_ERR_MBEDTLS_SSL_WRITE_FAILED

mbedtls api returned error

ESP_ERR_MBEDTLS_PK_PARSE_KEY_FAILED

mbedtls api returned failed

ESP_ERR_MBEDTLS_SSL_HANDSHAKE_FAILED

mbedtls api returned failed

ESP_ERR_MBEDTLS_SSL_CONF_PSK_FAILED

mbedtls api returned failed

ESP_ERR_MBEDTLS_SSL_TICKET_SETUP_FAILED

mbedtls api returned failed

ESP_ERR_WOLFSSL_SSL_SET_HOSTNAME_FAILED

wolfSSL api returned error

ESP_ERR_WOLFSSL_SSL_CONF_ALPN_PROTOCOLS_FAILED

wolfSSL api returned error

ESP_ERR_WOLFSSL_CERT_VERIFY_SETUP_FAILED

wolfSSL api returned error

ESP_ERR_WOLFSSL_KEY_VERIFY_SETUP_FAILED

wolfSSL api returned error

ESP_ERR_WOLFSSL_SSL_HANDSHAKE_FAILED

wolfSSL api returned failed

ESP_ERR_WOLFSSL_CTX_SETUP_FAILED

wolfSSL api returned failed

ESP_ERR_WOLFSSL_SSL_SETUP_FAILED

wolfSSL api returned failed

ESP_ERR_WOLFSSL_SSL_WRITE_FAILED

wolfSSL api returned failed

ESP_TLS_ERR_SSL_WANT_READ

Definition of errors reported from IO API (potentially non-blocking) in case of error:

- `esp_tls_conn_read()`
- `esp_tls_conn_write()`

ESP_TLS_ERR_SSL_WANT_WRITE**ESP_TLS_ERR_SSL_TIMEOUT**

Type Definitions

typedef struct *esp_tls_last_error* ***esp_tls_error_handle_t**

typedef struct *esp_tls_last_error* **esp_tls_last_error_t**

Error structure containing relevant errors in case tls error occurred.

Enumerations

enum **esp_tls_error_type_t**

Definition of different types/sources of error codes reported from different components

Values:

enumerator **ESP_TLS_ERR_TYPE_UNKNOWN**

enumerator **ESP_TLS_ERR_TYPE_SYSTEM**

System error — errno

enumerator **ESP_TLS_ERR_TYPE_MBEDTLS**

Error code from mbedtls library

enumerator **ESP_TLS_ERR_TYPE_MBEDTLS_CERT_FLAGS**

Certificate flags defined in mbedtls

enumerator **ESP_TLS_ERR_TYPE_ESP**

ESP-IDF error type — esp_err_t

enumerator **ESP_TLS_ERR_TYPE_WOLFSSL**

Error code from wolfSSL library

enumerator **ESP_TLS_ERR_TYPE_WOLFSSL_CERT_FLAGS**

Certificate flags defined in wolfSSL

enumerator **ESP_TLS_ERR_TYPE_MAX**

Last err type — invalid entry

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.2.5 ESP HTTP Client

Overview

esp_http_client component provides a set of APIs for making HTTP/S requests from ESP-IDF applications. The steps to use these APIs are as follows:

- [esp_http_client_init\(\)](#): Creates an [esp_http_client_handle_t](#) instance, i.e., an HTTP client handle based on the given [esp_http_client_config_t](#) configuration. This function must be the first to be called; default values are assumed for the configuration values that are not explicitly defined by the user.
- [esp_http_client_perform\(\)](#): Performs all operations of the esp_http_client - opening the connection, exchanging data, and closing the connection (as required), while blocking the current task before its completion. All related events are invoked through the event handler (as specified in [esp_http_client_config_t](#)).
- [esp_http_client_cleanup\(\)](#): Closes the connection (if any) and frees up all the memory allocated to the HTTP client instance. This must be the last function to be called after the completion of operations.

Application Example

Simple example that uses ESP HTTP Client to make HTTP/S requests can be found at [protocols/esp_http_client](#).

Basic HTTP Request

Check out the example functions `http_rest_with_url` and `http_rest_with_hostname_path` in the application example for implementation details.

Persistent Connections

Persistent connection means that the HTTP client can re-use the same connection for several exchanges. If the server does not request to close the connection with the `Connection: close` header, the connection is not dropped but is instead kept open and used for further requests.

To allow ESP HTTP client to take full advantage of persistent connections, one should make as many requests as possible using the same handle instance. Check out the example functions `http_rest_with_url` and `http_rest_with_hostname_path` in the application example. Here, once the connection is created, multiple requests (GET, POST, PUT, etc.) are made before the connection is closed.

HTTPS Request

ESP HTTP client supports SSL connections using **mbedTLS**, with the `url` configuration starting with `https` scheme or `transport_type` set to `HTTP_TRANSPORT_OVER_SSL`. HTTPS support can be configured via `CONFIG_ESP_HTTP_CLIENT_ENABLE_HTTPS` (enabled by default).

Note: While making HTTPS requests, if server verification is needed, an additional root certificate (in PEM format) needs to be provided to the `cert_pem` member in the `esp_http_client_config_t` configuration. Users can also use the ESP x509 Certificate Bundle for server verification using the `crt_bundle_attach` member of the `esp_http_client_config_t` configuration.

Check out the example functions `https_with_url` and `https_with_hostname_path` in the application example for implementation details of the above note.

HTTP Stream

Some applications need to open the connection and control the exchange of data actively (data streaming). In such cases, the application flow is different from regular requests. Example flow is given below:

- `esp_http_client_init()`: Create a HTTP client handle.
- `esp_http_client_set_*` or `esp_http_client_delete_*`: Modify the HTTP connection parameters (optional).
- `esp_http_client_open()`: Open the HTTP connection with `write_len` parameter (content length that needs to be written to server), set `write_len=0` for read-only connection.
- `esp_http_client_write()`: Write data to server with a maximum length equal to `write_len` of `esp_http_client_open()` function; no need to call this function for `write_len=0`.
- `esp_http_client_fetch_headers()`: Read the HTTP Server response headers, after sending the request headers and server data (if any). Returns the `content-length` from the server and can be succeeded by `esp_http_client_get_status_code()` for getting the HTTP status of the connection.
- `esp_http_client_read()`: Read the HTTP stream.
- `esp_http_client_close()`: Close the connection.
- `esp_http_client_cleanup()`: Release allocated resources.

Check out the example function `http_perform_as_stream_reader` in the application example for implementation details.

HTTP Authentication

ESP HTTP client supports both Basic and Digest Authentication.

- Users can provide the username and password in the `url` or the `username` and `password` members of the `esp_http_client_config_t` configuration. For `auth_type = HTTP_AUTH_TYPE_BASIC`, the HTTP client takes only one perform operation to pass the authentication process.
- If `auth_type = HTTP_AUTH_TYPE_NONE`, but the `username` and `password` fields are present in the configuration, the HTTP client takes two perform operations. The client will receive the 401 Unauthorized header in its first attempt to connect to the server. Based on this information, it decides which authentication method to choose and performs it in the second operation.
- Check out the example functions `http_auth_basic`, `http_auth_basic_redirect` (for Basic authentication) and `http_auth_digest` (for Digest authentication) in the application example for implementation details.

Examples of Authentication Configuration

- Authentication with URI

```
esp_http_client_config_t config = {
    .url = "http://user:passwd@httpbin.org/basic-auth/user/passwd",
    .auth_type = HTTP_AUTH_TYPE_BASIC,
};
```

- Authentication with username and password entry

```
esp_http_client_config_t config = {
    .url = "http://httpbin.org/basic-auth/user/passwd",
    .username = "user",
    .password = "passwd",
    .auth_type = HTTP_AUTH_TYPE_BASIC,
};
```

Event Handling

ESP HTTP Client supports event handling by triggering an event handler corresponding to the event which takes place. [esp_http_client_event_id_t](#) contains all the events which could occur while performing an HTTP request using the ESP HTTP Client.

To enable event handling, you just need to set a callback function using the [esp_http_client_config_t::event_handler](#) member.

ESP HTTP Client Diagnostic Information

Diagnostic information could be helpful to gain insights into a problem. In the case of ESP HTTP Client, the diagnostic information can be collected by registering an event handler with [the Event Loop library](#). This feature has been added by keeping in mind the [ESP Insights](#) framework which collects the diagnostic information. However, this feature can also be used without any dependency on the ESP Insights framework for the diagnostic purpose. Event handler can be registered to the event loop using the [esp_event_handler_register\(\)](#) function.

Expected data types for different HTTP Client events in the event loop are as follows:

- `HTTP_EVENT_ERROR`: [esp_http_client_handle_t](#)
- `HTTP_EVENT_ON_CONNECTED`: [esp_http_client_handle_t](#)
- `HTTP_EVENT_HEADERS_SENT`: [esp_http_client_handle_t](#)
- `HTTP_EVENT_ON_HEADER`: [esp_http_client_handle_t](#)
- `HTTP_EVENT_ON_DATA`: [esp_http_client_on_data_t](#)
- `HTTP_EVENT_ON_FINISH`: [esp_http_client_handle_t](#)
- `HTTP_EVENT_DISCONNECTED`: [esp_http_client_handle_t](#)
- `HTTP_EVENT_REDIRECT`: [esp_http_client_redirect_event_data_t](#)

The [esp_http_client_handle_t](#) received along with the event data will be valid until [HTTP_EVENT_DISCONNECTED](#) is not received. This handle has been sent primarily to differentiate between different client connections and must not be used for any other purpose, as it may change based on client connection state.

API Reference

Header File

- [components/esp_http_client/include/esp_http_client.h](#)
- This header file can be included with:

```
#include "esp_http_client.h"
```

- This header file is a part of the API provided by the `esp_http_client` component. To declare that your component depends on `esp_http_client`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_http_client
```

or

```
PRIV_REQUIRES esp_http_client
```

Functions

esp_http_client_handle_t **esp_http_client_init** (const *esp_http_client_config_t* *config)

Start a HTTP session This function must be the first function to call, and it returns a `esp_http_client_handle_t` that you must use as input to other functions in the interface. This call **MUST** have a corresponding call to `esp_http_client_cleanup` when the operation is complete.

Parameters `config` -- [in] The configurations, see `http_client_config_t`

Returns

- `esp_http_client_handle_t`
- NULL if any errors

esp_err_t **esp_http_client_perform** (*esp_http_client_handle_t* client)

Invoke this function after `esp_http_client_init` and all the options calls are made, and will perform the transfer as described in the options. It must be called with the same `esp_http_client_handle_t` as input as the `esp_http_client_init` call returned. `esp_http_client_perform` performs the entire request in either blocking or non-blocking manner. By default, the API performs request in a blocking manner and returns when done, or if it failed, and in non-blocking manner, it returns if `EAGAIN/EWOULDBLOCK` or `EINPROGRESS` is encountered, or if it failed. And in case of non-blocking request, the user may call this API multiple times unless request & response is complete or there is a failure. To enable non-blocking `esp_http_client_perform()`, `is_async` member of *esp_http_client_config_t* must be set while making a call to `esp_http_client_init()` API. You can do any amount of calls to `esp_http_client_perform` while using the same `esp_http_client_handle_t`. The underlying connection may be kept open if the server allows it. If you intend to transfer more than one file, you are even encouraged to do so. `esp_http_client` will then attempt to reuse the same connection for the following transfers, thus making the operations faster, less CPU intense and using less network resources. Just note that you will have to use `esp_http_client_set_*` between the invokes to set options for the following `esp_http_client_perform`.

Note: You must never call this function simultaneously from two places using the same client handle. Let the function return first before invoking it another time. If you want parallel transfers, you must use several `esp_http_client_handle_t`. This function include `esp_http_client_open` -> `esp_http_client_write` -> `esp_http_client_fetch_headers` -> `esp_http_client_read` (and option) `esp_http_client_close`.

Parameters `client` -- The `esp_http_client` handle

Returns

- `ESP_OK` on successful
- `ESP_FAIL` on error

esp_err_t **esp_http_client_cancel_request** (*esp_http_client_handle_t* client)

Cancel an ongoing HTTP request. This API closes the current socket and opens a new socket with the same `esp_http_client` context.

Parameters `client` -- The `esp_http_client` handle

Returns

- `ESP_OK` on successful
- `ESP_FAIL` on error
- `ESP_ERR_INVALID_ARG`
- `ESP_ERR_INVALID_STATE`

esp_err_t **esp_http_client_set_url** (*esp_http_client_handle_t* client, const char *url)

Set URL for client, when performing this behavior, the options in the URL will replace the old ones.

Parameters

- **client** -- [in] The esp_http_client handle
- **url** -- [in] The url

Returns

- ESP_OK
- ESP_FAIL

esp_err_t **esp_http_client_set_post_field** (*esp_http_client_handle_t* client, const char *data, int len)

Set post data, this function must be called before esp_http_client_perform. Note: The data parameter passed to this function is a pointer and this function will not copy the data.

Parameters

- **client** -- [in] The esp_http_client handle
- **data** -- [in] post data pointer
- **len** -- [in] post length

Returns

- ESP_OK
- ESP_FAIL

int **esp_http_client_get_post_field** (*esp_http_client_handle_t* client, char **data)

Get current post field information.

Parameters

- **client** -- [in] The client
- **data** -- [out] Point to post data pointer

Returns Size of post data

esp_err_t **esp_http_client_set_header** (*esp_http_client_handle_t* client, const char *key, const char *value)

Set http request header, this function must be called after esp_http_client_init and before any perform function.

Parameters

- **client** -- [in] The esp_http_client handle
- **key** -- [in] The header key
- **value** -- [in] The header value

Returns

- ESP_OK
- ESP_FAIL

esp_err_t **esp_http_client_get_header** (*esp_http_client_handle_t* client, const char *key, char **value)

Get http request header. The value parameter will be set to NULL if there is no header which is same as the key specified, otherwise the address of header value will be assigned to value parameter. This function must be called after esp_http_client_init.

Parameters

- **client** -- [in] The esp_http_client handle
- **key** -- [in] The header key
- **value** -- [out] The header value

Returns

- ESP_OK
- ESP_FAIL

esp_err_t **esp_http_client_get_username** (*esp_http_client_handle_t* client, char **value)

Get http request username. The address of username buffer will be assigned to value parameter. This function must be called after esp_http_client_init.

Parameters

- **client** -- [in] The esp_http_client handle
- **value** -- [out] The username value

Returns

- ESP_OK
- ESP_ERR_INVALID_ARG

esp_err_t **esp_http_client_set_username** (*esp_http_client_handle_t* client, const char *username)

Set http request username. The value of username parameter will be assigned to username buffer. If the username parameter is NULL then username buffer will be freed.

Parameters

- **client** -- [in] The esp_http_client handle
- **username** -- [in] The username value

Returns

- ESP_OK
- ESP_ERR_INVALID_ARG

esp_err_t **esp_http_client_get_password** (*esp_http_client_handle_t* client, char **value)

Get http request password. The address of password buffer will be assigned to value parameter. This function must be called after `esp_http_client_init`.

Parameters

- **client** -- [in] The esp_http_client handle
- **value** -- [out] The password value

Returns

- ESP_OK
- ESP_ERR_INVALID_ARG

esp_err_t **esp_http_client_set_password** (*esp_http_client_handle_t* client, const char *password)

Set http request password. The value of password parameter will be assigned to password buffer. If the password parameter is NULL then password buffer will be freed.

Parameters

- **client** -- [in] The esp_http_client handle
- **password** -- [in] The password value

Returns

- ESP_OK
- ESP_ERR_INVALID_ARG

esp_err_t **esp_http_client_set_auth_type** (*esp_http_client_handle_t* client, *esp_http_client_auth_type_t* auth_type)

Set http request auth_type.

Parameters

- **client** -- [in] The esp_http_client handle
- **auth_type** -- [in] The esp_http_client auth type

Returns

- ESP_OK
- ESP_ERR_INVALID_ARG

esp_err_t **esp_http_client_get_user_data** (*esp_http_client_handle_t* client, void **data)

Get http request user_data. The value stored from the *esp_http_client_config_t* will be written to the address passed into data.

Parameters

- **client** -- [in] The esp_http_client handle
- **data** -- [out] A pointer to the pointer that will be set to user_data.

Returns

- ESP_OK
- ESP_ERR_INVALID_ARG

esp_err_t **esp_http_client_set_user_data** (*esp_http_client_handle_t* client, void *data)

Set http request user_data. The value passed in +data+ will be available during event callbacks. No memory management will be performed on the user's behalf.

Parameters

- **client** -- [in] The esp_http_client handle
- **data** -- [in] The pointer to the user data

Returns

- ESP_OK
- ESP_ERR_INVALID_ARG

int **esp_http_client_get_errno** (*esp_http_client_handle_t* client)

Get HTTP client session errno.

Parameters **client** -- [in] The esp_http_client handle

Returns

- (-1) if invalid argument
- errno

esp_err_t **esp_http_client_set_method** (*esp_http_client_handle_t* client, *esp_http_client_method_t* method)

Set http request method.

Parameters

- **client** -- [in] The esp_http_client handle
- **method** -- [in] The method

Returns

- ESP_OK
- ESP_ERR_INVALID_ARG

esp_err_t **esp_http_client_set_timeout_ms** (*esp_http_client_handle_t* client, int timeout_ms)

Set http request timeout.

Parameters

- **client** -- [in] The esp_http_client handle
- **timeout_ms** -- [in] The timeout value

Returns

- ESP_OK
- ESP_ERR_INVALID_ARG

esp_err_t **esp_http_client_delete_header** (*esp_http_client_handle_t* client, const char *key)

Delete http request header.

Parameters

- **client** -- [in] The esp_http_client handle
- **key** -- [in] The key

Returns

- ESP_OK
- ESP_FAIL

esp_err_t **esp_http_client_open** (*esp_http_client_handle_t* client, int write_len)

This function will be open the connection, write all header strings and return.

Parameters

- **client** -- [in] The esp_http_client handle
- **write_len** -- [in] HTTP Content length need to write to the server

Returns

- ESP_OK
- ESP_FAIL

int **esp_http_client_write** (*esp_http_client_handle_t* client, const char *buffer, int len)

This function will write data to the HTTP connection previously opened by esp_http_client_open()

Parameters

- **client** -- [in] The esp_http_client handle
- **buffer** -- The buffer

- **len** -- **[in]** This value must not be larger than the `write_len` parameter provided to `esp_http_client_open()`

Returns

- (-1) if any errors
- Length of data written

int64_t **esp_http_client_fetch_headers** (*esp_http_client_handle_t* client)

This function need to call after `esp_http_client_open`, it will read from http stream, process all receive headers.

Parameters **client** -- **[in]** The `esp_http_client` handle

Returns

- (0) if stream doesn't contain content-length header, or chunked encoding (checked by `esp_http_client_is_chunked` response)
- (-1: `ESP_FAIL`) if any errors
- (`-ESP_ERR_HTTP_EAGAIN = -0x7007`) if call is timed-out before any data was ready
- Download data length defined by content-length header

bool **esp_http_client_is_chunked_response** (*esp_http_client_handle_t* client)

Check response data is chunked.

Parameters **client** -- **[in]** The `esp_http_client` handle

Returns true or false

int **esp_http_client_read** (*esp_http_client_handle_t* client, char *buffer, int len)

Read data from http stream.

Note: (`-ESP_ERR_HTTP_EAGAIN = -0x7007`) is returned when call is timed-out before any data was ready

Parameters

- **client** -- **[in]** The `esp_http_client` handle
- **buffer** -- The buffer
- **len** -- **[in]** The length

Returns

- (-1) if any errors
- Length of data was read

int **esp_http_client_get_status_code** (*esp_http_client_handle_t* client)

Get http response status code, the valid value if this function invoke after `esp_http_client_perform`

Parameters **client** -- **[in]** The `esp_http_client` handle

Returns Status code

int64_t **esp_http_client_get_content_length** (*esp_http_client_handle_t* client)

Get http response content length (from header Content-Length) the valid value if this function invoke after `esp_http_client_perform`

Parameters **client** -- **[in]** The `esp_http_client` handle

Returns

- (-1) Chunked transfer
- Content-Length value as bytes

esp_err_t **esp_http_client_close** (*esp_http_client_handle_t* client)

Close http connection, still kept all http request resources.

Parameters **client** -- **[in]** The `esp_http_client` handle

Returns

- `ESP_OK`
- `ESP_FAIL`

esp_err_t **esp_http_client_cleanup** (*esp_http_client_handle_t* client)

This function must be the last function to call for an session. It is the opposite of the `esp_http_client_init` function and must be called with the same handle as input that a `esp_http_client_init` call returned. This might close all connections this handle has used and possibly has kept open until now. Don't call this function if you intend to transfer more files, re-using handles is a key to good performance with `esp_http_client`.

Parameters `client` -- [in] The `esp_http_client` handle

Returns

- ESP_OK
- ESP_FAIL

esp_http_client_transport_t **esp_http_client_get_transport_type** (*esp_http_client_handle_t* client)

Get transport type.

Parameters `client` -- [in] The `esp_http_client` handle

Returns

- HTTP_TRANSPORT_UNKNOWN
- HTTP_TRANSPORT_OVER_TCP
- HTTP_TRANSPORT_OVER_SSL

esp_err_t **esp_http_client_set_redirection** (*esp_http_client_handle_t* client)

Set redirection URL. When received the 30x code from the server, the client stores the redirect URL provided by the server. This function will set the current URL to redirect to enable client to execute the redirection request. When `disable_auto_redirect` is set, the client will not call this function but the event HTTP_EVENT_REDIRECT will be dispatched giving the user control over the redirection event.

Parameters `client` -- [in] The `esp_http_client` handle

Returns

- ESP_OK
- ESP_FAIL

esp_err_t **esp_http_client_reset_redirect_counter** (*esp_http_client_handle_t* client)

Reset the redirection counter. This is useful to reset redirect counter in cases where the same handle is used for multiple requests.

Parameters `client` -- [in] The `esp_http_client` handle

Returns

- ESP_OK
- ESP_ERR_INVALID_ARG

esp_err_t **esp_http_client_set_auth_data** (*esp_http_client_handle_t* client, const char *auth_data, int len)

On receiving a custom authentication header, this API can be invoked to set the authentication information from the header. This API can be called from the event handler.

Parameters

- `client` -- [in] The `esp_http_client` handle
- `auth_data` -- [in] The authentication data received in the header
- `len` -- [in] length of `auth_data`.

Returns

- ESP_ERR_INVALID_ARG
- ESP_OK

void **esp_http_client_add_auth** (*esp_http_client_handle_t* client)

On receiving HTTP Status code 401, this API can be invoked to add authorization information.

Note: There is a possibility of receiving body message with redirection status codes, thus make sure to flush off body data after calling this API.

Parameters `client` -- [in] The `esp_http_client` handle

bool **esp_http_client_is_complete_data_received** (*esp_http_client_handle_t* client)

Checks if entire data in the response has been read without any error.

Parameters **client** -- [in] The esp_http_client handle

Returns

- true
- false

int **esp_http_client_read_response** (*esp_http_client_handle_t* client, char *buffer, int len)

Helper API to read larger data chunks This is a helper API which internally calls esp_http_client_read multiple times till the end of data is reached or till the buffer gets full.

Parameters

- **client** -- [in] The esp_http_client handle
- **buffer** -- The buffer
- **len** -- [in] The buffer length

Returns

- Length of data was read

esp_err_t **esp_http_client_flush_response** (*esp_http_client_handle_t* client, int *len)

Process all remaining response data This uses an internal buffer to repeatedly receive, parse, and discard response data until complete data is processed. As no additional user-supplied buffer is required, this may be preferable to esp_http_client_read_response in situations where the content of the response may be ignored.

Parameters

- **client** -- [in] The esp_http_client handle
- **len** -- Length of data discarded

Returns

- ESP_OK If successful, len will have discarded length
- ESP_FAIL If failed to read response
- ESP_ERR_INVALID_ARG If the client is NULL

esp_err_t **esp_http_client_get_url** (*esp_http_client_handle_t* client, char *url, const int len)

Get URL from client.

Parameters

- **client** -- [in] The esp_http_client handle
- **url** -- [inout] The buffer to store URL
- **len** -- [in] The buffer length

Returns

- ESP_OK
- ESP_FAIL

esp_err_t **esp_http_client_get_chunk_length** (*esp_http_client_handle_t* client, int *len)

Get Chunk-Length from client.

Parameters

- **client** -- [in] The esp_http_client handle
- **len** -- [out] Variable to store length

Returns

- ESP_OK If successful, len will have length of current chunk
- ESP_FAIL If the server is not a chunked server
- ESP_ERR_INVALID_ARG If the client or len are NULL

Structures

struct **esp_http_client_event**

HTTP Client events data.

Public Members

esp_http_client_event_id_t **event_id**

event_id, to know the cause of the event

esp_http_client_handle_t **client**

esp_http_client_handle_t context

void ***data**

data of the event

int **data_len**

data length of data

void ***user_data**

user_data context, from *esp_http_client_config_t* user_data

char ***header_key**

For HTTP_EVENT_ON_HEADER event_id, it's store current http header key

char ***header_value**

For HTTP_EVENT_ON_HEADER event_id, it's store current http header value

struct **esp_http_client_on_data**

Argument structure for HTTP_EVENT_ON_DATA event.

Public Members

esp_http_client_handle_t **client**

Client handle

int64_t **data_process**

Total data processed

struct **esp_http_client_redirect_event_data**

Argument structure for HTTP_EVENT_REDIRECT event.

Public Members

esp_http_client_handle_t **client**

Client handle

int **status_code**

Status Code

struct **esp_http_client_config_t**

HTTP configuration.

Public Members

const char ***url**

HTTP URL, the information on the URL is most important, it overrides the other fields below, if any

const char ***host**

Domain or IP as string

int **port**

Port to connect, default depend on esp_http_client_transport_t (80 or 443)

const char ***username**

Using for Http authentication

const char ***password**

Using for Http authentication

esp_http_client_auth_type_t **auth_type**

Http authentication type, see esp_http_client_auth_type_t

const char ***path**

HTTP Path, if not set, default is /

const char ***query**

HTTP query

const char ***cert_pem**

SSL server certification, PEM format as string, if the client requires to verify server

size_t **cert_len**

Length of the buffer pointed to by cert_pem. May be 0 for null-terminated pem

const char ***client_cert_pem**

SSL client certification, PEM format as string, if the server requires to verify client

size_t **client_cert_len**

Length of the buffer pointed to by client_cert_pem. May be 0 for null-terminated pem

const char ***client_key_pem**

SSL client key, PEM format as string, if the server requires to verify client

size_t **client_key_len**

Length of the buffer pointed to by client_key_pem. May be 0 for null-terminated pem

const char ***client_key_password**

Client key decryption password string

size_t **client_key_password_len**

String length of the password pointed to by client_key_password

esp_http_client_proto_ver_t **tls_version**

TLS protocol version of the connection, e.g., TLS 1.2, TLS 1.3 (default - no preference)

const char ***user_agent**

The User Agent string to send with HTTP requests

esp_http_client_method_t **method**

HTTP Method

int **timeout_ms**

Network timeout in milliseconds

bool **disable_auto_redirect**

Disable HTTP automatic redirects

int **max_redirection_count**

Max number of redirections on receiving HTTP redirect status code, using default value if zero

int **max_authorization_retries**

Max connection retries on receiving HTTP unauthorized status code, using default value if zero. Disables authorization retry if -1

http_event_handle_cb **event_handler**

HTTP Event Handle

esp_http_client_transport_t **transport_type**

HTTP transport type, see `esp_http_client_transport_t`

int **buffer_size**

HTTP receive buffer size

int **buffer_size_tx**

HTTP transmit buffer size

void ***user_data**

HTTP user_data context

bool **is_async**

Set asynchronous mode, only supported with HTTPS for now

bool **use_global_ca_store**

Use a global ca_store for all the connections in which this bool is set.

bool **skip_cert_common_name_check**

Skip any validation of server certificate CN field

const char ***common_name**

Pointer to the string containing server certificate common name. If non-NULL, server certificate CN must match this name, If NULL, server certificate CN must match hostname.

esp_err_t (***crt_bundle_attach**)(void *conf)

Function pointer to esp_crt_bundle_attach. Enables the use of certification bundle for server verification, must be enabled in menuconfig

bool **keep_alive_enable**

Enable keep-alive timeout

int **keep_alive_idle**

Keep-alive idle time. Default is 5 (second)

int **keep_alive_interval**

Keep-alive interval time. Default is 5 (second)

int **keep_alive_count**

Keep-alive packet retry send count. Default is 3 counts

struct ifreq ***if_name**

The name of interface for data to go through. Use the default interface without setting

void ***ds_data**

Pointer for digital signature peripheral context, see ESP-TLS Documentation for more details

esp_http_client_addr_type_t **addr_type**

Address type used in http client configurations

Macros

DEFAULT_HTTP_BUF_SIZE

ESP_ERR_HTTP_BASE

Starting number of HTTP error codes

ESP_ERR_HTTP_MAX_REDIRECT

The error exceeds the number of HTTP redirects

ESP_ERR_HTTP_CONNECT

Error open the HTTP connection

ESP_ERR_HTTP_WRITE_DATA

Error write HTTP data

ESP_ERR_HTTP_FETCH_HEADER

Error read HTTP header from server

ESP_ERR_HTTP_INVALID_TRANSPORT

There are no transport support for the input scheme

ESP_ERR_HTTP_CONNECTING

HTTP connection hasn't been established yet

ESP_ERR_HTTP_EAGAIN

Mapping of errno EAGAIN to esp_err_t

ESP_ERR_HTTP_CONNECTION_CLOSED

Read FIN from peer and the connection closed

Type Definitions

typedef struct esp_http_client ***esp_http_client_handle_t**

typedef struct *esp_http_client_event* ***esp_http_client_event_handle_t**

typedef struct *esp_http_client_event* **esp_http_client_event_t**

HTTP Client events data.

typedef struct *esp_http_client_on_data* **esp_http_client_on_data_t**

Argument structure for HTTP_EVENT_ON_DATA event.

typedef struct *esp_http_client_redirect_event_data* **esp_http_client_redirect_event_data_t**

Argument structure for HTTP_EVENT_REDIRECT event.

typedef *esp_err_t* (***http_event_handle_cb**)(*esp_http_client_event_t* *evt)

Enumerations

enum **esp_http_client_event_id_t**

HTTP Client events id.

Values:

enumerator **HTTP_EVENT_ERROR**

This event occurs when there are any errors during execution

enumerator **HTTP_EVENT_ON_CONNECTED**

Once the HTTP has been connected to the server, no data exchange has been performed

enumerator **HTTP_EVENT_HEADERS_SENT**

After sending all the headers to the server

enumerator **HTTP_EVENT_HEADER_SENT**

This header has been kept for backward compatibility and will be deprecated in future versions esp-idf

enumerator **HTTP_EVENT_ON_HEADER**

Occurs when receiving each header sent from the server

enumerator **HTTP_EVENT_ON_DATA**

Occurs when receiving data from the server, possibly multiple portions of the packet

enumerator **HTTP_EVENT_ON_FINISH**

Occurs when finish a HTTP session

enumerator **HTTP_EVENT_DISCONNECTED**

The connection has been disconnected

enumerator **HTTP_EVENT_REDIRECT**

Intercepting HTTP redirects to handle them manually

enum **esp_http_client_transport_t**

HTTP Client transport.

Values:

enumerator **HTTP_TRANSPORT_UNKNOWN**

Unknown

enumerator **HTTP_TRANSPORT_OVER_TCP**

Transport over tcp

enumerator **HTTP_TRANSPORT_OVER_SSL**

Transport over ssl

enum **esp_http_client_proto_ver_t**

Values:

enumerator **ESP_HTTP_CLIENT_TLS_VER_ANY**

enumerator **ESP_HTTP_CLIENT_TLS_VER_TLS_1_2**

enumerator **ESP_HTTP_CLIENT_TLS_VER_TLS_1_3**

enumerator **ESP_HTTP_CLIENT_TLS_VER_MAX**

enum **esp_http_client_method_t**

HTTP method.

Values:

enumerator **HTTP_METHOD_GET**

HTTP GET Method

enumerator **HTTP_METHOD_POST**

HTTP POST Method

enumerator **HTTP_METHOD_PUT**

HTTP PUT Method

enumerator **HTTP_METHOD_PATCH**

HTTP PATCH Method

enumerator **HTTP_METHOD_DELETE**

HTTP DELETE Method

enumerator **HTTP_METHOD_HEAD**
HTTP HEAD Method

enumerator **HTTP_METHOD_NOTIFY**
HTTP NOTIFY Method

enumerator **HTTP_METHOD_SUBSCRIBE**
HTTP SUBSCRIBE Method

enumerator **HTTP_METHOD_UNSUBSCRIBE**
HTTP UNSUBSCRIBE Method

enumerator **HTTP_METHOD_OPTIONS**
HTTP OPTIONS Method

enumerator **HTTP_METHOD_COPY**
HTTP COPY Method

enumerator **HTTP_METHOD_MOVE**
HTTP MOVE Method

enumerator **HTTP_METHOD_LOCK**
HTTP LOCK Method

enumerator **HTTP_METHOD_UNLOCK**
HTTP UNLOCK Method

enumerator **HTTP_METHOD_PROPFIND**
HTTP PROPFIND Method

enumerator **HTTP_METHOD_PROPPATCH**
HTTP PROPPATCH Method

enumerator **HTTP_METHOD_MKCOL**
HTTP MKCOL Method

enumerator **HTTP_METHOD_MAX**

enum **esp_http_client_auth_type_t**
HTTP Authentication type.

Values:

enumerator **HTTP_AUTH_TYPE_NONE**
No authentication

enumerator **HTTP_AUTH_TYPE_BASIC**
HTTP Basic authentication

enumerator **HTTP_AUTH_TYPE_DIGEST**

HTTP Digest authentication

enum **esp_http_client_addr_type_t**

Values:

enumerator **HTTP_ADDR_TYPE_UNSPEC**

Unspecified address family.

enumerator **HTTP_ADDR_TYPE_INET**

IPv4 address family.

enumerator **HTTP_ADDR_TYPE_INET6**

IPv6 address family.

enum **HttpStatus_Code**

Enum for the HTTP status codes.

Values:

enumerator **HttpStatus_Ok**

enumerator **HttpStatus_MultipleChoices**

enumerator **HttpStatus_MovedPermanently**

enumerator **HttpStatus_Found**

enumerator **HttpStatus_SeeOther**

enumerator **HttpStatus_TemporaryRedirect**

enumerator **HttpStatus_PermanentRedirect**

enumerator **HttpStatus_BadRequest**

enumerator **HttpStatus_Unauthorized**

enumerator **HttpStatus_Forbidden**

enumerator **HttpStatus_NotFound**

enumerator **HttpStatus_InternalError**

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.2.6 ESP Local Control

Overview

ESP Local Control (**esp_local_ctrl**) component in ESP-IDF provides capability to control an ESP device over HTTPS or Bluetooth® Low Energy. It provides access to application defined **properties** that are available for reading/writing via a set of configurable handlers.

Initialization of the **esp_local_ctrl** service over Bluetooth Low Energy transport is performed as follows:

```
esp_local_ctrl_config_t config = {
    .transport = ESP_LOCAL_CTRL_TRANSPORT_BLE,
    .transport_config = {
        .ble = & (protocomm_ble_config_t) {
            .device_name = SERVICE_NAME,
            .service_uuid = {
                /* LSB <-----> MSB */
                0x21, 0xd5, 0x3b, 0x8d, 0xbd, 0x75, 0x68, 0x8a,
                0xb4, 0x42, 0xeb, 0x31, 0x4a, 0x1e, 0x98, 0x3d
            }
        },
    },
    .proto_sec = {
        .version = PROTOCOM_SEC0,
        .custom_handle = NULL,
        .sec_params = NULL,
    },
    .handlers = {
        /* User defined handler functions */
        .get_prop_values = get_property_values,
        .set_prop_values = set_property_values,
        .usr_ctx = NULL,
        .usr_ctx_free_fn = NULL
    },
    /* Maximum number of properties that may be set */
    .max_properties = 10
};

/* Start esp_local_ctrl service */
ESP_ERROR_CHECK(esp_local_ctrl_start(&config));
```

Similarly for HTTPS transport:

```
/* Set the configuration */
httpd_ssl_config_t https_conf = HTTPD_SSL_CONFIG_DEFAULT();

/* Load server certificate */
extern const unsigned char servercert_start[] asm("_binary_servercert_pem_
↪start");
extern const unsigned char servercert_end[] asm("_binary_servercert_pem_
↪end");
https_conf.servercert = servercert_start;
https_conf.servercert_len = servercert_end - servercert_start;

/* Load server private key */
extern const unsigned char prvtkey_pem_start[] asm("_binary_prvtkey_pem_
↪start");
extern const unsigned char prvtkey_pem_end[] asm("_binary_prvtkey_pem_
↪end");
```

(continues on next page)

(continued from previous page)

```

https_conf.prvtkkey_pem = prvtkey_pem_start;
https_conf.prvtkkey_len = prvtkey_pem_end - prvtkey_pem_start;

esp_local_ctrl_config_t config = {
    .transport = ESP_LOCAL_CTRL_TRANSPORT_HTTPD,
    .transport_config = {
        .httpd = &https_conf
    },
    .proto_sec = {
        .version = PROTOCOL_SEC0,
        .custom_handle = NULL,
        .sec_params = NULL,
    },
    .handlers = {
        /* User defined handler functions */
        .get_prop_values = get_property_values,
        .set_prop_values = set_property_values,
        .usr_ctx         = NULL,
        .usr_ctx_free_fn = NULL
    },
    /* Maximum number of properties that may be set */
    .max_properties = 10
};

/* Start esp_local_ctrl service */
ESP_ERROR_CHECK(esp_local_ctrl_start(&config));

```

You may set security for transport in ESP local control using following options:

1. PROTOCOL_SEC2: specifies that SRP6a-based key exchange and end-to-end encryption based on AES-GCM are used. This is the most preferred option as it adds a robust security with Augmented PAKE protocol, i.e., SRP6a.
2. PROTOCOL_SEC1: specifies that Curve25519-based key exchange and end-to-end encryption based on AES-CTR are used.
3. PROTOCOL_SEC0: specifies that data will be exchanged as a plain text (no security).
4. PROTOCOL_SEC_CUSTOM: you can define your own security requirement. Please note that you will also have to provide `custom_handle` of type `protocomm_security_t *` in this context.

Note: The respective security schemes need to be enabled through the project configuration menu. Please refer to the Enabling protocol security version section in [Protocol Communication](#) for more details.

Creating a Property

Now that we know how to start the `esp_local_ctrl` service, let's add a property to it. Each property must have a unique `name` (string), a `type` (e.g., enum), `flags` (bit fields) and `size`.

The `size` is to be kept 0, if we want our property value to be of variable length (e.g., if it is a string or bytestream). For data types with fixed-length property value, like int, float, etc., setting the `size` field to the right value helps `esp_local_ctrl` to perform internal checks on arguments received with write requests.

The interpretation of `type` and `flags` fields is totally upto the application, hence they may be used as enumerations, bitfields, or even simple integers. One way is to use `type` values to classify properties, while `flags` to specify characteristics of a property.

Here is an example property which is to function as a timestamp. It is assumed that the application defines `TYPE_TIMESTAMP` and `READONLY`, which are used for setting the `type` and `flags` fields here.

```

/* Create a timestamp property */
esp_local_ctrl_prop_t timestamp = {
    .name      = "timestamp",
    .type      = TYPE_TIMESTAMP,
    .size      = sizeof(int32_t),
    .flags     = READONLY,
    .ctx       = func_get_time,
    .ctx_free_fn = NULL
};

/* Now register the property */
esp_local_ctrl_add_property(&timestamp);

```

Also notice that there is a `ctx` field, which is set to point to some custom `func_get_time()`. This can be used inside the property get/set handlers to retrieve timestamp.

Here is an example of `get_prop_values()` handler, which is used for retrieving the timestamp.

```

static esp_err_t get_property_values(size_t props_count,
                                     const esp_local_ctrl_prop_t *props,
                                     esp_local_ctrl_prop_val_t *prop_
                                     ↪values,
                                     void *usr_ctx)
{
    for (uint32_t i = 0; i < props_count; i++) {
        ESP_LOGI(TAG, "Reading %s", props[i].name);
        if (props[i].type == TYPE_TIMESTAMP) {
            /* Obtain the timer function from ctx */
            int32_t (*func_get_time)(void) = props[i].ctx;

            /* Use static variable for saving the value. This is_
            ↪essential because the value has to be valid even after this function_
            ↪returns. Alternative is to use dynamic allocation and set the free_fn_
            ↪field */
            static int32_t ts = func_get_time();
            prop_values[i].data = &ts;
        }
    }
    return ESP_OK;
}

```

Here is an example of `set_prop_values()` handler. Notice how we restrict from writing to read-only properties.

```

static esp_err_t set_property_values(size_t props_count,
                                     const esp_local_ctrl_prop_t *props,
                                     const esp_local_ctrl_prop_val_t_
                                     ↪*prop_values,
                                     void *usr_ctx)
{
    for (uint32_t i = 0; i < props_count; i++) {
        if (props[i].flags & READONLY) {
            ESP_LOGE(TAG, "Cannot write to read-only property %s",_
            ↪props[i].name);
            return ESP_ERR_INVALID_ARG;
        } else {
            ESP_LOGI(TAG, "Setting %s", props[i].name);

            /* For keeping it simple, lets only log the incoming data */
            ESP_LOG_BUFFER_HEX_LEVEL(TAG, prop_values[i].data,
                                     prop_values[i].size, ESP_LOG_INFO);
        }
    }
}

```

(continues on next page)

(continued from previous page)

```

    return ESP_OK;
}

```

For complete example see [protocols/esp_local_ctrl](#).

Client Side Implementation

The client side implementation establishes a protocomm session with the device first, over the supported mode of transport, and then send and receive protobuf messages understood by the **esp_local_ctrl** service. The service translates these messages into requests and then call the appropriate handlers (set/get). Then, the generated response for each handler is again packed into a protobuf message and transmitted back to the client.

See below the various protobuf messages understood by the **esp_local_ctrl** service:

1. `get_prop_count` : This should simply return the total number of properties supported by the service.
2. `get_prop_values` : This accepts an array of indices and should return the information (name, type, flags) and values of the properties corresponding to those indices.
3. `set_prop_values` : This accepts an array of indices and an array of new values, which are used for setting the values of the properties corresponding to the indices.

Note that indices may or may not be the same for a property, across multiple sessions. Therefore, the client must only use the names of the properties to uniquely identify them. So, every time a new session is established, the client should first call `get_prop_count` and then `get_prop_values`, hence form an index-to-name mapping for all properties. Now when calling `set_prop_values` for a set of properties, it must first convert the names to indexes, using the created mapping. As emphasized earlier, the client must refresh the index-to-name mapping every time a new session is established with the same device.

The various protocomm endpoints provided by **esp_local_ctrl** are listed below:

Table 1: Endpoints provided by ESP Local Control

Endpoint Name (Blue-tooth Low Energy + GATT Server)	URI (HTTPS Server + mDNS)	Description
<code>esp_local_ctrl</code>	<code>https://<mdns-hostname>.local/esp_local_ctrl/version</code>	Endpoint used for retrieving version string
<code>esp_local_ctrl</code>	<code>https://<mdns-hostname>.local/esp_local_ctrl/control</code>	Endpoint used for sending or receiving control messages

API Reference

Header File

- [components/esp_local_ctrl/include/esp_local_ctrl.h](#)
- This header file can be included with:

```
#include "esp_local_ctrl.h"
```

- This header file is a part of the API provided by the `esp_local_ctrl` component. To declare that your component depends on `esp_local_ctrl`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_local_ctrl
```

or

`PRIV_REQUIRES esp_local_ctrl`

Functions

`const esp_local_ctrl_transport_t *esp_local_ctrl_get_transport_ble` (void)

Function for obtaining BLE transport mode.

`const esp_local_ctrl_transport_t *esp_local_ctrl_get_transport_httpd` (void)

Function for obtaining HTTPD transport mode.

`esp_err_t esp_local_ctrl_start` (const *esp_local_ctrl_config_t* *config)

Start local control service.

Parameters **config** -- [in] Pointer to configuration structure

Returns

- ESP_OK : Success
- ESP_FAIL : Failure

`esp_err_t esp_local_ctrl_stop` (void)

Stop local control service.

`esp_err_t esp_local_ctrl_add_property` (const *esp_local_ctrl_prop_t* *prop)

Add a new property.

This adds a new property and allocates internal resources for it. The total number of properties that could be added is limited by configuration option `max_properties`

Parameters **prop** -- [in] Property description structure

Returns

- ESP_OK : Success
- ESP_FAIL : Failure

`esp_err_t esp_local_ctrl_remove_property` (const char *name)

Remove a property.

This finds a property by name, and releases the internal resources which are associated with it.

Parameters **name** -- [in] Name of the property to remove

Returns

- ESP_OK : Success
- ESP_ERR_NOT_FOUND : Failure

`const esp_local_ctrl_prop_t *esp_local_ctrl_get_property` (const char *name)

Get property description structure by name.

This API may be used to get a property's context structure `esp_local_ctrl_prop_t` when its name is known

Parameters **name** -- [in] Name of the property to find

Returns

- Pointer to property
- NULL if not found

`esp_err_t esp_local_ctrl_set_handler` (const char *ep_name, *protocomm_req_handler_t* handler, void *user_ctx)

Register protocomm handler for a custom endpoint.

This API can be called by the application to register a protocomm handler for an endpoint after the local control service has started.

Note: In case of BLE transport the names and uuids of all custom endpoints must be provided beforehand as a part of the `protocomm_ble_config_t` structure set in `esp_local_ctrl_config_t`, and passed

to `esp_local_ctrl_start()`.

Parameters

- **ep_name** -- [in] Name of the endpoint
- **handler** -- [in] Endpoint handler function
- **user_ctx** -- [in] User data

Returns

- ESP_OK : Success
- ESP_FAIL : Failure

Unions

union **esp_local_ctrl_transport_config_t**

#include <esp_local_ctrl.h> Transport mode (BLE / HTTPD) configuration.

Public Members

esp_local_ctrl_transport_config_ble_t *ble

This is same as `protocomm_ble_config_t`. See `protocomm_ble.h` for available configuration parameters.

esp_local_ctrl_transport_config_httpd_t *httpd

This is same as `httpd_ssl_config_t`. See `esp_https_server.h` for available configuration parameters.

Structures

struct **esp_local_ctrl_prop**

Property description data structure, which is to be populated and passed to the `esp_local_ctrl_add_property()` function.

Once a property is added, its structure is available for read-only access inside `get_prop_values()` and `set_prop_values()` handlers.

Public Members

char ***name**

Unique name of property

uint32_t **type**

Type of property. This may be set to application defined enums

size_t **size**

Size of the property value, which:

- if zero, the property can have values of variable size
- if non-zero, the property can have values of fixed size only, therefore, checks are performed internally by `esp_local_ctrl` when setting the value of such a property

uint32_t flags

Flags set for this property. This could be a bit field. A flag may indicate property behavior, e.g. read-only / constant

void *ctx

Pointer to some context data relevant for this property. This will be available for use inside the `get_prop_values` and `set_prop_values` handlers as a part of this property structure. When set, this is valid throughout the lifetime of a property, till either the property is removed or the `esp_local_ctrl` service is stopped.

void (*ctx_free_fn)(void *ctx)

Function used by `esp_local_ctrl` to internally free the property context when `esp_local_ctrl_remove_property()` or `esp_local_ctrl_stop()` is called.

struct esp_local_ctrl_prop_val

Property value data structure. This gets passed to the `get_prop_values()` and `set_prop_values()` handlers for the purpose of retrieving or setting the present value of a property.

Public Members

void *data

Pointer to memory holding property value

size_t size

Size of property value

void (*free_fn)(void *data)

This may be set by the application in `get_prop_values()` handler to tell `esp_local_ctrl` to call this function on the data pointer above, for freeing its resources after sending the `get_prop_values` response.

struct esp_local_ctrl_handlers

Handlers for receiving and responding to local control commands for getting and setting properties.

Public Members

esp_err_t (***get_prop_values**)(size_t props_count, const *esp_local_ctrl_prop_t* props[], *esp_local_ctrl_prop_val_t* prop_values[], void *usr_ctx)

Handler function to be implemented for retrieving current values of properties.

Note: If any of the properties have fixed sizes, the size field of corresponding element in `prop_values` need to be set

Param props_count [in] Total elements in the props array

Param props [in] Array of properties, the current values for which have been requested by the client

Param prop_values [out] Array of empty property values, the elements of which need to be populated with the current values of those properties specified by props argument

Param usr_ctx [in] This provides value of the `usr_ctx` field of `esp_local_ctrl_handlers_t` structure

Return Returning different error codes will convey the corresponding protocol level errors to the client :

- ESP_OK : Success
- ESP_ERR_INVALID_ARG : InvalidArgument
- ESP_ERR_INVALID_STATE : InvalidProto
- All other error codes : InternalError

esp_err_t (***set_prop_values**)(size_t props_count, const *esp_local_ctrl_prop_t* props[], const *esp_local_ctrl_prop_val_t* prop_values[], void *usr_ctx)

Handler function to be implemented for changing values of properties.

Note: If any of the properties have variable sizes, the size field of the corresponding element in `prop_values` must be checked explicitly before making any assumptions on the size.

Param props_count [in] Total elements in the props array

Param props [in] Array of properties, the values for which the client requests to change

Param prop_values [in] Array of property values, the elements of which need to be used for updating those properties specified by props argument

Param usr_ctx [in] This provides value of the `usr_ctx` field of `esp_local_ctrl_handlers_t` structure

Return Returning different error codes will convey the corresponding protocol level errors to the client :

- ESP_OK : Success
- ESP_ERR_INVALID_ARG : InvalidArgument
- ESP_ERR_INVALID_STATE : InvalidProto
- All other error codes : InternalError

void ***usr_ctx**

Context pointer to be passed to above handler functions upon invocation. This is different from the property level context, as this is valid throughout the lifetime of the `esp_local_ctrl` service, and freed only when the service is stopped.

void (***usr_ctx_free_fn**)(void *usr_ctx)

Pointer to function which will be internally invoked on `usr_ctx` for freeing the context resources when `esp_local_ctrl_stop()` is called.

struct **esp_local_ctrl_proto_sec_cfg**

Protocom security configs

Public Members

esp_local_ctrl_proto_sec_t **version**

This sets protocom security version, sec0/sec1 or custom. If custom, user must provide handle via `proto_sec_custom_handle` below

void ***custom_handle**

Custom security handle if security is set custom via `proto_sec` above. This handle must follow `protocomm_security_t` signature

const void ***pop**

Proof of possession to be used for local control. Could be NULL.

const void ***sec_params**

Pointer to security params (NULL if not needed). This is not needed for protocomm security 0. This pointer should hold the struct of type `esp_local_ctrl_security1_params_t` for protocomm security 1 and `esp_local_ctrl_security2_params_t` for protocomm security 2 respectively. Could be NULL.

struct **esp_local_ctrl_config**

Configuration structure to pass to `esp_local_ctrl_start()`

Public Members

const *esp_local_ctrl_transport_t* ***transport**

Transport layer over which service will be provided

esp_local_ctrl_transport_config_t **transport_config**

Transport layer over which service will be provided

esp_local_ctrl_proto_sec_cfg_t **proto_sec**

Security version and POP

esp_local_ctrl_handlers_t **handlers**

Register handlers for responding to get/set requests on properties

size_t **max_properties**

This limits the number of properties that are available at a time

Macros

ESP_LOCAL_CTRL_TRANSPORT_BLE

ESP_LOCAL_CTRL_TRANSPORT_HTTPD

Type Definitions

typedef struct *esp_local_ctrl_prop* **esp_local_ctrl_prop_t**

Property description data structure, which is to be populated and passed to the `esp_local_ctrl_add_property()` function.

Once a property is added, its structure is available for read-only access inside `get_prop_values()` and `set_prop_values()` handlers.

typedef struct *esp_local_ctrl_prop_val* **esp_local_ctrl_prop_val_t**

Property value data structure. This gets passed to the `get_prop_values()` and `set_prop_values()` handlers for the purpose of retrieving or setting the present value of a property.

typedef struct *esp_local_ctrl_handlers* **esp_local_ctrl_handlers_t**

Handlers for receiving and responding to local control commands for getting and setting properties.

typedef struct *esp_local_ctrl_transport* **esp_local_ctrl_transport_t**

Transport mode (BLE / HTTPD) over which the service will be provided.

This is forward declaration of a private structure, implemented internally by `esp_local_ctrl`.

typedef struct *protocomm_ble_config* **esp_local_ctrl_transport_config_ble_t**

Configuration for transport mode BLE.

This is a forward declaration for `protocomm_ble_config_t`. To use this, application must set `CONFIG_BT_BLUEDROID_ENABLED` and include `protocomm_ble.h`.

typedef struct *httpd_config* **esp_local_ctrl_transport_config_httpd_t**

Configuration for transport mode HTTPD.

This is a forward declaration for `httpd_ssl_config_t` (for HTTPS) or `httpd_config_t` (for HTTP)

typedef enum *esp_local_ctrl_proto_sec* **esp_local_ctrl_proto_sec_t**

Security types for `esp_local_control`.

typedef *protocomm_security1_params_t* **esp_local_ctrl_security1_params_t**

typedef *protocomm_security2_params_t* **esp_local_ctrl_security2_params_t**

typedef struct *esp_local_ctrl_proto_sec_cfg* **esp_local_ctrl_proto_sec_cfg_t**

Protocom security configs

typedef struct *esp_local_ctrl_config* **esp_local_ctrl_config_t**

Configuration structure to pass to `esp_local_ctrl_start()`

Enumerations

enum **esp_local_ctrl_proto_sec**

Security types for `esp_local_control`.

Values:

enumerator **PROTOCOLCOM_SEC0**

enumerator **PROTOCOLCOM_SEC1**

enumerator **PROTOCOLCOM_SEC2**

enumerator **PROTOCOLCOM_SEC_CUSTOM**

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the <code>add_warnings_pages</code> list.

2.2.7 ESP Serial Slave Link

Overview

Espressif provides several chips that can work as slaves. These slave devices rely on some common buses, and have their own communication protocols over those buses. The `esp_serial_slave_link` component is designed for the master to communicate with ESP slave devices through those protocols over the bus drivers.

After an `esp_serial_slave_link` device is initialized properly, the application can use it to communicate with the ESP slave devices conveniently.

Note: The ESP-IDF component `esp_serial_slave_link` has been moved from ESP-IDF since version v5.0 to a separate repository:

- [ESSL component on GitHub](#)

To add ESSL component in your project, please run `idf.py add-dependency espressif/esp_serial_slave_link`.

Espressif Device Protocols

For more details about Espressif device protocols, see the following documents.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

ESP SPI Slave HD (Half Duplex) Mode Protocol

SPI Slave Capabilities of Espressif Chips

	ESP32	ESP32-S2	ESP32-C3	ESP32-S3	ESP32-C2	ESP32-C6	ESP32-H2
SPI Slave HD	N	Y (v2)	Y (v2)	Y (v2)	Y (v2)	Y (v2)	Y (v2)
Tohost intr		N	N	N	N	N	N
Frhost intr		2 *	2 *	2 *	2 *	2 *	2 *
TX DMA		Y	Y	Y	Y	Y	Y
RX DMA		Y	Y	Y	Y	Y	Y
Shared registers		72	64	64	64	64	64

Introduction In the half duplex mode, the master has to use the protocol defined by the slave to communicate with the slave. Each transaction may consist of the following phases (listed by the order they should exist):

- Command: 8-bit, master to slave
This phase determines the rest phases of the transactions. See [Supported Commands](#).
- Address: 8-bit, master to slave, optional
For some commands (WRBUF, RDBUF), this phase specifies the address of the shared register to write to/read from. For other commands with this phase, they are meaningless but still have to exist in the transaction.
- Dummy: 8-bit, floating, optional
This phase is the turnaround time between the master and the slave on the bus, and also provides enough time for the slave to prepare the data to send to the master.
- Data: variable length, the direction is also determined by the command.
This may be a data OUT phase, in which the direction is slave to master, or a data IN phase, in which the direction is master to slave.

The **direction** means which side (master or slave) controls the MOSI, MISO, WP, and HD pins.

Data IO Modes In some IO modes, more data wires can be used to send the data. As a result, the SPI clock cycles required for the same amount of data will be less than in the 1-bit mode. For example, in QIO mode, address and data (IN and OUT) should be sent on all 4 data wires (MOSI, MISO, WP, and HD). Here are the modes supported by the ESP32-S2 SPI slave and the wire number (WN) used in corresponding modes.

Mode	Command WN	Address WN	Dummy cycles	Data WN
1-bit	1	1	1	1
DOUT	1	1	4	2
DIO	1	2	4	2
QOUT	1	1	4	4
QIO	1	4	4	4
QPI	4	4	4	4

Normally, which mode is used is determined by the command sent by the master (See [Supported Commands](#)), except the QPI mode.

QPI Mode The QPI mode is a special state of the SPI Slave. The master can send the ENQPI command to put the slave into the QPI mode state. In the QPI mode, the command is also sent in 4-bit, thus it is not compatible with the normal modes. The master should only send QPI commands when the slave is in QPI mode. To exit from the QPI mode, master can send the EXQPI command.

Supported Commands

Note: The command name is in a master-oriented direction. For example, WRBUF means master writes the buffer of slave.

Name	Description	Command	Address	Data
WRBUF	Write buffer	0x01	Buf addr	master to slave, no longer than buffer size
RDBUF	Read buffer	0x02	Buf addr	slave to master, no longer than buffer size
WRDMA	Write DMA	0x03	8 bits	master to slave, no longer than length provided by slave
RDDMA	Read DMA	0x04	8 bits	slave to master, no longer than length provided by slave
SEG_DONE	Segments done	0x05	•	•
ENQPI	Enter QPI mode	0x06	•	•
WR_DONE	Write segments done	0x07	•	•
CMD8	Interrupt	0x08	•	•
CMD9	Interrupt	0x09	•	•
CMDA	Interrupt	0x0A	•	•
EXQPI	Exit QPI mode	0xDD	•	•

Moreover, WRBUF, RDBUF, WRDMA, and RDDMA commands have their 2-bit and 4-bit version. To do transactions in 2-bit or 4-bit mode, send the original command ORed by the corresponding command mask below. For example, command 0xA1 means WRBUF in QIO mode.

Mode	Mask
1-bit	0x00
DOUT	0x10
DIO	0x50
QOUT	0x20
QIO	0xA0
QPI	0xA0

Segment Transaction Mode Segment transaction mode is the only mode supported by the SPI Slave HD driver for now. In this mode, for a transaction the slave loads onto the DMA, the master is allowed to read or write in segments. In this way, the master does not have to prepare a large buffer as the size of data provided by the slave. After the master finishes reading/writing a buffer, it has to send the corresponding termination command to the slave as a synchronization signal. The slave driver will update new data (if exist) onto the DMA upon seeing the termination command.

The termination command is WR_DONE (0x07) for WRDMA and CMD8 (0x08) for RDDMA.

Here is an example for the flow the master read data from the slave DMA:

1. The slave loads 4092 bytes of data onto the RDDMA.
2. The master do seven RDDMA transactions, each of them is 512 bytes long, and reads the first 3584 bytes from the slave.
3. The master do the last RDDMA transaction of 512 bytes (equal, longer, or shorter than the total length loaded by the slave are all allowed). The first 508 bytes are valid data from the slave, while the last 4 bytes are meaningless bytes.
4. The master sends CMD8 to the slave.
5. The slave loads another 4092 bytes of data onto the RDDMA.
6. The master can start new reading transactions after it sends the CMD8.

Terminology

- ESSL: Abbreviation for ESP Serial Slave Link, the component described by this document.
- Master: The device running the `esp_serial_slave_link` component.
- ESSL device: a virtual device on the master associated with an ESP slave device. The device context has the knowledge of the slave protocol above the bus, relying on some bus drivers to communicate with the slave.
- ESSL device handle: a handle to ESSL device context containing the configuration, status and data required by the ESSL component. The context stores the driver configurations, communication state, data shared by master and slave, etc.
 - The context should be initialized before it is used, and get deinitialized if not used any more. The master application operates on the ESSL device through this handle.
- ESP slave: the slave device connected to the bus, which ESSL component is designed to communicate with.
- Bus: The bus over which the master and the slave communicate with each other.
- Slave protocol: The special communication protocol specified by Espressif HW/SW over the bus.
- TX buffer num: a counter, which is on the slave and can be read by the master, indicates the accumulated buffer numbers that the slave has loaded to the hardware to receive data from the master.
- RX data size: a counter, which is on the slave and can be read by the master, indicates the accumulated data size that the slave has loaded to the hardware to send to the master.

Services Provided by ESP Slave

There are some common services provided by the Espressif slaves:

1. Tohost Interrupts: The slave can inform the master about certain events by the interrupt line. (optional)

2. Frhost Interrupts: The master can inform the slave about certain events.
3. TX FIFO (master to slave): The slave can receive data from the master in units of receiving buffers. The slave updates the TX buffer num to inform the master how much data it can receive, and the master then read the TX buffer num, and take off the used buffer number to know how many buffers are remaining.
4. RX FIFO (slave to master): The slave can send data in stream to the master. The SDIO slave can also indicate it has new data to send to master by the interrupt line. The slave updates the RX data size to inform the master how much data it has prepared to send, and then the master read the data size, and take off the data length it has already received to know how many data is remaining.
5. Shared registers: The master can read some part of the registers on the slave, and also write these registers to let the slave read.

The services provided by the slave depends on the slave's model. See [SPI Slave Capabilities of Espressif Chips](#) for more details.

Initialization of ESP Serial Slave Link

ESP SDIO Slave The ESP SDIO slave link (ESSL SDIO) devices relies on the SDMMC component. It includes the usage of communicating with ESP SDIO Slave device via the SDMMC Host or SDSPI Host feature. The ESSL device should be initialized as below:

1. Initialize a SDMMC card (see `:doc:` Document of SDMMC driver </api-reference/storage/sdmmc>``) structure.
2. Call `sdmmc_card_init()` to initialize the card.
3. Initialize the ESSL device with `essl_sdio_config_t`. The `card` member should be the `sdmmc_card_t` got in step 2, and the `recv_buffer_size` member should be filled correctly according to pre-negotiated value.
4. Call `essl_init()` to do initialization of the SDIO part.
5. Call `essl_wait_for_ready()` to wait for the slave to be ready.

ESP SPI Slave

Note: If you are communicating with the ESP SDIO Slave device through SPI interface, you should use the [SDIO interface](#) instead.

Has not been supported yet.

APIs

After the initialization process above is performed, you can call the APIs below to make use of the services provided by the slave:

Tohost Interrupts (Optional)

1. Call `essl_get_intr_ena()` to know which events trigger the interrupts to the master.
2. Call `essl_set_intr_ena()` to set the events that trigger interrupts to the master.
3. Call `essl_wait_int()` to wait until interrupt from the slave, or timeout.
4. When interrupt is triggered, call `essl_get_intr()` to know which events are active, and call `essl_clear_intr()` to clear them.

Frhost Interrupts

1. Call `essl_send_slave_intr()` to trigger general purpose interrupt of the slave.

TX FIFO

1. Call `essl_get_tx_buffer_num()` to know how many buffers the slave has prepared to receive data from the master. This is optional. The master will poll `tx_buffer_num` when it tries to send packets to the slave, until the slave has enough buffer or timeout.
2. Call `essl_send_packet()` to send data to the slave.

RX FIFO

1. Call `essl_get_rx_data_size()` to know how many data the slave has prepared to send to the master. This is optional. When the master tries to receive data from the slave, it updates the `rx_data_size` for once, if the current `rx_data_size` is shorter than the buffer size the master prepared to receive. And it may poll the `rx_data_size` if the `rx_data_size` keeps 0, until timeout.
2. Call `essl_get_packet()` to receive data from the slave.

Reset Counters (Optional) Call `essl_reset_cnt()` to reset the internal counter if you find the slave has reset its counter.

Application Example

The example below shows how ESP32-P4 SDIO host and slave communicate with each other. The host uses the ESSL SDIO:

[peripherals/sdio](#)

Please refer to the specific example README.md for details.

API Reference

Header File

- [components/driver/test_apps/components/esp_serial_slave_link/include/esp_serial_slave_link/essl.h](#)

Functions

`esp_err_t` **essl_init** (`essl_handle_t` handle, `uint32_t` wait_ms)

Initialize the slave.

Parameters

- **handle** -- Handle of an ESSL device.
- **wait_ms** -- Millisecond to wait before timeout, will not wait at all if set to 0-9.

Returns

- ESP_OK: If success
- ESP_ERR_NOT_SUPPORTED: Current device does not support this function.
- Other value returned from lower layer `init`.

`esp_err_t` **essl_wait_for_ready** (`essl_handle_t` handle, `uint32_t` wait_ms)

Wait for interrupt of an ESSL slave device.

Parameters

- **handle** -- Handle of an ESSL device.
- **wait_ms** -- Millisecond to wait before timeout, will not wait at all if set to 0-9.

Returns

- ESP_OK: If success
- ESP_ERR_NOT_SUPPORTED: Current device does not support this function.
- One of the error codes from SDMMC host controller

esp_err_t **essl_get_tx_buffer_num** (*essl_handle_t* handle, uint32_t *out_tx_num, uint32_t wait_ms)

Get buffer num for the host to send data to the slave. The buffers are size of `buffer_size`.

Parameters

- **handle** -- Handle of a ESSL device.
- **out_tx_num** -- Output of buffer num that host can send data to ESSL slave.
- **wait_ms** -- Millisecond to wait before timeout, will not wait at all if set to 0-9.

Returns

- ESP_OK: Success
- ESP_ERR_NOT_SUPPORTED: This API is not supported in this mode
- One of the error codes from SDMMC/SPI host controller

esp_err_t **essl_get_rx_data_size** (*essl_handle_t* handle, uint32_t *out_rx_size, uint32_t wait_ms)

Get the size, in bytes, of the data that the ESSL slave is ready to send

Parameters

- **handle** -- Handle of an ESSL device.
- **out_rx_size** -- Output of data size to read from slave, in bytes
- **wait_ms** -- Millisecond to wait before timeout, will not wait at all if set to 0-9.

Returns

- ESP_OK: Success
- ESP_ERR_NOT_SUPPORTED: This API is not supported in this mode
- One of the error codes from SDMMC/SPI host controller

esp_err_t **essl_reset_cnt** (*essl_handle_t* handle)

Reset the counters of this component. Usually you don't need to do this unless you know the slave is reset.

Parameters **handle** -- Handle of an ESSL device.

Returns

- ESP_OK: Success
- ESP_ERR_NOT_SUPPORTED: This API is not supported in this mode
- ESP_ERR_INVALID_ARG: Invalid argument, handle is not init.

esp_err_t **essl_send_packet** (*essl_handle_t* handle, const void *start, size_t length, uint32_t wait_ms)

Send a packet to the ESSL Slave. The Slave receives the packet into buffers whose size is `buffer_size` (configured during initialization).

Parameters

- **handle** -- Handle of an ESSL device.
- **start** -- Start address of the packet to send
- **length** -- Length of data to send, if the packet is over-size, the it will be divided into blocks and hold into different buffers automatically.
- **wait_ms** -- Millisecond to wait before timeout, will not wait at all if set to 0-9.

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG: Invalid argument, handle is not init or other argument is not valid.
- ESP_ERR_TIMEOUT: No buffer to use, or error from SDMMC host controller.
- ESP_ERR_NOT_FOUND: Slave is not ready for receiving.
- ESP_ERR_NOT_SUPPORTED: This API is not supported in this mode
- One of the error codes from SDMMC/SPI host controller.

esp_err_t **essl_get_packet** (*essl_handle_t* handle, void *out_data, size_t size, size_t *out_length, uint32_t wait_ms)

Get a packet from ESSL slave.

Parameters

- **handle** -- Handle of an ESSL device.
- **out_data** -- [out] Data output address
- **size** -- The size of the output buffer, if the buffer is smaller than the size of data to receive from slave, the driver returns ESP_ERR_NOT_FINISHED

- **out_length** -- [out] Output of length the data actually received from slave.
- **wait_ms** -- Millisecond to wait before timeout, will not wait at all if set to 0-9.

Returns

- ESP_OK Success: All the data has been read from the slave.
- ESP_ERR_INVALID_ARG: Invalid argument, The handle is not initialized or the other arguments are invalid.
- ESP_ERR_NOT_FINISHED: Read was successful, but there is still data remaining.
- ESP_ERR_NOT_FOUND: Slave is not ready to send data.
- ESP_ERR_NOT_SUPPORTED: This API is not supported in this mode
- One of the error codes from SDMMC/SPI host controller.

esp_err_t **essl_write_reg** (*essl_handle_t* handle, uint8_t addr, uint8_t value, uint8_t *value_o, uint32_t wait_ms)

Write general purpose R/W registers (8-bit) of ESSL slave.

Note: sdio 28-31 are reserved, the lower API helps to skip.

Parameters

- **handle** -- Handle of an ESSL device.
- **addr** -- Address of register to write. For SDIO, valid address: 0-59. For SPI, see `essl_spi.h`
- **value** -- Value to write to the register.
- **value_o** -- Output of the returned written value.
- **wait_ms** -- Millisecond to wait before timeout, will not wait at all if set to 0-9.

Returns

- ESP_OK Success
- One of the error codes from SDMMC/SPI host controller

esp_err_t **essl_read_reg** (*essl_handle_t* handle, uint8_t addr, uint8_t *value_o, uint32_t wait_ms)

Read general purpose R/W registers (8-bit) of ESSL slave.

Parameters

- **handle** -- Handle of a `essl` device.
- **addr** -- Address of register to read. For SDIO, Valid address: 0-27, 32-63 (28-31 reserved, return interrupt bits on read). For SPI, see `essl_spi.h`
- **value_o** -- Output value read from the register.
- **wait_ms** -- Millisecond to wait before timeout, will not wait at all if set to 0-9.

Returns

- ESP_OK Success
- One of the error codes from SDMMC/SPI host controller

esp_err_t **essl_wait_int** (*essl_handle_t* handle, uint32_t wait_ms)

wait for an interrupt of the slave

Parameters

- **handle** -- Handle of an ESSL device.
- **wait_ms** -- Millisecond to wait before timeout, will not wait at all if set to 0-9.

Returns

- ESP_OK: If interrupt is triggered.
- ESP_ERR_NOT_SUPPORTED: Current device does not support this function.
- ESP_ERR_TIMEOUT: No interrupts before timeout.

esp_err_t **essl_clear_intr** (*essl_handle_t* handle, uint32_t intr_mask, uint32_t wait_ms)

Clear interrupt bits of ESSL slave. All the bits set in the mask will be cleared, while other bits will stay the same.

Parameters

- **handle** -- Handle of an ESSL device.
- **intr_mask** -- Mask of interrupt bits to clear.

- **wait_ms** -- Millisecond to wait before timeout, will not wait at all if set to 0-9.

Returns

- ESP_OK: Success
- ESP_ERR_NOT_SUPPORTED: Current device does not support this function.
- One of the error codes from SDMMC host controller

esp_err_t **essl_get_intr** (*essl_handle_t* handle, uint32_t *intr_raw, uint32_t *intr_st, uint32_t wait_ms)

Get interrupt bits of ESSL slave.

Parameters

- **handle** -- Handle of an ESSL device.
- **intr_raw** -- Output of the raw interrupt bits. Set to NULL if only masked bits are read.
- **intr_st** -- Output of the masked interrupt bits. set to NULL if only raw bits are read.
- **wait_ms** -- Millisecond to wait before timeout, will not wait at all if set to 0-9.

Returns

- ESP_OK: Success
- ESP_INVALID_ARG: If both *intr_raw* and *intr_st* are NULL.
- ESP_ERR_NOT_SUPPORTED: Current device does not support this function.
- One of the error codes from SDMMC host controller

esp_err_t **essl_set_intr_ena** (*essl_handle_t* handle, uint32_t ena_mask, uint32_t wait_ms)

Set interrupt enable bits of ESSL slave. The slave only sends interrupt on the line when there is a bit both the raw status and the enable are set.

Parameters

- **handle** -- Handle of an ESSL device.
- **ena_mask** -- Mask of the interrupt bits to enable.
- **wait_ms** -- Millisecond to wait before timeout, will not wait at all if set to 0-9.

Returns

- ESP_OK: Success
- ESP_ERR_NOT_SUPPORTED: Current device does not support this function.
- One of the error codes from SDMMC host controller

esp_err_t **essl_get_intr_ena** (*essl_handle_t* handle, uint32_t *ena_mask_o, uint32_t wait_ms)

Get interrupt enable bits of ESSL slave.

Parameters

- **handle** -- Handle of an ESSL device.
- **ena_mask_o** -- Output of interrupt bit enable mask.
- **wait_ms** -- Millisecond to wait before timeout, will not wait at all if set to 0-9.

Returns

- ESP_OK Success
- One of the error codes from SDMMC host controller

esp_err_t **essl_send_slave_intr** (*essl_handle_t* handle, uint32_t intr_mask, uint32_t wait_ms)

Send interrupts to slave. Each bit of the interrupt will be triggered.

Parameters

- **handle** -- Handle of an ESSL device.
- **intr_mask** -- Mask of interrupt bits to send to slave.
- **wait_ms** -- Millisecond to wait before timeout, will not wait at all if set to 0-9.

Returns

- ESP_OK: Success
- ESP_ERR_NOT_SUPPORTED: Current device does not support this function.
- One of the error codes from SDMMC host controller

Type Definitions

```
typedef struct essl_dev_t *essl_handle_t
```

Handle of an ESSL device.

Header File

- [components/driver/test_apps/components/esp_serial_slave_link/include/esp_serial_slave_link/essl_sdio.h](#)

Functions

esp_err_t **essl_sdio_init_dev** (*essl_handle_t* *out_handle, const *essl_sdio_config_t* *config)

Initialize the ESSL SDIO device and get its handle.

Parameters

- **out_handle** -- Output of the handle.
- **config** -- Configuration for the ESSL SDIO device.

Returns

- ESP_OK: on success
- ESP_ERR_NO_MEM: memory exhausted.

esp_err_t **essl_sdio_deinit_dev** (*essl_handle_t* handle)

Deinitialize and free the space used by the ESSL SDIO device.

Parameters **handle** -- Handle of the ESSL SDIO device to deinit.

Returns

- ESP_OK: on success
- ESP_ERR_INVALID_ARG: wrong handle passed

Structures

struct **essl_sdio_config_t**

Configuration for the ESSL SDIO device.

Public Members

sdmmc_card_t ***card**

The initialized sdmmc card pointer of the slave.

int **recv_buffer_size**

The pre-negotiated recv buffer size used by both the host and the slave.

Header File

- [components/driver/test_apps/components/esp_serial_slave_link/include/esp_serial_slave_link/essl_spi.h](#)

Functions

esp_err_t **essl_spi_init_dev** (*essl_handle_t* *out_handle, const *essl_spi_config_t* *init_config)

Initialize the ESSL SPI device function list and get its handle.

Parameters

- **out_handle** -- [out] Output of the handle
- **init_config** -- Configuration for the ESSL SPI device

Returns

- ESP_OK: On success
- ESP_ERR_NO_MEM: Memory exhausted
- ESP_ERR_INVALID_STATE: SPI driver is not initialized
- ESP_ERR_INVALID_ARG: Wrong register ID

esp_err_t **essl_spi_deinit_dev** (*essl_handle_t* handle)

Deinitialize the ESSL SPI device and free the memory used by the device.

Parameters **handle** -- Handle of the ESSL SPI device

Returns

- `ESP_OK`: On success
- `ESP_ERR_INVALID_STATE`: ESSL SPI is not in use

esp_err_t **essl_spi_read_reg** (void *arg, uint8_t addr, uint8_t *out_value, uint32_t wait_ms)

Read from the shared registers.

Note: The registers for Master/Slave synchronization are reserved. Do not use them. (see `rx_sync_reg` in `essl_spi_config_t`)

Parameters

- **arg** -- Context of the component. (Member `arg` from `essl_handle_t`)
- **addr** -- Address of the shared registers. (Valid: 0 ~ `SOC_SPI_MAXIMUM_BUFFER_SIZE`, registers for M/S sync are reserved, see note1).
- **out_value** -- [out] Read buffer for the shared registers.
- **wait_ms** -- Time to wait before timeout (reserved for future use, user should set this to 0).

Returns

- `ESP_OK`: success
- `ESP_ERR_INVALID_STATE`: ESSL SPI has not been initialized.
- `ESP_ERR_INVALID_ARG`: The address argument is not valid. See note 1.
- or other return value from `:cpp:func:spi_device_transmit`.

esp_err_t **essl_spi_get_packet** (void *arg, void *out_data, size_t size, uint32_t wait_ms)

Get a packet from Slave.

Parameters

- **arg** -- Context of the component. (Member `arg` from `essl_handle_t`)
- **out_data** -- [out] Output data address
- **size** -- The size of the output data.
- **wait_ms** -- Time to wait before timeout (reserved for future use, user should set this to 0).

Returns

- `ESP_OK`: On Success
- `ESP_ERR_INVALID_STATE`: ESSL SPI has not been initialized.
- `ESP_ERR_INVALID_ARG`: The output data address is neither DMA capable nor 4 byte-aligned
- `ESP_ERR_INVALID_SIZE`: Master requires `size` bytes of data but Slave did not load enough bytes.

esp_err_t **essl_spi_write_reg** (void *arg, uint8_t addr, uint8_t value, uint8_t *out_value, uint32_t wait_ms)

Write to the shared registers.

Note: The registers for Master/Slave synchronization are reserved. Do not use them. (see `tx_sync_reg` in `essl_spi_config_t`)

Note: Feature of checking the actual written value (`out_value`) is not supported.

Parameters

- **arg** -- Context of the component. (Member `arg` from `essl_handle_t`)

- **addr** -- Address of the shared registers. (Valid: 0 ~ SOC_SPI_MAXIMUM_BUFFER_SIZE, registers for M/S sync are reserved, see note1)
- **value** -- Buffer for data to send, should be align to 4.
- **out_value** -- [out] Not supported, should be set to NULL.
- **wait_ms** -- Time to wait before timeout (reserved for future use, user should set this to 0).

Returns

- ESP_OK: success
- ESP_ERR_INVALID_STATE: ESSL SPI has not been initialized.
- ESP_ERR_INVALID_ARG: The address argument is not valid. See note 1.
- ESP_ERR_NOT_SUPPORTED: Should set `out_value` to NULL. See note 2.
- or other return value from `:cpp:func:spi_device_transmit`.

esp_err_t **essl_spi_send_packet** (void *arg, const void *data, size_t size, uint32_t wait_ms)

Send a packet to Slave.

Parameters

- **arg** -- Context of the component. (Member `arg` from `essl_handle_t`)
- **data** -- Address of the data to send
- **size** -- Size of the data to send.
- **wait_ms** -- Time to wait before timeout (reserved for future use, user should set this to 0).

Returns

- ESP_OK: On success
- ESP_ERR_INVALID_STATE: ESSL SPI has not been initialized.
- ESP_ERR_INVALID_ARG: The data address is not DMA capable
- ESP_ERR_INVALID_SIZE: Master will send `size` bytes of data but Slave did not load enough RX buffer

void **essl_spi_reset_cnt** (void *arg)

Reset the counter in Master context.

Note: Shall only be called if the slave has reset its counter. Else, Slave and Master would be desynchronized

Parameters **arg** -- Context of the component. (Member `arg` from `essl_handle_t`)

esp_err_t **essl_spi_rdbuf** (*spi_device_handle_t* spi, uint8_t *out_data, int addr, int len, uint32_t flags)

Read the shared buffer from the slave in ISR way.

Note: The slave's HW doesn't guarantee the data in one SPI transaction is consistent. It sends data in unit of byte. In other words, if the slave SW attempts to update the shared register when a `rdbuf` SPI transaction is in-flight, the data got by the master will be the combination of bytes of different writes of slave SW.

Note: `out_data` should be prepared in words and in the DRAM. The buffer may be written in words by the DMA. When a byte is written, the remaining bytes in the same word will also be overwritten, even the `len` is shorter than a word.

Parameters

- **spi** -- SPI device handle representing the slave
- **out_data** -- [out] Buffer for read data, strongly suggested to be in the DRAM and aligned to 4
- **addr** -- Address of the slave shared buffer
- **len** -- Length to read

- **flags** -- SPI_TRANS_* flags to control the transaction mode of the transaction to send.

Returns

- ESP_OK: on success
- or other return value from :cpp:func:spi_device_transmit.

esp_err_t **essl_spi_rdbuf_polling** (*spi_device_handle_t* spi, uint8_t *out_data, int addr, int len, uint32_t flags)

Read the shared buffer from the slave in polling way.

Note: out_data should be prepared in words and in the DRAM. The buffer may be written in words by the DMA. When a byte is written, the remaining bytes in the same word will also be overwritten, even the len is shorter than a word.

Parameters

- **spi** -- SPI device handle representing the slave
- **out_data** -- [out] Buffer for read data, strongly suggested to be in the DRAM and aligned to 4
- **addr** -- Address of the slave shared buffer
- **len** -- Length to read
- **flags** -- SPI_TRANS_* flags to control the transaction mode of the transaction to send.

Returns

- ESP_OK: on success
- or other return value from :cpp:func:spi_device_transmit.

esp_err_t **essl_spi_wrbuf** (*spi_device_handle_t* spi, const uint8_t *data, int addr, int len, uint32_t flags)

Write the shared buffer of the slave in ISR way.

Note: out_data should be prepared in words and in the DRAM. The buffer may be written in words by the DMA. When a byte is written, the remaining bytes in the same word will also be overwritten, even the len is shorter than a word.

Parameters

- **spi** -- SPI device handle representing the slave
- **data** -- Buffer for data to send, strongly suggested to be in the DRAM
- **addr** -- Address of the slave shared buffer,
- **len** -- Length to write
- **flags** -- SPI_TRANS_* flags to control the transaction mode of the transaction to send.

Returns

- ESP_OK: success
- or other return value from :cpp:func:spi_device_transmit.

esp_err_t **essl_spi_wrbuf_polling** (*spi_device_handle_t* spi, const uint8_t *data, int addr, int len, uint32_t flags)

Write the shared buffer of the slave in polling way.

Note: out_data should be prepared in words and in the DRAM. The buffer may be written in words by the DMA. When a byte is written, the remaining bytes in the same word will also be overwritten, even the len is shorter than a word.

Parameters

- **spi** -- SPI device handle representing the slave
- **data** -- Buffer for data to send, strongly suggested to be in the DRAM
- **addr** -- Address of the slave shared buffer,

- **len** -- Length to write
- **flags** -- SPI_TRANS_* flags to control the transaction mode of the transaction to send.

Returns

- ESP_OK: success
- or other return value from :cpp:func:spi_device_polling_transmit.

esp_err_t **essl_spi_rddma** (*spi_device_handle_t* spi, uint8_t *out_data, int len, int seg_len, uint32_t flags)

Receive long buffer in segments from the slave through its DMA.

Note: This function combines several :cpp:func:essl_spi_rddma_seg and one :cpp:func:essl_spi_rddma_done at the end. Used when the slave is working in segment mode.

Parameters

- **spi** -- SPI device handle representing the slave
- **out_data** -- [out] Buffer to hold the received data, strongly suggested to be in the DRAM and aligned to 4
- **len** -- Total length of data to receive.
- **seg_len** -- Length of each segment, which is not larger than the maximum transaction length allowed for the spi device. Suggested to be multiples of 4. When set < 0, means send all data in one segment (the rddma_done will still be sent.)
- **flags** -- SPI_TRANS_* flags to control the transaction mode of the transaction to send.

Returns

- ESP_OK: success
- or other return value from :cpp:func:spi_device_transmit.

esp_err_t **essl_spi_rddma_seg** (*spi_device_handle_t* spi, uint8_t *out_data, int seg_len, uint32_t flags)

Read one data segment from the slave through its DMA.

Note: To read long buffer, call :cpp:func:essl_spi_rddma instead.

Parameters

- **spi** -- SPI device handle representing the slave
- **out_data** -- [out] Buffer to hold the received data. strongly suggested to be in the DRAM and aligned to 4
- **seg_len** -- Length of this segment
- **flags** -- SPI_TRANS_* flags to control the transaction mode of the transaction to send.

Returns

- ESP_OK: success
- or other return value from :cpp:func:spi_device_transmit.

esp_err_t **essl_spi_rddma_done** (*spi_device_handle_t* spi, uint32_t flags)

Send the rddma_done command to the slave. Upon receiving this command, the slave will stop sending the current buffer even there are data unsent, and maybe prepare the next buffer to send.

Note: This is required only when the slave is working in segment mode.

Parameters

- **spi** -- SPI device handle representing the slave
- **flags** -- SPI_TRANS_* flags to control the transaction mode of the transaction to send.

Returns

- ESP_OK: success
- or other return value from :cpp:func:spi_device_transmit.

esp_err_t **essl_spi_wr dma** (*spi_device_handle_t* spi, const uint8_t *data, int len, int seg_len, uint32_t flags)

Send long buffer in segments to the slave through its DMA.

Note: This function combines several :cpp:func:essl_spi_wr dma_seg and one :cpp:func:essl_spi_wr dma_done at the end. Used when the slave is working in segment mode.

Parameters

- **spi** -- SPI device handle representing the slave
- **data** -- Buffer for data to send, strongly suggested to be in the DRAM
- **len** -- Total length of data to send.
- **seg_len** -- Length of each segment, which is not larger than the maximum transaction length allowed for the spi device. Suggested to be multiples of 4. When set < 0, means send all data in one segment (the wr dma_done will still be sent.)
- **flags** -- SPI_TRANS_* flags to control the transaction mode of the transaction to send.

Returns

- ESP_OK: success
- or other return value from :cpp:func:spi_device_transmit.

esp_err_t **essl_spi_wr dma_seg** (*spi_device_handle_t* spi, const uint8_t *data, int seg_len, uint32_t flags)

Send one data segment to the slave through its DMA.

Note: To send long buffer, call :cpp:func:essl_spi_wr dma instead.

Parameters

- **spi** -- SPI device handle representing the slave
- **data** -- Buffer for data to send, strongly suggested to be in the DRAM
- **seg_len** -- Length of this segment
- **flags** -- SPI_TRANS_* flags to control the transaction mode of the transaction to send.

Returns

- ESP_OK: success
- or other return value from :cpp:func:spi_device_transmit.

esp_err_t **essl_spi_wr dma_done** (*spi_device_handle_t* spi, uint32_t flags)

Send the wr dma_done command to the slave. Upon receiving this command, the slave will stop receiving, process the received data, and maybe prepare the next buffer to receive.

Note: This is required only when the slave is working in segment mode.

Parameters

- **spi** -- SPI device handle representing the slave
- **flags** -- SPI_TRANS_* flags to control the transaction mode of the transaction to send.

Returns

- ESP_OK: success
- or other return value from :cpp:func:spi_device_transmit.

Structures

struct **essl_spi_config_t**

Configuration of ESSL SPI device.

Public Members

spi_device_handle_t ***spi**

Pointer to SPI device handle.

uint32_t **tx_buf_size**

The pre-negotiated Master TX buffer size used by both the host and the slave.

uint8_t **tx_sync_reg**

The pre-negotiated register ID for Master-TX-SLAVE-RX synchronization. 1 word (4 Bytes) will be reserved for the synchronization.

uint8_t **rx_sync_reg**

The pre-negotiated register ID for Master-RX-Slave-TX synchronization. 1 word (4 Bytes) will be reserved for the synchronization.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.2.8 ESP x509 Certificate Bundle

Overview

The ESP x509 Certificate Bundle API provides an easy way to include a bundle of custom x509 root certificates for TLS server verification.

Note: The bundle is currently not available when using WolfSSL.

The bundle comes with the complete list of root certificates from Mozilla's NSS root certificate store. Using the *gen_cert_bundle.py* python utility, the certificates' subject name and public key are stored in a file and embedded in the ESP32-P4 binary.

When generating the bundle you may choose between:

- The full root certificate bundle from Mozilla, containing more than 130 certificates. The current bundle was updated Tue Jan 10 04:12:06 2023 GMT.
- A pre-selected filter list of the name of the most commonly used root certificates, reducing the amount of certificates to around 41 while still having around 90% absolute usage coverage and 99% market share coverage according to SSL certificate authorities statistics.

In addition, it is possible to specify a path to a certificate file or a directory containing certificates which then will be added to the generated bundle.

Note: Trusting all root certificates means the list will have to be updated if any of the certificates are retracted. This includes removing them from *cacrt_all.pem*.

Configuration

Most configuration is done through menuconfig. CMake generates the bundle according to the configuration and embed it.

- `CONFIG_MBEDTLS_CERTIFICATE_BUNDLE`: automatically build and attach the bundle.
- `CONFIG_MBEDTLS_DEFAULT_CERTIFICATE_BUNDLE`: decide which certificates to include from the complete root certificate list.
- `CONFIG_MBEDTLS_CUSTOM_CERTIFICATE_BUNDLE_PATH`: specify the path of any additional certificates to embed in the bundle.

To enable the bundle when using ESP-TLS simply pass the function pointer to the bundle attach function:

```
esp_tls_cfg_t cfg = {  
    .crt_bundle_attach = esp_cert_bundle_attach,  
};
```

This is done to avoid embedding the certificate bundle unless activated by the user.

If using mbedTLS directly then the bundle may be activated by directly calling the attach function during the setup process:

```
mbedtls_ssl_config conf;  
mbedtls_ssl_config_init(&conf);  
  
esp_cert_bundle_attach(&conf);
```

Generating the List of Root Certificates

The list of root certificates comes from Mozilla's NSS root certificate store, which can be found [here](#)

The list can be downloaded and created by running the script `mk-ca-bundle.pl` that is distributed as a part of [curl](#).

Another alternative would be to download the finished list directly from the curl website: [CA certificates extracted from Mozilla](#)

The common certificates bundle were made by selecting the authorities with a market share of more than 1% from w3tech's [SSL Survey](#).

These authorities were then used to pick the names of the certificates for the filter list, `cmn_cert_authorities.csv`, from Mozilla's [Included CA Certificate List](#).

Updating the Certificate Bundle

The bundle is embedded into the app and can be updated along with the app by an OTA update. If you want to include a more up-to-date bundle than the bundle currently included in ESP-IDF, then the certificate list can be downloaded from Mozilla as described in [Generating the List of Root Certificates](#).

Application Examples

Simple HTTPS example that uses ESP-TLS to establish a secure socket connection using the certificate bundle with two custom certificates added for verification: [protocols/https_x509_bundle](#).

HTTPS example that uses ESP-TLS and the default bundle: [protocols/https_request](#).

HTTPS example that uses mbedTLS and the default bundle: [protocols/https_mbedtls](#).

API Reference

Header File

- [components/mbedtls/esp_crt_bundle/include/esp_crt_bundle.h](#)
- This header file can be included with:

```
#include "esp_crt_bundle.h"
```

- This header file is a part of the API provided by the `mbedtls` component. To declare that your component depends on `mbedtls`, add the following to your `CMakeLists.txt`:

```
REQUIRES mbedtls
```

or

```
PRIV_REQUIRES mbedtls
```

Functions

esp_err_t **esp_crt_bundle_attach** (void *conf)

Attach and enable use of a bundle for certificate verification.

Attach and enable use of a bundle for certificate verification through a verification callback. If no specific bundle has been set through `esp_crt_bundle_set()` it will default to the bundle defined in `menuconfig` and embedded in the binary.

Parameters **conf** -- [in] The config struct for the SSL connection.

Returns

- `ESP_OK` if adding certificates was successful.
- Other if an error occurred or an action must be taken by the calling process.

void **esp_crt_bundle_detach** (mbedtls_ssl_config *conf)

Disable and dealloc the certification bundle.

Removes the certificate verification callback and deallocates used resources

Parameters **conf** -- [in] The config struct for the SSL connection.

esp_err_t **esp_crt_bundle_set** (const uint8_t *x509_bundle, size_t bundle_size)

Set the default certificate bundle used for verification.

Overrides the default certificate bundle only in case of successful initialization. In most use cases the bundle should be set through `menuconfig`. The bundle needs to be sorted by subject name since binary search is used to find certificates.

Parameters

- **x509_bundle** -- [in] A pointer to the certificate bundle.
- **bundle_size** -- [in] Size of the certificate bundle in bytes.

Returns

- `ESP_OK` if adding certificates was successful.
- Other if an error occurred or an action must be taken by the calling process.

bool **esp_crt_bundle_in_use** (const mbedtls_x509_crt *ca_chain)

Check if the given CA certificate chain is the default "dummy" certificate chain attached by the `esp_crt_bundle`.

Parameters **ca_chain** -- A pointer to the CA chain.

Returns true if the `ca_chain` is the dummy CA chain attached by `esp_crt_bundle`

Returns false otherwise

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.2.9 HTTP Server

Overview

The HTTP Server component provides an ability for running a lightweight web server on ESP32-P4. Following are detailed steps to use the API exposed by HTTP Server:

- `httpd_start()`: Creates an instance of HTTP server, allocate memory/resources for it depending upon the specified configuration and outputs a handle to the server instance. The server has both, a listening socket (TCP) for HTTP traffic, and a control socket (UDP) for control signals, which are selected in a round robin fashion in the server task loop. The task priority and stack size are configurable during server instance creation by passing `httpd_config_t` structure to `httpd_start()`. TCP traffic is parsed as HTTP requests and, depending on the requested URI, user registered handlers are invoked which are supposed to send back HTTP response packets.
- `httpd_stop()`: This stops the server with the provided handle and frees up any associated memory/resources. This is a blocking function that first signals a halt to the server task and then waits for the task to terminate. While stopping, the task closes all open connections, removes registered URI handlers and resets all session context data to empty.
- `httpd_register_uri_handler()`: A URI handler is registered by passing object of type `httpd_uri_t` structure which has members including uri name, method type (eg. `HTTPD_GET/HTTPD_POST/HTTPD_PUT` etc.), function pointer of type `esp_err_t *handler (httpd_req_t *req)` and `user_ctx` pointer to user context data.

Application Example

```
/* Our URI handler function to be called during GET /uri request */
esp_err_t get_handler(httpd_req_t *req)
{
    /* Send a simple response */
    const char resp[] = "URI GET Response";
    httpd_resp_send(req, resp, HTTPD_RESP_USE_STRLEN);
    return ESP_OK;
}

/* Our URI handler function to be called during POST /uri request */
esp_err_t post_handler(httpd_req_t *req)
{
    /* Destination buffer for content of HTTP POST request.
     * httpd_req_recv() accepts char* only, but content could
     * as well be any binary data (needs type casting).
     * In case of string data, null termination will be absent, and
     * content length would give length of string */
    char content[100];

    /* Truncate if content length larger than the buffer */
    size_t recv_size = MIN(req->content_len, sizeof(content));

    int ret = httpd_req_recv(req, content, recv_size);
    if (ret <= 0) { /* 0 return value indicates connection closed */
        /* Check if timeout occurred */
        if (ret == HTTPD_SOCK_ERR_TIMEOUT) {
            /* In case of timeout one can choose to retry calling
             * httpd_req_recv(), but to keep it simple, here we
             * respond with an HTTP 408 (Request Timeout) error */
            httpd_resp_send_408(req);
        }
        /* In case of error, returning ESP_FAIL will
```

(continues on next page)

(continued from previous page)

```

        * ensure that the underlying socket is closed */
        return ESP_FAIL;
    }

    /* Send a simple response */
    const char resp[] = "URI POST Response";
    httpd_resp_send(req, resp, HTTPD_RESP_USE_STRLEN);
    return ESP_OK;
}

/* URI handler structure for GET /uri */
httpd_uri_t uri_get = {
    .uri      = "/uri",
    .method   = HTTP_GET,
    .handler  = get_handler,
    .user_ctx = NULL
};

/* URI handler structure for POST /uri */
httpd_uri_t uri_post = {
    .uri      = "/uri",
    .method   = HTTP_POST,
    .handler  = post_handler,
    .user_ctx = NULL
};

/* Function for starting the webserver */
httpd_handle_t start_webserver(void)
{
    /* Generate default configuration */
    httpd_config_t config = HTTPD_DEFAULT_CONFIG();

    /* Empty handle to esp_http_server */
    httpd_handle_t server = NULL;

    /* Start the httpd server */
    if (httpd_start(&server, &config) == ESP_OK) {
        /* Register URI handlers */
        httpd_register_uri_handler(server, &uri_get);
        httpd_register_uri_handler(server, &uri_post);
    }
    /* If server failed to start, handle will be NULL */
    return server;
}

/* Function for stopping the webserver */
void stop_webserver(httpd_handle_t server)
{
    if (server) {
        /* Stop the httpd server */
        httpd_stop(server);
    }
}

```

Simple HTTP Server Example Check HTTP server example under [protocols/http_server/simple](#) where handling of arbitrary content lengths, reading request headers and URL query parameters, and setting response headers is demonstrated.

Persistent Connections

HTTP server features persistent connections, allowing for the re-use of the same connection (session) for several transfers, all the while maintaining context specific data for the session. Context data may be allocated dynamically by the handler in which case a custom function may need to be specified for freeing this data when the connection/session is closed.

Persistent Connections Example

```
/* Custom function to free context */
void free_ctx_func(void *ctx)
{
    /* Could be something other than free */
    free(ctx);
}

esp_err_t adder_post_handler(httpd_req_t *req)
{
    /* Create session's context if not already available */
    if (! req->sess_ctx) {
        req->sess_ctx = malloc(sizeof(ANY_DATA_TYPE)); /*!< Pointer to context_
↪data */
        req->free_ctx = free_ctx_func; /*!< Function to free_
↪context data */
    }

    /* Access context data */
    ANY_DATA_TYPE *ctx_data = (ANY_DATA_TYPE *) req->sess_ctx;

    /* Respond */
    .....
    .....
    .....

    return ESP_OK;
}
```

Check the example under [protocols/http_server/persistent_sockets](#).

Websocket Server

The HTTP server component provides websocket support. The websocket feature can be enabled in menuconfig using the `CONFIG_HTTPD_WS_SUPPORT` option. Please refer to the [protocols/http_server/ws_echo_server](#) example which demonstrates usage of the websocket feature.

Event Handling

ESP HTTP server has various events for which a handler can be triggered by *the Event Loop library* when the particular event occurs. The handler has to be registered using `esp_event_handler_register()`. This helps in event handling for ESP HTTP server.

`esp_http_server_event_id_t` has all the events which can happen for ESP HTTP server.

Expected data type for different ESP HTTP server events in event loop:

- `HTTP_SERVER_EVENT_ERROR`: `httpd_err_code_t`
- `HTTP_SERVER_EVENT_START`: `NULL`
- `HTTP_SERVER_EVENT_ON_CONNECTED`: `int`
- `HTTP_SERVER_EVENT_ON_HEADER`: `int`
- `HTTP_SERVER_EVENT_HEADERS_SENT`: `int`

- `HTTP_SERVER_EVENT_ON_DATA` : `esp_http_server_event_data`
- `HTTP_SERVER_EVENT_SENT_DATA` : `esp_http_server_event_data`
- `HTTP_SERVER_EVENT_DISCONNECTED` : `int`
- `HTTP_SERVER_EVENT_STOP` : `NULL`

API Reference

Header File

- `components/esp_http_server/include/esp_http_server.h`
- This header file can be included with:

```
#include "esp_http_server.h"
```

- This header file is a part of the API provided by the `esp_http_server` component. To declare that your component depends on `esp_http_server`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_http_server
```

or

```
PRIV_REQUIRES esp_http_server
```

Functions

`esp_err_t httpd_register_uri_handler` (`httpd_handle_t` handle, const `httpd_uri_t` *uri_handler)

Registers a URI handler.

Example usage:

```
esp_err_t my_uri_handler(httpd_req_t* req)
{
    // Recv , Process and Send
    ....
    ....
    ....

    // Fail condition
    if (....) {
        // Return fail to close session //
        return ESP_FAIL;
    }

    // On success
    return ESP_OK;
}

// URI handler structure
httpd_uri_t my_uri {
    .uri      = "/my_uri/path/xyz",
    .method   = HTTPD_GET,
    .handler  = my_uri_handler,
    .user_ctx = NULL
};

// Register handler
if (httpd_register_uri_handler(server_handle, &my_uri) != ESP_OK) {
    // If failed to register handler
    ....
}
```

Note: URI handlers can be registered in real time as long as the server handle is valid.

Parameters

- **handle** -- **[in]** handle to HTTPD server instance
- **uri_handler** -- **[in]** pointer to handler that needs to be registered

Returns

- ESP_OK : On successfully registering the handler
- ESP_ERR_INVALID_ARG : Null arguments
- ESP_ERR_HTTPD_HANDLERS_FULL : If no slots left for new handler
- ESP_ERR_HTTPD_HANDLER_EXISTS : If handler with same URI and method is already registered

esp_err_t **httpd_unregister_uri_handler** (*httpd_handle_t* handle, const char *uri, *httpd_method_t* method)

Unregister a URI handler.

Parameters

- **handle** -- **[in]** handle to HTTPD server instance
- **uri** -- **[in]** URI string
- **method** -- **[in]** HTTP method

Returns

- ESP_OK : On successfully deregistering the handler
- ESP_ERR_INVALID_ARG : Null arguments
- ESP_ERR_NOT_FOUND : Handler with specified URI and method not found

esp_err_t **httpd_unregister_uri** (*httpd_handle_t* handle, const char *uri)

Unregister all URI handlers with the specified uri string.

Parameters

- **handle** -- **[in]** handle to HTTPD server instance
- **uri** -- **[in]** uri string specifying all handlers that need to be deregisterd

Returns

- ESP_OK : On successfully deregistering all such handlers
- ESP_ERR_INVALID_ARG : Null arguments
- ESP_ERR_NOT_FOUND : No handler registered with specified uri string

esp_err_t **httpd_sess_set_recv_override** (*httpd_handle_t* hd, int sockfd, *httpd_recv_func_t* recv_func)

Override web server's receive function (by session FD)

This function overrides the web server's receive function. This same function is used to read HTTP request packets.

Note: This API is supposed to be called either from the context of

- an http session APIs where sockfd is a valid parameter
 - a URI handler where sockfd is obtained using httpd_req_to_sockfd()
-

Parameters

- **hd** -- **[in]** HTTPD instance handle
- **sockfd** -- **[in]** Session socket FD
- **recv_func** -- **[in]** The receive function to be set for this session

Returns

- ESP_OK : On successfully registering override
- ESP_ERR_INVALID_ARG : Null arguments

esp_err_t **httpd_sess_set_send_override** (*httpd_handle_t* hd, int sockfd, *httpd_send_func_t* send_func)

Override web server's send function (by session FD)

This function overrides the web server's send function. This same function is used to send out any response to any HTTP request.

Note: This API is supposed to be called either from the context of

- an http session APIs where sockfd is a valid parameter
 - a URI handler where sockfd is obtained using httpd_req_to_sockfd()
-

Parameters

- **hd** -- [in] HTTPD instance handle
- **sockfd** -- [in] Session socket FD
- **send_func** -- [in] The send function to be set for this session

Returns

- ESP_OK : On successfully registering override
- ESP_ERR_INVALID_ARG : Null arguments

esp_err_t **httpd_sess_set_pending_override** (*httpd_handle_t* hd, int sockfd, *httpd_pending_func_t* pending_func)

Override web server's pending function (by session FD)

This function overrides the web server's pending function. This function is used to test for pending bytes in a socket.

Note: This API is supposed to be called either from the context of

- an http session APIs where sockfd is a valid parameter
 - a URI handler where sockfd is obtained using httpd_req_to_sockfd()
-

Parameters

- **hd** -- [in] HTTPD instance handle
- **sockfd** -- [in] Session socket FD
- **pending_func** -- [in] The receive function to be set for this session

Returns

- ESP_OK : On successfully registering override
- ESP_ERR_INVALID_ARG : Null arguments

esp_err_t **httpd_req_async_handler_begin** (*httpd_req_t* *r, *httpd_req_t* **out)

Start an asynchronous request. This function can be called in a request handler to get a request copy that can be used on an async thread.

Note:

- This function is necessary in order to handle multiple requests simultaneously. See examples/async_requests for example usage.
 - You must call httpd_req_async_handler_complete() when you are done with the request and also on any error conditions.
-

Parameters

- **r** -- [in] The request to create an async copy of
- **out** -- [out] A newly allocated request which can be used on an async thread

Returns

- ESP_OK : async request object created

esp_err_t **httpd_req_async_handler_complete** (*httpd_req_t* *r)

Mark an asynchronous request as completed. This will.

- free the request memory
- relinquish ownership of the underlying socket, so it can be reused.
- allow the http server to close our socket if needed (lru_purge_enable)

Note: If async requests are not marked completed, eventually the server will no longer accept incoming connections. The server will log a "httpd_accept_conn: error in accept (23)" message if this happens.

Parameters **r** -- [in] The request to mark async work as completed

Returns

- ESP_OK : async request was marked completed

int **httpd_req_to_sockfd** (*httpd_req_t* *r)

Get the Socket Descriptor from the HTTP request.

This API will return the socket descriptor of the session for which URI handler was executed on reception of HTTP request. This is useful when user wants to call functions that require session socket fd, from within a URI handler, ie. : `httpd_sess_get_ctx()`, `httpd_sess_trigger_close()`, `httpd_sess_update_lru_counter()`.

Note: This API is supposed to be called only from the context of a URI handler where `httpd_req_t*` request pointer is valid.

Parameters **r** -- [in] The request whose socket descriptor should be found

Returns

- Socket descriptor : The socket descriptor for this request
- -1 : Invalid/NULL request pointer

int **httpd_req_recv** (*httpd_req_t* *r, char *buf, size_t buf_len)

API to read content data from the HTTP request.

This API will read HTTP content data from the HTTP request into provided buffer. Use `content_len` provided in `httpd_req_t` structure to know the length of data to be fetched. If `content_len` is too large for the buffer then user may have to make multiple calls to this function, each time fetching 'buf_len' number of bytes, while the pointer to content data is incremented internally by the same number.

Note:

- This API is supposed to be called only from the context of a URI handler where `httpd_req_t*` request pointer is valid.
 - If an error is returned, the URI handler must further return an error. This will ensure that the erroneous socket is closed and cleaned up by the web server.
 - Presently Chunked Encoding is not supported
-

Parameters

- **r** -- [in] The request being responded to
- **buf** -- [in] Pointer to a buffer that the data will be read into
- **buf_len** -- [in] Length of the buffer

Returns

- Bytes : Number of bytes read into the buffer successfully
- 0 : Buffer length parameter is zero / connection closed by peer
- HTTPD SOCK_ERR_INVALID : Invalid arguments

- HTTPD SOCK_ERR_TIMEOUT : Timeout/interrupted while calling socket recv()
- HTTPD SOCK_ERR_FAIL : Unrecoverable error while calling socket recv()

size_t **httpd_req_get_hdr_value_len** (*httpd_req_t* *r, const char *field)

Search for a field in request headers and return the string length of it's value.

Note:

- This API is supposed to be called only from the context of a URI handler where httpd_req_t* request pointer is valid.
- Once httpd_resp_send() API is called all request headers are purged, so request headers need be copied into separate buffers if they are required later.

Parameters

- **r** -- [in] The request being responded to
- **field** -- [in] The header field to be searched in the request

Returns

- Length : If field is found in the request URL
- Zero : Field not found / Invalid request / Null arguments

esp_err_t **httpd_req_get_hdr_value_str** (*httpd_req_t* *r, const char *field, char *val, size_t val_size)

Get the value string of a field from the request headers.

Note:

- This API is supposed to be called only from the context of a URI handler where httpd_req_t* request pointer is valid.
- Once httpd_resp_send() API is called all request headers are purged, so request headers need be copied into separate buffers if they are required later.
- If output size is greater than input, then the value is truncated, accompanied by truncation error as return value.
- Use httpd_req_get_hdr_value_len() to know the right buffer length

Parameters

- **r** -- [in] The request being responded to
- **field** -- [in] The field to be searched in the header
- **val** -- [out] Pointer to the buffer into which the value will be copied if the field is found
- **val_size** -- [in] Size of the user buffer "val"

Returns

- ESP_OK : Field found in the request header and value string copied
- ESP_ERR_NOT_FOUND : Key not found
- ESP_ERR_INVALID_ARG : Null arguments
- ESP_ERR_HTTPD_INVALID_REQ : Invalid HTTP request pointer
- ESP_ERR_HTTPD_RESULT_TRUNC : Value string truncated

size_t **httpd_req_get_url_query_len** (*httpd_req_t* *r)

Get Query string length from the request URL.

Note: This API is supposed to be called only from the context of a URI handler where httpd_req_t* request pointer is valid

Parameters **r** -- [in] The request being responded to

Returns

- Length : Query is found in the request URL

- Zero : Query not found / Null arguments / Invalid request / uri is empty

esp_err_t **httpd_req_get_url_query_str** (*httpd_req_t* *r, char *buf, size_t buf_len)

Get Query string from the request URL.

Note:

- Presently, the user can fetch the full URL query string, but decoding will have to be performed by the user. Request headers can be read using `httpd_req_get_hdr_value_str()` to know the 'Content-Type' (eg. Content-Type: application/x-www-form-urlencoded) and then the appropriate decoding algorithm needs to be applied.
 - This API is supposed to be called only from the context of a URI handler where `httpd_req_t*` request pointer is valid
 - If output size is greater than input, then the value is truncated, accompanied by truncation error as return value
 - Prior to calling this function, one can use `httpd_req_get_url_query_len()` to know the query string length beforehand and hence allocate the buffer of right size (usually query string length + 1 for null termination) for storing the query string
-

Parameters

- **r** -- **[in]** The request being responded to
- **buf** -- **[out]** Pointer to the buffer into which the query string will be copied (if found)
- **buf_len** -- **[in]** Length of output buffer

Returns

- **ESP_OK** : Query is found in the request URL and copied to buffer
- **ESP_FAIL** : uri is empty
- **ESP_ERR_NOT_FOUND** : Query not found
- **ESP_ERR_INVALID_ARG** : Null arguments
- **ESP_ERR_HTTPD_INVALID_REQ** : Invalid HTTP request pointer
- **ESP_ERR_HTTPD_RESULT_TRUNC** : Query string truncated

esp_err_t **httpd_query_key_value** (const char *qry, const char *key, char *val, size_t val_size)

Helper function to get a URL query tag from a query string of the type param1=val1&param2=val2.

Note:

- The components of URL query string (keys and values) are not URLdecoded. The user must check for 'Content-Type' field in the request headers and then depending upon the specified encoding (URLencoded or otherwise) apply the appropriate decoding algorithm.
 - If actual value size is greater than `val_size`, then the value is truncated, accompanied by truncation error as return value.
-

Parameters

- **qry** -- **[in]** Pointer to query string
- **key** -- **[in]** The key to be searched in the query string
- **val** -- **[out]** Pointer to the buffer into which the value will be copied if the key is found
- **val_size** -- **[in]** Size of the user buffer "val"

Returns

- **ESP_OK** : Key is found in the URL query string and copied to buffer
- **ESP_ERR_NOT_FOUND** : Key not found
- **ESP_ERR_INVALID_ARG** : Null arguments
- **ESP_ERR_HTTPD_RESULT_TRUNC** : Value string truncated

esp_err_t **httpd_req_get_cookie_val** (*httpd_req_t* *req, const char *cookie_name, char *val, size_t *val_size)

Get the value string of a cookie value from the "Cookie" request headers by cookie name.

Parameters

- **req** -- [in] Pointer to the HTTP request
- **cookie_name** -- [in] The cookie name to be searched in the request
- **val** -- [out] Pointer to the buffer into which the value of cookie will be copied if the cookie is found
- **val_size** -- [inout] Pointer to size of the user buffer "val". This variable will contain cookie length if ESP_OK is returned and required buffer length incase ESP_ERR_HTTPD_RESULT_TRUNC is returned.

Returns

- ESP_OK : Key is found in the cookie string and copied to buffer
- ESP_ERR_NOT_FOUND : Key not found
- ESP_ERR_INVALID_ARG : Null arguments
- ESP_ERR_HTTPD_RESULT_TRUNC : Value string truncated
- ESP_ERR_NO_MEM : Memory allocation failure

bool **httpd_uri_match_wildcard** (const char *uri_template, const char *uri_to_match, size_t match_upto)

Test if a URI matches the given wildcard template.

Template may end with "?" to make the previous character optional (typically a slash), "*" for a wildcard match, and "?*" to make the previous character optional, and if present, allow anything to follow.

Example:

- * matches everything
- /foo/? matches /foo and /foo/
- /foo/* (sans the backslash) matches /foo/ and /foo/bar, but not /foo or /fo
- /foo/?* or /foo/*? (sans the backslash) matches /foo/, /foo/bar, and also /foo, but not /foox or /fo

The special characters "?" and "*" anywhere else in the template will be taken literally.

Parameters

- **uri_template** -- [in] URI template (pattern)
- **uri_to_match** -- [in] URI to be matched
- **match_upto** -- [in] how many characters of the URI buffer to test (there may be trailing query string etc.)

Returns true if a match was found

esp_err_t **httpd_resp_send** (*httpd_req_t* *r, const char *buf, ssize_t buf_len)

API to send a complete HTTP response.

This API will send the data as an HTTP response to the request. This assumes that you have the entire response ready in a single buffer. If you wish to send response in incremental chunks use `httpd_resp_send_chunk()` instead.

If no status code and content-type were set, by default this will send 200 OK status code and content type as text/html. You may call the following functions before this API to configure the response headers : `httpd_resp_set_status()` - for setting the HTTP status string, `httpd_resp_set_type()` - for setting the Content Type, `httpd_resp_set_hdr()` - for appending any additional field value entries in the response header

Note:

- This API is supposed to be called only from the context of a URI handler where `httpd_req_t*` request pointer is valid.
 - Once this API is called, the request has been responded to.
 - No additional data can then be sent for the request.
 - Once this API is called, all request headers are purged, so request headers need be copied into separate buffers if they are required later.
-

Parameters

- **r** -- [in] The request being responded to
- **buf** -- [in] Buffer from where the content is to be fetched
- **buf_len** -- [in] Length of the buffer, HTTPD_RESP_USE_STRLEN to use strlen()

Returns

- ESP_OK : On successfully sending the response packet
- ESP_ERR_INVALID_ARG : Null request pointer
- ESP_ERR_HTTPD_RESP_HDR : Essential headers are too large for internal buffer
- ESP_ERR_HTTPD_RESP_SEND : Error in raw send
- ESP_ERR_HTTPD_INVALID_REQ : Invalid request

esp_err_t **httpd_resp_send_chunk** (*httpd_req_t* *r, const char *buf, ssize_t buf_len)

API to send one HTTP chunk.

This API will send the data as an HTTP response to the request. This API will use chunked-encoding and send the response in the form of chunks. If you have the entire response contained in a single buffer, please use `httpd_resp_send()` instead.

If no status code and content-type were set, by default this will send 200 OK status code and content type as text/html. You may call the following functions before this API to configure the response headers `httpd_resp_set_status()` - for setting the HTTP status string, `httpd_resp_set_type()` - for setting the Content Type, `httpd_resp_set_hdr()` - for appending any additional field value entries in the response header

Note:

- This API is supposed to be called only from the context of a URI handler where `httpd_req_t*` request pointer is valid.
 - When you are finished sending all your chunks, you must call this function with `buf_len` as 0.
 - Once this API is called, all request headers are purged, so request headers need be copied into separate buffers if they are required later.
-

Parameters

- **r** -- [in] The request being responded to
- **buf** -- [in] Pointer to a buffer that stores the data
- **buf_len** -- [in] Length of the buffer, HTTPD_RESP_USE_STRLEN to use strlen()

Returns

- ESP_OK : On successfully sending the response packet chunk
- ESP_ERR_INVALID_ARG : Null request pointer
- ESP_ERR_HTTPD_RESP_HDR : Essential headers are too large for internal buffer
- ESP_ERR_HTTPD_RESP_SEND : Error in raw send
- ESP_ERR_HTTPD_INVALID_REQ : Invalid request pointer

static inline *esp_err_t* **httpd_resp_sendstr** (*httpd_req_t* *r, const char *str)

API to send a complete string as HTTP response.

This API simply calls `httpd_resp_send` with buffer length set to string length assuming the buffer contains a null terminated string

Parameters

- **r** -- [in] The request being responded to
- **str** -- [in] String to be sent as response body

Returns

- ESP_OK : On successfully sending the response packet
- ESP_ERR_INVALID_ARG : Null request pointer
- ESP_ERR_HTTPD_RESP_HDR : Essential headers are too large for internal buffer
- ESP_ERR_HTTPD_RESP_SEND : Error in raw send
- ESP_ERR_HTTPD_INVALID_REQ : Invalid request

static inline [esp_err_t](#) **httpd_resp_sendstr_chunk** ([httpd_req_t](#) *r, const char *str)

API to send a string as an HTTP response chunk.

This API simply calls `http_resp_send_chunk` with buffer length set to string length assuming the buffer contains a null terminated string

Parameters

- **r** -- **[in]** The request being responded to
- **str** -- **[in]** String to be sent as response body (NULL to finish response packet)

Returns

- `ESP_OK` : On successfully sending the response packet
- `ESP_ERR_INVALID_ARG` : Null request pointer
- `ESP_ERR_HTTPD_RESP_HDR` : Essential headers are too large for internal buffer
- `ESP_ERR_HTTPD_RESP_SEND` : Error in raw send
- `ESP_ERR_HTTPD_INVALID_REQ` : Invalid request

[esp_err_t](#) **httpd_resp_set_status** ([httpd_req_t](#) *r, const char *status)

API to set the HTTP status code.

This API sets the status of the HTTP response to the value specified. By default, the '200 OK' response is sent as the response.

Note:

- This API is supposed to be called only from the context of a URI handler where `httpd_req_t*` request pointer is valid.
 - This API only sets the status to this value. The status isn't sent out until any of the send APIs is executed.
 - Make sure that the lifetime of the status string is valid till send function is called.
-

Parameters

- **r** -- **[in]** The request being responded to
- **status** -- **[in]** The HTTP status code of this response

Returns

- `ESP_OK` : On success
- `ESP_ERR_INVALID_ARG` : Null arguments
- `ESP_ERR_HTTPD_INVALID_REQ` : Invalid request pointer

[esp_err_t](#) **httpd_resp_set_type** ([httpd_req_t](#) *r, const char *type)

API to set the HTTP content type.

This API sets the 'Content Type' field of the response. The default content type is 'text/html'.

Note:

- This API is supposed to be called only from the context of a URI handler where `httpd_req_t*` request pointer is valid.
 - This API only sets the content type to this value. The type isn't sent out until any of the send APIs is executed.
 - Make sure that the lifetime of the type string is valid till send function is called.
-

Parameters

- **r** -- **[in]** The request being responded to
- **type** -- **[in]** The Content Type of the response

Returns

- `ESP_OK` : On success
- `ESP_ERR_INVALID_ARG` : Null arguments
- `ESP_ERR_HTTPD_INVALID_REQ` : Invalid request pointer

esp_err_t **httpd_resp_set_hdr** (*httpd_req_t* *r, const char *field, const char *value)

API to append any additional headers.

This API sets any additional header fields that need to be sent in the response.

Note:

- This API is supposed to be called only from the context of a URI handler where *httpd_req_t** request pointer is valid.
 - The header isn't sent out until any of the send APIs is executed.
 - The maximum allowed number of additional headers is limited to value of *max_resp_headers* in config structure.
 - Make sure that the lifetime of the field value strings are valid till send function is called.
-

Parameters

- **r** -- [in] The request being responded to
- **field** -- [in] The field name of the HTTP header
- **value** -- [in] The value of this HTTP header

Returns

- ESP_OK : On successfully appending new header
- ESP_ERR_INVALID_ARG : Null arguments
- ESP_ERR_HTTPD_RESP_HDR : Total additional headers exceed max allowed
- ESP_ERR_HTTPD_INVALID_REQ : Invalid request pointer

esp_err_t **httpd_resp_send_err** (*httpd_req_t* *req, *httpd_err_code_t* error, const char *msg)

For sending out error code in response to HTTP request.

Note:

- This API is supposed to be called only from the context of a URI handler where *httpd_req_t** request pointer is valid.
 - Once this API is called, all request headers are purged, so request headers need be copied into separate buffers if they are required later.
 - If you wish to send additional data in the body of the response, please use the lower-level functions directly.
-

Parameters

- **req** -- [in] Pointer to the HTTP request for which the response needs to be sent
- **error** -- [in] Error type to send
- **msg** -- [in] Error message string (pass NULL for default message)

Returns

- ESP_OK : On successfully sending the response packet
- ESP_ERR_INVALID_ARG : Null arguments
- ESP_ERR_HTTPD_RESP_SEND : Error in raw send
- ESP_ERR_HTTPD_INVALID_REQ : Invalid request pointer

static inline *esp_err_t* **httpd_resp_send_404** (*httpd_req_t* *r)

Helper function for HTTP 404.

Send HTTP 404 message. If you wish to send additional data in the body of the response, please use the lower-level functions directly.

Note:

- This API is supposed to be called only from the context of a URI handler where *httpd_req_t** request pointer is valid.

- Once this API is called, all request headers are purged, so request headers need be copied into separate buffers if they are required later.
-

Parameters *r* -- [in] The request being responded to

Returns

- ESP_OK : On successfully sending the response packet
- ESP_ERR_INVALID_ARG : Null arguments
- ESP_ERR_HTTPD_RESP_SEND : Error in raw send
- ESP_ERR_HTTPD_INVALID_REQ : Invalid request pointer

static inline [*esp_err_t*](#) **httpd_resp_send_408** ([*httpd_req_t*](#) *r)

Helper function for HTTP 408.

Send HTTP 408 message. If you wish to send additional data in the body of the response, please use the lower-level functions directly.

Note:

- This API is supposed to be called only from the context of a URI handler where *httpd_req_t** request pointer is valid.
 - Once this API is called, all request headers are purged, so request headers need be copied into separate buffers if they are required later.
-

Parameters *r* -- [in] The request being responded to

Returns

- ESP_OK : On successfully sending the response packet
- ESP_ERR_INVALID_ARG : Null arguments
- ESP_ERR_HTTPD_RESP_SEND : Error in raw send
- ESP_ERR_HTTPD_INVALID_REQ : Invalid request pointer

static inline [*esp_err_t*](#) **httpd_resp_send_500** ([*httpd_req_t*](#) *r)

Helper function for HTTP 500.

Send HTTP 500 message. If you wish to send additional data in the body of the response, please use the lower-level functions directly.

Note:

- This API is supposed to be called only from the context of a URI handler where *httpd_req_t** request pointer is valid.
 - Once this API is called, all request headers are purged, so request headers need be copied into separate buffers if they are required later.
-

Parameters *r* -- [in] The request being responded to

Returns

- ESP_OK : On successfully sending the response packet
- ESP_ERR_INVALID_ARG : Null arguments
- ESP_ERR_HTTPD_RESP_SEND : Error in raw send
- ESP_ERR_HTTPD_INVALID_REQ : Invalid request pointer

int **httpd_send** ([*httpd_req_t*](#) *r, const char *buf, size_t buf_len)

Raw HTTP send.

Call this API if you wish to construct your custom response packet. When using this, all essential header, eg. HTTP version, Status Code, Content Type and Length, Encoding, etc. will have to be constructed manually,

and HTTP delimiters (CRLF) will need to be placed correctly for separating sub-sections of the HTTP response packet.

If the send override function is set, this API will end up calling that function eventually to send data out.

Note:

- This API is supposed to be called only from the context of a URI handler where `httpd_req_t*` request pointer is valid.
 - Unless the response has the correct HTTP structure (which the user must now ensure) it is not guaranteed that it will be recognized by the client. For most cases, you wouldn't have to call this API, but you would rather use either of : `httpd_resp_send()`, `httpd_resp_send_chunk()`
-

Parameters

- **r** -- **[in]** The request being responded to
- **buf** -- **[in]** Buffer from where the fully constructed packet is to be read
- **buf_len** -- **[in]** Length of the buffer

Returns

- Bytes : Number of bytes that were sent successfully
- HTTPD_SOCK_ERR_INVALID : Invalid arguments
- HTTPD_SOCK_ERR_TIMEOUT : Timeout/interrupted while calling socket send()
- HTTPD_SOCK_ERR_FAIL : Unrecoverable error while calling socket send()

int **httpd_socket_send** ([*httpd_handle_t*](#) hd, int sockfd, const char *buf, size_t buf_len, int flags)

A low level API to send data on a given socket

This internally calls the default send function, or the function registered by `httpd_sess_set_send_override()`.

Note: This API is not recommended to be used in any request handler. Use this only for advanced use cases, wherein some asynchronous data is to be sent over a socket.

Parameters

- **hd** -- **[in]** server instance
- **sockfd** -- **[in]** session socket file descriptor
- **buf** -- **[in]** buffer with bytes to send
- **buf_len** -- **[in]** data size
- **flags** -- **[in]** flags for the send() function

Returns

- Bytes : The number of bytes sent successfully
- HTTPD_SOCK_ERR_INVALID : Invalid arguments
- HTTPD_SOCK_ERR_TIMEOUT : Timeout/interrupted while calling socket send()
- HTTPD_SOCK_ERR_FAIL : Unrecoverable error while calling socket send()

int **httpd_socket_recv** ([*httpd_handle_t*](#) hd, int sockfd, char *buf, size_t buf_len, int flags)

A low level API to receive data from a given socket

This internally calls the default recv function, or the function registered by `httpd_sess_set_recv_override()`.

Note: This API is not recommended to be used in any request handler. Use this only for advanced use cases, wherein some asynchronous communication is required.

Parameters

- **hd** -- **[in]** server instance
- **sockfd** -- **[in]** session socket file descriptor
- **buf** -- **[in]** buffer with bytes to send
- **buf_len** -- **[in]** data size
- **flags** -- **[in]** flags for the send() function

Returns

- Bytes : The number of bytes received successfully
- 0 : Buffer length parameter is zero / connection closed by peer
- HTTPD SOCK_ERR_INVALID : Invalid arguments
- HTTPD SOCK_ERR_TIMEOUT : Timeout/interrupted while calling socket recv()
- HTTPD SOCK_ERR_FAIL : Unrecoverable error while calling socket recv()

esp_err_t **httpd_register_err_handler** (*httpd_handle_t* handle, *httpd_err_code_t* error, *httpd_err_handler_func_t* handler_fn)

Function for registering HTTP error handlers.

This function maps a handler function to any supported error code given by *httpd_err_code_t*. See prototype *httpd_err_handler_func_t* above for details.

Parameters

- **handle** -- **[in]** HTTP server handle
- **error** -- **[in]** Error type
- **handler_fn** -- **[in]** User implemented handler function (Pass NULL to unset any previously set handler)

Returns

- ESP_OK : handler registered successfully
- ESP_ERR_INVALID_ARG : invalid error code or server handle

esp_err_t **httpd_start** (*httpd_handle_t* *handle, const *httpd_config_t* *config)

Starts the web server.

Create an instance of HTTP server and allocate memory/resources for it depending upon the specified configuration.

Example usage:

```
//Function for starting the webserver
httpd_handle_t start_webserver(void)
{
    // Generate default configuration
    httpd_config_t config = HTTPD_DEFAULT_CONFIG();

    // Empty handle to http_server
    httpd_handle_t server = NULL;

    // Start the httpd server
    if (httpd_start(&server, &config) == ESP_OK) {
        // Register URI handlers
        httpd_register_uri_handler(server, &uri_get);
        httpd_register_uri_handler(server, &uri_post);
    }
    // If server failed to start, handle will be NULL
    return server;
}
```

Parameters

- **config** -- **[in]** Configuration for new instance of the server
- **handle** -- **[out]** Handle to newly created instance of the server. NULL on error

Returns

- ESP_OK : Instance created successfully
- ESP_ERR_INVALID_ARG : Null argument(s)

- ESP_ERR_HTTPD_ALLOC_MEM : Failed to allocate memory for instance
- ESP_ERR_HTTPD_TASK : Failed to launch server task

esp_err_t **httpd_stop** (*httpd_handle_t* handle)

Stops the web server.

Deallocates memory/resources used by an HTTP server instance and deletes it. Once deleted the handle can no longer be used for accessing the instance.

Example usage:

```
// Function for stopping the webserver
void stop_webserver(httpd_handle_t server)
{
    // Ensure handle is non NULL
    if (server != NULL) {
        // Stop the httpd server
        httpd_stop(server);
    }
}
```

Parameters **handle** -- [in] Handle to server returned by httpd_start

Returns

- ESP_OK : Server stopped successfully
- ESP_ERR_INVALID_ARG : Handle argument is Null

esp_err_t **httpd_queue_work** (*httpd_handle_t* handle, *httpd_work_fn_t* work, void *arg)

Queue execution of a function in HTTPD's context.

This API queues a work function for asynchronous execution

Note: Some protocols require that the web server generate some asynchronous data and send it to the persistently opened connection. This facility is for use by such protocols.

Parameters

- **handle** -- [in] Handle to server returned by httpd_start
- **work** -- [in] Pointer to the function to be executed in the HTTPD's context
- **arg** -- [in] Pointer to the arguments that should be passed to this function

Returns

- ESP_OK : On successfully queueing the work
- ESP_FAIL : Failure in ctrl socket
- ESP_ERR_INVALID_ARG : Null arguments

void ***httpd_sess_get_ctx** (*httpd_handle_t* handle, int sockfd)

Get session context from socket descriptor.

Typically if a session context is created, it is available to URI handlers through the httpd_req_t structure. But, there are cases where the web server's send/receive functions may require the context (for example, for accessing keying information etc). Since the send/receive function only have the socket descriptor at their disposal, this API provides them with a way to retrieve the session context.

Parameters

- **handle** -- [in] Handle to server returned by httpd_start
- **sockfd** -- [in] The socket descriptor for which the context should be extracted.

Returns

- void* : Pointer to the context associated with this session
- NULL : Empty context / Invalid handle / Invalid socket fd

void **httpd_sess_set_ctx** (*httpd_handle_t* handle, int sockfd, void *ctx, *httpd_free_ctx_fn_t* free_fn)

Set session context by socket descriptor.

Parameters

- **handle** -- [in] Handle to server returned by httpd_start
- **sockfd** -- [in] The socket descriptor for which the context should be extracted.
- **ctx** -- [in] Context object to assign to the session
- **free_fn** -- [in] Function that should be called to free the context

void ***httpd_sess_get_transport_ctx** (*httpd_handle_t* handle, int sockfd)

Get session 'transport' context by socket descriptor.

This context is used by the send/receive functions, for example to manage SSL context.

See also:

httpd_sess_get_ctx()

Parameters

- **handle** -- [in] Handle to server returned by httpd_start
- **sockfd** -- [in] The socket descriptor for which the context should be extracted.

Returns

- void* : Pointer to the transport context associated with this session
- NULL : Empty context / Invalid handle / Invalid socket fd

void **httpd_sess_set_transport_ctx** (*httpd_handle_t* handle, int sockfd, void *ctx, *httpd_free_ctx_fn_t* free_fn)

Set session 'transport' context by socket descriptor.

See also:

httpd_sess_set_ctx()

Parameters

- **handle** -- [in] Handle to server returned by httpd_start
- **sockfd** -- [in] The socket descriptor for which the context should be extracted.
- **ctx** -- [in] Transport context object to assign to the session
- **free_fn** -- [in] Function that should be called to free the transport context

void ***httpd_get_global_user_ctx** (*httpd_handle_t* handle)

Get HTTPD global user context (it was set in the server config struct)

Parameters **handle** -- [in] Handle to server returned by httpd_start

Returns global user context

void ***httpd_get_global_transport_ctx** (*httpd_handle_t* handle)

Get HTTPD global transport context (it was set in the server config struct)

Parameters **handle** -- [in] Handle to server returned by httpd_start

Returns global transport context

esp_err_t **httpd_sess_trigger_close** (*httpd_handle_t* handle, int sockfd)

Trigger an httpd session close externally.

Note: Calling this API is only required in special circumstances wherein some application requires to close an httpd client session asynchronously.

Parameters

- **handle** -- **[in]** Handle to server returned by `httpd_start`
- **sockfd** -- **[in]** The socket descriptor of the session to be closed

Returns

- `ESP_OK` : On successfully initiating closure
- `ESP_FAIL` : Failure to queue work
- `ESP_ERR_NOT_FOUND` : Socket fd not found
- `ESP_ERR_INVALID_ARG` : Null arguments

esp_err_t **httpd_sess_update_lru_counter** (*httpd_handle_t* handle, int sockfd)

Update LRU counter for a given socket.

LRU Counters are internally associated with each session to monitor how recently a session exchanged traffic. When LRU purge is enabled, if a client is requesting for connection but maximum number of sockets/sessions is reached, then the session having the earliest LRU counter is closed automatically.

Updating the LRU counter manually prevents the socket from being purged due to the Least Recently Used (LRU) logic, even though it might not have received traffic for some time. This is useful when all open sockets/session are frequently exchanging traffic but the user specifically wants one of the sessions to be kept open, irrespective of when it last exchanged a packet.

Note: Calling this API is only necessary if the LRU Purge Enable option is enabled.

Parameters

- **handle** -- **[in]** Handle to server returned by `httpd_start`
- **sockfd** -- **[in]** The socket descriptor of the session for which LRU counter is to be updated

Returns

- `ESP_OK` : Socket found and LRU counter updated
- `ESP_ERR_NOT_FOUND` : Socket not found
- `ESP_ERR_INVALID_ARG` : Null arguments

esp_err_t **httpd_get_client_list** (*httpd_handle_t* handle, size_t *fds, int *client_fds)

Returns list of current socket descriptors of active sessions.

Note: Size of provided array has to be equal or greater then maximum number of opened sockets, configured upon initialization with `max_open_sockets` field in `httpd_config_t` structure.

Parameters

- **handle** -- **[in]** Handle to server returned by `httpd_start`
- **fds** -- **[inout]** In: Size of provided `client_fds` array Out: Number of valid client fds returned in `client_fds`,
- **client_fds** -- **[out]** Array of client fds

Returns

- `ESP_OK` : Successfully retrieved session list
- `ESP_ERR_INVALID_ARG` : Wrong arguments or list is longer than provided array

Structures

struct **esp_http_server_event_data**

Argument structure for `HTTP_SERVER_EVENT_ON_DATA` and `HTTP_SERVER_EVENT_SENT_DATA` event

Public Members

int **fd**

Session socket file descriptor

int **data_len**

Data length

struct **httpd_config**

HTTP Server Configuration Structure.

Note: Use HTTPD_DEFAULT_CONFIG() to initialize the configuration to a default value and then modify only those fields that are specifically determined by the use case.

Public Members

unsigned **task_priority**

Priority of FreeRTOS task which runs the server

size_t **stack_size**

The maximum stack size allowed for the server task

BaseType_t **core_id**

The core the HTTP server task will run on

uint16_t **server_port**

TCP Port number for receiving and transmitting HTTP traffic

uint16_t **ctrl_port**

UDP Port number for asynchronously exchanging control signals between various components of the server

uint16_t **max_open_sockets**

Max number of sockets/clients connected at any time (3 sockets are reserved for internal working of the HTTP server)

uint16_t **max_uri_handlers**

Maximum allowed uri handlers

uint16_t **max_resp_headers**

Maximum allowed additional headers in HTTP response

uint16_t **backlog_conn**

Number of backlog connections

bool **lru_purge_enable**

Purge "Least Recently Used" connection

uint16_t **recv_wait_timeout**

Timeout for recv function (in seconds)

uint16_t **send_wait_timeout**

Timeout for send function (in seconds)

void ***global_user_ctx**

Global user context.

This field can be used to store arbitrary user data within the server context. The value can be retrieved using the server handle, available e.g. in the `httpd_req_t` struct.

When shutting down, the server frees up the user context by calling `free()` on the `global_user_ctx` field. If you wish to use a custom function for freeing the global user context, please specify that here.

[*httpd_free_ctx_fn_t*](#) **global_user_ctx_free_fn**

Free function for global user context

void ***global_transport_ctx**

Global transport context.

Similar to `global_user_ctx`, but used for session encoding or encryption (e.g. to hold the SSL context). It will be freed using `free()`, unless `global_transport_ctx_free_fn` is specified.

[*httpd_free_ctx_fn_t*](#) **global_transport_ctx_free_fn**

Free function for global transport context

bool **enable_so_linger**

bool to enable/disable linger

int **linger_timeout**

linger timeout (in seconds)

bool **keep_alive_enable**

Enable keep-alive timeout

int **keep_alive_idle**

Keep-alive idle time. Default is 5 (second)

int **keep_alive_interval**

Keep-alive interval time. Default is 5 (second)

int **keep_alive_count**

Keep-alive packet retry send count. Default is 3 counts

[*httpd_open_func_t*](#) **open_fn**

Custom session opening callback.

Called on a new session socket just after `accept()`, but before reading any data.

This is an opportunity to set up e.g. SSL encryption using `global_transport_ctx` and the `send/recv/pending` session overrides.

If a context needs to be maintained between these functions, store it in the session using `httpd_sess_set_transport_ctx()` and retrieve it later with `httpd_sess_get_transport_ctx()`

Returning a value other than `ESP_OK` will immediately close the new socket.

***httpd_close_func_t* close_fn**

Custom session closing callback.

Called when a session is deleted, before freeing user and transport contexts and before closing the socket. This is a place for custom de-init code common to all sockets.

The server will only close the socket if no custom session closing callback is set. If a custom callback is used, `close(sockfd)` should be called in here for most cases.

Set the user or transport context to NULL if it was freed here, so the server does not try to free it again.

This function is run for all terminated sessions, including sessions where the socket was closed by the network stack - that is, the file descriptor may not be valid anymore.

***httpd_uri_match_func_t* uri_match_fn**

URI matcher function.

Called when searching for a matching URI: 1) whose request handler is to be executed right after an HTTP request is successfully parsed 2) in order to prevent duplication while registering a new URI handler using `httpd_register_uri_handler()`

Available options are: 1) NULL : Internally do basic matching using `strcmp()` 2) `httpd_uri_match_wildcard()` : URI wildcard matcher

Users can implement their own matching functions (See description of the `httpd_uri_match_func_t` function prototype)

struct httpd_req

HTTP Request Data Structure.

Public Members***httpd_handle_t* handle**

Handle to server instance

int method

The type of HTTP request, -1 if unsupported method

const char uri[HTTPD_MAX_URI_LEN + 1]

The URI of this request (1 byte extra for null termination)

size_t content_len

Length of the request body

void *aux

Internally used members

void *user_ctx

User context pointer passed during URI registration.

void *sess_ctx

Session Context Pointer

A session context. Contexts are maintained across 'sessions' for a given open TCP connection. One session could have multiple request responses. The web server will ensure that the context persists across all these request and responses.

By default, this is NULL. URI Handlers can set this to any meaningful value.

If the underlying socket gets closed, and this pointer is non-NULL, the web server will free up the context by calling `free()`, unless `free_ctx` function is set.

`httpd_free_ctx_fn_t` **`free_ctx`**

Pointer to free context hook

Function to free session context

If the web server's socket closes, it frees up the session context by calling `free()` on the `sess_ctx` member. If you wish to use a custom function for freeing the session context, please specify that here.

bool **`ignore_sess_ctx_changes`**

Flag indicating if Session Context changes should be ignored

By default, if you change the `sess_ctx` in some URI handler, the http server will internally free the earlier context (if non NULL), after the URI handler returns. If you want to manage the allocation/reallocation/freeing of `sess_ctx` yourself, set this flag to true, so that the server will not perform any checks on it. The context will be cleared by the server (by calling `free_ctx` or `free()`) only if the socket gets closed.

struct **`httpd_uri`**

Structure for URI handler.

Public Members

const char ***`uri`**

The URI to handle

`httpd_method_t` **`method`**

Method supported by the URI

`esp_err_t` (***`handler`**)(*`httpd_req_t`* *r)

Handler to call for supported request method. This must return `ESP_OK`, or else the underlying socket will be closed.

void ***`user_ctx`**

Pointer to user context data which will be available to handler

Macros

`HTTPD_MAX_REQ_HDR_LEN`

`HTTPD_MAX_URI_LEN`

`HTTPD SOCK_ERR_FAIL`

`HTTPD SOCK_ERR_INVALID`

`HTTPD SOCK_ERR_TIMEOUT`

HTTPD_200

HTTP Response 200

HTTPD_204

HTTP Response 204

HTTPD_207

HTTP Response 207

HTTPD_400

HTTP Response 400

HTTPD_404

HTTP Response 404

HTTPD_408

HTTP Response 408

HTTPD_500

HTTP Response 500

HTTPD_TYPE_JSON

HTTP Content type JSON

HTTPD_TYPE_TEXT

HTTP Content type text/HTML

HTTPD_TYPE_OCTET

HTTP Content type octext-stream

ESP_HTTPD_DEF_CTRL_PORT

HTTP Server control socket port

HTTPD_DEFAULT_CONFIG()

ESP_ERR_HTTPD_BASE

Starting number of HTTPD error codes

ESP_ERR_HTTPD_HANDLERS_FULL

All slots for registering URI handlers have been consumed

ESP_ERR_HTTPD_HANDLER_EXISTS

URI handler with same method and target URI already registered

ESP_ERR_HTTPD_INVALID_REQ

Invalid request pointer

ESP_ERR_HTTPD_RESULT_TRUNC

Result string truncated

ESP_ERR_HTTPD_RESP_HDR

Response header field larger than supported

ESP_ERR_HTTPD_RESP_SEND

Error occurred while sending response packet

ESP_ERR_HTTPD_ALLOC_MEM

Failed to dynamically allocate memory for resource

ESP_ERR_HTTPD_TASK

Failed to launch server task/thread

HTTPD_RESP_USE_STRLEN

Type Definitions

typedef struct [httpd_req](#) **httpd_req_t**

HTTP Request Data Structure.

typedef struct [httpd_uri](#) **httpd_uri_t**

Structure for URI handler.

typedef int (***httpd_send_func_t**)([httpd_handle_t](#) hd, int sockfd, const char *buf, size_t buf_len, int flags)

Prototype for HTTPDs low-level send function.

Note: User specified send function must handle errors internally, depending upon the set value of `errno`, and return specific `HTTPD_SOCK_ERR_` codes, which will eventually be conveyed as return value of `httpd_send()` function

Param hd [in] server instance

Param sockfd [in] session socket file descriptor

Param buf [in] buffer with bytes to send

Param buf_len [in] data size

Param flags [in] flags for the `send()` function

Return

- Bytes : The number of bytes sent successfully
- `HTTPD_SOCK_ERR_INVALID` : Invalid arguments
- `HTTPD_SOCK_ERR_TIMEOUT` : Timeout/interrupted while calling socket `send()`
- `HTTPD_SOCK_ERR_FAIL` : Unrecoverable error while calling socket `send()`

typedef int (***httpd_recv_func_t**)([httpd_handle_t](#) hd, int sockfd, char *buf, size_t buf_len, int flags)

Prototype for HTTPDs low-level recv function.

Note: User specified recv function must handle errors internally, depending upon the set value of `errno`, and return specific `HTTPD_SOCK_ERR_` codes, which will eventually be conveyed as return value of `httpd_req_recv()` function

Param hd [in] server instance

Param sockfd [in] session socket file descriptor

Param buf [in] buffer with bytes to send

Param buf_len [in] data size

Param flags [in] flags for the send() function

Return

- Bytes : The number of bytes received successfully
- 0 : Buffer length parameter is zero / connection closed by peer
- HTTPD_SOCK_ERR_INVALID : Invalid arguments
- HTTPD_SOCK_ERR_TIMEOUT : Timeout/interrupted while calling socket recv()
- HTTPD_SOCK_ERR_FAIL : Unrecoverable error while calling socket recv()

```
typedef int (*httpd_pending_func_t)(httpd\_handle\_t hd, int sockfd)
```

Prototype for HTTPDs low-level "get pending bytes" function.

Note: User specified pending function must handle errors internally, depending upon the set value of errno, and return specific HTTPD_SOCK_ERR_ codes, which will be handled accordingly in the server task.

Param hd [in] server instance

Param sockfd [in] session socket file descriptor

Return

- Bytes : The number of bytes waiting to be received
- HTTPD_SOCK_ERR_INVALID : Invalid arguments
- HTTPD_SOCK_ERR_TIMEOUT : Timeout/interrupted while calling socket pending()
- HTTPD_SOCK_ERR_FAIL : Unrecoverable error while calling socket pending()

```
typedef esp\_err\_t (*httpd_err_handler_func_t)(httpd\_req\_t *req, httpd\_err\_code\_t error)
```

Function prototype for HTTP error handling.

This function is executed upon HTTP errors generated during internal processing of an HTTP request. This is used to override the default behavior on error, which is to send HTTP error response and close the underlying socket.

Note:

- If implemented, the server will not automatically send out HTTP error response codes, therefore, httpd_resp_send_err() must be invoked inside this function if user wishes to generate HTTP error responses.
 - When invoked, the validity of uri, method, content_len and user_ctx fields of the httpd_req_t parameter is not guaranteed as the HTTP request may be partially received/parsed.
 - The function must return ESP_OK if underlying socket needs to be kept open. Any other value will ensure that the socket is closed. The return value is ignored when error is of type HTTPD_500_INTERNAL_SERVER_ERROR and the socket closed anyway.
-

Param req [in] HTTP request for which the error needs to be handled

Param error [in] Error type

Return

- ESP_OK : error handled successful
- ESP_FAIL : failure indicates that the underlying socket needs to be closed

```
typedef void *httpd_handle_t
```

HTTP Server Instance Handle.

Every instance of the server will have a unique handle.

typedef enum http_method **httpd_method_t**

HTTP Method Type wrapper over "enum http_method" available in "http_parser" library.

typedef void (***httpd_free_ctx_fn_t**)(void *ctx)

Prototype for freeing context data (if any)

Param ctx [in] object to free

typedef *esp_err_t* (***httpd_open_func_t**)(*httpd_handle_t* hd, int sockfd)

Function prototype for opening a session.

Called immediately after the socket was opened to set up the send/recv functions and other parameters of the socket.

Param hd [in] server instance

Param sockfd [in] session socket file descriptor

Return

- ESP_OK : On success
- Any value other than ESP_OK will signal the server to close the socket immediately

typedef void (***httpd_close_func_t**)(*httpd_handle_t* hd, int sockfd)

Function prototype for closing a session.

Note: It's possible that the socket descriptor is invalid at this point, the function is called for all terminated sessions. Ensure proper handling of return codes.

Param hd [in] server instance

Param sockfd [in] session socket file descriptor

typedef bool (***httpd_uri_match_func_t**)(const char *reference_uri, const char *uri_to_match, size_t match_upto)

Function prototype for URI matching.

Param reference_uri [in] URI/template with respect to which the other URI is matched

Param uri_to_match [in] URI/template being matched to the reference URI/template

Param match_upto [in] For specifying the actual length of uri_to_match up to which the matching algorithm is to be applied (The maximum value is strlen(uri_to_match), independent of the length of reference_uri)

Return true on match

typedef struct *httpd_config* **httpd_config_t**

HTTP Server Configuration Structure.

Note: Use HTTPD_DEFAULT_CONFIG() to initialize the configuration to a default value and then modify only those fields that are specifically determined by the use case.

typedef void (***httpd_work_fn_t**)(void *arg)

Prototype of the HTTPD work function Please refer to httpd_queue_work() for more details.

Param arg [in] The arguments for this work function

Enumerations

enum **httpd_err_code_t**

Error codes sent as HTTP response in case of errors encountered during processing of an HTTP request.

Values:

enumerator **HTTPD_500_INTERNAL_SERVER_ERROR**

enumerator **HTTPD_501_METHOD_NOT_IMPLEMENTED**

enumerator **HTTPD_505_VERSION_NOT_SUPPORTED**

enumerator **HTTPD_400_BAD_REQUEST**

enumerator **HTTPD_401_UNAUTHORIZED**

enumerator **HTTPD_403_FORBIDDEN**

enumerator **HTTPD_404_NOT_FOUND**

enumerator **HTTPD_405_METHOD_NOT_ALLOWED**

enumerator **HTTPD_408_REQ_TIMEOUT**

enumerator **HTTPD_411_LENGTH_REQUIRED**

enumerator **HTTPD_414_URI_TOO_LONG**

enumerator **HTTPD_431_REQ_HDR_FIELDS_TOO_LARGE**

enumerator **HTTPD_ERR_CODE_MAX**

enum **esp_http_server_event_id_t**

HTTP Server events id.

Values:

enumerator **HTTP_SERVER_EVENT_ERROR**

This event occurs when there are any errors during execution

enumerator **HTTP_SERVER_EVENT_START**

This event occurs when HTTP Server is started

enumerator **HTTP_SERVER_EVENT_ON_CONNECTED**

Once the HTTP Server has been connected to the client, no data exchange has been performed

enumerator **HTTP_SERVER_EVENT_ON_HEADER**

Occurs when receiving each header sent from the client

enumerator **HTTP_SERVER_EVENT_HEADERS_SENT**

After sending all the headers to the client

enumerator **HTTP_SERVER_EVENT_ON_DATA**

Occurs when receiving data from the client

enumerator **HTTP_SERVER_EVENT_SENT_DATA**

Occurs when an ESP HTTP server session is finished

enumerator **HTTP_SERVER_EVENT_DISCONNECTED**

The connection has been disconnected

enumerator **HTTP_SERVER_EVENT_STOP**

This event occurs when HTTP Server is stopped

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.2.10 HTTPS Server

Overview

This component is built on top of [HTTP Server](#). The HTTPS server takes advantage of hook registration functions in the regular HTTP server to provide callback function for SSL session.

All documentation for [HTTP Server](#) applies also to a server you create this way.

Used APIs

The following APIs of [HTTP Server](#) should not be used with [HTTPS Server](#), as they are used internally to handle secure sessions and to maintain internal state:

- "send", "receive" and "pending" callback registration functions - secure socket handling
 - [httpd_sess_set_send_override\(\)](#)
 - [httpd_sess_set_recv_override\(\)](#)
 - [httpd_sess_set_pending_override\(\)](#)
- "transport context" - both global and session
 - [httpd_sess_get_transport_ctx\(\)](#) - returns SSL used for the session
 - [httpd_sess_set_transport_ctx\(\)](#)
 - [httpd_get_global_transport_ctx\(\)](#) - returns the shared SSL context
 - [httpd_config::global_transport_ctx](#)
 - [httpd_config::global_transport_ctx_free_fn](#)
 - [httpd_config::open_fn](#) - used to set up secure sockets

Everything else can be used without limitations.

Usage

Please see the example [protocols/https_server](#) to learn how to set up a secure server.

Basically, all you need is to generate a certificate, embed it into the firmware, and pass the init struct into the start function after the certificate address and lengths are correctly configured in the init struct.

The server can be started with or without SSL by changing a flag in the init struct - [httpd_ssl_config::transport_mode](#). This could be used, e.g., for testing or in trusted environments where you prefer speed over security.

Performance

The initial session setup can take about two seconds, or more with slower clock speed or more verbose logging. Subsequent requests through the open secure socket are much faster (down to under 100 ms).

API Reference

Header File

- [components/esp_https_server/include/esp_https_server.h](#)
- This header file can be included with:

```
#include "esp_https_server.h"
```

- This header file is a part of the API provided by the `esp_https_server` component. To declare that your component depends on `esp_https_server`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_https_server
```

or

```
PRIV_REQUIRES esp_https_server
```

Functions

[esp_err_t](#) **httpd_ssl_start** ([httpd_handle_t](#) *handle, [httpd_ssl_config_t](#) *config)

Create a SSL capable HTTP server (secure mode may be disabled in config)

Parameters

- **config** -- [inout] - server config, must not be const. Does not have to stay valid after calling this function.
- **handle** -- [out] - storage for the server handle, must be a valid pointer

Returns

 success

[esp_err_t](#) **httpd_ssl_stop** ([httpd_handle_t](#) handle)

Stop the server. Blocks until the server is shut down.

Parameters

handle -- [in]

Returns

- ESP_OK: Server stopped successfully
- ESP_ERR_INVALID_ARG: Invalid argument
- ESP_FAIL: Failure to shut down server

Structures

struct **esp_https_server_user_cb_arg**

Callback data struct, contains the ESP-TLS connection handle and the connection state at which the callback is executed.

Public Members

httpd_ssl_user_cb_state_t **user_cb_state**

State of user callback

esp_tls_t ***tls**

ESP-TLS connection handle

struct **httpd_ssl_config**

HTTPS server config struct

Please use HTTPD_SSL_CONFIG_DEFAULT() to initialize it.

Public Members

httpd_config_t **httpd**

Underlying HTTPD server config

Parameters like task stack size and priority can be adjusted here.

const uint8_t ***servercert**

Server certificate

size_t **servercert_len**

Server certificate byte length

const uint8_t ***cacert_pem**

CA certificate ((CA used to sign clients, or client cert itself)

size_t **cacert_len**

CA certificate byte length

const uint8_t ***prvtkey_pem**

Private key

size_t **prvtkey_len**

Private key byte length

bool **use_ecdsa_peripheral**

Use ECDSA peripheral to use private key

uint8_t **ecdsa_key_efuse_blk**

The efuse block where ECDSA key is stored

httpd_ssl_transport_mode_t **transport_mode**

Transport Mode (default secure)

uint16_t **port_secure**

Port used when transport mode is secure (default 443)

uint16_t **port_insecure**

Port used when transport mode is insecure (default 80)

bool **session_tickets**

Enable tls session tickets

bool **use_secure_element**

Enable secure element for server session

esp_https_server_user_cb ***user_cb**

User callback for esp_https_server

void ***ssl_userdata**

user data to add to the ssl context

esp_tls_handshake_callback **cert_select_cb**

Certificate selection callback to use

const char ****alpn_protos**

Application protocols the server supports in order of preference. Used for negotiating during the TLS handshake, first one the client supports is selected. The data structure must live as long as the https server itself!

Macros

HTTPD_SSL_CONFIG_DEFAULT ()

Default config struct init

(http_server default config had to be copied for customization)

Notes:

- port is set when starting the server, according to 'transport_mode'
- one socket uses ~ 40kB RAM with SSL, we reduce the default socket count to 4
- SSL sockets are usually long-lived, closing LRU prevents pool exhaustion DOS
- Stack size may need adjustments depending on the user application

Type Definitions

typedef struct *esp_https_server_user_cb_arg* **esp_https_server_user_cb_arg_t**

Callback data struct, contains the ESP-TLS connection handle and the connection state at which the callback is executed.

typedef void **esp_https_server_user_cb** (*esp_https_server_user_cb_arg_t* *user_cb)

Callback function prototype Can be used to get connection or client information (SSL context) E.g. Client certificate, Socket FD, Connection state, etc.

Param user_cb Callback data struct

typedef struct *httpd_ssl_config* **httpd_ssl_config_t**

Enumerations

enum **httpd_ssl_transport_mode_t**

Values:

enumerator `HTTPD_SSL_TRANSPORT_SECURE`

enumerator `HTTPD_SSL_TRANSPORT_INSECURE`

enum `httpd_ssl_user_cb_state_t`

Indicates the state at which the user callback is executed, i.e at session creation or session close.

Values:

enumerator `HTTPD_SSL_USER_CB_SESS_CREATE`

enumerator `HTTPD_SSL_USER_CB_SESS_CLOSE`

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.2.11 ICMP Echo

Overview

ICMP (Internet Control Message Protocol) is used for diagnostic or control purposes or generated in response to errors in IP operations. The common network util `ping` is implemented based on the ICMP packets with the type field value of 0, also called Echo Reply.

During a ping session, the source host firstly sends out an ICMP echo request packet and wait for an ICMP echo reply with specific times. In this way, it also measures the round-trip time for the messages. After receiving a valid ICMP echo reply, the source host will generate statistics about the IP link layer (e.g., packet loss, elapsed time, etc).

It is common that IoT device needs to check whether a remote server is alive or not. The device should show the warnings to users when it got offline. It can be achieved by creating a ping session and sending or parsing ICMP echo packets periodically.

To make this internal procedure much easier for users, ESP-IDF provides some out-of-box APIs.

Create a New Ping Session To create a ping session, you need to fill in the `esp_ping_config_t` configuration structure firstly, specifying target IP address, interval times, and etc. Optionally, you can also register some callback functions with the `esp_ping_callbacks_t` structure.

Example method to create a new ping session and register callbacks:

```
static void test_on_ping_success(esp_ping_handle_t hdl, void *args)
{
    // optionally, get callback arguments
    // const char* str = (const char*) args;
    // printf("%s\r\n", str); // "foo"
    uint8_t ttl;
    uint16_t seqno;
    uint32_t elapsed_time, recv_len;
    ip_addr_t target_addr;
    esp_ping_get_profile(hdl, ESP_PING_PROF_SEQNO, &seqno, sizeof(seqno));
    esp_ping_get_profile(hdl, ESP_PING_PROF_TTL, &ttl, sizeof(ttl));
    esp_ping_get_profile(hdl, ESP_PING_PROF_IPADDR, &target_addr, sizeof(target_
↵addr));
```

(continues on next page)

(continued from previous page)

```

    esp_ping_get_profile(hdl, ESP_PING_PROF_SIZE, &recv_len, sizeof(recv_len));
    esp_ping_get_profile(hdl, ESP_PING_PROF_TIMEGAP, &elapsed_time, sizeof(elapsed_
↪time));
    printf("%d bytes from %s icmp_seq=%d ttl=%d time=%d ms\n",
           recv_len, inet_ntoa(target_addr.u_addr.ip4), seqno, ttl, elapsed_time);
}

static void test_on_ping_timeout(esp_ping_handle_t hdl, void *args)
{
    uint16_t seqno;
    ip_addr_t target_addr;
    esp_ping_get_profile(hdl, ESP_PING_PROF_SEQNO, &seqno, sizeof(seqno));
    esp_ping_get_profile(hdl, ESP_PING_PROF_IPADDR, &target_addr, sizeof(target_
↪addr));
    printf("From %s icmp_seq=%d timeout\n", inet_ntoa(target_addr.u_addr.ip4), ↪
↪seqno);
}

static void test_on_ping_end(esp_ping_handle_t hdl, void *args)
{
    uint32_t transmitted;
    uint32_t received;
    uint32_t total_time_ms;

    esp_ping_get_profile(hdl, ESP_PING_PROF_REQUEST, &transmitted, ↪
↪sizeof(transmitted));
    esp_ping_get_profile(hdl, ESP_PING_PROF_REPLY, &received, sizeof(received));
    esp_ping_get_profile(hdl, ESP_PING_PROF_DURATION, &total_time_ms, sizeof(total_
↪time_ms));
    printf("%d packets transmitted, %d received, time %dms\n", transmitted, ↪
↪received, total_time_ms);
}

void initialize_ping()
{
    /* convert URL to IP address */
    ip_addr_t target_addr;
    struct addrinfo hint;
    struct addrinfo *res = NULL;
    memset(&hint, 0, sizeof(hint));
    memset(&target_addr, 0, sizeof(target_addr));
    getaddrinfo("www.espressif.com", NULL, &hint, &res);
    struct in_addr addr4 = ((struct sockaddr_in *) (res->ai_addr))->sin_addr;
    inet_addr_to_ip4addr(ip_2_ip4(&target_addr), &addr4);
    freeaddrinfo(res);

    esp_ping_config_t ping_config = ESP_PING_DEFAULT_CONFIG();
    ping_config.target_addr = target_addr;           // target IP address
    ping_config.count = ESP_PING_COUNT_INFINITE;     // ping in infinite mode, esp_
↪ping_stop can stop it

    /* set callback functions */
    esp_ping_callbacks_t cbs;
    cbs.on_ping_success = test_on_ping_success;
    cbs.on_ping_timeout = test_on_ping_timeout;
    cbs.on_ping_end = test_on_ping_end;
    cbs.cb_args = "foo"; // arguments that feeds to all callback functions, can ↪
↪be NULL
    cbs.cb_args = eth_event_group;

    esp_ping_handle_t ping;

```

(continues on next page)

(continued from previous page)

```

    esp_ping_new_session(&ping_config, &cbs, &ping);
}

```

Start and Stop Ping Session You can start and stop ping session with the handle returned by `esp_ping_new_session`. Note that, the ping session does not start automatically after creation. If the ping session is stopped, and restart again, the sequence number in ICMP packets will recount from zero again.

Delete a Ping Session If a ping session will not be used any more, you can delete it with `esp_ping_delete_session`. Please make sure the ping session is in stop state (i.e., you have called `esp_ping_stop` before or the ping session has finished all the procedures) when you call this function.

Get Runtime Statistics As the example code above, you can call `esp_ping_get_profile` to get different runtime statistics of ping session in the callback function.

Application Example

ICMP echo example: [protocols/icmp_echo](#)

API Reference

Header File

- [components/lwip/include/apps/ping/ping_sock.h](#)
- This header file can be included with:

```
#include "ping/ping_sock.h"
```

- This header file is a part of the API provided by the `lwip` component. To declare that your component depends on `lwip`, add the following to your `CMakeLists.txt`:

```
REQUIRES lwip
```

or

```
PRIV_REQUIRES lwip
```

Functions

esp_err_t **esp_ping_new_session** (const *esp_ping_config_t* *config, const *esp_ping_callbacks_t* *cbs, *esp_ping_handle_t* *hdl_out)

Create a ping session.

Parameters

- **config** -- ping configuration
- **cbs** -- a bunch of callback functions invoked by internal ping task
- **hdl_out** -- handle of ping session

Returns

- `ESP_ERR_INVALID_ARG`: invalid parameters (e.g. configuration is null, etc)
- `ESP_ERR_NO_MEM`: out of memory
- `ESP_FAIL`: other internal error (e.g. socket error)
- `ESP_OK`: create ping session successfully, user can take the ping handle to do follow-on jobs

esp_err_t **esp_ping_delete_session** (*esp_ping_handle_t* hdl)

Delete a ping session.

Parameters **hdl** -- handle of ping session

Returns

- ESP_ERR_INVALID_ARG: invalid parameters (e.g. ping handle is null, etc)
- ESP_OK: delete ping session successfully

esp_err_t **esp_ping_start** (*esp_ping_handle_t* hdl)

Start the ping session.

Parameters **hdl** -- handle of ping session

Returns

- ESP_ERR_INVALID_ARG: invalid parameters (e.g. ping handle is null, etc)
- ESP_OK: start ping session successfully

esp_err_t **esp_ping_stop** (*esp_ping_handle_t* hdl)

Stop the ping session.

Parameters **hdl** -- handle of ping session

Returns

- ESP_ERR_INVALID_ARG: invalid parameters (e.g. ping handle is null, etc)
- ESP_OK: stop ping session successfully

esp_err_t **esp_ping_get_profile** (*esp_ping_handle_t* hdl, *esp_ping_profile_t* profile, void *data, uint32_t size)

Get runtime profile of ping session.

Parameters

- **hdl** -- handle of ping session
- **profile** -- type of profile
- **data** -- profile data
- **size** -- profile data size

Returns

- ESP_ERR_INVALID_ARG: invalid parameters (e.g. ping handle is null, etc)
- ESP_ERR_INVALID_SIZE: the actual profile data size doesn't match the "size" parameter
- ESP_OK: get profile successfully

Structures

struct **esp_ping_callbacks_t**

Type of "ping" callback functions.

Public Members

void ***cb_args**

arguments for callback functions

void (***on_ping_success**)(*esp_ping_handle_t* hdl, void *args)

Invoked by internal ping thread when received ICMP echo reply packet.

void (***on_ping_timeout**)(*esp_ping_handle_t* hdl, void *args)

Invoked by internal ping thread when receive ICMP echo reply packet timeout.

void (***on_ping_end**)(*esp_ping_handle_t* hdl, void *args)

Invoked by internal ping thread when a ping session is finished.

struct **esp_ping_config_t**

Type of "ping" configuration.

Public Members

uint32_t **count**

A "ping" session contains count procedures

uint32_t **interval_ms**

Milliseconds between each ping procedure

uint32_t **timeout_ms**

Timeout value (in milliseconds) of each ping procedure

uint32_t **data_size**

Size of the data next to ICMP packet header

int **tos**

Type of Service, a field specified in the IP header

int **ttl**

Time to Live, a field specified in the IP header

ip_addr_t **target_addr**

Target IP address, either IPv4 or IPv6

uint32_t **task_stack_size**

Stack size of internal ping task

uint32_t **task_prio**

Priority of internal ping task

uint32_t **interface**

Netif index, interface=0 means NETIF_NO_INDEX

Macros

ESP_PING_DEFAULT_CONFIG()

Default ping configuration.

ESP_PING_COUNT_INFINITE

Set ping count to zero will ping target infinitely

Type Definitions

typedef void ***esp_ping_handle_t**

Type of "ping" session handle.

Enumerations

enum **esp_ping_profile_t**

Profile of ping session.

Values:

enumerator **ESP_PING_PROF_SEQNO**

Sequence number of a ping procedure

enumerator **ESP_PING_PROF_TOS**

Type of service of a ping procedure

enumerator **ESP_PING_PROF_TTL**

Time to live of a ping procedure

enumerator **ESP_PING_PROF_REQUEST**

Number of request packets sent out

enumerator **ESP_PING_PROF_REPLY**

Number of reply packets received

enumerator **ESP_PING_PROF_IPADDR**

IP address of replied target

enumerator **ESP_PING_PROF_SIZE**

Size of received packet

enumerator **ESP_PING_PROF_TIMEGAP**

Elapsed time between request and reply packet

enumerator **ESP_PING_PROF_DURATION**

Elapsed time of the whole ping session

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.2.12 mDNS Service

mDNS is a multicast UDP service that is used to provide local network service and host discovery.

The ESP-IDF component mDNS has been moved from ESP-IDF since version v5.0 to a separate repository:

- [mDNS component on GitHub](#)

To add mDNS component in your project, please run `idf.py add-dependency espressif/mdns`.

Hosted Documentation

The documentation can be found on the link below:

- [mDNS documentation](#)

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.2.13 Mbed TLS

[Mbed TLS](#) is a C library that implements cryptographic primitives, X.509 certificate manipulation and the SSL/TLS and DTLS protocols. Its small code footprint makes it suitable for embedded systems.

Note: ESP-IDF uses a [fork](#) of Mbed TLS which includes a few patches (related to hardware routines of certain modules like `bignum` (MPI) and ECC) over vanilla Mbed TLS.

Mbed TLS supports TLS 1.2, TLS 1.3 and DTLS 1.2 communication by providing the following:

- TCP/IP communication functions: listen, connect, accept, read/write.
- SSL/TLS communication functions: init, handshake, read/write.
- X.509 functions: CRT, CRL and key handling
- Random number generation
- Hashing
- Encryption/decryption

Note: Mbed TLS v3.x.x series supports only TLS 1.2 and TLS 1.3 protocols. Support for SSL 3.0, TLS 1.0/1.1 and DTLS 1.0 has been removed (deprecated). TLS 1.3 is fully supported starting Mbed TLS v3.6.0 release, before this release some features were still in experimental state. Please refer to [Mbed TLS ChangeLog](#) for more details.

Mbed TLS Documentation

For Mbed TLS documentation please refer to the following (upstream) pointers:

- [API Reference](#)
- [Knowledge Base](#)

Mbed TLS Support in ESP-IDF

Please find the information about the Mbed TLS versions presented in different branches of ESP-IDF [here](#).

Note: Please refer the [Mbed TLS](#) to migrate from Mbed TLS version 2.x to version 3.0 or greater.

Application Examples

Examples in ESP-IDF use [ESP-TLS](#) which provides a simplified API interface for accessing the commonly used TLS functionality.

Refer to the examples [protocols/https_server/simple](#) (Simple HTTPS server) and [protocols/https_request](#) (Make HTTPS requests) for more information.

If the Mbed TLS API is to be used directly, refer to the example [protocols/https_mbedtls](#).

Alternatives

ESP-TLS acts as an abstraction layer over the underlying SSL/TLS library and thus has an option to use Mbed TLS or wolfSSL as the underlying library. By default, only Mbed TLS is available and used in ESP-IDF whereas wolfSSL is available publicly at <<https://github.com/espressif/esp-wolfSSL>> with the upstream submodule pointer.

Please refer to *ESP-TLS: Underlying SSL/TLS Library Options* docs for more information on this and comparison of Mbed TLS and wolfSSL.

Important Config Options

Following is a brief list of important config options accessible at `Component Config -> mbedTLS`. The full list of config options can be found [here](#).

- *CONFIG_MBEDTLS_SSL_PROTO_TLS1_2*: Support for TLS 1.2
- *CONFIG_MBEDTLS_SSL_PROTO_TLS1_3*: Support for TLS 1.3
- *CONFIG_MBEDTLS_CERTIFICATE_BUNDLE*: Support for trusted root certificate bundle (more about this: *ESP x509 Certificate Bundle*)
- *CONFIG_MBEDTLS_CLIENT_SSL_SESSION_TICKETS*: Support for TLS Session Resumption: Client session tickets
- *CONFIG_MBEDTLS_SERVER_SSL_SESSION_TICKETS*: Support for TLS Session Resumption: Server session tickets
- *CONFIG_MBEDTLS_HARDWARE_SHA*: Support for hardware SHA acceleration
- *CONFIG_MBEDTLS_HARDWARE_MPI*: Support for hardware MPI (bignum) acceleration
- *CONFIG_MBEDTLS_HARDWARE_ECC*: Support for hardware ECC acceleration

Performance and Memory Tweaks

Reducing Heap Usage The following table shows typical memory usage with different configs when the [protocols/https_request](#) example (with Server Validation enabled) was run with Mbed TLS as the SSL/TLS library.

Mbed TLS Test	Related Configs	Heap Usage (approx.)
Default	NA	42196 B
Enable SSL Variable Length	<i>CONFIG_MBEDTLS_SSL_VARIABLE_BUFFER_LENGTH</i>	42120 B
Disable Keep Peer Certificate	<i>CONFIG_MBEDTLS_SSL_KEEP_PEER_CERTIFICATE</i>	38533 B
Enable Dynamic TX/RX Buffer	<i>CONFIG_MBEDTLS_DYNAMIC_BUFFER</i> <i>CONFIG_MBEDTLS_DYNAMIC_FREE_CONFIG_DATA</i> <i>CONFIG_MBEDTLS_DYNAMIC_FREE_CA_CERT</i>	22013 B

Note: These values are subject to change with change in configuration options and versions of Mbed TLS.

Reducing Binary Size Under `Component Config -> mbedTLS`, there are multiple Mbed TLS features which are enabled by default but can be disabled if not needed to save code size. More information can be about this can be found in *Minimizing Binary Size* docs.

Code examples for this API section are provided in the [protocols](#) directory of ESP-IDF examples.

2.2.14 IP Network Layer

Documentation for IP Network Layer protocols (below the Application Protocol layer) are provided in *Networking APIs*.

2.3 Error Codes Reference

This section lists various error code constants defined in ESP-IDF.

For general information about error codes in ESP-IDF, see [Error Handling](#).

[ESP_FAIL](#) (-1): Generic esp_err_t code indicating failure

[ESP_OK](#) (0): esp_err_t value indicating success (no error)

[ESP_ERR_NO_MEM](#) (0x101): Out of memory

[ESP_ERR_INVALID_ARG](#) (0x102): Invalid argument

[ESP_ERR_INVALID_STATE](#) (0x103): Invalid state

[ESP_ERR_INVALID_SIZE](#) (0x104): Invalid size

[ESP_ERR_NOT_FOUND](#) (0x105): Requested resource not found

[ESP_ERR_NOT_SUPPORTED](#) (0x106): Operation or feature not supported

[ESP_ERR_TIMEOUT](#) (0x107): Operation timed out

[ESP_ERR_INVALID_RESPONSE](#) (0x108): Received response was invalid

[ESP_ERR_INVALID_CRC](#) (0x109): CRC or checksum was invalid

[ESP_ERR_INVALID_VERSION](#) (0x10a): Version was invalid

[ESP_ERR_INVALID_MAC](#) (0x10b): MAC address was invalid

[ESP_ERR_NOT_FINISHED](#) (0x10c): Operation has not fully completed

[ESP_ERR_NOT_ALLOWED](#) (0x10d): Operation is not allowed

[ESP_ERR_NVS_BASE](#) (0x1100): Starting number of error codes

[ESP_ERR_NVS_NOT_INITIALIZED](#) (0x1101): The storage driver is not initialized

[ESP_ERR_NVS_NOT_FOUND](#) (0x1102): A requested entry couldn't be found or namespace doesn't exist yet and mode is NVS_READONLY

[ESP_ERR_NVS_TYPE_MISMATCH](#) (0x1103): The type of set or get operation doesn't match the type of value stored in NVS

[ESP_ERR_NVS_READ_ONLY](#) (0x1104): Storage handle was opened as read only

[ESP_ERR_NVS_NOT_ENOUGH_SPACE](#) (0x1105): There is not enough space in the underlying storage to save the value

[ESP_ERR_NVS_INVALID_NAME](#) (0x1106): Namespace name doesn't satisfy constraints

[ESP_ERR_NVS_INVALID_HANDLE](#) (0x1107): Handle has been closed or is NULL

[ESP_ERR_NVS_REMOVE_FAILED](#) (0x1108): The value wasn't updated because flash write operation has failed. The value was written however, and update will be finished after re-initialization of nvs, provided that flash operation doesn't fail again.

[ESP_ERR_NVS_KEY_TOO_LONG](#) (0x1109): Key name is too long

[ESP_ERR_NVS_PAGE_FULL](#) (0x110a): Internal error; never returned by nvs API functions

[ESP_ERR_NVS_INVALID_STATE](#) (0x110b): NVS is in an inconsistent state due to a previous error. Call nvs_flash_init and nvs_open again, then retry.

[ESP_ERR_NVS_INVALID_LENGTH](#) (0x110c): String or blob length is not sufficient to store data

[ESP_ERR_NVS_NO_FREE_PAGES](#) (0x110d): NVS partition doesn't contain any empty pages. This may happen if NVS partition was truncated. Erase the whole partition and call nvs_flash_init again.

ESP_ERR_NVS_VALUE_TOO_LONG (0x110e): Value doesn't fit into the entry or string or blob length is longer than supported by the implementation

ESP_ERR_NVS_PART_NOT_FOUND (0x110f): Partition with specified name is not found in the partition table

ESP_ERR_NVS_NEW_VERSION_FOUND (0x1110): NVS partition contains data in new format and cannot be recognized by this version of code

ESP_ERR_NVS_XTS_ENCR_FAILED (0x1111): XTS encryption failed while writing NVS entry

ESP_ERR_NVS_XTS_DECR_FAILED (0x1112): XTS decryption failed while reading NVS entry

ESP_ERR_NVS_XTS_CFG_FAILED (0x1113): XTS configuration setting failed

ESP_ERR_NVS_XTS_CFG_NOT_FOUND (0x1114): XTS configuration not found

ESP_ERR_NVS_ENCR_NOT_SUPPORTED (0x1115): NVS encryption is not supported in this version

ESP_ERR_NVS_KEYS_NOT_INITIALIZED (0x1116): NVS key partition is uninitialized

ESP_ERR_NVS_CORRUPT_KEY_PART (0x1117): NVS key partition is corrupt

ESP_ERR_NVS_CONTENT_DIFFERS (0x1118): Internal error; never returned by nvs API functions. NVS key is different in comparison

ESP_ERR_NVS_WRONG_ENCRYPTION (0x1119): NVS partition is marked as encrypted with generic flash encryption. This is forbidden since the NVS encryption works differently.

ESP_ERR_ULP_BASE (0x1200): Offset for ULP-related error codes

ESP_ERR_ULP_SIZE_TOO_BIG (0x1201): Program doesn't fit into RTC memory reserved for the ULP

ESP_ERR_ULP_INVALID_LOAD_ADDR (0x1202): Load address is outside of RTC memory reserved for the ULP

ESP_ERR_ULP_DUPLICATE_LABEL (0x1203): More than one label with the same number was defined

ESP_ERR_ULP_UNDEFINED_LABEL (0x1204): Branch instructions references an undefined label

ESP_ERR_ULP_BRANCH_OUT_OF_RANGE (0x1205): Branch target is out of range of B instruction (try replacing with BX)

ESP_ERR_OTA_BASE (0x1500): Base error code for ota_ops api

ESP_ERR_OTA_PARTITION_CONFLICT (0x1501): Error if request was to write or erase the current running partition

ESP_ERR_OTA_SELECT_INFO_INVALID (0x1502): Error if OTA data partition contains invalid content

ESP_ERR_OTA_VALIDATE_FAILED (0x1503): Error if OTA app image is invalid

ESP_ERR_OTA_SMALL_SEC_VER (0x1504): Error if the firmware has a secure version less than the running firmware.

ESP_ERR_OTA_ROLLBACK_FAILED (0x1505): Error if flash does not have valid firmware in passive partition and hence rollback is not possible

ESP_ERR_OTA_ROLLBACK_INVALID_STATE (0x1506): Error if current active firmware is still marked in pending validation state (ESP_OTA_IMG_PENDING_VERIFY), essentially first boot of firmware image post upgrade and hence firmware upgrade is not possible

ESP_ERR_EFUSE (0x1600): Base error code for efuse api.

ESP_OK_EFUSE_CNT (0x1601): OK the required number of bits is set.

ESP_ERR_EFUSE_CNT_IS_FULL (0x1602): Error field is full.

ESP_ERR_EFUSE_REPEATED_PROG (0x1603): Error repeated programming of programmed bits is strictly forbidden.

ESP_ERR_CODING (0x1604): Error while a encoding operation.

ESP_ERR_NOT_ENOUGH_UNUSED_KEY_BLOCKS (0x1605): Error not enough unused key blocks available

ESP_ERR_DAMAGED_READING (0x1606): Error. Burn or reset was done during a reading operation leads to damage read data. This error is internal to the efuse component and not returned by any public API.

ESP_ERR_IMAGE_BASE (0x2000)

ESP_ERR_IMAGE_FLASH_FAIL (0x2001)

ESP_ERR_IMAGE_INVALID (0x2002)

ESP_ERR_WIFI_BASE (0x3000): Starting number of WiFi error codes

ESP_ERR_WIFI_NOT_INIT (0x3001): WiFi driver was not installed by esp_wifi_init

ESP_ERR_WIFI_NOT_STARTED (0x3002): WiFi driver was not started by esp_wifi_start

ESP_ERR_WIFI_NOT_STOPPED (0x3003): WiFi driver was not stopped by esp_wifi_stop

ESP_ERR_WIFI_IF (0x3004): WiFi interface error

ESP_ERR_WIFI_MODE (0x3005): WiFi mode error

ESP_ERR_WIFI_STATE (0x3006): WiFi internal state error

ESP_ERR_WIFI_CONN (0x3007): WiFi internal control block of station or soft-AP error

ESP_ERR_WIFI_NVS (0x3008): WiFi internal NVS module error

ESP_ERR_WIFI_MAC (0x3009): MAC address is invalid

ESP_ERR_WIFI_SSID (0x300a): SSID is invalid

ESP_ERR_WIFI_PASSWORD (0x300b): Password is invalid

ESP_ERR_WIFI_TIMEOUT (0x300c): Timeout error

ESP_ERR_WIFI_WAKE_FAIL (0x300d): WiFi is in sleep state(RF closed) and wakeup fail

ESP_ERR_WIFI_WOULD_BLOCK (0x300e): The caller would block

ESP_ERR_WIFI_NOT_CONNECT (0x300f): Station still in disconnect status

ESP_ERR_WIFI_POST (0x3012): Failed to post the event to WiFi task

ESP_ERR_WIFI_INIT_STATE (0x3013): Invalid WiFi state when init/deinit is called

ESP_ERR_WIFI_STOP_STATE (0x3014): Returned when WiFi is stopping

ESP_ERR_WIFI_NOT_ASSOC (0x3015): The WiFi connection is not associated

ESP_ERR_WIFI_TX_DISALLOW (0x3016): The WiFi TX is disallowed

ESP_ERR_WIFI_TWT_FULL (0x3017): no available flow id

ESP_ERR_WIFI_TWT_SETUP_TIMEOUT (0x3018): Timeout of receiving twt setup response frame, timeout times can be set during twt setup

ESP_ERR_WIFI_TWT_SETUP_TXFAIL (0x3019): TWT setup frame tx failed

ESP_ERR_WIFI_TWT_SETUP_REJECT (0x301a): The twt setup request was rejected by the AP

ESP_ERR_WIFI_DISCARD (0x301b): Discard frame

ESP_ERR_WIFI_ROC_IN_PROGRESS (0x301c): ROC op is in progress

ESP_ERR_WIFI_REGISTRAR (0x3033): WPS registrar is not supported

ESP_ERR_WIFI_WPS_TYPE (0x3034): WPS type error

ESP_ERR_WIFI_WPS_SM (0x3035): WPS state machine is not initialized

ESP_ERR_ESPNOW_BASE (0x3064): ESPNOW error number base.

ESP_ERR_ESPNOW_NOT_INIT (0x3065): ESPNOW is not initialized.

ESP_ERR_ESPNOW_ARG (0x3066): Invalid argument

ESP_ERR_ESPNOW_NO_MEM (0x3067): Out of memory

ESP_ERR_ESPNOW_FULL (**0x3068**): ESPNOW peer list is full

ESP_ERR_ESPNOW_NOT_FOUND (**0x3069**): ESPNOW peer is not found

ESP_ERR_ESPNOW_INTERNAL (**0x306a**): Internal error

ESP_ERR_ESPNOW_EXIST (**0x306b**): ESPNOW peer has existed

ESP_ERR_ESPNOW_IF (**0x306c**): Interface error

ESP_ERR_ESPNOW_CHAN (**0x306d**): Channel error

ESP_ERR_DPP_FAILURE (**0x3097**): Generic failure during DPP Operation

ESP_ERR_DPP_TX_FAILURE (**0x3098**): DPP Frame Tx failed OR not Acked

ESP_ERR_DPP_INVALID_ATTR (**0x3099**): Encountered invalid DPP Attribute

ESP_ERR_DPP_AUTH_TIMEOUT (**0x309a**): DPP Auth response was not recieved in time

[ESP_ERR_MESH_BASE](#) (**0x4000**): Starting number of MESH error codes

ESP_ERR_MESH_WIFI_NOT_START (**0x4001**)

ESP_ERR_MESH_NOT_INIT (**0x4002**)

ESP_ERR_MESH_NOT_CONFIG (**0x4003**)

ESP_ERR_MESH_NOT_START (**0x4004**)

ESP_ERR_MESH_NOT_SUPPORT (**0x4005**)

ESP_ERR_MESH_NOT_ALLOWED (**0x4006**)

ESP_ERR_MESH_NO_MEMORY (**0x4007**)

ESP_ERR_MESH_ARGUMENT (**0x4008**)

ESP_ERR_MESH_EXCEED_MTU (**0x4009**)

ESP_ERR_MESH_TIMEOUT (**0x400a**)

ESP_ERR_MESH_DISCONNECTED (**0x400b**)

ESP_ERR_MESH_QUEUE_FAIL (**0x400c**)

ESP_ERR_MESH_QUEUE_FULL (**0x400d**)

ESP_ERR_MESH_NO_PARENT_FOUND (**0x400e**)

ESP_ERR_MESH_NO_ROUTE_FOUND (**0x400f**)

ESP_ERR_MESH_OPTION_NULL (**0x4010**)

ESP_ERR_MESH_OPTION_UNKNOWN (**0x4011**)

ESP_ERR_MESH_XON_NO_WINDOW (**0x4012**)

ESP_ERR_MESH_INTERFACE (**0x4013**)

ESP_ERR_MESH_DISCARD_DUPLICATE (**0x4014**)

ESP_ERR_MESH_DISCARD (**0x4015**)

ESP_ERR_MESH_VOTING (**0x4016**)

ESP_ERR_MESH_XMIT (**0x4017**)

ESP_ERR_MESH_QUEUE_READ (**0x4018**)

ESP_ERR_MESH_PS (**0x4019**)

ESP_ERR_MESH_RECV_RELEASE (**0x401a**)

[ESP_ERR_ESP_NETIF_BASE](#) (**0x5000**)

[ESP_ERR_ESP_NETIF_INVALID_PARAMS](#) (**0x5001**)

ESP_ERR_ESP_NETIF_IF_NOT_READY (0x5002)

ESP_ERR_ESP_NETIF_DHCP_START_FAILED (0x5003)

ESP_ERR_ESP_NETIF_DHCP_ALREADY_STARTED (0x5004)

ESP_ERR_ESP_NETIF_DHCP_ALREADY_STOPPED (0x5005)

ESP_ERR_ESP_NETIF_NO_MEM (0x5006)

ESP_ERR_ESP_NETIF_DHCP_NOT_STOPPED (0x5007)

ESP_ERR_ESP_NETIF_DRIVER_ATTACH_FAILED (0x5008)

ESP_ERR_ESP_NETIF_INIT_FAILED (0x5009)

ESP_ERR_ESP_NETIF_DNS_NOT_CONFIGURED (0x500a)

ESP_ERR_ESP_NETIF_MLD6_FAILED (0x500b)

ESP_ERR_ESP_NETIF_IP6_ADDR_FAILED (0x500c)

ESP_ERR_ESP_NETIF_DHCP_START_FAILED (0x500d)

ESP_ERR_FLASH_BASE (0x6000): Starting number of flash error codes

ESP_ERR_FLASH_OP_FAIL (0x6001)

ESP_ERR_FLASH_OP_TIMEOUT (0x6002)

ESP_ERR_FLASH_NOT_INITIALISED (0x6003)

ESP_ERR_FLASH_UNSUPPORTED_HOST (0x6004)

ESP_ERR_FLASH_UNSUPPORTED_CHIP (0x6005)

ESP_ERR_FLASH_PROTECTED (0x6006)

ESP_ERR_HTTP_BASE (0x7000): Starting number of HTTP error codes

ESP_ERR_HTTP_MAX_REDIRECT (0x7001): The error exceeds the number of HTTP redirects

ESP_ERR_HTTP_CONNECT (0x7002): Error open the HTTP connection

ESP_ERR_HTTP_WRITE_DATA (0x7003): Error write HTTP data

ESP_ERR_HTTP_FETCH_HEADER (0x7004): Error read HTTP header from server

ESP_ERR_HTTP_INVALID_TRANSPORT (0x7005): There are no transport support for the input scheme

ESP_ERR_HTTP_CONNECTING (0x7006): HTTP connection hasn't been established yet

ESP_ERR_HTTP_EAGAIN (0x7007): Mapping of errno EAGAIN to esp_err_t

ESP_ERR_HTTP_CONNECTION_CLOSED (0x7008): Read FIN from peer and the connection closed

ESP_ERR_ESP_TLS_BASE (0x8000): Starting number of ESP-TLS error codes

ESP_ERR_ESP_TLS_CANNOT_RESOLVE_HOSTNAME (0x8001): Error if hostname couldn't be resolved upon tls connection

ESP_ERR_ESP_TLS_CANNOT_CREATE_SOCKET (0x8002): Failed to create socket

ESP_ERR_ESP_TLS_UNSUPPORTED_PROTOCOL_FAMILY (0x8003): Unsupported protocol family

ESP_ERR_ESP_TLS_FAILED_CONNECT_TO_HOST (0x8004): Failed to connect to host

ESP_ERR_ESP_TLS_SOCKET_SETOPT_FAILED (0x8005): failed to set/get socket option

ESP_ERR_ESP_TLS_CONNECTION_TIMEOUT (0x8006): new connection in esp_tls_low_level_conn connection timed out

ESP_ERR_ESP_TLS_SE_FAILED (0x8007)

ESP_ERR_ESP_TLS_TCP_CLOSED_FIN (0x8008)

ESP_ERR_MBEDTLS_CERT_PARTLY_OK (0x8010): mbedtls parse certificates was partly successful

ESP_ERR_MBEDTLS_CTR_DRBG_SEED_FAILED (0x8011): mbedtls api returned error

ESP_ERR_MBEDTLS_SSL_SET_HOSTNAME_FAILED (0x8012): mbedtls api returned error

ESP_ERR_MBEDTLS_SSL_CONFIG_DEFAULTS_FAILED (0x8013): mbedtls api returned error

ESP_ERR_MBEDTLS_SSL_CONF_ALPN_PROTOCOLS_FAILED (0x8014): mbedtls api returned error

ESP_ERR_MBEDTLS_X509_CRT_PARSE_FAILED (0x8015): mbedtls api returned error

ESP_ERR_MBEDTLS_SSL_CONF_OWN_CERT_FAILED (0x8016): mbedtls api returned error

ESP_ERR_MBEDTLS_SSL_SETUP_FAILED (0x8017): mbedtls api returned error

ESP_ERR_MBEDTLS_SSL_WRITE_FAILED (0x8018): mbedtls api returned error

ESP_ERR_MBEDTLS_PK_PARSE_KEY_FAILED (0x8019): mbedtls api returned failed

ESP_ERR_MBEDTLS_SSL_HANDSHAKE_FAILED (0x801a): mbedtls api returned failed

ESP_ERR_MBEDTLS_SSL_CONF_PSK_FAILED (0x801b): mbedtls api returned failed

ESP_ERR_MBEDTLS_SSL_TICKET_SETUP_FAILED (0x801c): mbedtls api returned failed

ESP_ERR_WOLFSSL_SSL_SET_HOSTNAME_FAILED (0x8031): wolfSSL api returned error

ESP_ERR_WOLFSSL_SSL_CONF_ALPN_PROTOCOLS_FAILED (0x8032): wolfSSL api returned error

ESP_ERR_WOLFSSL_CERT_VERIFY_SETUP_FAILED (0x8033): wolfSSL api returned error

ESP_ERR_WOLFSSL_KEY_VERIFY_SETUP_FAILED (0x8034): wolfSSL api returned error

ESP_ERR_WOLFSSL_SSL_HANDSHAKE_FAILED (0x8035): wolfSSL api returned failed

ESP_ERR_WOLFSSL_CTX_SETUP_FAILED (0x8036): wolfSSL api returned failed

ESP_ERR_WOLFSSL_SSL_SETUP_FAILED (0x8037): wolfSSL api returned failed

ESP_ERR_WOLFSSL_SSL_WRITE_FAILED (0x8038): wolfSSL api returned failed

ESP_ERR_HTTPS_OTA_BASE (0x9000)

ESP_ERR_HTTPS_OTA_IN_PROGRESS (0x9001)

ESP_ERR_PING_BASE (0xa000)

ESP_ERR_PING_INVALID_PARAMS (0xa001)

ESP_ERR_PING_NO_MEM (0xa002)

ESP_ERR_HTTPD_BASE (0xb000): Starting number of HTTPD error codes

ESP_ERR_HTTPD_HANDLERS_FULL (0xb001): All slots for registering URI handlers have been consumed

ESP_ERR_HTTPD_HANDLER_EXISTS (0xb002): URI handler with same method and target URI already registered

ESP_ERR_HTTPD_INVALID_REQ (0xb003): Invalid request pointer

ESP_ERR_HTTPD_RESULT_TRUNC (0xb004): Result string truncated

ESP_ERR_HTTPD_RESP_HDR (0xb005): Response header field larger than supported

ESP_ERR_HTTPD_RESP_SEND (0xb006): Error occurred while sending response packet

ESP_ERR_HTTPD_ALLOC_MEM (0xb007): Failed to dynamically allocate memory for resource

ESP_ERR_HTTPD_TASK (0xb008): Failed to launch server task/thread

ESP_ERR_HW_CRYPTO_BASE (0xc000): Starting number of HW cryptography module error codes

ESP_ERR_HW_CRYPTO_DS_HMAC_FAIL (0xc001): HMAC peripheral problem

ESP_ERR_HW_CRYPTO_DS_INVALID_KEY (0xc002)

ESP_ERR_HW_CRYPTO_DS_INVALID_DIGEST (0xc004)

ESP_ERR_HW_CRYPTODS_INVALID_PADDING (0xc005)

[ESP_ERR_MEMPROT_BASE \(0xd000\)](#): Starting number of Memory Protection API error codes

ESP_ERR_MEMPROT_MEMORY_TYPE_INVALID (0xd001)

ESP_ERR_MEMPROT_SPLIT_ADDR_INVALID (0xd002)

ESP_ERR_MEMPROT_SPLIT_ADDR_OUT_OF_RANGE (0xd003)

ESP_ERR_MEMPROT_SPLIT_ADDR_UNALIGNED (0xd004)

ESP_ERR_MEMPROT_UNIMGMT_BLOCK_INVALID (0xd005)

ESP_ERR_MEMPROT_WORLD_INVALID (0xd006)

ESP_ERR_MEMPROT_AREA_INVALID (0xd007)

ESP_ERR_MEMPROT_CPUID_INVALID (0xd008)

ESP_ERR_TCP_TRANSPORT_BASE (0xe000): Starting number of TCP Transport error codes

ESP_ERR_TCP_TRANSPORT_CONNECTION_TIMEOUT (0xe001): Connection has timed out

ESP_ERR_TCP_TRANSPORT_CONNECTION_CLOSED_BY_FIN (0xe002): Read FIN from peer and the connection has closed (in a clean way)

ESP_ERR_TCP_TRANSPORT_CONNECTION_FAILED (0xe003): Failed to connect to the peer

ESP_ERR_TCP_TRANSPORT_NO_MEM (0xe004): Memory allocation failed

[ESP_ERR_NVS_SEC_BASE \(0xf000\)](#): Starting number of error codes

[ESP_ERR_NVS_SEC_HMAC_KEY_NOT_FOUND \(0xf001\)](#): HMAC Key required to generate the NVS encryption keys not found

[ESP_ERR_NVS_SEC_HMAC_KEY_BLK_ALREADY_USED \(0xf002\)](#): Provided eFuse block for HMAC key generation is already in use

[ESP_ERR_NVS_SEC_HMAC_KEY_GENERATION_FAILED \(0xf003\)](#): Failed to generate/write the HMAC key to eFuse

[ESP_ERR_NVS_SEC_HMAC_XTS_KEYS_DERIV_FAILED \(0xf004\)](#): Failed to derive the NVS encryption keys based on the HMAC-based scheme

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.4 Networking APIs

2.4.1 Ethernet

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Ethernet

Overview ESP-IDF provides a set of consistent and flexible APIs to support both internal Ethernet MAC (EMAC) controller and external SPI-Ethernet modules.

This programming guide is split into the following sections:

1. *Basic Ethernet Concepts*
2. *Configure MAC and PHY*
3. *Connect Driver to TCP/IP Stack*
4. *Misc Control of Ethernet Driver*

Basic Ethernet Concepts Ethernet is an asynchronous Carrier Sense Multiple Access with Collision Detect (CSMA/CD) protocol/interface. It is generally not well suited for low-power applications. However, with ubiquitous deployment, internet connectivity, high data rates, and limitless-range expandability, Ethernet can accommodate nearly all wired communications.

Normal IEEE 802.3 compliant Ethernet frames are between 64 and 1518 bytes in length. They are made up of five or six different fields: a destination MAC address (DA), a source MAC address (SA), a type/length field, a data payload, an optional padding field and a Cyclic Redundancy Check (CRC). Additionally, when transmitted on the Ethernet medium, a 7-byte preamble field and Start-of-Frame (SOF) delimiter byte are appended to the beginning of the Ethernet packet.

Thus the traffic on the twist-pair cabling appears as shown below:




Fig. 1: Ethernet Data Frame Format

Preamble and Start-of-Frame Delimiter The preamble contains seven bytes of 55H. It allows the receiver to lock onto the stream of data before the actual frame arrives.

The Start-of-Frame Delimiter (SFD) is a binary sequence 10101011 (as seen on the physical medium). It is sometimes considered to be part of the preamble.

When transmitting and receiving data, the preamble and SFD bytes will be automatically generated or stripped from the packets.

Destination Address The destination address field contains a 6-byte length MAC address of the device that the packet is directed to. If the Least Significant bit in the first byte of the MAC address is set, the address is a multicast destination. For example, 01-00-00-00-F0-00 and 33-45-67-89-AB-CD are multi-cast addresses, while 00-00-00-00-F0-00 and 32-45-67-89-AB-CD are not.

Packets with multi-cast destination addresses are designed to arrive and be important to a selected group of Ethernet nodes. If the destination address field is the reserved multicast address, i.e., FF-FF-FF-FF-FF-FF, the packet is a broadcast packet and it will be directed to everyone sharing the network. If the Least Significant bit in the first byte of the MAC address is clear, the address is a unicast address and will be designed for usage by only the addressed node.

Normally the EMAC controller incorporates receive filters which can be used to discard or accept packets with multi-cast, broadcast and/or unicast destination addresses. When transmitting packets, the host controller is responsible for writing the desired destination address into the transmit buffer.

Source Address The source address field contains a 6-byte length MAC address of the node which created the Ethernet packet. Users of Ethernet must generate a unique MAC address for each controller used. MAC addresses consist of two portions. The first three bytes are known as the Organizationally Unique Identifier (OUI). OUIs are distributed by the IEEE. The last three bytes are address bytes at the discretion of the company that purchased the OUI. For more information about MAC Address used in ESP-IDF, please see [MAC Address Allocation](#).

When transmitting packets, the assigned source MAC address must be written into the transmit buffer by the host controller.

Type/Length The type/length field is a 2-byte field. If the value in this field is ≤ 1500 (decimal), it is considered a length field and it specifies the amount of non-padding data which follows in the data field. If the value is ≥ 1536 , it represents the protocol the following packet data belongs to. The followings are the most common type values:

- IPv4 = 0800H
- IPv6 = 86DDH
- ARP = 0806H

Users implementing proprietary networks may choose to treat this field as a length field, while applications implementing protocols such as the Internet Protocol (IP) or Address Resolution Protocol (ARP), should program this field with the appropriate type defined by the protocol's specification when transmitting packets.

Payload The payload field is a variable length field, anywhere from 0 to 1500 bytes. Larger data packets violates Ethernet standards and will be dropped by most Ethernet nodes.

This field contains the client data, such as an IP datagram.

Padding and FCS The padding field is a variable length field added to meet the IEEE 802.3 specification requirements when small data payloads are used.

The DA, SA, type, payload, and padding of an Ethernet packet must be no smaller than 60 bytes in total. If the required 4-byte FCS field is added, packets must be no smaller than 64 bytes. If the payload field is less than 46-byte long, a padding field is required.

The FCS field is a 4-byte field that contains an industry-standard 32-bit CRC calculated with the data from the DA, SA, type, payload, and padding fields. Given the complexity of calculating a CRC, the hardware normally automatically generates a valid CRC and transmit it. Otherwise, the host controller must generate the CRC and place it in the transmit buffer.

Normally, the host controller does not need to concern itself with padding and the CRC which the hardware EMAC will also be able to automatically generate when transmitting and verify when receiving. However, the padding and

CRC fields will be written into the receive buffer when packets arrive, so they may be evaluated by the host controller if needed.

Note: Besides the basic data frame described above, there are two other common frame types in 10/100 Mbps Ethernet: control frames and VLAN-tagged frames. They are not supported in ESP-IDF.

Configure MAC and PHY The Ethernet driver is composed of two parts: MAC and PHY.

You need to set up the necessary parameters for MAC and PHY respectively based on your Ethernet board design, and then combine the two together to complete the driver installation.

Configuration for MAC is described in `eth_mac_config_t`, including:

- `eth_mac_config_t::sw_reset_timeout_ms`: software reset timeout value, in milliseconds. Typically, MAC reset should be finished within 100 ms.
- `eth_mac_config_t::rx_task_stack_size` and `eth_mac_config_t::rx_task_prio`: the MAC driver creates a dedicated task to process incoming packets. These two parameters are used to set the stack size and priority of the task.
- `eth_mac_config_t::flags`: specifying extra features that the MAC driver should have, it could be useful in some special situations. The value of this field can be OR'd with macros prefixed with `ETH_MAC_FLAG_`. For example, if the MAC driver should work when the cache is disabled, then you should configure this field with `ETH_MAC_FLAG_WORK_WITH_CACHE_DISABLE`.

Configuration for PHY is described in `eth_phy_config_t`, including:

- `eth_phy_config_t::phy_addr`: multiple PHY devices can share the same SMI bus, so each PHY needs a unique address. Usually, this address is configured during hardware design by pulling up/down some PHY strapping pins. You can set the value from 0 to 15 based on your Ethernet board. Especially, if the SMI bus is shared by only one PHY device, setting this value to -1 can enable the driver to detect the PHY address automatically.
- `eth_phy_config_t::reset_timeout_ms`: reset timeout value, in milliseconds. Typically, PHY reset should be finished within 100 ms.
- `eth_phy_config_t::autonego_timeout_ms`: auto-negotiation timeout value, in milliseconds. The Ethernet driver starts negotiation with the peer Ethernet node automatically, to determine to duplex and speed mode. This value usually depends on the ability of the PHY device on your board.
- `eth_phy_config_t::reset_gpio_num`: if your board also connects the PHY reset pin to one of the GPIO, then set it here. Otherwise, set this field to -1.

ESP-IDF provides a default configuration for MAC and PHY in macro `ETH_MAC_DEFAULT_CONFIG` and `ETH_PHY_DEFAULT_CONFIG`.

Create MAC and PHY Instance The Ethernet driver is implemented in an Object-Oriented style. Any operation on MAC and PHY should be based on the instance of the two.

SPI-Ethernet Module

```
eth_mac_config_t mac_config = ETH_MAC_DEFAULT_CONFIG();           // apply default_
↪common MAC configuration
eth_phy_config_t phy_config = ETH_PHY_DEFAULT_CONFIG();           // apply default PHY_
↪configuration
phy_config.phy_addr = CONFIG_EXAMPLE_ETH_PHY_ADDR;                 // alter the PHY_
↪address according to your board design
phy_config.reset_gpio_num = CONFIG_EXAMPLE_ETH_PHY_RST_GPIO;      // alter the GPIO_
↪used for PHY reset
// Install GPIO interrupt service (as the SPI-Ethernet module is interrupt-driven)
```

(continues on next page)

(continued from previous page)

```

gpio_install_isr_service(0);
// SPI bus configuration
spi_device_handle_t spi_handle = NULL;
spi_bus_config_t buscfg = {
    .miso_io_num = CONFIG_EXAMPLE_ETH_SPI_MISO_GPIO,
    .mosi_io_num = CONFIG_EXAMPLE_ETH_SPI_MOSI_GPIO,
    .sclk_io_num = CONFIG_EXAMPLE_ETH_SPI_SCLK_GPIO,
    .quadwp_io_num = -1,
    .quadhd_io_num = -1,
};
ESP_ERROR_CHECK(spi_bus_initialize(CONFIG_EXAMPLE_ETH_SPI_HOST, &buscfg, 1));
// Configure SPI device
spi_device_interface_config_t spi_devcfg = {
    .mode = 0,
    .clock_speed_hz = CONFIG_EXAMPLE_ETH_SPI_CLOCK_MHZ * 1000 * 1000,
    .spics_io_num = CONFIG_EXAMPLE_ETH_SPI_CS_GPIO,
    .queue_size = 20
};
/* dm9051 ethernet driver is based on spi driver */
eth_dm9051_config_t dm9051_config = ETH_DM9051_DEFAULT_CONFIG(CONFIG_EXAMPLE_ETH_
↳SPI_HOST, &spi_devcfg);
dm9051_config.int_gpio_num = CONFIG_EXAMPLE_ETH_SPI_INT_GPIO;
esp_eth_mac_t *mac = esp_eth_mac_new_dm9051(&dm9051_config, &mac_config);
esp_eth_phy_t *phy = esp_eth_phy_new_dm9051(&phy_config);

```

Note:

- When creating MAC and PHY instances for SPI-Ethernet modules (e.g., DM9051), the constructor function must have the same suffix (e.g., `esp_eth_mac_new_dm9051` and `esp_eth_phy_new_dm9051`). This is because we don't have other choices but the integrated PHY.
- The SPI device configuration (i.e., `spi_device_interface_config_t`) may slightly differ for other Ethernet modules or to meet SPI timing on specific PCB. Please check out your module's specs and the examples in ESP-IDF.

Install Driver To install the Ethernet driver, we need to combine the instance of MAC and PHY and set some additional high-level configurations (i.e., not specific to either MAC or PHY) in `esp_eth_config_t`:

- `esp_eth_config_t::mac`: instance that created from MAC generator (e.g., `esp_eth_mac_new_esp32()`).
- `esp_eth_config_t::phy`: instance that created from PHY generator (e.g., `esp_eth_phy_new_ip101()`).
- `esp_eth_config_t::check_link_period_ms`: Ethernet driver starts an OS timer to check the link status periodically, this field is used to set the interval, in milliseconds.
- `esp_eth_config_t::stack_input`: In most Ethernet IoT applications, any Ethernet frame received by a driver should be passed to the upper layer (e.g., TCP/IP stack). This field is set to a function that is responsible to deal with the incoming frames. You can even update this field at runtime via function `esp_eth_update_input_path()` after driver installation.
- `esp_eth_config_t::on_lowlevel_init_done` and `esp_eth_config_t::on_lowlevel_deinit_done`: These two fields are used to specify the hooks which get invoked when low-level hardware has been initialized or de-initialized.

ESP-IDF provides a default configuration for driver installation in macro `ETH_DEFAULT_CONFIG`.

```

esp_eth_config_t config = ETH_DEFAULT_CONFIG(mac, phy); // apply default driver_
↳configuration
esp_eth_handle_t eth_handle = NULL; // after the driver is installed, we will get_
↳the handle of the driver
esp_eth_driver_install(&config, &eth_handle); // install driver

```

The Ethernet driver also includes an event-driven model, which sends useful and important events to user space. We need to initialize the event loop before installing the Ethernet driver. For more information about event-driven programming, please refer to [ESP Event](#).

```

/** Event handler for Ethernet events */
static void eth_event_handler(void *arg, esp_event_base_t event_base,
                             int32_t event_id, void *event_data)
{
    uint8_t mac_addr[6] = {0};
    /* we can get the ethernet driver handle from event data */
    esp_eth_handle_t eth_handle = *(esp_eth_handle_t *)event_data;

    switch (event_id) {
        case ETHERNET_EVENT_CONNECTED:
            esp_eth_ioctl(eth_handle, ETH_CMD_G_MAC_ADDR, mac_addr);
            ESP_LOGI(TAG, "Ethernet Link Up");
            ESP_LOGI(TAG, "Ethernet HW Addr %02x:%02x:%02x:%02x:%02x:%02x",
                     mac_addr[0], mac_addr[1], mac_addr[2], mac_addr[3], mac_
↳addr[4], mac_addr[5]);
            break;
        case ETHERNET_EVENT_DISCONNECTED:
            ESP_LOGI(TAG, "Ethernet Link Down");
            break;
        case ETHERNET_EVENT_START:
            ESP_LOGI(TAG, "Ethernet Started");
            break;
        case ETHERNET_EVENT_STOP:
            ESP_LOGI(TAG, "Ethernet Stopped");
            break;
        default:
            break;
    }
}

esp_event_loop_create_default(); // create a default event loop that runs in the_
↳background
esp_event_handler_register(ETH_EVENT, ESP_EVENT_ANY_ID, &eth_event_handler, NULL);_
↳// register Ethernet event handler (to deal with user-specific stuff when events_
↳like link up/down happened)

```

Start Ethernet Driver After driver installation, we can start Ethernet immediately.

```
esp_eth_start(eth_handle); // start Ethernet driver state machine
```

Connect Driver to TCP/IP Stack Up until now, we have installed the Ethernet driver. From the view of OSI (Open System Interconnection), we are still on level 2 (i.e., Data Link Layer). While we can detect link up and down events and gain MAC address in user space, it is infeasible to obtain the IP address, let alone send an HTTP request. The TCP/IP stack used in ESP-IDF is called LwIP. For more information about it, please refer to [LwIP](#).

To connect the Ethernet driver to TCP/IP stack, follow these three steps:

1. Create a network interface for the Ethernet driver
2. Attach the network interface to the Ethernet driver
3. Register IP event handlers

For more information about the network interface, please refer to [Network Interface](#).

```

/** Event handler for IP_EVENT_ETH_GOT_IP */
static void got_ip_event_handler(void *arg, esp_event_base_t event_base,
                                int32_t event_id, void *event_data)

```

(continues on next page)

(continued from previous page)

```

{
    ip_event_got_ip_t *event = (ip_event_got_ip_t *) event_data;
    const esp_netif_ip_info_t *ip_info = &event->ip_info;

    ESP_LOGI(TAG, "Ethernet Got IP Address");
    ESP_LOGI(TAG, "~~~~~");
    ESP_LOGI(TAG, "ETHIP:" IPSTR, IP2STR(&ip_info->ip));
    ESP_LOGI(TAG, "ETHMASK:" IPSTR, IP2STR(&ip_info->netmask));
    ESP_LOGI(TAG, "ETHGW:" IPSTR, IP2STR(&ip_info->gw));
    ESP_LOGI(TAG, "~~~~~");
}

esp_netif_init(); // Initialize TCP/IP network interface (should be called only
↳once in application)
esp_netif_config_t cfg = ESP_NETIF_DEFAULT_ETH(); // apply default network
↳interface configuration for Ethernet
esp_netif_t *eth_netif = esp_netif_new(&cfg); // create network interface for
↳Ethernet driver

esp_netif_attach(eth_netif, esp_eth_new_netif_glue(eth_handle)); // attach
↳Ethernet driver to TCP/IP stack
esp_event_handler_register(IP_EVENT, IP_EVENT_ETH_GOT_IP, &got_ip_event_handler,
↳NULL); // register user defined IP event handlers
esp_eth_start(eth_handle); // start Ethernet driver state machine

```

Warning: It is recommended to fully initialize the Ethernet driver and network interface before registering the user's Ethernet/IP event handlers, i.e., register the event handlers as the last thing prior to starting the Ethernet driver. Such an approach ensures that Ethernet/IP events get executed first by the Ethernet driver or network interface so the system is in the expected state when executing the user's handlers.

Misc Control of Ethernet Driver The following functions should only be invoked after the Ethernet driver has been installed.

- Stop Ethernet driver: `esp_eth_stop()`
- Update Ethernet data input path: `esp_eth_update_input_path()`
- Misc get/set of Ethernet driver attributes: `esp_eth_ioctl()`

```

/* get MAC address */
uint8_t mac_addr[6];
memset(mac_addr, 0, sizeof(mac_addr));
esp_eth_ioctl(eth_handle, ETH_CMD_G_MAC_ADDR, mac_addr);
ESP_LOGI(TAG, "Ethernet MAC Address: %02x:%02x:%02x:%02x:%02x:%02x",
    mac_addr[0], mac_addr[1], mac_addr[2], mac_addr[3], mac_addr[4], mac_
↳addr[5]);

/* get PHY address */
int phy_addr = -1;
esp_eth_ioctl(eth_handle, ETH_CMD_G_PHY_ADDR, &phy_addr);
ESP_LOGI(TAG, "Ethernet PHY Address: %d", phy_addr);

```

Flow Control Ethernet on MCU usually has a limitation in the number of frames it can handle during network congestion, because of the limitation in RAM size. A sending station might be transmitting data faster than the peer end can accept it. The ethernet flow control mechanism allows the receiving node to signal the sender requesting the suspension of transmissions until the receiver catches up. The magic behind that is the pause frame, which was defined in IEEE 802.3x.

Pause frame is a special Ethernet frame used to carry the pause command, whose EtherType field is 0x8808, with

the Control opcode set to 0x0001. Only stations configured for full-duplex operation may send pause frames. When a station wishes to pause the other end of a link, it sends a pause frame to the 48-bit reserved multicast address of 01-80-C2-00-00-01. The pause frame also includes the period of pause time being requested, in the form of a two-byte integer, ranging from 0 to 65535.

After the Ethernet driver installation, the flow control feature is disabled by default. You can enable it by:

```
bool flow_ctrl_enable = true;
esp_eth_ioctl(eth_handle, ETH_CMD_S_FLOW_CTRL, &flow_ctrl_enable);
```

One thing that should be kept in mind is that the pause frame ability is advertised to the peer end by PHY during auto-negotiation. The Ethernet driver sends a pause frame only when both sides of the link support it.

Application Examples

- Ethernet basic example: [ethernet/basic](#)
- Ethernet iperf example: [ethernet/iperf](#)
- Ethernet to Wi-Fi AP "router": [network/eth2ap](#)
- Wi-Fi station to Ethernet "bridge": [network/sta2eth](#)
- Most protocol examples should also work for Ethernet: [protocols](#)

Advanced Topics

Custom PHY Driver There are multiple PHY manufacturers with wide portfolios of chips available. The ESP-IDF already supports several PHY chips however one can easily get to a point where none of them satisfies the user's actual needs due to price, features, stock availability, etc.

Luckily, a management interface between EMAC and PHY is standardized by IEEE 802.3 in Section 22.2.4 Management Functions. It defines provisions of the so-called "MII Management Interface" to control the PHY and gather status from the PHY. A set of management registers is defined to control chip behavior, link properties, auto-negotiation configuration, etc. This basic management functionality is addressed by [esp_eth/src/esp_eth_phy_802_3.c](#) in ESP-IDF and so it makes the creation of a new custom PHY chip driver quite a simple task.

Note: Always consult with PHY datasheet since some PHY chips may not comply with IEEE 802.3, Section 22.2.4. It does not mean you are not able to create a custom PHY driver, but it just requires more effort. You will have to define all PHY management functions.

The majority of PHY management functionality required by the ESP-IDF Ethernet driver is covered by the [esp_eth/src/esp_eth_phy_802_3.c](#). However, the following may require developing chip-specific management functions:

- Link status which is almost always chip-specific
- Chip initialization, even though not strictly required, should be customized to at least ensure that the expected chip is used
- Chip-specific features configuration

Steps to create a custom PHY driver:

1. Define vendor-specific registry layout based on the PHY datasheet. See [esp_eth/src/esp_eth_phy_ip101.c](#) as an example.
2. Prepare derived PHY management object info structure which:
 - must contain at least parent IEEE 802.3 [phy_802_3_t](#) object
 - optionally contain additional variables needed to support non-IEEE 802.3 or customized functionality. See [esp_eth/src/esp_eth_phy_ksz80xx.c](#) as an example.
3. Define chip-specific management call-back functions.
4. Initialize parent IEEE 802.3 object and re-assign chip-specific management call-back functions.

Once you finish the new custom PHY driver implementation, consider sharing it among other users via [IDF Component Registry](#).

API Reference

Header File

- [components/esp_eth/include/esp_eth.h](#)
- This header file can be included with:

```
#include "esp_eth.h"
```

- This header file is a part of the API provided by the `esp_eth` component. To declare that your component depends on `esp_eth`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_eth
```

or

```
PRIV_REQUIRES esp_eth
```

Header File

- [components/esp_eth/include/esp_eth_driver.h](#)
- This header file can be included with:

```
#include "esp_eth_driver.h"
```

- This header file is a part of the API provided by the `esp_eth` component. To declare that your component depends on `esp_eth`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_eth
```

or

```
PRIV_REQUIRES esp_eth
```

Functions

[esp_err_t](#) **esp_eth_driver_install** (const [esp_eth_config_t](#) *config, [esp_eth_handle_t](#) *out_hdl)

Install Ethernet driver.

Parameters

- **config** -- [in] configuration of the Ethernet driver
- **out_hdl** -- [out] handle of Ethernet driver

Returns

- `ESP_OK`: install `esp_eth` driver successfully
- `ESP_ERR_INVALID_ARG`: install `esp_eth` driver failed because of some invalid argument
- `ESP_ERR_NO_MEM`: install `esp_eth` driver failed because there's no memory for driver
- `ESP_FAIL`: install `esp_eth` driver failed because some other error occurred

[esp_err_t](#) **esp_eth_driver_uninstall** ([esp_eth_handle_t](#) hdl)

Uninstall Ethernet driver.

Note: It's not recommended to uninstall Ethernet driver unless it won't get used any more in application code. To uninstall Ethernet driver, you have to make sure, all references to the driver are released. Ethernet driver can only be uninstalled successfully when reference counter equals to one.

Parameters **hdl** -- [in] handle of Ethernet driver

Returns

- `ESP_OK`: uninstall `esp_eth` driver successfully

- `ESP_ERR_INVALID_ARG`: uninstall `esp_eth` driver failed because of some invalid argument
- `ESP_ERR_INVALID_STATE`: uninstall `esp_eth` driver failed because it has more than one reference
- `ESP_FAIL`: uninstall `esp_eth` driver failed because some other error occurred

esp_err_t **esp_eth_start** (*esp_eth_handle_t* hdl)

Start Ethernet driver **ONLY** in standalone mode (i.e. without TCP/IP stack)

Note: This API will start driver state machine and internal software timer (for checking link status).

Parameters `hdl` -- [in] handle of Ethernet driver

Returns

- `ESP_OK`: start `esp_eth` driver successfully
- `ESP_ERR_INVALID_ARG`: start `esp_eth` driver failed because of some invalid argument
- `ESP_ERR_INVALID_STATE`: start `esp_eth` driver failed because driver has started already
- `ESP_FAIL`: start `esp_eth` driver failed because some other error occurred

esp_err_t **esp_eth_stop** (*esp_eth_handle_t* hdl)

Stop Ethernet driver.

Note: This function does the oppsite operation of `esp_eth_start`.

Parameters `hdl` -- [in] handle of Ethernet driver

Returns

- `ESP_OK`: stop `esp_eth` driver successfully
- `ESP_ERR_INVALID_ARG`: stop `esp_eth` driver failed because of some invalid argument
- `ESP_ERR_INVALID_STATE`: stop `esp_eth` driver failed because driver has not started yet
- `ESP_FAIL`: stop `esp_eth` driver failed because some other error occurred

esp_err_t **esp_eth_update_input_path** (*esp_eth_handle_t* hdl, *esp_err_t* (*stack_input)(*esp_eth_handle_t* hdl, uint8_t *buffer, uint32_t length, void *priv), void *priv)

Update Ethernet data input path (i.e. specify where to pass the input buffer)

Note: After install driver, Ethernet still don't know where to deliver the input buffer. In fact, this API registers a callback function which get invoked when Ethernet received new packets.

Parameters

- `hdl` -- [in] handle of Ethernet driver
- `stack_input` -- [in] function pointer, which does the actual process on incoming packets
- `priv` -- [in] private resource, which gets passed to `stack_input` callback without any modification

Returns

- `ESP_OK`: update input path successfully
- `ESP_ERR_INVALID_ARG`: update input path failed because of some invalid argument
- `ESP_FAIL`: update input path failed because some other error occurred

esp_err_t **esp_eth_transmit** (*esp_eth_handle_t* hdl, void *buf, size_t length)

General Transmit.

Parameters

- **hdl** -- [in] handle of Ethernet driver
- **buf** -- [in] buffer of the packet to transfer
- **length** -- [in] length of the buffer to transfer

Returns

- ESP_OK: transmit frame buffer successfully
- ESP_ERR_INVALID_ARG: transmit frame buffer failed because of some invalid argument
- ESP_ERR_INVALID_STATE: invalid driver state (e.i. driver is not started)
- ESP_ERR_TIMEOUT: transmit frame buffer failed because HW was not get available in predefined period
- ESP_FAIL: transmit frame buffer failed because some other error occurred

esp_err_t **esp_eth_transmit_vargs** (*esp_eth_handle_t* hdl, uint32_t argc, ...)

Special Transmit with variable number of arguments.

Parameters

- **hdl** -- [in] handle of Ethernet driver
- **argc** -- [in] number variable arguments
- ... -- variable arguments

Returns

- ESP_OK: transmit successfull
- ESP_ERR_INVALID_STATE: invalid driver state (e.i. driver is not started)
- ESP_ERR_TIMEOUT: transmit frame buffer failed because HW was not get available in predefined period
- ESP_FAIL: transmit frame buffer failed because some other error occurred

esp_err_t **esp_eth_ioctl** (*esp_eth_handle_t* hdl, *esp_eth_io_cmd_t* cmd, void *data)

Misc IO function of Etherent driver.

The following common IO control commands are supported:

- ETH_CMD_S_MAC_ADDR sets Ethernet interface MAC address. *data* argument is pointer to MAC address buffer with expected size of 6 bytes.
- ETH_CMD_G_MAC_ADDR gets Ethernet interface MAC address. *data* argument is pointer to a buffer to which MAC address is to be copied. The buffer size must be at least 6 bytes.
- ETH_CMD_S_PHY_ADDR sets PHY address in range of <0-31>. *data* argument is pointer to memory of uint32_t datatype from where the configuration option is read.
- ETH_CMD_G_PHY_ADDR gets PHY address. *data* argument is pointer to memory of uint32_t datatype to which the PHY address is to be stored.
- ETH_CMD_S_AUTONEGO enables or disables Ethernet link speed and duplex mode autonegotiation. *data* argument is pointer to memory of bool datatype from which the configuration option is read. Preconditions: Ethernet driver needs to be stopped.
- ETH_CMD_G_AUTONEGO gets current configuration of the Ethernet link speed and duplex mode autonegotiation. *data* argument is pointer to memory of bool datatype to which the current configuration is to be stored.
- ETH_CMD_S_SPEED sets the Ethernet link speed. *data* argument is pointer to memory of eth_speed_t datatype from which the configuration option is read. Preconditions: Ethernet driver needs to be stopped and auto-negotiation disabled.
- ETH_CMD_G_SPEED gets current Ethernet link speed. *data* argument is pointer to memory of eth_speed_t datatype to which the speed is to be stored.
- ETH_CMD_S_PROMISCUOUS sets/resets Ethernet interface promiscuous mode. *data* argument is pointer to memory of bool datatype from which the configuration option is read.
- ETH_CMD_S_FLOW_CTRL sets/resets Ethernet interface flow control. *data* argument is pointer to memory of bool datatype from which the configuration option is read.
- ETH_CMD_S_DUPLEX_MODE sets the Ethernet duplex mode. *data* argument is pointer to memory of eth_duplex_t datatype from which the configuration option is read. Preconditions: Ethernet driver needs to be stopped and auto-negotiation disabled.

- `ETH_CMD_G_DUPLEX_MODE` gets current Ethernet link duplex mode. `data` argument is pointer to memory of `eth_duplex_t` datatype to which the duplex mode is to be stored.
- `ETH_CMD_S_PHY_LOOPBACK` sets/resets PHY to/from loopback mode. `data` argument is pointer to memory of `bool` datatype from which the configuration option is read.
- Note that additional control commands may be available for specific MAC or PHY chips. Please consult specific MAC or PHY documentation or driver code.

Parameters

- **hdl** -- **[in]** handle of Ethernet driver
- **cmd** -- **[in]** IO control command
- **data** -- **[inout]** address of data for `set` command or address where to store the data when used with `get` command

Returns

- `ESP_OK`: process io command successfully
- `ESP_ERR_INVALID_ARG`: process io command failed because of some invalid argument
- `ESP_FAIL`: process io command failed because some other error occurred
- `ESP_ERR_NOT_SUPPORTED`: requested feature is not supported

esp_err_t **esp_eth_increase_reference** (*esp_eth_handle_t* hdl)

Increase Ethernet driver reference.

Note: Ethernet driver handle can be obtained by `os timer`, `netif`, etc. It's dangerous when thread A is using Ethernet but thread B uninstall the driver. Using reference counter can prevent such risk, but care should be taken, when you obtain Ethernet driver, this API must be invoked so that the driver won't be uninstalled during your using time.

Parameters **hdl** -- **[in]** handle of Ethernet driver

Returns

- `ESP_OK`: increase reference successfully
- `ESP_ERR_INVALID_ARG`: increase reference failed because of some invalid argument

esp_err_t **esp_eth_decrease_reference** (*esp_eth_handle_t* hdl)

Decrease Ethernet driver reference.

Parameters **hdl** -- **[in]** handle of Ethernet driver

Returns

- `ESP_OK`: increase reference successfully
- `ESP_ERR_INVALID_ARG`: increase reference failed because of some invalid argument

Structures

struct **esp_eth_config_t**

Configuration of Ethernet driver.

Public Members

esp_eth_mac_t ***mac**

Ethernet MAC object.

esp_eth_phy_t ***phy**

Ethernet PHY object.

uint32_t **check_link_period_ms**

Period time of checking Ethernet link status.

esp_err_t (***stack_input**)(*esp_eth_handle_t* eth_handle, uint8_t *buffer, uint32_t length, void *priv)

Input frame buffer to user's stack.

Param eth_handle [in] handle of Ethernet driver

Param buffer [in] frame buffer that will get input to upper stack

Param length [in] length of the frame buffer

Return

- ESP_OK: input frame buffer to upper stack successfully
- ESP_FAIL: error occurred when inputting buffer to upper stack

esp_err_t (***on_lowlevel_init_done**)(*esp_eth_handle_t* eth_handle)

Callback function invoked when lowlevel initialization is finished.

Param eth_handle [in] handle of Ethernet driver

Return

- ESP_OK: process extra lowlevel initialization successfully
- ESP_FAIL: error occurred when processing extra lowlevel initialization

esp_err_t (***on_lowlevel_deinit_done**)(*esp_eth_handle_t* eth_handle)

Callback function invoked when lowlevel deinitialization is finished.

Param eth_handle [in] handle of Ethernet driver

Return

- ESP_OK: process extra lowlevel deinitialization successfully
- ESP_FAIL: error occurred when processing extra lowlevel deinitialization

esp_err_t (***read_phy_reg**)(*esp_eth_handle_t* eth_handle, uint32_t phy_addr, uint32_t phy_reg, uint32_t *reg_value)

Read PHY register.

Note: Usually the PHY register read/write function is provided by MAC (SMI interface), but if the PHY device is managed by other interface (e.g. I2C), then user needs to implement the corresponding read/write. Setting this to NULL means your PHY device is managed by MAC's SMI interface.

Param eth_handle [in] handle of Ethernet driver

Param phy_addr [in] PHY chip address (0~31)

Param phy_reg [in] PHY register index code

Param reg_value [out] PHY register value

Return

- ESP_OK: read PHY register successfully
- ESP_ERR_INVALID_ARG: read PHY register failed because of invalid argument
- ESP_ERR_TIMEOUT: read PHY register failed because of timeout
- ESP_FAIL: read PHY register failed because some other error occurred

esp_err_t (***write_phy_reg**)(*esp_eth_handle_t* eth_handle, uint32_t phy_addr, uint32_t phy_reg, uint32_t reg_value)

Write PHY register.

Note: Usually the PHY register read/write function is provided by MAC (SMI interface), but if the PHY device is managed by other interface (e.g. I2C), then user needs to implement the corresponding

read/write. Setting this to NULL means your PHY device is managed by MAC's SMI interface.

Param eth_handle [in] handle of Ethernet driver

Param phy_addr [in] PHY chip address (0~31)

Param phy_reg [in] PHY register index code

Param reg_value [in] PHY register value

Return

- ESP_OK: write PHY register successfully
- ESP_ERR_INVALID_ARG: read PHY register failed because of invalid argument
- ESP_ERR_TIMEOUT: write PHY register failed because of timeout
- ESP_FAIL: write PHY register failed because some other error occurred

struct **esp_eth_phy_reg_rw_data_t**

Data structure to Read/Write PHY register via ioctl API.

Public Members

uint32_t **reg_addr**

PHY register address

uint32_t ***reg_value_p**

Pointer to a memory where the register value is read/written

Macros

ETH_DEFAULT_CONFIG (emac, ephy)

Default configuration for Ethernet driver.

Type Definitions

typedef void ***esp_eth_handle_t**

Handle of Ethernet driver.

Enumerations

enum **esp_eth_io_cmd_t**

Command list for ioctl API.

Values:

enumerator **ETH_CMD_G_MAC_ADDR**

Get MAC address

enumerator **ETH_CMD_S_MAC_ADDR**

Set MAC address

enumerator **ETH_CMD_G_PHY_ADDR**

Get PHY address

enumerator **ETH_CMD_S_PHY_ADDR**

Set PHY address

enumerator **ETH_CMD_G_AUTONEGO**

Get PHY Auto Negotiation

enumerator **ETH_CMD_S_AUTONEGO**

Set PHY Auto Negotiation

enumerator **ETH_CMD_G_SPEED**

Get Speed

enumerator **ETH_CMD_S_SPEED**

Set Speed

enumerator **ETH_CMD_S_PROMISCUOUS**

Set promiscuous mode

enumerator **ETH_CMD_S_FLOW_CTRL**

Set flow control

enumerator **ETH_CMD_G_DUPLEX_MODE**

Get Duplex mode

enumerator **ETH_CMD_S_DUPLEX_MODE**

Set Duplex mode

enumerator **ETH_CMD_S_PHY_LOOPBACK**

Set PHY loopback

enumerator **ETH_CMD_READ_PHY_REG**

Read PHY register

enumerator **ETH_CMD_WRITE_PHY_REG**

Write PHY register

enumerator **ETH_CMD_CUSTOM_MAC_CMDS**

enumerator **ETH_CMD_CUSTOM_PHY_CMDS**

Header File

- [components/esp_eth/include/esp_eth_com.h](#)
- This header file can be included with:

```
#include "esp_eth_com.h"
```

- This header file is a part of the API provided by the `esp_eth` component. To declare that your component depends on `esp_eth`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_eth
```

or

PRIV_REQUIRES esp_eth

Structures

struct **esp_eth_mediator_s**

Ethernet mediator.

Public Members

esp_err_t (***phy_reg_read**)(*esp_eth_mediator_t* *eth, uint32_t phy_addr, uint32_t phy_reg, uint32_t *reg_value)

Read PHY register.

Param eth [in] mediator of Ethernet driver

Param phy_addr [in] PHY Chip address (0~31)

Param phy_reg [in] PHY register index code

Param reg_value [out] PHY register value

Return

- ESP_OK: read PHY register successfully
- ESP_FAIL: read PHY register failed because some error occurred

esp_err_t (***phy_reg_write**)(*esp_eth_mediator_t* *eth, uint32_t phy_addr, uint32_t phy_reg, uint32_t reg_value)

Write PHY register.

Param eth [in] mediator of Ethernet driver

Param phy_addr [in] PHY Chip address (0~31)

Param phy_reg [in] PHY register index code

Param reg_value [in] PHY register value

Return

- ESP_OK: write PHY register successfully
- ESP_FAIL: write PHY register failed because some error occurred

esp_err_t (***stack_input**)(*esp_eth_mediator_t* *eth, uint8_t *buffer, uint32_t length)

Deliver packet to upper stack.

Param eth [in] mediator of Ethernet driver

Param buffer [in] packet buffer

Param length [in] length of the packet

Return

- ESP_OK: deliver packet to upper stack successfully
- ESP_FAIL: deliver packet failed because some error occurred

esp_err_t (***on_state_changed**)(*esp_eth_mediator_t* *eth, *esp_eth_state_t* state, void *args)

Callback on Ethernet state changed.

Param eth [in] mediator of Ethernet driver

Param state [in] new state

Param args [in] optional argument for the new state

Return

- ESP_OK: process the new state successfully
- ESP_FAIL: process the new state failed because some error occurred

Type Definitions

typedef struct *esp_eth_mediator_s* **esp_eth_mediator_t**
Ethernet mediator.

Enumerations

enum **esp_eth_state_t**

Ethernet driver state.

Values:

enumerator **ETH_STATE_LLINIT**

Lowlevel init done

enumerator **ETH_STATE_DEINIT**

Deinit done

enumerator **ETH_STATE_LINK**

Link status changed

enumerator **ETH_STATE_SPEED**

Speed updated

enumerator **ETH_STATE_DUPLEX**

Duplex updated

enumerator **ETH_STATE_PAUSE**

Pause ability updated

enum **eth_event_t**

Ethernet event declarations.

Values:

enumerator **ETHERNET_EVENT_START**

Ethernet driver start

enumerator **ETHERNET_EVENT_STOP**

Ethernet driver stop

enumerator **ETHERNET_EVENT_CONNECTED**

Ethernet got a valid link

enumerator **ETHERNET_EVENT_DISCONNECTED**

Ethernet lost a valid link

Header File

- [components/esp_eth/include/esp_eth_mac.h](#)
- This header file can be included with:

```
#include "esp_eth_mac.h"
```

- This header file is a part of the API provided by the `esp_eth` component. To declare that your component depends on `esp_eth`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_eth
```

or

```
PRIV_REQUIRES esp_eth
```

Unions

union **eth_mac_clock_config_t**

#include <esp_eth_mac.h> Ethernet MAC Clock Configuration.

Public Members

struct *eth_mac_clock_config_t*::[anonymous] **mii**
EMAC MII Clock Configuration

emac_rmii_clock_mode_t **clock_mode**
RMII Clock Mode Configuration

emac_rmii_clock_gpio_t **clock_gpio**
RMII Clock GPIO Configuration

struct *eth_mac_clock_config_t*::[anonymous] **rmii**
EMAC RMII Clock Configuration

Structures

struct **esp_eth_mac_s**
Ethernet MAC.

Public Members

esp_err_t (***set_mediator**)(*esp_eth_mac_t* *mac, *esp_eth_mediator_t* *eth)
Set mediator for Ethernet MAC.

Param mac [in] Ethernet MAC instance

Param eth [in] Ethernet mediator

Return

- ESP_OK: set mediator for Ethernet MAC successfully
- ESP_ERR_INVALID_ARG: set mediator for Ethernet MAC failed because of invalid argument

esp_err_t (***init**)(*esp_eth_mac_t* *mac)
Initialize Ethernet MAC.

Param mac [in] Ethernet MAC instance

Return

- ESP_OK: initialize Ethernet MAC successfully

- ESP_ERR_TIMEOUT: initialize Ethernet MAC failed because of timeout
- ESP_FAIL: initialize Ethernet MAC failed because some other error occurred

esp_err_t (***deinit**)(*esp_eth_mac_t* *mac)

Deinitialize Ethernet MAC.

Param mac [in] Ethernet MAC instance

Return

- ESP_OK: deinitialize Ethernet MAC successfully
- ESP_FAIL: deinitialize Ethernet MAC failed because some error occurred

esp_err_t (***start**)(*esp_eth_mac_t* *mac)

Start Ethernet MAC.

Param mac [in] Ethernet MAC instance

Return

- ESP_OK: start Ethernet MAC successfully
- ESP_FAIL: start Ethernet MAC failed because some other error occurred

esp_err_t (***stop**)(*esp_eth_mac_t* *mac)

Stop Ethernet MAC.

Param mac [in] Ethernet MAC instance

Return

- ESP_OK: stop Ethernet MAC successfully
- ESP_FAIL: stop Ethernet MAC failed because some error occurred

esp_err_t (***transmit**)(*esp_eth_mac_t* *mac, uint8_t *buf, uint32_t length)

Transmit packet from Ethernet MAC.

Note: Returned error codes may differ for each specific MAC chip.

Param mac [in] Ethernet MAC instance

Param buf [in] packet buffer to transmit

Param length [in] length of packet

Return

- ESP_OK: transmit packet successfully
- ESP_ERR_INVALID_SIZE: number of actually sent bytes differs to expected
- ESP_FAIL: transmit packet failed because some other error occurred

esp_err_t (***transmit_vargs**)(*esp_eth_mac_t* *mac, uint32_t argc, va_list args)

Transmit packet from Ethernet MAC constructed with special parameters at Layer2.

Note: Typical intended use case is to make possible to construct a frame from multiple higher layer buffers without a need of buffer reallocations. However, other use cases are not limited.

Note: Returned error codes may differ for each specific MAC chip.

Param mac [in] Ethernet MAC instance

Param argc [in] number variable arguments

Param args [in] variable arguments

Return

- ESP_OK: transmit packet successfully
- ESP_ERR_INVALID_SIZE: number of actually sent bytes differs to expected
- ESP_FAIL: transmit packet failed because some other error occurred

esp_err_t (***receive**)(*esp_eth_mac_t* *mac, uint8_t *buf, uint32_t *length)

Receive packet from Ethernet MAC.

Note: Memory of buf is allocated in the Layer2, make sure it get free after process.

Note: Before this function got invoked, the value of "length" should set by user, equals the size of buffer. After the function returned, the value of "length" means the real length of received data.

Param mac [in] Ethernet MAC instance

Param buf [out] packet buffer which will preserve the received frame

Param length [out] length of the received packet

Return

- ESP_OK: receive packet successfully
- ESP_ERR_INVALID_ARG: receive packet failed because of invalid argument
- ESP_ERR_INVALID_SIZE: input buffer size is not enough to hold the incoming data.
in this case, value of returned "length" indicates the real size of incoming data.
- ESP_FAIL: receive packet failed because some other error occurred

esp_err_t (***read_phy_reg**)(*esp_eth_mac_t* *mac, uint32_t phy_addr, uint32_t phy_reg, uint32_t *reg_value)

Read PHY register.

Param mac [in] Ethernet MAC instance

Param phy_addr [in] PHY chip address (0~31)

Param phy_reg [in] PHY register index code

Param reg_value [out] PHY register value

Return

- ESP_OK: read PHY register successfully
- ESP_ERR_INVALID_ARG: read PHY register failed because of invalid argument
- ESP_ERR_INVALID_STATE: read PHY register failed because of wrong state of MAC
- ESP_ERR_TIMEOUT: read PHY register failed because of timeout
- ESP_FAIL: read PHY register failed because some other error occurred

esp_err_t (***write_phy_reg**)(*esp_eth_mac_t* *mac, uint32_t phy_addr, uint32_t phy_reg, uint32_t reg_value)

Write PHY register.

Param mac [in] Ethernet MAC instance

Param phy_addr [in] PHY chip address (0~31)

Param phy_reg [in] PHY register index code

Param reg_value [in] PHY register value

Return

- ESP_OK: write PHY register successfully
- ESP_ERR_INVALID_STATE: write PHY register failed because of wrong state of MAC
- ESP_ERR_TIMEOUT: write PHY register failed because of timeout
- ESP_FAIL: write PHY register failed because some other error occurred

esp_err_t (***set_addr**)(*esp_eth_mac_t* *mac, uint8_t *addr)

Set MAC address.

Param mac [in] Ethernet MAC instance

Param addr [in] MAC address

Return

- ESP_OK: set MAC address successfully
- ESP_ERR_INVALID_ARG: set MAC address failed because of invalid argument
- ESP_FAIL: set MAC address failed because some other error occurred

esp_err_t (***get_addr**)(*esp_eth_mac_t* *mac, uint8_t *addr)

Get MAC address.

Param mac [in] Ethernet MAC instance

Param addr [out] MAC address

Return

- ESP_OK: get MAC address successfully
- ESP_ERR_INVALID_ARG: get MAC address failed because of invalid argument
- ESP_FAIL: get MAC address failed because some other error occurred

esp_err_t (***set_speed**)(*esp_eth_mac_t* *mac, eth_speed_t speed)

Set speed of MAC.

Param mac [in] Ethernet MAC instance

Param speed [in] MAC speed

Return

- ESP_OK: set MAC speed successfully
- ESP_ERR_INVALID_ARG: set MAC speed failed because of invalid argument
- ESP_FAIL: set MAC speed failed because some other error occurred

esp_err_t (***set_duplex**)(*esp_eth_mac_t* *mac, eth_duplex_t duplex)

Set duplex mode of MAC.

Param mac [in] Ethernet MAC instance

Param duplex [in] MAC duplex

Return

- ESP_OK: set MAC duplex mode successfully
- ESP_ERR_INVALID_ARG: set MAC duplex failed because of invalid argument
- ESP_FAIL: set MAC duplex failed because some other error occurred

esp_err_t (***set_link**)(*esp_eth_mac_t* *mac, eth_link_t link)

Set link status of MAC.

Param mac [in] Ethernet MAC instance

Param link [in] Link status

Return

- ESP_OK: set link status successfully
- ESP_ERR_INVALID_ARG: set link status failed because of invalid argument
- ESP_FAIL: set link status failed because some other error occurred

esp_err_t (***set_promiscuous**)(*esp_eth_mac_t* *mac, bool enable)

Set promiscuous of MAC.

Param mac [in] Ethernet MAC instance

Param enable [in] set true to enable promiscuous mode; set false to disable promiscuous mode

Return

- ESP_OK: set promiscuous mode successfully
- ESP_FAIL: set promiscuous mode failed because some error occurred

esp_err_t (***enable_flow_ctrl**)(*esp_eth_mac_t* *mac, bool enable)

Enable flow control on MAC layer or not.

Param mac [in] Ethernet MAC instance

Param enable [in] set true to enable flow control; set false to disable flow control

Return

- ESP_OK: set flow control successfully
- ESP_FAIL: set flow control failed because some error occurred

esp_err_t (***set_peer_pause_ability**)(*esp_eth_mac_t* *mac, uint32_t ability)

Set the PAUSE ability of peer node.

Param mac [in] Ethernet MAC instance

Param ability [in] zero indicates that pause function is supported by link partner; non-zero indicates that pause function is not supported by link partner

Return

- ESP_OK: set peer pause ability successfully
- ESP_FAIL: set peer pause ability failed because some error occurred

esp_err_t (***custom_ioctl**)(*esp_eth_mac_t* *mac, uint32_t cmd, void *data)

Custom IO function of MAC driver. This function is intended to extend common options of *esp_eth_ioctl* to cover specifics of MAC chip.

Note: This function may not be assigned when the MAC chip supports only most common set of configuration options.

Param mac [in] Ethernet MAC instance

Param cmd [in] IO control command

Param data [inout] address of data for *set* command or address where to store the data when used with *get* command

Return

- ESP_OK: process io command successfully
- ESP_ERR_INVALID_ARG: process io command failed because of some invalid argument
- ESP_FAIL: process io command failed because some other error occurred
- ESP_ERR_NOT_SUPPORTED: requested feature is not supported

esp_err_t (***del**)(*esp_eth_mac_t* *mac)

Free memory of Ethernet MAC.

Param mac [in] Ethernet MAC instance

Return

- ESP_OK: free Ethernet MAC instance successfully
- ESP_FAIL: free Ethernet MAC instance failed because some error occurred

struct **eth_mac_config_t**

Configuration of Ethernet MAC object.

Public Members

uint32_t **sw_reset_timeout_ms**

Software reset timeout value (Unit: ms)

uint32_t **rx_task_stack_size**

Stack size of the receive task

uint32_t **rx_task_prio**

Priority of the receive task

uint32_t **flags**

Flags that specify extra capability for mac driver

struct **eth_spi_custom_driver_config_t**

Custom SPI Driver Configuration. This structure declares configuration and callback functions to access Ethernet SPI module via user's custom SPI driver.

Public Members

void ***config**

Custom driver specific configuration data used by *init()* function.

Note: Type and its content is fully under user's control

void **(*init)**(const void *spi_config)

Custom driver SPI Initialization.

Note: return type and its content is fully under user's control

Param spi_config [in] Custom driver specific configuration

Return

- spi_ctx: when initialization is successful, a pointer to context structure holding all variables needed for subsequent SPI access operations (e.g. SPI bus identification, mutexes, etc.)
- NULL: driver initialization failed

esp_err_t **(*deinit)**(void *spi_ctx)

Custom driver De-initialization.

Param spi_ctx [in] a pointer to driver specific context structure

Return

- ESP_OK: driver de-initialization was successful
- ESP_FAIL: driver de-initialization failed
- any other failure codes are allowed to be used to provide failure isolation

esp_err_t **(*read)**(void *spi_ctx, uint32_t cmd, uint32_t addr, void *data, uint32_t data_len)

Custom driver SPI read.

Note: The read function is responsible to construct command, address and data fields of the SPI frame in format expected by particular SPI Ethernet module

Param spi_ctx [in] a pointer to driver specific context structure

Param cmd [in] command
Param addr [in] register address
Param data [out] read data
Param data_len [in] read data length in bytes
Return

- ESP_OK: read was successful
- ESP_FAIL: read failed
- any other failure codes are allowed to be used to provide failure isolation

esp_err_t (***write**)(void *spi_ctx, uint32_t cmd, uint32_t addr, const void *data, uint32_t data_len)

Custom driver SPI write.

Note: The write function is responsible to construct command, address and data fields of the SPI frame in format expected by particular SPI Ethernet module

Param spi_ctx [in] a pointer to driver specific context structure
Param cmd [in] command
Param addr [in] register address
Param data [in] data to write
Param data_len [in] length of data to write in bytes
Return

- ESP_OK: write was successful
- ESP_FAIL: write failed
- any other failure codes are allowed to be used to provide failure isolation

Macros

ETH_MAC_FLAG_WORK_WITH_CACHE_DISABLE

MAC driver can work when cache is disabled

ETH_MAC_FLAG_PIN_TO_CORE

Pin MAC task to the CPU core where driver installation happened

ETH_MAC_DEFAULT_CONFIG()

Default configuration for Ethernet MAC object.

ETH_DEFAULT_SPI

Default configuration of the custom SPI driver. Internal ESP-IDF SPI Master driver is used by default.

Type Definitions

typedef struct *esp_eth_mac_s* **esp_eth_mac_t**

Ethernet MAC.

Enumerations

enum **emac_rmii_clock_mode_t**

RMII Clock Mode Options.

Values:

enumerator **EMAC_CLK_DEFAULT**

Default values configured using Kconfig are going to be used when "Default" selected.

enumerator **EMAC_CLK_EXT_IN**

Input RMII Clock from external. EMAC Clock GPIO number needs to be configured when this option is selected.

Note: MAC will get RMII clock from outside. Note that ESP32 only supports GPIO0 to input the RMII clock.

enumerator **EMAC_CLK_OUT**

Output RMII Clock from internal APLL Clock. EMAC Clock GPIO number needs to be configured when this option is selected.

enum **emac_rmii_clock_gpio_t**

RMII Clock GPIO number Options.

Warning: If you want the Ethernet to work with WiFi, don't select ESP32 as RMII CLK output as it would result in clock instability.

Values:

enumerator **EMAC_CLK_IN_GPIO**

MAC will get RMII clock from outside at this GPIO.

Note: ESP32 only supports GPIO0 to input the RMII clock.

enumerator **EMAC_APPL_CLK_OUT_GPIO**

Output RMII Clock from internal APLL Clock available at GPIO0.

Note: GPIO0 can be set to output a pre-divided PLL clock. Enabling this option will configure GPIO0 to output a 50MHz clock. In fact this clock doesn't have directly relationship with EMAC peripheral. Sometimes this clock may not work well with your PHY chip.

enumerator **EMAC_CLK_OUT_GPIO**

Output RMII Clock from internal APLL Clock available at GPIO16.

enumerator **EMAC_CLK_OUT_180_GPIO**

Inverted Output RMII Clock from internal APLL Clock available at GPIO17.

Header File

- [components/esp_eth/include/esp_eth_phy.h](#)
- This header file can be included with:

```
#include "esp_eth_phy.h"
```

- This header file is a part of the API provided by the `esp_eth` component. To declare that your component depends on `esp_eth`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_eth
```

or

PRIV_REQUIRES esp_eth

Functions

esp_eth_phy_t ***esp_eth_phy_new_ip101** (const *eth_phy_config_t* *config)

Create a PHY instance of IP101.

Parameters **config** -- [in] configuration of PHY

Returns

- instance: create PHY instance successfully
- NULL: create PHY instance failed because some error occurred

esp_eth_phy_t ***esp_eth_phy_new_rt18201** (const *eth_phy_config_t* *config)

Create a PHY instance of RTL8201.

Parameters **config** -- [in] configuration of PHY

Returns

- instance: create PHY instance successfully
- NULL: create PHY instance failed because some error occurred

esp_eth_phy_t ***esp_eth_phy_new_lan87xx** (const *eth_phy_config_t* *config)

Create a PHY instance of LAN87xx.

Parameters **config** -- [in] configuration of PHY

Returns

- instance: create PHY instance successfully
- NULL: create PHY instance failed because some error occurred

esp_eth_phy_t ***esp_eth_phy_new_dp83848** (const *eth_phy_config_t* *config)

Create a PHY instance of DP83848.

Parameters **config** -- [in] configuration of PHY

Returns

- instance: create PHY instance successfully
- NULL: create PHY instance failed because some error occurred

esp_eth_phy_t ***esp_eth_phy_new_ksz80xx** (const *eth_phy_config_t* *config)

Create a PHY instance of KSZ80xx.

The phy model from the KSZ80xx series is detected automatically. If the driver is unable to detect a supported model, NULL is returned.

Currently, the following models are supported: KSZ8001, KSZ8021, KSZ8031, KSZ8041, KSZ8051, KSZ8061, KSZ8081, KSZ8091

Parameters **config** -- [in] configuration of PHY

Returns

- instance: create PHY instance successfully
- NULL: create PHY instance failed because some error occurred

Structures

struct **esp_eth_phy_s**

Ethernet PHY.

Public Members

esp_err_t (***set_mediator**)(*esp_eth_phy_t* *phy, *esp_eth_mediator_t* *mediator)

Set mediator for PHY.

Param phy [in] Ethernet PHY instance

Param mediator [in] mediator of Ethernet driver

Return

- ESP_OK: set mediator for Ethernet PHY instance successfully
- ESP_ERR_INVALID_ARG: set mediator for Ethernet PHY instance failed because of some invalid arguments

esp_err_t (***reset**)(*esp_eth_phy_t* *phy)

Software Reset Ethernet PHY.

Param phy [in] Ethernet PHY instance

Return

- ESP_OK: reset Ethernet PHY successfully
- ESP_FAIL: reset Ethernet PHY failed because some error occurred

esp_err_t (***reset_hw**)(*esp_eth_phy_t* *phy)

Hardware Reset Ethernet PHY.

Note: Hardware reset is mostly done by pull down and up PHY's nRST pin

Param phy [in] Ethernet PHY instance

Return

- ESP_OK: reset Ethernet PHY successfully
- ESP_FAIL: reset Ethernet PHY failed because some error occurred

esp_err_t (***init**)(*esp_eth_phy_t* *phy)

Initialize Ethernet PHY.

Param phy [in] Ethernet PHY instance

Return

- ESP_OK: initialize Ethernet PHY successfully
- ESP_FAIL: initialize Ethernet PHY failed because some error occurred

esp_err_t (***deinit**)(*esp_eth_phy_t* *phy)

Deinitialize Ethernet PHY.

Param phy [in] Ethernet PHY instance

Return

- ESP_OK: deinitialize Ethernet PHY successfully
- ESP_FAIL: deinitialize Ethernet PHY failed because some error occurred

esp_err_t (***autonego_ctrl**)(*esp_eth_phy_t* *phy, *eth_phy_autoneg_cmd_t* cmd, bool *autonego_en_stat)

Configure auto negotiation.

Param phy [in] Ethernet PHY instance

Param cmd [in] Configuration command, it is possible to Enable (restart), Disable or get current status of PHY auto negotiation

Param autonego_en_stat [out] Address where to store current status of auto negotiation configuration

Return

- ESP_OK: restart auto negotiation successfully
- ESP_FAIL: restart auto negotiation failed because some error occurred
- ESP_ERR_INVALID_ARG: invalid command

esp_err_t (***get_link**)(*esp_eth_phy_t* *phy)

Get Ethernet PHY link status.

Param phy [in] Ethernet PHY instance

Return

- ESP_OK: get Ethernet PHY link status successfully
- ESP_FAIL: get Ethernet PHY link status failed because some error occurred

esp_err_t (***set_link**)(*esp_eth_phy_t* *phy, *eth_link_t* link)

Set Ethernet PHY link status.

Param phy [in] Ethernet PHY instance

Param link [in] new link status

Return

- ESP_OK: set Ethernet PHY link status successfully
- ESP_FAIL: set Ethernet PHY link status failed because some error occurred

esp_err_t (***pwrcctl**)(*esp_eth_phy_t* *phy, bool enable)

Power control of Ethernet PHY.

Param phy [in] Ethernet PHY instance

Param enable [in] set true to power on Ethernet PHY; set false to power off Ethernet PHY

Return

- ESP_OK: control Ethernet PHY power successfully
- ESP_FAIL: control Ethernet PHY power failed because some error occurred

esp_err_t (***set_addr**)(*esp_eth_phy_t* *phy, *uint32_t* addr)

Set PHY chip address.

Param phy [in] Ethernet PHY instance

Param addr [in] PHY chip address

Return

- ESP_OK: set Ethernet PHY address successfully
- ESP_FAIL: set Ethernet PHY address failed because some error occurred

esp_err_t (***get_addr**)(*esp_eth_phy_t* *phy, *uint32_t* *addr)

Get PHY chip address.

Param phy [in] Ethernet PHY instance

Param addr [out] PHY chip address

Return

- ESP_OK: get Ethernet PHY address successfully
- ESP_ERR_INVALID_ARG: get Ethernet PHY address failed because of invalid argument

esp_err_t (***advertise_pause_ability**)(*esp_eth_phy_t* *phy, *uint32_t* ability)

Advertise pause function supported by MAC layer.

Param phy [in] Ethernet PHY instance

Param addr [out] Pause ability

Return

- ESP_OK: Advertise pause ability successfully
- ESP_ERR_INVALID_ARG: Advertise pause ability failed because of invalid argument

esp_err_t (***loopback**)(*esp_eth_phy_t* *phy, bool enable)

Sets the PHY to loopback mode.

Param phy [in] Ethernet PHY instance

Param enable [in] enables or disables PHY loopback

Return

- ESP_OK: PHY instance loopback mode has been configured successfully
- ESP_FAIL: PHY instance loopback configuration failed because some error occurred

esp_err_t (***set_speed**)(*esp_eth_phy_t* *phy, eth_speed_t speed)

Sets PHY speed mode.

Note: Autonegotiation feature needs to be disabled prior to calling this function for the new setting to be applied

Param phy [in] Ethernet PHY instance

Param speed [in] Speed mode to be set

Return

- ESP_OK: PHY instance speed mode has been configured successfully
- ESP_FAIL: PHY instance speed mode configuration failed because some error occurred

esp_err_t (***set_duplex**)(*esp_eth_phy_t* *phy, eth_duplex_t duplex)

Sets PHY duplex mode.

Note: Autonegotiation feature needs to be disabled prior to calling this function for the new setting to be applied

Param phy [in] Ethernet PHY instance

Param duplex [in] Duplex mode to be set

Return

- ESP_OK: PHY instance duplex mode has been configured successfully
- ESP_FAIL: PHY instance duplex mode configuration failed because some error occurred

esp_err_t (***custom_ioctl**)(*esp_eth_phy_t* *phy, uint32_t cmd, void *data)

Custom IO function of PHY driver. This function is intended to extend common options of *esp_eth_ioctl* to cover specifics of PHY chip.

Note: This function may not be assigned when the PHY chip supports only most common set of configuration options.

Param phy [in] Ethernet PHY instance

Param cmd [in] IO control command

Param data [inout] address of data for *set* command or address where to store the data when used with *get* command

Return

- ESP_OK: process io command successfully
- ESP_ERR_INVALID_ARG: process io command failed because of some invalid argument
- ESP_FAIL: process io command failed because some other error occurred
- ESP_ERR_NOT_SUPPORTED: requested feature is not supported

esp_err_t (***del**)(*esp_eth_phy_t* *phy)

Free memory of Ethernet PHY instance.

Param phy [in] Ethernet PHY instance

Return

- ESP_OK: free PHY instance successfully
- ESP_FAIL: free PHY instance failed because some error occurred

struct **eth_phy_config_t**

Ethernet PHY configuration.

Public Members

int32_t **phy_addr**

PHY address, set -1 to enable PHY address detection at initialization stage

uint32_t **reset_timeout_ms**

Reset timeout value (Unit: ms)

uint32_t **autonego_timeout_ms**

Auto-negotiation timeout value (Unit: ms)

int **reset_gpio_num**

Reset GPIO number, -1 means no hardware reset

Macros

ESP_ETH_PHY_ADDR_AUTO

ETH_PHY_DEFAULT_CONFIG()

Default configuration for Ethernet PHY object.

Type Definitions

typedef struct *esp_eth_phy_s* **esp_eth_phy_t**

Ethernet PHY.

Enumerations

enum **eth_phy_autoneg_cmd_t**

Auto-negotiation controll commands.

Values:

enumerator **ESP_ETH_PHY_AUTONEGO_RESTART**

enumerator **ESP_ETH_PHY_AUTONEGO_EN**

enumerator **ESP_ETH_PHY_AUTONEGO_DIS**

enumerator **ESP_ETH_PHY_AUTONEGO_G_STAT**

Header File

- [components/esp_eth/include/esp_eth_phy_802_3.h](#)
- This header file can be included with:

```
#include "esp_eth_phy_802_3.h"
```

- This header file is a part of the API provided by the `esp_eth` component. To declare that your component depends on `esp_eth`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_eth
```

or

```
PRIV_REQUIRES esp_eth
```

Functions

[esp_err_t](#) **esp_eth_phy_802_3_set_mediator** ([phy_802_3_t](#) *phy_802_3, [esp_eth_mediator_t](#) *eth)

Set Ethernet mediator.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **eth** -- Ethernet mediator pointer

Returns

- ESP_OK: Ethernet mediator set successfully
- ESP_ERR_INVALID_ARG: if `eth` is NULL

[esp_err_t](#) **esp_eth_phy_802_3_reset** ([phy_802_3_t](#) *phy_802_3)

Reset PHY.

Parameters **phy_802_3** -- IEEE 802.3 PHY object infostructure

Returns

- ESP_OK: Ethernet PHY reset successfully
- ESP_FAIL: reset Ethernet PHY failed because some error occurred

[esp_err_t](#) **esp_eth_phy_802_3_autonego_ctrl** ([phy_802_3_t](#) *phy_802_3, [eth_phy_autoneg_cmd_t](#) cmd, bool *autonego_en_stat)

Control autonegotiation mode of Ethernet PHY.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **cmd** -- autonegotiation command enumeration
- **autonego_en_stat** -- [out] autonegotiation enabled flag

Returns

- ESP_OK: Ethernet PHY autonegotiation configured successfully
- ESP_FAIL: Ethernet PHY autonegotiation configuration fail because some error occurred
- ESP_ERR_INVALID_ARG: invalid value of `cmd`

[esp_err_t](#) **esp_eth_phy_802_3_pwrctl** ([phy_802_3_t](#) *phy_802_3, bool enable)

Power control of Ethernet PHY.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **enable** -- set true to power ON Ethernet PHY; set false to power OFF Ethernet PHY

Returns

- ESP_OK: Ethernet PHY power down mode set successfully
- ESP_FAIL: Ethernet PHY power up or power down failed because some error occurred

[esp_err_t](#) **esp_eth_phy_802_3_set_addr** ([phy_802_3_t](#) *phy_802_3, uint32_t addr)

Set Ethernet PHY address.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure

- **addr** -- new PHY address

Returns

- ESP_OK: Ethernet PHY address set

esp_err_t **esp_eth_phy_802_3_get_addr** (*phy_802_3_t* *phy_802_3, uint32_t *addr)

Get Ethernet PHY address.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **addr** -- [out] Ethernet PHY address

Returns

- ESP_OK: Ethernet PHY address read successfully
- ESP_ERR_INVALID_ARG: addr pointer is NULL

esp_err_t **esp_eth_phy_802_3_advertise_pause_ability** (*phy_802_3_t* *phy_802_3, uint32_t ability)

Advertise pause function ability.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **ability** -- enable or disable pause ability

Returns

- ESP_OK: pause ability set successfully
- ESP_FAIL: Advertise pause function ability failed because some error occurred

esp_err_t **esp_eth_phy_802_3_loopback** (*phy_802_3_t* *phy_802_3, bool enable)

Set Ethernet PHY loopback mode.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **enable** -- set true to enable loopback; set false to disable loopback

Returns

- ESP_OK: Ethernet PHY loopback mode set successfully
- ESP_FAIL: Ethernet PHY loopback configuration failed because some error occurred

esp_err_t **esp_eth_phy_802_3_set_speed** (*phy_802_3_t* *phy_802_3, eth_speed_t speed)

Set Ethernet PHY speed.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **speed** -- new speed of Ethernet PHY link

Returns

- ESP_OK: Ethernet PHY speed set successfully
- ESP_FAIL: Set Ethernet PHY speed failed because some error occurred

esp_err_t **esp_eth_phy_802_3_set_duplex** (*phy_802_3_t* *phy_802_3, eth_duplex_t duplex)

Set Ethernet PHY duplex mode.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **duplex** -- new duplex mode for Ethernet PHY link

Returns

- ESP_OK: Ethernet PHY duplex mode set successfully
- ESP_ERR_INVALID_STATE: unable to set duplex mode to Half if loopback is enabled
- ESP_FAIL: Set Ethernet PHY duplex mode failed because some error occurred

esp_err_t **esp_eth_phy_802_3_set_link** (*phy_802_3_t* *phy_802_3, eth_link_t link)

Set Ethernet PHY link status.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **link** -- new link status

Returns

- ESP_OK: Ethernet PHY link set successfully

esp_err_t **esp_eth_phy_802_3_init** (*phy_802_3_t* *phy_802_3)

Initialize Ethernet PHY.

Parameters **phy_802_3** -- IEEE 802.3 PHY object infostructure

Returns

- ESP_OK: Ethernet PHY initialized successfully

esp_err_t **esp_eth_phy_802_3_deinit** (*phy_802_3_t* *phy_802_3)

Power off Ethernet PHY.

Parameters **phy_802_3** -- IEEE 802.3 PHY object infostructure

Returns

- ESP_OK: Ethernet PHY powered off successfully

esp_err_t **esp_eth_phy_802_3_del** (*phy_802_3_t* *phy_802_3)

Delete Ethernet PHY infostructure.

Parameters **phy_802_3** -- IEEE 802.3 PHY object infostructure

Returns

- ESP_OK: Ethernet PHY infostructure deleted

esp_err_t **esp_eth_phy_802_3_reset_hw** (*phy_802_3_t* *phy_802_3, uint32_t reset_assert_us)

Performs hardware reset with specific reset pin assertion time.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **reset_assert_us** -- Hardware reset pin assertion time

Returns

- ESP_OK: reset Ethernet PHY successfully

esp_err_t **esp_eth_phy_802_3_detect_phy_addr** (*esp_eth_mediator_t* *eth, int *detected_addr)

Detect PHY address.

Parameters

- **eth** -- Mediator of Ethernet driver
- **detected_addr** -- [out] a valid address after detection

Returns

- ESP_OK: detect phy address successfully
- ESP_ERR_INVALID_ARG: invalid parameter
- ESP_ERR_NOT_FOUND: can't detect any PHY device
- ESP_FAIL: detect phy address failed because some error occurred

esp_err_t **esp_eth_phy_802_3_basic_phy_init** (*phy_802_3_t* *phy_802_3)

Performs basic PHY chip initialization.

Note: It should be called as the first function in PHY specific driver instance

Parameters **phy_802_3** -- IEEE 802.3 PHY object infostructure

Returns

- ESP_OK: initialized Ethernet PHY successfully
- ESP_FAIL: initialization of Ethernet PHY failed because some error occurred
- ESP_ERR_INVALID_ARG: invalid argument
- ESP_ERR_NOT_FOUND: PHY device not detected
- ESP_ERR_TIMEOUT: MII Management read/write operation timeout
- ESP_ERR_INVALID_STATE: PHY is in invalid state to perform requested operation

esp_err_t **esp_eth_phy_802_3_basic_phy_deinit** (*phy_802_3_t* *phy_802_3)

Performs basic PHY chip de-initialization.

Note: It should be called as the last function in PHY specific driver instance

Parameters **phy_802_3** -- IEEE 802.3 PHY object infostructure

Returns

- ESP_OK: de-initialized Ethernet PHY successfully
- ESP_FAIL: de-initialization of Ethernet PHY failed because some error occurred
- ESP_ERR_TIMEOUT: MII Management read/write operation timeout
- ESP_ERR_INVALID_STATE: PHY is in invalid state to perform requested operation

esp_err_t **esp_eth_phy_802_3_read_oui** (*phy_802_3_t* *phy_802_3, uint32_t *oui)

Reads raw content of OUI field.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **oui** -- [out] OUI value

Returns

- ESP_OK: OUI field read successfully
- ESP_FAIL: OUI field read failed because some error occurred
- ESP_ERR_INVALID_ARG: invalid `oui` argument
- ESP_ERR_TIMEOUT: MII Management read/write operation timeout
- ESP_ERR_INVALID_STATE: PHY is in invalid state to perform requested operation

esp_err_t **esp_eth_phy_802_3_read_manufac_info** (*phy_802_3_t* *phy_802_3, uint8_t *model, uint8_t *rev)

Reads manufacturer's model and revision number.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **model** -- [out] Manufacturer's model number (can be NULL when not required)
- **rev** -- [out] Manufacturer's revision number (can be NULL when not required)

Returns

- ESP_OK: Manufacturer's info read successfully
- ESP_FAIL: Manufacturer's info read failed because some error occurred
- ESP_ERR_TIMEOUT: MII Management read/write operation timeout
- ESP_ERR_INVALID_STATE: PHY is in invalid state to perform requested operation

esp_err_t **esp_eth_phy_802_3_get_mmd_addr** (*phy_802_3_t* *phy_802_3, uint8_t devaddr, uint16_t *mmd_addr)

Reads MDIO device's internal address register.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **devaddr** -- Address of MDIO device
- **mmd_addr** -- [out] Current address stored in device's register

Returns

- ESP_OK: Address register read successfully
- ESP_FAIL: Address register read failed because of some error occurred
- ESP_ERR_INVALID_ARG: Device address provided is out of range (hardware limits device address to 5 bits)

esp_err_t **esp_eth_phy_802_3_set_mmd_addr** (*phy_802_3_t* *phy_802_3, uint8_t devaddr, uint16_t mmd_addr)

Write to DIO device's internal address register.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure

- **devaddr** -- Address of MDIO device
- **mmd_addr** -- [out] New value of MDIO device's address register value

Returns

- ESP_OK: Address register written to successfully
- ESP_FAIL: Address register write failed because of some error occurred
- ESP_ERR_INVALID_ARG: Device address provided is out of range (hardware limits device address to 5 bits)

esp_err_t **esp_eth_phy_802_3_read_mmd_data** (*phy_802_3_t* *phy_802_3, uint8_t devaddr, *esp_eth_phy_802_3_mmd_func_t* function, uint32_t *data)

Read data of MDIO device's memory at address register.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **devaddr** -- Address of MDIO device
- **function** -- MMD function
- **data** -- [out] Data read from the device's memory

Returns

- ESP_OK: Memory read successfully
- ESP_FAIL: Memory read failed because of some error occurred
- ESP_ERR_INVALID_ARG: Device address provided is out of range (hardware limits device address to 5 bits) or MMD access function is invalid

esp_err_t **esp_eth_phy_802_3_write_mmd_data** (*phy_802_3_t* *phy_802_3, uint8_t devaddr, *esp_eth_phy_802_3_mmd_func_t* function, uint32_t data)

Write data to MDIO device's memory at address register.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **devaddr** -- Address of MDIO device
- **function** -- MMD function
- **data** -- [out] Data to write to the device's memory

Returns

- ESP_OK: Memory written successfully
- ESP_FAIL: Memory write failed because of some error occurred
- ESP_ERR_INVALID_ARG: Device address provided is out of range (hardware limits device address to 5 bits) or MMD access function is invalid

esp_err_t **esp_eth_phy_802_3_read_mmd_register** (*phy_802_3_t* *phy_802_3, uint8_t devaddr, uint16_t mmd_addr, uint32_t *data)

Set MMD address to mmd_addr with function MMD_FUNC_NOINCR and read contents to *data.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure
- **devaddr** -- Address of MDIO device
- **mmd_addr** -- Address of MDIO device register
- **data** -- [out] Data read from the device's memory

Returns

- ESP_OK: Memory read successfully
- ESP_FAIL: Memory read failed because of some error occurred
- ESP_ERR_INVALID_ARG: Device address provided is out of range (hardware limits device address to 5 bits)

esp_err_t **esp_eth_phy_802_3_write_mmd_register** (*phy_802_3_t* *phy_802_3, uint8_t devaddr, uint16_t mmd_addr, uint32_t data)

Set MMD address to mmd_addr with function MMD_FUNC_NOINCR and write data.

Parameters

- **phy_802_3** -- IEEE 802.3 PHY object infostructure

- **devaddr** -- Address of MDIO device
- **mmd_addr** -- Address of MDIO device register
- **data** -- [out] Data to write to the device's memory

Returns

- ESP_OK: Memory written to successfully
- ESP_FAIL: Memory write failed because of some error occurred
- ESP_ERR_INVALID_ARG: Device address provided is out of range (hardware limits device address to 5 bits)

inline [phy_802_3_t](#) ***esp_eth_phy_into_phy_802_3** ([esp_eth_phy_t](#) *phy)

Returns address to parent IEEE 802.3 PHY object infostructure.

Parameters **phy** -- Ethernet PHY instance

Returns [phy_802_3_t](#)*

- address to parent IEEE 802.3 PHY object infostructure

[esp_err_t](#) **esp_eth_phy_802_3_obj_config_init** ([phy_802_3_t](#) *phy_802_3, const [eth_phy_config_t](#) *config)

Initializes configuration of parent IEEE 802.3 PHY object infostructure.

Parameters

- **phy_802_3** -- Address to IEEE 802.3 PHY object infostructure
- **config** -- Configuration of the IEEE 802.3 PHY object

Returns

- ESP_OK: configuration initialized successfully
- ESP_ERR_INVALID_ARG: invalid `config` argument

Structures

struct **phy_802_3_t**

IEEE 802.3 PHY object infostructure.

Public Members

[esp_eth_phy_t](#) **parent**

Parent Ethernet PHY instance

[esp_eth_mediator_t](#) ***eth**

Mediator of Ethernet driver

int **addr**

PHY address

uint32_t **reset_timeout_ms**

Reset timeout value (Unit: ms)

uint32_t **autonego_timeout_ms**

Auto-negotiation timeout value (Unit: ms)

[eth_link_t](#) **link_status**

Current Link status

int **reset_gpio_num**

Reset GPIO number, -1 means no hardware reset

Enumerations

enum **esp_eth_phy_802_3_mmd_func_t**

IEEE 802.3 MMD modes enumeration.

Values:

enumerator **MMD_FUNC_ADDRESS**

enumerator **MMD_FUNC_DATA_NOINCR**

enumerator **MMD_FUNC_DATA_RWINCR**

enumerator **MMD_FUNC_DATA_WINCR**

Header File

- [components/esp_eth/include/esp_eth_netif_glue.h](#)
- This header file can be included with:

```
#include "esp_eth_netif_glue.h"
```

- This header file is a part of the API provided by the `esp_eth` component. To declare that your component depends on `esp_eth`, add the following to your CMakeLists.txt:

```
REQUIRES esp_eth
```

or

```
PRIV_REQUIRES esp_eth
```

Functions

esp_eth_netif_glue_handle_t **esp_eth_new_netif_glue** (*esp_eth_handle_t* eth_hdl)

Create a netif glue for Ethernet driver.

Note: netif glue is used to attach io driver to TCP/IP netif

Parameters **eth_hdl** -- Ethernet driver handle

Returns glue object, which inherits `esp_netif_driver_base_t`

esp_err_t **esp_eth_del_netif_glue** (*esp_eth_netif_glue_handle_t* eth_netif_glue)

Delete netif glue of Ethernet driver.

Parameters **eth_netif_glue** -- netif glue

Returns -ESP_OK: delete netif glue successfully

Type Definitions

typedef struct esp_eth_netif_glue_t ***esp_eth_netif_glue_handle_t**

Handle of netif glue - an intermediate layer between netif and Ethernet driver.

Code examples for the Ethernet API are provided in the [ethernet](#) directory of ESP-IDF examples.

2.4.2 Thread

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Thread

Introduction `Thread` is an IP-based mesh networking protocol. It is based on the 802.15.4 physical and MAC layer.

Application Examples The `openthread` directory of ESP-IDF examples contains the following applications:

- The OpenThread interactive shell `openthread/ot_cli`
- The Thread Border Router `openthread/ot_br`
- The Thread Radio Co-Processor `openthread/ot_rcp`
- `openthread/ot_trel` demonstrates Thread Radio Encapsulation Link (TREL) function. This requires a board equipped with a Wi-Fi module.
- `openthread/ot_sleepy_device/deep_sleep` demonstrates Thread Deep-sleep function.
- `openthread/ot_sleepy_device/light_sleep` demonstrates Thread Light-sleep function.

API Reference For manipulating the Thread network, the OpenThread API shall be used. The OpenThread API docs can be found at the [OpenThread API docs](#).

ESP-IDF provides extra APIs for launching and managing the OpenThread stack, binding to network interfaces and border routing features.

Header File

- `components/openthread/include/esp_openthread.h`
- This header file can be included with:

```
#include "esp_openthread.h"
```

- This header file is a part of the API provided by the `openthread` component. To declare that your component depends on `openthread`, add the following to your `CMakeLists.txt`:

```
REQUIRES openthread
```

or

```
PRIV_REQUIRES openthread
```

Functions

`esp_err_t esp_openthread_init` (const `esp_openthread_platform_config_t` *init_config)

Initializes the full OpenThread stack.

Note: The OpenThread instance will also be initialized in this function.

Parameters `init_config` -- [in] The initialization configuration.

Returns

- `ESP_OK` on success

- ESP_ERR_NO_MEM if allocation has failed
- ESP_ERR_INVALID_ARG if radio or host connection mode not supported
- ESP_ERR_INVALID_STATE if already initialized

esp_err_t **esp_openthread_auto_start** (otOperationalDatasetTlvs *datasetTlvs)

Starts the Thread protocol operation and attaches to a Thread network.

Parameters **datasetTlvs** -- [in] The operational dataset (TLV encoded), if it's NULL, the function will generate the dataset based on the configurations from kconfig.

Returns

- ESP_OK on success
- ESP_FAIL on failures

esp_err_t **esp_openthread_launch_mainloop** (void)

Launches the OpenThread main loop.

Note: This function will not return unless error happens when running the OpenThread stack.

Returns

- ESP_OK on success
- ESP_ERR_NO_MEM if allocation has failed
- ESP_FAIL on other failures

esp_err_t **esp_openthread_deinit** (void)

This function performs OpenThread stack and platform driver deinitialization.

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE if not initialized

otInstance ***esp_openthread_get_instance** (void)

This function acquires the underlying OpenThread instance.

Note: This function can be called on other tasks without lock.

Returns The OpenThread instance pointer

esp_err_t **esp_openthread_mainloop_exit** (void)

Signals the OpenThread main loop to exit.

Returns

- ESP_OK on success
- ESP_FAIL on failures

esp_err_t **esp_openthread_start** (const *esp_openthread_config_t* *config)

Starts the full OpenThread stack and create a handle task.

Note: The OpenThread instance will also be initialized in this function.

Parameters **config** -- [in] The OpenThread platform configuration.

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE if already initialized

esp_err_t **esp_openthread_stop** (void)

This function performs OpenThread stack and platform driver deinitialization and delete the handle task.

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE if Thread is already active

Header File

- [components/openthread/include/esp_openthread_types.h](#)
- This header file can be included with:

```
#include "esp_openthread_types.h"
```

- This header file is a part of the API provided by the `openthread` component. To declare that your component depends on `openthread`, add the following to your `CMakeLists.txt`:

```
REQUIRES openthread
```

or

```
PRIV_REQUIRES openthread
```

Structures

struct **esp_openthread_role_changed_event_t**

OpenThread role changed event data.

Public Members

otDeviceRole **previous_role**

Previous Thread role

otDeviceRole **current_role**

Current Thread role

struct **esp_openthread_dataset_changed_event_t**

OpenThread dataset changed event data.

Public Members

esp_openthread_dataset_type_t **type**

Dataset type

otOperationalDataset **new_dataset**

New dataset

struct **esp_openthread_mainloop_context_t**

This structure represents a context for a `select()` based mainloop.

Public Members

fd_set **read_fds**

The read file descriptors

fd_set **write_fds**

The write file descriptors

fd_set **error_fds**

The error file descriptors

int **max_fd**

The max file descriptor

struct timeval **timeout**

The timeout

struct **esp_openthread_uart_config_t**

The uart port config for OpenThread.

Public Members

[*uart_port_t*](#) **port**

UART port number

[*uart_config_t*](#) **uart_config**

UART configuration, see [*uart_config_t*](#) docs

gpio_num_t **rx_pin**

UART RX pin

gpio_num_t **tx_pin**

UART TX pin

struct **esp_openthread_spi_host_config_t**

The spi port config for OpenThread.

Public Members

[*spi_host_device_t*](#) **host_device**

SPI host device

[*spi_dma_chan_t*](#) **dma_channel**

DMA channel

[*spi_bus_config_t*](#) **spi_interface**

SPI bus

spi_device_interface_config_t **spi_device**

SPI peripheral device

gpio_num_t **intr_pin**

SPI interrupt pin

struct **esp_openthread_spi_slave_config_t**

The spi slave config for OpenThread.

Public Members

spi_host_device_t **host_device**

SPI host device

spi_bus_config_t **bus_config**

SPI bus config

spi_slave_interface_config_t **slave_config**

SPI slave config

gpio_num_t **intr_pin**

SPI interrupt pin

struct **esp_openthread_radio_config_t**

The OpenThread radio configuration.

Public Members

esp_openthread_radio_mode_t **radio_mode**

The radio mode

esp_openthread_uart_config_t **radio_uart_config**

The uart configuration to RCP

esp_openthread_spi_host_config_t **radio_spi_config**

The spi configuration to RCP

struct **esp_openthread_host_connection_config_t**

The OpenThread host connection configuration.

Public Members

esp_openthread_host_connection_mode_t **host_connection_mode**

The host connection mode

esp_openthread_uart_config_t **host_uart_config**

The uart configuration to host

usb_serial_jtag_driver_config_t **host_usb_config**

The usb configuration to host

esp_openthread_spi_slave_config_t **spi_slave_config**

The spi configuration to host

struct **esp_openthread_port_config_t**

The OpenThread port specific configuration.

Public Members

const char ***storage_partition_name**

The partition for storing OpenThread dataset

uint8_t **netif_queue_size**

The packet queue size for the network interface

uint8_t **task_queue_size**

The task queue size

struct **esp_openthread_platform_config_t**

The OpenThread platform configuration.

Public Members

esp_openthread_radio_config_t **radio_config**

The radio configuration

esp_openthread_host_connection_config_t **host_config**

The host connection configuration

esp_openthread_port_config_t **port_config**

The port configuration

struct **esp_openthread_config_t**

The OpenThread configuration.

Public Members

esp_netif_config_t **netif_config**

The netif configuration

esp_openthread_platform_config_t **platform_config**

The platform configuration

Type Definitions

typedef void (***esp_openthread_rcp_failure_handler**)(void)

The OpenThread rcp failure handler.

typedef void (***esp_openthread_compatibility_error_callback**)(void)

The OpenThread compatibility error callback.

typedef void (***esp_openthread_coprocessor_reset_failure_callback**)(void)

The OpenThread co-processor reset failure callback.

Enumerations

enum **esp_openthread_event_t**

OpenThread event declarations.

Values:

enumerator **OPENTHREAD_EVENT_START**

OpenThread stack start

enumerator **OPENTHREAD_EVENT_STOP**

OpenThread stack stop

enumerator **OPENTHREAD_EVENT_DETACHED**

OpenThread detached

enumerator **OPENTHREAD_EVENT_ATTACHED**

OpenThread attached

enumerator **OPENTHREAD_EVENT_ROLE_CHANGED**

OpenThread role changed

enumerator **OPENTHREAD_EVENT_IF_UP**

OpenThread network interface up

enumerator **OPENTHREAD_EVENT_IF_DOWN**

OpenThread network interface down

enumerator **OPENTHREAD_EVENT_GOT_IP6**

OpenThread stack added IPv6 address

enumerator **OPENTHREAD_EVENT_LOST_IP6**

OpenThread stack removed IPv6 address

enumerator **OPENTHREAD_EVENT_MULTICAST_GROUP_JOIN**

OpenThread stack joined IPv6 multicast group

enumerator **OPENTHREAD_EVENT_MULTICAST_GROUP_LEAVE**

OpenThread stack left IPv6 multicast group

enumerator **OPENTHREAD_EVENT_TREL_ADD_IP6**

OpenThread stack added TREL IPv6 address

enumerator **OPENTHREAD_EVENT_TREL_REMOVE_IP6**

OpenThread stack removed TREL IPv6 address

enumerator **OPENTHREAD_EVENT_TREL_MULTICAST_GROUP_JOIN**

OpenThread stack joined TREL IPv6 multicast group

enumerator **OPENTHREAD_EVENT_SET_DNS_SERVER**

OpenThread stack set DNS server >

enumerator **OPENTHREAD_EVENT_PUBLISH_MESHCOPE**

OpenThread stack start to publish meshcop-e service >

enumerator **OPENTHREAD_EVENT_REMOVE_MESHCOPE**

OpenThread stack start to remove meshcop-e service >

enumerator **OPENTHREAD_EVENT_DATASET_CHANGED**

OpenThread dataset changed >

enum **esp_openthread_dataset_type_t**

OpenThread dataset type.

Values:

enumerator **OPENTHREAD_ACTIVE_DATASET**

Active dataset

enumerator **OPENTHREAD_PENDING_DATASET**

Pending dataset

enum **esp_openthread_radio_mode_t**

The radio mode of OpenThread.

Values:

enumerator **RADIO_MODE_NATIVE**

Use the native 15.4 radio

enumerator **RADIO_MODE_UART_RCP**

UART connection to a 15.4 capable radio co-processor (RCP)

enumerator **RADIO_MODE_SPI_RCP**

SPI connection to a 15.4 capable radio co-processor (RCP)

enumerator **RADIO_MODE_TREL**

Use the Thread Radio Encapsulation Link (TREL)

enumerator **RADIO_MODE_MAX**

Using for parameter check

enum **esp_openthread_host_connection_mode_t**

How OpenThread connects to the host.

Values:

enumerator **HOST_CONNECTION_MODE_NONE**

Disable host connection

enumerator **HOST_CONNECTION_MODE_CLI_UART**

CLI UART connection to the host

enumerator **HOST_CONNECTION_MODE_CLI_USB**

CLI USB connection to the host

enumerator **HOST_CONNECTION_MODE_RCP_UART**

RCP UART connection to the host

enumerator **HOST_CONNECTION_MODE_RCP_SPI**

RCP SPI connection to the host

enumerator **HOST_CONNECTION_MODE_MAX**

Using for parameter check

Header File

- [components/openthread/include/esp_openthread_lock.h](#)
- This header file can be included with:

```
#include "esp_openthread_lock.h"
```

- This header file is a part of the API provided by the `openthread` component. To declare that your component depends on `openthread`, add the following to your `CMakeLists.txt`:

```
REQUIRES openthread
```

or

```
PRIV_REQUIRES openthread
```

Functions

esp_err_t **esp_openthread_lock_init** (void)

This function initializes the OpenThread API lock.

Returns

- ESP_OK on success
- ESP_ERR_NO_MEM if allocation has failed
- ESP_ERR_INVALID_STATE if already initialized

void **esp_openthread_lock_deinit** (void)

This function deinitializes the OpenThread API lock.

bool **esp_openthread_lock_acquire** (TickType_t block_ticks)

This function acquires the OpenThread API lock.

Note: Every OpenThread APIs that takes an `otInstance` argument **MUST** be protected with this API lock except that the call site is in OpenThread callbacks.

Parameters **block_ticks** -- [in] The maximum number of RTOS ticks to wait for the lock.

Returns

- True on lock acquired
- False on failing to acquire the lock with the timeout.

void **esp_openthread_lock_release** (void)

This function releases the OpenThread API lock.

bool **esp_openthread_task_switching_lock_acquire** (TickType_t block_ticks)

This function acquires the OpenThread API task switching lock.

Note: In OpenThread API context, it waits for some actions to be done in other tasks (like lwip), after task switching, it needs to call OpenThread API again. Normally it's not allowed, since the previous OpenThread API lock is not released yet. This task_switching lock allows the OpenThread API can be called in this case.

Note: Please use esp_openthread_lock_acquire() for normal cases.

Parameters **block_ticks** -- [in] The maximum number of RTOS ticks to wait for the lock.

Returns

- True on lock acquired
- False on failing to acquire the lock with the timeout.

void **esp_openthread_task_switching_lock_release** (void)

This function releases the OpenThread API task switching lock.

Note: This API must be called after esp_openthread_task_switching_lock_acquire or esp_openthread_lock_acquire and will cause a crash if the current task is not the task switching lock holder. This error could be caused by calling OpenThread APIs without locking OpenThread stack.

Header File

- [components/openthread/include/esp_openthread_netif_glue.h](#)
- This header file can be included with:

```
#include "esp_openthread_netif_glue.h"
```

- This header file is a part of the API provided by the openthread component. To declare that your component depends on openthread, add the following to your CMakeLists.txt:

```
REQUIRES openthread
```

or

```
PRIV_REQUIRES openthread
```

Functions

void ***esp_openthread_netif_glue_init** (const *esp_openthread_platform_config_t* *config)

This function initializes the OpenThread network interface glue.

Parameters **config** -- [in] The platform configuration.

Returns

- glue pointer on success
- NULL on failure

void **esp_openthread_netif_glue_deinit** (void)

This function deinitializes the OpenThread network interface glue.

esp_netif_t ***esp_openthread_get_netif** (void)

This function acquires the OpenThread netif.

Returns The OpenThread netif or NULL if not initialized.

void **esp_openthread_register_meshcop_e_handler** (*esp_event_handler_t* handler, bool for_publish)

This function register a handler for meshcop-e service publish event and remove event.

Parameters

- **handler** -- [in] The handler.
- **for_publish** -- [in] The usage of handler, true for publish event and false for remove event.

bool **is_openthread_internal_mesh_local_addr** (const otIp6Address *address)

This function judges the target address is openthread mesh local or not.

Parameters **address** -- [in] The address.

Returns

- True if the address is openthread mesh local, otherwise false

Macros

ESP_NETIF_INHERENT_DEFAULT_OPENTHREAD ()

Default configuration reference of OpenThread esp-netif.

ESP_NETIF_DEFAULT_OPENTHREAD ()

Header File

- [components/openthread/include/esp_openthread_border_router.h](#)
- This header file can be included with:

```
#include "esp_openthread_border_router.h"
```

- This header file is a part of the API provided by the `openthread` component. To declare that your component depends on `openthread`, add the following to your CMakeLists.txt:

```
REQUIRES openthread
```

or

```
PRIV_REQUIRES openthread
```

Functions

void **esp_openthread_set_backbone_netif** (*esp_netif_t* *backbone_netif)

Sets the backbone interface used for border routing.

Note: This function must be called before `esp_openthread_init`

Parameters **backbone_netif** -- [in] The backbone network interface (WiFi or ethernet)

esp_err_t **esp_openthread_border_router_init** (void)

Initializes the border router features of OpenThread.

Note: Calling this function will make the device behave as an OpenThread border router. Kconfig option `CONFIG_OPENTHREAD_BORDER_ROUTER` is required.

Returns

- `ESP_OK` on success
- `ESP_ERR_NOT_SUPPORTED` if feature not supported
- `ESP_ERR_INVALID_STATE` if already initialized
- `ESP_FIAL` on other failures

esp_err_t **esp_openthread_border_router_deinit** (void)

Deinitializes the border router features of OpenThread.

Returns

- `ESP_OK` on success
- `ESP_ERR_INVALID_STATE` if not initialized
- `ESP_FIAL` on other failures

esp_netif_t ***esp_openthread_get_backbone_netif** (void)

Gets the backbone interface of OpenThread border router.

Returns The backbone interface or NULL if border router not initialized.

esp_err_t **esp_openthread_set_meshcop_instance_name** (const char *instance_name)

Sets the meshcop(e) instance name.

Note: This function can only be called before `esp_openthread_border_router_init`. If `instance_name` is NULL, then the service will use the hostname as instance name.

Parameters `instance_name` -- [in] The instance name, can be NULL.

Returns

- `ESP_OK` on success
- `ESP_FAIL` if fail to initialize RCP

const char ***esp_openthread_get_meshcop_instance_name** (void)

Gets the meshcop(e) instance name.

Returns The instance name.

Thread is an IPv6-based mesh networking technology for IoT.

Code examples for the Thread API are provided in the [openthread](#) directory of ESP-IDF examples.

2.4.3 ESP-NETIF

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

ESP-NETIF

The purpose of the ESP-NETIF library is twofold:

- It provides an abstraction layer for the application on top of the TCP/IP stack. This allows applications to choose between IP stacks in the future.

- The APIs it provides are thread-safe, even if the underlying TCP/IP stack APIs are not.


ESP-IDF currently implements ESP-NETIF for the lwIP TCP/IP stack only. However, the adapter itself is TCP/IP implementation-agnostic and allows different implementations.

It is also possible to use a custom TCP/IP stack with ESP-IDF, provided it implements BSD API. For more information on building ESP-IDF without lwIP, please refer to [components/esp_netif_stack/README.md](#).

Some ESP-NETIF API functions are intended to be called by application code, for example, to get or set interface IP addresses, and configure DHCP. Other functions are intended for internal ESP-IDF use by the network driver layer.


In many cases, applications do not need to call ESP-NETIF APIs directly as they are called by the default network event handlers.

ESP-NETIF Architecture



(continues on next page)

(continued from previous page)



Data and Event Flow in the Diagram

- Initialization line from user code to ESP-NETIF and communication driver
- --<--->-- Data packets going from communication media to TCP/IP stack and back
- ***** Events aggregated in ESP-NETIF propagate to the driver, user code, and network stack
- | User settings and runtime configuration

ESP-NETIF Interaction

A) User Code, Boilerplate Overall application interaction with a specific IO driver for communication media and configured TCP/IP network stack is abstracted using ESP-NETIF APIs and is outlined as below:

A) Initialization code

- 1) Initializes IO driver
- 2) Creates a new instance of ESP-NETIF and configure it with
 - ESP-NETIF specific options (flags, behavior, name)
 - Network stack options (netif init and input functions, not publicly available)
 - IO driver specific options (transmit, free rx buffer functions, IO driver handle)
- 3) Attaches the IO driver handle to the ESP-NETIF instance created in the above steps
- 4) Configures event handlers
 - Use default handlers for common interfaces defined in IO drivers; or define a specific handler for customized behavior or new interfaces
 - Register handlers for app-related events (such as IP lost or acquired)

B) Interaction with network interfaces using ESP-NETIF API

- 1) Gets and sets TCP/IP-related parameters (DHCP, IP, etc)
- 2) Receives IP events (connect or disconnect)
- 3) Controls application lifecycle (set interface up or down)

B) Communication Driver, IO Driver, and Media Driver Communication driver plays these two important roles in relation to ESP-NETIF:

- 1) Event handlers: Defines behavior patterns of interaction with ESP-NETIF (e.g., ethernet link-up -> turn netif on)
- 2) Glue IO layer: Adapts the input or output functions to use ESP-NETIF transmit, receive, and free receive buffer
 - Installs driver_transmit to the appropriate ESP-NETIF object so that outgoing packets from the network stack are passed to the IO driver
 - Calls `esp_netif_receive()` to pass incoming data to the network stack

C) ESP-NETIF ESP-NETIF serves as an intermediary between an IO driver and a network stack, connecting the packet data path between the two. It provides a set of interfaces for attaching a driver to an ESP-NETIF object at runtime and configures a network stack during compiling. Additionally, a set of APIs is provided to control the network interface lifecycle and its TCP/IP properties. As an overview, the ESP-NETIF public interface can be divided into six groups:

- 1) Initialization APIs (to create and configure ESP-NETIF instance)
- 2) Input or Output API (for passing data between IO driver and network stack)
- 3) Event or Action API
 - Used for network interface lifecycle management
 - ESP-NETIF provides building blocks for designing event handlers
- 4) Setters and Getters API for basic network interface properties
- 5) Network stack abstraction API: enabling user interaction with TCP/IP stack
 - Set interface up or down
 - DHCP server and client API
 - DNS API
 - [SNTP API](#)
- 6) Driver conversion utilities API

D) Network Stack The network stack has no public interaction with application code with regard to public interfaces and shall be fully abstracted by ESP-NETIF API.

E) ESP-NETIF L2 TAP Interface The ESP-NETIF L2 TAP interface is a mechanism in ESP-IDF used to access Data Link Layer (L2 per OSI/ISO) for frame reception and transmission from the user application. Its typical usage in the embedded world might be the implementation of non-IP-related protocols, e.g., PTP, Wake on LAN. Note that only Ethernet (IEEE 802.3) is currently supported.

From a user perspective, the ESP-NETIF L2 TAP interface is accessed using file descriptors of VFS, which provides file-like interfacing (using functions like `open()`, `read()`, `write()`, etc). To learn more, refer to [Virtual Filesystem Component](#).

There is only one ESP-NETIF L2 TAP interface device (path name) available. However multiple file descriptors with different configurations can be opened at a time since the ESP-NETIF L2 TAP interface can be understood as a generic entry point to the Layer 2 infrastructure. What is important is then the specific configuration of the particular file descriptor. It can be configured to give access to a specific Network Interface identified by `if_key` (e.g., `ETH_DEF`) and to filter only specific frames based on their type (e.g., Ethernet type in the case of IEEE 802.3). Filtering only specific frames is crucial since the ESP-NETIF L2 TAP needs to exist along with the IP stack and so the IP-related traffic (IP, ARP, etc.) should not be passed directly to the user application. Even though this option is still configurable, it is not recommended in standard use cases. Filtering is also advantageous from the perspective of the user's application, as it only gets access to the frame types it is interested in, and the remaining traffic is either passed to other L2 TAP file descriptors or to the IP stack.

ESP-NETIF L2 TAP Interface Usage Manual

Initialization To be able to use the ESP-NETIF L2 TAP interface, it needs to be enabled in Kconfig by [CONFIG_ESP_NETIF_L2_TAP](#) first and then registered by `esp_vfs_l2tap_intf_register()` prior usage of any VFS function.

`open()` Once the ESP-NETIF L2 TAP is registered, it can be opened at path name `"/dev/net/tap"`. The same path name can be opened multiple times up to [CONFIG_ESP_NETIF_L2_TAP_MAX_FDS](#) and multiple file descriptors with a different configuration may access the Data Link Layer frames.

The ESP-NETIF L2 TAP can be opened with the `O_NONBLOCK` file status flag to make sure the `read()` does not block. Note that the `write()` may block in the current implementation when accessing a Network interface since

it is a shared resource among multiple ESP-NETIF L2 TAP file descriptors and IP stack, and there is currently no queuing mechanism deployed. The file status flag can be retrieved and modified using `fcntl()`.

On success, `open()` returns the new file descriptor (a nonnegative integer). On error, -1 is returned, and `errno` is set to indicate the error.

`ioctl()` The newly opened ESP-NETIF L2 TAP file descriptor needs to be configured prior to its usage since it is not bounded to any specific Network Interface and no frame type filter is configured. The following configuration options are available to do so:

- `L2TAP_S_INTF_DEVICE` - bounds the file descriptor to a specific Network Interface that is identified by its `if_key`. ESP-NETIF Network Interface `if_key` is passed to `ioctl()` as the third parameter. Note that default Network Interfaces `if_key`'s used in ESP-IDF can be found in [esp_netif/include/esp_netif_defaults.h](#).
- `L2TAP_S_DEVICE_DRV_HNDL` - is another way to bound the file descriptor to a specific Network Interface. In this case, the Network interface is identified directly by IO Driver handle (e.g., `esp_eth_handle_t` in case of Ethernet). The IO Driver handle is passed to `ioctl()` as the third parameter.
- `L2TAP_S_RCV_FILTER` - sets the filter to frames with the type to be passed to the file descriptor. In the case of Ethernet frames, the frames are to be filtered based on the Length and Ethernet type field. In case the filter value is set less than or equal to 0x05DC, the Ethernet type field is considered to represent IEEE802.3 Length Field, and all frames with values in interval $[0, 0x05DC]$ at that field are passed to the file descriptor. The IEEE802.2 logical link control (LLC) resolution is then expected to be performed by the user's application. In case the filter value is set greater than 0x05DC, the Ethernet type field is considered to represent protocol identification and only frames that are equal to the set value are to be passed to the file descriptor.

All above-set configuration options have a getter counterpart option to read the current settings.

Warning: The file descriptor needs to be firstly bounded to a specific Network Interface by `L2TAP_S_INTF_DEVICE` or `L2TAP_S_DEVICE_DRV_HNDL` to make `L2TAP_S_RCV_FILTER` option available.

Note: VLAN-tagged frames are currently not recognized. If the user needs to process VLAN-tagged frames, they need a set filter to be equal to the VLAN tag (i.e., 0x8100 or 0x88A8) and process the VLAN-tagged frames in the user application.

Note: `L2TAP_S_DEVICE_DRV_HNDL` is particularly useful when the user's application does not require the usage of an IP stack and so ESP-NETIF is not required to be initialized too. As a result, Network Interface cannot be identified by its `if_key` and hence it needs to be identified directly by its IO Driver handle.

On success, `ioctl()` returns 0. On error, -1 is returned, and `errno` is set to indicate the error.

- * `EBADF` - not a valid file descriptor.
- * `EACCES` - options change is denied in this state (e.g., file descriptor has not been bounded to Network interface yet).
- * `EINVAL` - invalid configuration argument. Ethernet type filter is already used by other file descriptors on that same Network interface.
- * `ENODEV` - no such Network Interface which is tried to be assigned to the file descriptor exists.
- * `ENOSYS` - unsupported operation, passed configuration option does not exist.

`fcntl()` `fcntl()` is used to manipulate with properties of opened ESP-NETIF L2 TAP file descriptor.

The following commands manipulate the status flags associated with the file descriptor:

- `F_GETFD` - the function returns the file descriptor flags, and the third argument is ignored.

- `F_SETFD` - sets the file descriptor flags to the value specified by the third argument. Zero is returned.

On success, `ioctl()` returns 0. On error, -1 is returned, and `errno` is set to indicate the error.

* `EBADF` - not a valid file descriptor.

* `ENOSYS` - unsupported command.

read() Opened and configured ESP-NETIF L2 TAP file descriptor can be accessed by `read()` to get inbound frames. The read operation can be either blocking or non-blocking based on the actual state of the `O_NONBLOCK` file status flag. When the file status flag is set to blocking, the read operation waits until a frame is received and the context is switched to other tasks. When the file status flag is set to non-blocking, the read operation returns immediately. In such case, either a frame is returned if it was already queued or the function indicates the queue is empty. The number of queued frames associated with one file descriptor is limited by [`CONFIG\_ESP\_NETIF\_L2\_TAP\_RX\_QUEUE\_SIZE`](#) Kconfig option. Once the number of queued frames reached a configured threshold, the newly arrived frames are dropped until the queue has enough room to accept incoming traffic (Tail Drop queue management).

On success, `read()` returns the number of bytes read. Zero is returned when the size of the destination buffer is 0. On error, -1 is returned, and `errno` is set to indicate the error.

* `EBADF` - not a valid file descriptor.

* `EAGAIN` - the file descriptor has been marked non-blocking (`O_NONBLOCK`), and the read would block.

write() A raw Data Link Layer frame can be sent to Network Interface via opened and configured ESP-NETIF L2 TAP file descriptor. The user's application is responsible to construct the whole frame except for fields which are added automatically by the physical interface device. The following fields need to be constructed by the user's application in case of an Ethernet link: source/destination MAC addresses, Ethernet type, actual protocol header, and user data. The length of these fields is as follows:

Destination MAC	Source MAC	Type/Length	Payload (protocol header/data)
6 B	6 B	2 B	0-1486 B

In other words, there is no additional frame processing performed by the ESP-NETIF L2 TAP interface. It only checks the Ethernet type of the frame is the same as the filter configured in the file descriptor. If the Ethernet type is different, an error is returned and the frame is not sent. Note that the `write()` may block in the current implementation when accessing a Network interface since it is a shared resource among multiple ESP-NETIF L2 TAP file descriptors and IP stack, and there is currently no queuing mechanism deployed.

On success, `write()` returns the number of bytes written. Zero is returned when the size of the input buffer is 0. On error, -1 is returned, and `errno` is set to indicate the error.

* `EBADF` - not a valid file descriptor.

* `EBADMSG` - The Ethernet type of the frame is different from the file descriptor configured filter.

* `EIO` - Network interface not available or busy.

close() Opened ESP-NETIF L2 TAP file descriptor can be closed by the `close()` to free its allocated resources. The ESP-NETIF L2 TAP implementation of `close()` may block. On the other hand, it is thread-safe and can be called from a different task than the file descriptor is actually used. If such a situation occurs and one task is blocked in the I/O operation and another task tries to close the file descriptor, the first task is unblocked. The first's task read operation then ends with an error.

On success, `close()` returns zero. On error, -1 is returned, and `errno` is set to indicate the error.

* `EBADF` - not a valid file descriptor.

select () Select is used in a standard way, just `CONFIG_VFS_SUPPORT_SELECT` needs to be enabled to make the `select ()` function available.

SNTP API You can find a brief introduction to SNTP in general, its initialization code, and basic modes in Section *SNTP Time Synchronization* in *System Time*.

This section provides more details about specific use cases of the SNTP service, with statically configured servers, or use the DHCP-provided servers, or both. The workflow is usually very simple:

- 1) Initialize and configure the service using `esp_netif_sntp_init ()`. This operations can only be called once (unless the SNTP service has been destroyed by `esp_netif_sntp_deinit ()`)
- 2) Start the service via `esp_netif_sntp_start ()`. This step is not needed if we auto-started the service in the previous step (default). It is useful to start the service explicitly after connecting if we want to use the DHCP-obtained NTP servers. Please note, this option needs to be enabled before connecting, but the SNTP service should be started after.
- 3) Wait for the system time to synchronize using `esp_netif_sntp_sync_wait ()` (only if needed).
- 4) Stop and destroy the service using `esp_netif_sntp_deinit ()`.

Basic Mode with Statically Defined Server(s) Initialize the module with the default configuration after connecting to the network. Note that it is possible to provide multiple NTP servers in the configuration struct:

```
esp_sntp_config_t config = ESP_NETIF_SNTP_DEFAULT_CONFIG_MULTIPLE(2,  
                                                                    ESP_SNTP_SERVER_LIST("time.windows.com", "pool.ntp.org"  
↪) );  
esp_netif_sntp_init(&config);
```

Note: If we want to configure multiple SNTP servers, we have to update the lwIP configuration `CONFIG_LWIP_SNTP_MAX_SERVERS`.

Use DHCP-Obtained SNTP Server(s) First of all, we have to enable the lwIP configuration option `CONFIG_LWIP_DHCP_GET_NTP_SRV`. Then we have to initialize the SNTP module with the DHCP option and without the NTP server:

```
esp_sntp_config_t config = ESP_NETIF_SNTP_DEFAULT_CONFIG_MULTIPLE(0, {});  
config.start = false; // start the SNTP service explicitly  
config.server_from_dhcp = true; // accept the NTP offer from the DHCP_  
↪server  
esp_netif_sntp_init(&config);
```

Then, once we are connected, we could start the service using:

```
esp_netif_sntp_start();
```

Note: It is also possible to start the service during initialization (default `config.start=true`). This would likely to cause the initial SNTP request to fail (since we are not connected yet) and lead to some back-off time for subsequent requests.

Use Both Static and Dynamic Servers Very similar to the scenario above (DHCP provided SNTP server), but in this configuration, we need to make sure that the static server configuration is refreshed when obtaining NTP servers by DHCP. The underlying lwIP code cleans up the rest of the list of NTP servers when the DHCP-provided information gets accepted. Thus the ESP-NETIF SNTP module saves the statically configured server(s) and reconfigures them after obtaining a DHCP lease.

The typical configuration now looks as per below, providing the specific `IP_EVENT` to update the config and index of the first server to reconfigure (for example setting `config.index_of_first_server=1` would keep the DHCP provided server at index 0, and the statically configured server at index 1).

```
esp_sntp_config_t config = ESP_NETIF_SNTP_DEFAULT_CONFIG("pool.ntp.org");
config.start = false; // start the SNTP service explicitly
↳ (after connecting)
config.server_from_dhcp = true; // accept the NTP offers from DHCP
↳ server
config.renew_servers_after_new_IP = true; // let esp-netif update the configured
↳ SNTP server(s) after receiving the DHCP lease
config.index_of_first_server = 1; // updates from server num 1, leaving
↳ server 0 (from DHCP) intact
config.ip_event_to_renew = IP_EVENT_STA_GOT_IP; // IP event on which we refresh
↳ the configuration
```

Then we start the service normally with `esp_netif_sntp_start()`.

ESP-NETIF Programmer's Manual Please refer to the following example to understand the initialization process of the default interface:

- Ethernet: [ethernet/basic/main/ethernet_example_main.c](#)
- L2 TAP: [protocols/l2tap/main/l2tap_main.c](#)
- Wi-Fi Access Point: [wifi/getting_started/softAP/main/softap_example_main.c](#)

For more specific cases, please consult this guide: [ESP-NETIF Custom I/O Driver](#).

API Reference

Header File

- [components/esp_netif/include/esp_netif.h](#)
- This header file can be included with:

```
#include "esp_netif.h"
```

- This header file is a part of the API provided by the `esp_netif` component. To declare that your component depends on `esp_netif`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_netif
```

or

```
PRIV_REQUIRES esp_netif
```

Functions

[esp_err_t](#) **esp_netif_init** (void)

Initialize the underlying TCP/IP stack.

Note: This function should be called exactly once from application code, when the application starts up.

Returns

- `ESP_OK` on success
- `ESP_FAIL` if initializing failed

esp_err_t **esp_netif_deinit** (void)

Deinitialize the esp-netif component (and the underlying TCP/IP stack)

Note: Deinitialization **is not** supported yet

Returns

- ESP_ERR_INVALID_STATE if esp_netif not initialized
- ESP_ERR_NOT_SUPPORTED otherwise

esp_netif_t ***esp_netif_new** (const *esp_netif_config_t* *esp_netif_config)

Creates an instance of new esp-netif object based on provided config.

Parameters *esp_netif_config* -- pointer esp-netif configuration

Returns

- pointer to esp-netif object on success
- NULL otherwise

void **esp_netif_destroy** (*esp_netif_t* *esp_netif)

Destroys the esp_netif object.

Parameters *esp_netif* -- [in] pointer to the object to be deleted

esp_err_t **esp_netif_set_driver_config** (*esp_netif_t* *esp_netif, const *esp_netif_driver_ifconfig_t* *driver_config)

Configures driver related options of esp_netif object.

Parameters

- *esp_netif* -- [inout] pointer to the object to be configured
- *driver_config* -- [in] pointer esp-netif io driver related configuration

Returns

- ESP_OK on success
- ESP_ERR_ESP_NETIF_INVALID_PARAMS if invalid parameters provided

esp_err_t **esp_netif_attach** (*esp_netif_t* *esp_netif, *esp_netif_io_driver_handle* driver_handle)

Attaches esp_netif instance to the io driver handle.

Calling this function enables connecting specific esp_netif object with already initialized io driver to update esp_netif object with driver specific configuration (i.e. calls post_attach callback, which typically sets io driver callbacks to esp_netif instance and starts the driver)

Parameters

- *esp_netif* -- [inout] pointer to esp_netif object to be attached
- *driver_handle* -- [in] pointer to the driver handle

Returns

- ESP_OK on success
- ESP_ERR_ESP_NETIF_DRIVER_ATTACH_FAILED if driver's post_attach callback failed

esp_err_t **esp_netif_receive** (*esp_netif_t* *esp_netif, void *buffer, size_t len, void *eb)

Passes the raw packets from communication media to the appropriate TCP/IP stack.

This function is called from the configured (peripheral) driver layer. The data are then forwarded as frames to the TCP/IP stack.

Parameters

- *esp_netif* -- [in] Handle to esp-netif instance
- *buffer* -- [in] Received data
- *len* -- [in] Length of the data frame
- *eb* -- [in] Pointer to internal buffer (used in Wi-Fi driver)

Returns

- ESP_OK

void **esp_netif_action_start** (void *esp_netif, esp_event_base_t base, int32_t event_id, void *data)

Default building block for network interface action upon IO driver start event. Creates network interface, if AUTOUP enabled turns the interface on, if DHCP enabled starts dhcp server.

Note: This API can be directly used as event handler

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **base** -- The base type of the event
- **event_id** -- The specific ID of the event
- **data** -- Optional data associated with the event

void **esp_netif_action_stop** (void *esp_netif, esp_event_base_t base, int32_t event_id, void *data)

Default building block for network interface action upon IO driver stop event.

Note: This API can be directly used as event handler

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **base** -- The base type of the event
- **event_id** -- The specific ID of the event
- **data** -- Optional data associated with the event

void **esp_netif_action_connected** (void *esp_netif, esp_event_base_t base, int32_t event_id, void *data)

Default building block for network interface action upon IO driver connected event.

Note: This API can be directly used as event handler

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **base** -- The base type of the event
- **event_id** -- The specific ID of the event
- **data** -- Optional data associated with the event

void **esp_netif_action_disconnected** (void *esp_netif, esp_event_base_t base, int32_t event_id, void *data)

Default building block for network interface action upon IO driver disconnected event.

Note: This API can be directly used as event handler

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **base** -- The base type of the event
- **event_id** -- The specific ID of the event
- **data** -- Optional data associated with the event

void **esp_netif_action_got_ip** (void *esp_netif, esp_event_base_t base, int32_t event_id, void *data)

Default building block for network interface action upon network got IP event.

Note: This API can be directly used as event handler

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **base** -- The base type of the event
- **event_id** -- The specific ID of the event
- **data** -- Optional data associated with the event

void **esp_netif_action_join_ip6_multicast_group** (void *esp_netif, esp_event_base_t base, int32_t event_id, void *data)

Default building block for network interface action upon IPv6 multicast group join.

Note: This API can be directly used as event handler

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **base** -- The base type of the event
- **event_id** -- The specific ID of the event
- **data** -- Optional data associated with the event

void **esp_netif_action_leave_ip6_multicast_group** (void *esp_netif, esp_event_base_t base, int32_t event_id, void *data)

Default building block for network interface action upon IPv6 multicast group leave.

Note: This API can be directly used as event handler

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **base** -- The base type of the event
- **event_id** -- The specific ID of the event
- **data** -- Optional data associated with the event

void **esp_netif_action_add_ip6_address** (void *esp_netif, esp_event_base_t base, int32_t event_id, void *data)

Default building block for network interface action upon IPv6 address added by the underlying stack.

Note: This API can be directly used as event handler

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **base** -- The base type of the event
- **event_id** -- The specific ID of the event
- **data** -- Optional data associated with the event

void **esp_netif_action_remove_ip6_address** (void *esp_netif, esp_event_base_t base, int32_t event_id, void *data)

Default building block for network interface action upon IPv6 address removed by the underlying stack.

Note: This API can be directly used as event handler

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **base** -- The base type of the event
- **event_id** -- The specific ID of the event
- **data** -- Optional data associated with the event

esp_err_t **esp_netif_set_default_netif** (*esp_netif_t* *esp_netif)

Manual configuration of the default netif.

This API overrides the automatic configuration of the default interface based on the route_prio. If the selected netif is set default using this API, no other interface could be set-default disregarding its route_prio number (unless the selected netif gets destroyed).

Parameters **esp_netif** -- [in] Handle to esp-netif instance

Returns ESP_OK on success

esp_netif_t ***esp_netif_get_default_netif** (void)

Getter function of the default netif.

This API returns the selected default netif.

Returns Handle to esp-netif instance of the default netif.

esp_err_t **esp_netif_join_ip6_multicast_group** (*esp_netif_t* *esp_netif, const *esp_ip6_addr_t* *addr)

Cause the TCP/IP stack to join a IPv6 multicast group.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **addr** -- [in] The multicast group to join

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS
- ESP_ERR_ESP_NETIF_MLD6_FAILED
- ESP_ERR_NO_MEM

esp_err_t **esp_netif_leave_ip6_multicast_group** (*esp_netif_t* *esp_netif, const *esp_ip6_addr_t* *addr)

Cause the TCP/IP stack to leave a IPv6 multicast group.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **addr** -- [in] The multicast group to leave

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS
- ESP_ERR_ESP_NETIF_MLD6_FAILED
- ESP_ERR_NO_MEM

esp_err_t **esp_netif_set_mac** (*esp_netif_t* *esp_netif, uint8_t mac[])

Set the mac address for the interface instance.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **mac** -- [in] Desired mac address for the related network interface

Returns

- ESP_OK - success
- ESP_ERR_ESP_NETIF_IF_NOT_READY - interface status error
- ESP_ERR_NOT_SUPPORTED - mac not supported on this interface

esp_err_t **esp_netif_get_mac** (*esp_netif_t* *esp_netif, uint8_t mac[])

Get the mac address for the interface instance.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **mac** -- [out] Resultant mac address for the related network interface

Returns

- ESP_OK - success
- ESP_ERR_ESP_NETIF_IF_NOT_READY - interface status error
- ESP_ERR_NOT_SUPPORTED - mac not supported on this interface

esp_err_t **esp_netif_set_hostname** (*esp_netif_t* *esp_netif, const char *hostname)

Set the hostname of an interface.

The configured hostname overrides the default configuration value CONFIG_LWIP_LOCAL_HOSTNAME. Please note that when the hostname is altered after interface started/connected the changes would only be reflected once the interface restarts/reconnects

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **hostname** -- [in] New hostname for the interface. Maximum length 32 bytes.

Returns

- ESP_OK - success
- ESP_ERR_ESP_NETIF_IF_NOT_READY - interface status error
- ESP_ERR_ESP_NETIF_INVALID_PARAMS - parameter error

esp_err_t **esp_netif_get_hostname** (*esp_netif_t* *esp_netif, const char **hostname)

Get interface hostname.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **hostname** -- [out] Returns a pointer to the hostname. May be NULL if no hostname is set. If set non-NULL, pointer remains valid (and string may change if the hostname changes).

Returns

- ESP_OK - success
- ESP_ERR_ESP_NETIF_IF_NOT_READY - interface status error
- ESP_ERR_ESP_NETIF_INVALID_PARAMS - parameter error

bool **esp_netif_is_netif_up** (*esp_netif_t* *esp_netif)

Test if supplied interface is up or down.

Parameters **esp_netif** -- [in] Handle to esp-netif instance

Returns

- true - Interface is up
- false - Interface is down

esp_err_t **esp_netif_get_ip_info** (*esp_netif_t* *esp_netif, *esp_netif_ip_info_t* *ip_info)

Get interface's IP address information.

If the interface is up, IP information is read directly from the TCP/IP stack. If the interface is down, IP information is read from a copy kept in the ESP-NETIF instance

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **ip_info** -- [out] If successful, IP information will be returned in this argument.

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS

esp_err_t **esp_netif_get_old_ip_info** (*esp_netif_t* *esp_netif, *esp_netif_ip_info_t* *ip_info)

Get interface's old IP information.

Returns an "old" IP address previously stored for the interface when the valid IP changed.

If the IP lost timer has expired (meaning the interface was down for longer than the configured interval) then the old IP information will be zero.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **ip_info** -- [out] If successful, IP information will be returned in this argument.

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS

esp_err_t **esp_netif_set_ip_info** (*esp_netif_t* *esp_netif, const *esp_netif_ip_info_t* *ip_info)

Set interface's IP address information.

This function is mainly used to set a static IP on an interface.

If the interface is up, the new IP information is set directly in the TCP/IP stack.

The copy of IP information kept in the ESP-NETIF instance is also updated (this copy is returned if the IP is queried while the interface is still down.)

Note: DHCP client/server must be stopped (if enabled for this interface) before setting new IP information.

Note: Calling this interface for may generate a SYSTEM_EVENT_STA_GOT_IP or SYSTEM_EVENT_ETH_GOT_IP event.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **ip_info** -- [in] IP information to set on the specified interface

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS
- ESP_ERR_ESP_NETIF_DHCP_NOT_STOPPED If DHCP server or client is still running

esp_err_t **esp_netif_set_old_ip_info** (*esp_netif_t* *esp_netif, const *esp_netif_ip_info_t* *ip_info)

Set interface old IP information.

This function is called from the DHCP client (if enabled), before a new IP is set. It is also called from the default handlers for the SYSTEM_EVENT_STA_CONNECTED and SYSTEM_EVENT_ETH_CONNECTED events.

Calling this function stores the previously configured IP, which can be used to determine if the IP changes in the future.

If the interface is disconnected or down for too long, the "IP lost timer" will expire (after the configured interval) and set the old IP information to zero.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **ip_info** -- [in] Store the old IP information for the specified interface

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS

int **esp_netif_get_netif_impl_index** (*esp_netif_t* *esp_netif)

Get net interface index from network stack implementation.

Note: This index could be used in `setsockopt ()` to bind socket with multicast interface

Parameters **esp_netif** -- [in] Handle to esp-netif instance

Returns implementation specific index of interface represented with supplied esp_netif

esp_err_t **esp_netif_get_netif_impl_name** (*esp_netif_t* *esp_netif, char *name)

Get net interface name from network stack implementation.

Note: This name could be used in `setsockopt()` to bind socket with appropriate interface

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **name** -- [out] Interface name as specified in underlying TCP/IP stack. Note that the actual name will be copied to the specified buffer, which must be allocated to hold maximum interface name size (6 characters for lwIP)

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS

esp_err_t **esp_netif_napt_enable** (*esp_netif_t* *esp_netif)

Enable NAPT on an interface.

Note: Enable operation can be performed only on one interface at a time. NAPT cannot be enabled on multiple interfaces according to this implementation.

Parameters **esp_netif** -- [in] Handle to esp-netif instance

Returns

- ESP_OK
- ESP_FAIL
- ESP_ERR_NOT_SUPPORTED

esp_err_t **esp_netif_napt_disable** (*esp_netif_t* *esp_netif)

Disable NAPT on an interface.

Parameters **esp_netif** -- [in] Handle to esp-netif instance

Returns

- ESP_OK
- ESP_FAIL
- ESP_ERR_NOT_SUPPORTED

esp_err_t **esp_netif_dhcps_option** (*esp_netif_t* *esp_netif, *esp_netif_dhcp_option_mode_t* opt_op, *esp_netif_dhcp_option_id_t* opt_id, void *opt_val, uint32_t opt_len)

Set or Get DHCP server option.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **opt_op** -- [in] ESP_NETIF_OP_SET to set an option, ESP_NETIF_OP_GET to get an option.
- **opt_id** -- [in] Option index to get or set, must be one of the supported enum values.
- **opt_val** -- [inout] Pointer to the option parameter.
- **opt_len** -- [in] Length of the option parameter.

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS
- ESP_ERR_ESP_NETIF_DHCP_ALREADY_STOPPED
- ESP_ERR_ESP_NETIF_DHCP_ALREADY_STARTED

esp_err_t **esp_netif_dhcpc_option** (*esp_netif_t* *esp_netif, *esp_netif_dhcp_option_mode_t* opt_op, *esp_netif_dhcp_option_id_t* opt_id, void *opt_val, uint32_t opt_len)

Set or Get DHCP client option.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **opt_op** -- [in] ESP_NETIF_OP_SET to set an option, ESP_NETIF_OP_GET to get an option.
- **opt_id** -- [in] Option index to get or set, must be one of the supported enum values.
- **opt_val** -- [inout] Pointer to the option parameter.
- **opt_len** -- [in] Length of the option parameter.

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS
- ESP_ERR_ESP_NETIF_DHCP_ALREADY_STOPPED
- ESP_ERR_ESP_NETIF_DHCP_ALREADY_STARTED

esp_err_t **esp_netif_dhcpc_start** (*esp_netif_t* *esp_netif)

Start DHCP client (only if enabled in interface object)

Note: The default event handlers for the SYSTEM_EVENT_STA_CONNECTED and SYSTEM_EVENT_ETH_CONNECTED events call this function.

Parameters **esp_netif** -- [in] Handle to esp-netif instance

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS
- ESP_ERR_ESP_NETIF_DHCP_ALREADY_STARTED
- ESP_ERR_ESP_NETIF_DHCPC_START_FAILED

esp_err_t **esp_netif_dhcpc_stop** (*esp_netif_t* *esp_netif)

Stop DHCP client (only if enabled in interface object)

Note: Calling action_netif_stop() will also stop the DHCP Client if it is running.

Parameters **esp_netif** -- [in] Handle to esp-netif instance

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS
- ESP_ERR_ESP_NETIF_DHCP_ALREADY_STOPPED
- ESP_ERR_ESP_NETIF_IF_NOT_READY

esp_err_t **esp_netif_dhcpc_get_status** (*esp_netif_t* *esp_netif, *esp_netif_dhcp_status_t* *status)

Get DHCP client status.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **status** -- [out] If successful, the status of DHCP client will be returned in this argument.

Returns

- ESP_OK

esp_err_t **esp_netif_dhcps_get_status** (*esp_netif_t* *esp_netif, *esp_netif_dhcp_status_t* *status)

Get DHCP Server status.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance

- **status** -- [out] If successful, the status of the DHCP server will be returned in this argument.

Returns

- ESP_OK

esp_err_t **esp_netif_dhcps_start** (*esp_netif_t* *esp_netif)

Start DHCP server (only if enabled in interface object)

Parameters *esp_netif* -- [in] Handle to esp-netif instance

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS
- ESP_ERR_ESP_NETIF_DHCP_ALREADY_STARTED

esp_err_t **esp_netif_dhcps_stop** (*esp_netif_t* *esp_netif)

Stop DHCP server (only if enabled in interface object)

Parameters *esp_netif* -- [in] Handle to esp-netif instance

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS
- ESP_ERR_ESP_NETIF_DHCP_ALREADY_STOPPED
- ESP_ERR_ESP_NETIF_IF_NOT_READY

esp_err_t **esp_netif_dhcps_get_clients_by_mac** (*esp_netif_t* *esp_netif, int num, *esp_netif_pair_mac_ip_t* *mac_ip_pair)

Populate IP addresses of clients connected to DHCP server listed by their MAC addresses.

Parameters

- *esp_netif* -- [in] Handle to esp-netif instance
- *num* -- [in] Number of clients with specified MAC addresses in the array of pairs
- *mac_ip_pair* -- [inout] Array of pairs of MAC and IP addresses (MAC are inputs, IP outputs)

Returns

- ESP_OK on success
- ESP_ERR_ESP_NETIF_INVALID_PARAMS on invalid params
- ESP_ERR_NOT_SUPPORTED if DHCP server not enabled

esp_err_t **esp_netif_set_dns_info** (*esp_netif_t* *esp_netif, *esp_netif_dns_type_t* type, *esp_netif_dns_info_t* *dns)

Set DNS Server information.

This function behaves differently if DHCP server or client is enabled

If DHCP client is enabled, main and backup DNS servers will be updated automatically from the DHCP lease if the relevant DHCP options are set. Fallback DNS Server is never updated from the DHCP lease and is designed to be set via this API. If DHCP client is disabled, all DNS server types can be set via this API only.

Note that LWIP stores DNS server information globally, not per interface, so the first parameter is unused in the default LWIP configuration. If CONFIG_ESP_NETIF_SET_DNS_PER_DEFAULT_NETIF=1 this API sets internal DNS server information per netif. It's also possible to set the global DNS server info by supplying esp_netif=NULL

If DHCP server is enabled, the Main DNS Server setting is used by the DHCP server to provide a DNS Server option to DHCP clients (Wi-Fi stations).

- The default Main DNS server is typically the IP of the DHCP server itself.
- This function can override it by setting server type ESP_NETIF_DNS_MAIN.
- Other DNS Server types are not supported for the DHCP server.
- To propagate the DNS info to client, please stop the DHCP server before using this API.

Parameters

- *esp_netif* -- [in] Handle to esp-netif instance

- **type** -- **[in]** Type of DNS Server to set: ESP_NETIF_DNS_MAIN, ESP_NETIF_DNS_BACKUP, ESP_NETIF_DNS_FALLBACK
- **dns** -- **[in]** DNS Server address to set

Returns

- ESP_OK on success
- ESP_ERR_ESP_NETIF_INVALID_PARAMS invalid params

esp_err_t **esp_netif_get_dns_info** (*esp_netif_t* *esp_netif, *esp_netif_dns_type_t* type, *esp_netif_dns_info_t* *dns)

Get DNS Server information.

Return the currently configured DNS Server address for the specified interface and Server type.

This may be result of a previous call to *esp_netif_set_dns_info()*. If the interface's DHCP client is enabled, the Main or Backup DNS Server may be set by the current DHCP lease.

Note that LWIP stores DNS server information globally, not per interface, so the first parameter is unused in the default LWIP configuration. If CONFIG_ESP_NETIF_SET_DNS_PER_DEFAULT_NETIF=1 this API returns internally saved DNS server information per netif. It's also possible to ask for the global DNS server info by supplying esp_netif=NULL

Parameters

- **esp_netif** -- **[in]** Handle to esp-netif instance
- **type** -- **[in]** Type of DNS Server to get: ESP_NETIF_DNS_MAIN, ESP_NETIF_DNS_BACKUP, ESP_NETIF_DNS_FALLBACK
- **dns** -- **[out]** DNS Server result is written here on success

Returns

- ESP_OK on success
- ESP_ERR_ESP_NETIF_INVALID_PARAMS invalid params

esp_err_t **esp_netif_create_ip6_linklocal** (*esp_netif_t* *esp_netif)

Create interface link-local IPv6 address.

Cause the TCP/IP stack to create a link-local IPv6 address for the specified interface.

This function also registers a callback for the specified interface, so that if the link-local address becomes verified as the preferred address then a SYSTEM_EVENT_GOT_IP6 event will be sent.

Parameters **esp_netif** -- **[in]** Handle to esp-netif instance

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS

esp_err_t **esp_netif_get_ip6_linklocal** (*esp_netif_t* *esp_netif, *esp_ip6_addr_t* *if_ip6)

Get interface link-local IPv6 address.

If the specified interface is up and a preferred link-local IPv6 address has been created for the interface, return a copy of it.

Parameters

- **esp_netif** -- **[in]** Handle to esp-netif instance
- **if_ip6** -- **[out]** IPv6 information will be returned in this argument if successful.

Returns

- ESP_OK
- ESP_FAIL If interface is down, does not have a link-local IPv6 address, or the link-local IPv6 address is not a preferred address.

esp_err_t **esp_netif_get_ip6_global** (*esp_netif_t* *esp_netif, *esp_ip6_addr_t* *if_ip6)

Get interface global IPv6 address.

If the specified interface is up and a preferred global IPv6 address has been created for the interface, return a copy of it.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **if_ip6** -- [out] IPv6 information will be returned in this argument if successful.

Returns

- ESP_OK
- ESP_FAIL If interface is down, does not have a global IPv6 address, or the global IPv6 address is not a preferred address.

int **esp_netif_get_all_ip6** (*esp_netif_t* *esp_netif, *esp_ip6_addr_t* if_ip6[])

Get all IPv6 addresses of the specified interface.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **if_ip6** -- [out] Array of IPv6 addresses will be copied to the argument

Returns number of returned IPv6 addresses

int **esp_netif_get_all_preferred_ip6** (*esp_netif_t* *esp_netif, *esp_ip6_addr_t* if_ip6[])

Get all preferred IPv6 addresses of the specified interface.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **if_ip6** -- [out] Array of IPv6 addresses will be copied to the argument

Returns number of returned IPv6 addresses

esp_err_t **esp_netif_add_ip6_address** (*esp_netif_t* *esp_netif, const *esp_ip6_addr_t* addr, bool preferred)

Cause the TCP/IP stack to add an IPv6 address to the interface.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **addr** -- [in] The address to be added
- **preferred** -- [in] The preferred status of the address

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS
- ESP_ERR_ESP_NETIF_IP6_ADDR_FAILED
- ESP_ERR_NO_MEM

esp_err_t **esp_netif_remove_ip6_address** (*esp_netif_t* *esp_netif, const *esp_ip6_addr_t* *addr)

Cause the TCP/IP stack to remove an IPv6 address from the interface.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **addr** -- [in] The address to be removed

Returns

- ESP_OK
- ESP_ERR_ESP_NETIF_INVALID_PARAMS
- ESP_ERR_ESP_NETIF_IP6_ADDR_FAILED
- ESP_ERR_NO_MEM

void **esp_netif_set_ip4_addr** (*esp_ip4_addr_t* *addr, uint8_t a, uint8_t b, uint8_t c, uint8_t d)

Sets IPv4 address to the specified octets.

Parameters

- **addr** -- [out] IP address to be set
- **a** -- the first octet (127 for IP 127.0.0.1)
- **b** --
- **c** --
- **d** --

char ***esp_ip4addr_ntoa** (const *esp_ip4_addr_t* *addr, char *buf, int buflen)

Converts numeric IP address into decimal dotted ASCII representation.

Parameters

- **addr** -- ip address in network order to convert
- **buf** -- target buffer where the string is stored
- **buflen** -- length of buf

Returns either pointer to buf which now holds the ASCII representation of addr or NULL if buf was too small

uint32_t **esp_ip4addr_aton** (const char *addr)

Ascii internet address interpretation routine The value returned is in network order.

Parameters **addr** -- IP address in ascii representation (e.g. "127.0.0.1")

Returns ip address in network order

esp_err_t **esp_netif_str_to_ip4** (const char *src, *esp_ip4_addr_t* *dst)

Converts Ascii internet IPv4 address into esp_ip4_addr_t.

Parameters

- **src** -- **[in]** IPv4 address in ascii representation (e.g. "127.0.0.1")
- **dst** -- **[out]** Address of the target esp_ip4_addr_t structure to receive converted address

Returns

- ESP_OK on success
- ESP_FAIL if conversion failed
- ESP_ERR_INVALID_ARG if invalid parameter is passed into

esp_err_t **esp_netif_str_to_ip6** (const char *src, *esp_ip6_addr_t* *dst)

Converts Ascii internet IPv6 address into esp_ip4_addr_t Zeros in the IP address can be stripped or completely omitted: "2001:db8:85a3:0:0:0:2:1" or "2001:db8::2:1")

Parameters

- **src** -- **[in]** IPv6 address in ascii representation (e.g. ""2001:0db8:85a3:0000:0000:0002:0001")
- **dst** -- **[out]** Address of the target esp_ip6_addr_t structure to receive converted address

Returns

- ESP_OK on success
- ESP_FAIL if conversion failed
- ESP_ERR_INVALID_ARG if invalid parameter is passed into

esp_netif_iedriver_handle **esp_netif_get_io_driver** (*esp_netif_t* *esp_netif)

Gets media driver handle for this esp-netif instance.

Parameters **esp_netif** -- **[in]** Handle to esp-netif instance

Returns opaque pointer of related IO driver

esp_netif_t ***esp_netif_get_handle_from_ifkey** (const char *if_key)

Searches over a list of created objects to find an instance with supplied if key.

Parameters **if_key** -- Textual description of network interface

Returns Handle to esp-netif instance

esp_netif_flags_t **esp_netif_get_flags** (*esp_netif_t* *esp_netif)

Returns configured flags for this interface.

Parameters **esp_netif** -- **[in]** Handle to esp-netif instance

Returns Configuration flags

const char ***esp_netif_get_ifkey** (*esp_netif_t* *esp_netif)

Returns configured interface key for this esp-netif instance.

Parameters **esp_netif** -- **[in]** Handle to esp-netif instance

Returns Textual description of related interface

const char ***esp_netif_get_desc** (*esp_netif_t* *esp_netif)

Returns configured interface type for this esp-netif instance.

Parameters **esp_netif** -- **[in]** Handle to esp-netif instance

Returns Enumerated type of this interface, such as station, AP, ethernet

int **esp_netif_get_route_prio** (*esp_netif_t* *esp_netif)

Returns configured routing priority number.

Parameters **esp_netif** -- [in] Handle to esp-netif instance

Returns Integer representing the instance's route-prio, or -1 if invalid parameters

int32_t **esp_netif_get_event_id** (*esp_netif_t* *esp_netif, *esp_netif_ip_event_type_t* event_type)

Returns configured event for this esp-netif instance and supplied event type.

Parameters

- **esp_netif** -- [in] Handle to esp-netif instance
- **event_type** -- (either get or lost IP)

Returns specific event id which is configured to be raised if the interface lost or acquired IP address
-1 if supplied event_type is not known

esp_netif_t ***esp_netif_next** (*esp_netif_t* *esp_netif)

Iterates over list of interfaces. Returns first netif if NULL given as parameter.

Note: This API doesn't lock the list, nor the TCPIP context, as this it's usually required to get atomic access between iteration steps rather than within a single iteration. Therefore it is recommended to iterate over the interfaces inside *esp_netif_tcpip_exec()*

Note: This API is deprecated. Please use *esp_netif_next_unsafe()* directly if all the system interfaces are under your control and you can safely iterate over them. Otherwise, iterate over interfaces using *esp_netif_tcpip_exec()*, or use *esp_netif_find_if()* to search in the list of netifs with defined predicate.

Parameters **esp_netif** -- [in] Handle to esp-netif instance

Returns First netif from the list if supplied parameter is NULL, next one otherwise

esp_netif_t ***esp_netif_next_unsafe** (*esp_netif_t* *esp_netif)

Iterates over list of interfaces without list locking. Returns first netif if NULL given as parameter.

Used for bulk search loops within TCPIP context, e.g. using *esp_netif_tcpip_exec()*, or if we're sure that the iteration is safe from our application perspective (e.g. no interface is removed between iterations)

Parameters **esp_netif** -- [in] Handle to esp-netif instance

Returns First netif from the list if supplied parameter is NULL, next one otherwise

esp_netif_t ***esp_netif_find_if** (*esp_netif_find_predicate_t* fn, void *ctx)

Return a netif pointer for the first interface that meets criteria defined by the callback.

Parameters

- **fn** -- Predicate function returning true for the desired interface
- **ctx** -- Context pointer passed to the predicate, typically a descriptor to compare with

Returns valid netif pointer if found, NULL if not

size_t **esp_netif_get_nr_of_ifs** (void)

Returns number of registered esp_netif objects.

Returns Number of esp_netifs

void **esp_netif_netstack_buf_ref** (void *netstack_buf)

increase the reference counter of net stack buffer

Parameters **netstack_buf** -- [in] the net stack buffer

void **esp_netif_netstack_buf_free** (void *netstack_buf)

free the netstack buffer

Parameters **netstack_buf** -- [in] the net stack buffer

esp_err_t **esp_netif_tcpip_exec** (*esp_netif_callback_fn* fn, void *ctx)

Utility to execute the supplied callback in TCP/IP context.

Parameters

- **fn** -- Pointer to the callback
- **ctx** -- Parameter to the callback

Returns The error code (*esp_err_t*) returned by the callback

Type Definitions

typedef bool (***esp_netif_find_predicate_t**)(*esp_netif_t* *netif, void *ctx)

Predicate callback for *esp_netif_find_if()* used to find interface which meets defined criteria.

typedef *esp_err_t* (***esp_netif_callback_fn**)(void *ctx)

TCPIP thread safe callback used with *esp_netif_tcpip_exec()*

Header File

- [components/esp_netif/include/esp_netif_sntp.h](#)
- This header file can be included with:

```
#include "esp_netif_sntp.h"
```

- This header file is a part of the API provided by the *esp_netif* component. To declare that your component depends on *esp_netif*, add the following to your CMakeLists.txt:

```
REQUIRES esp_netif
```

or

```
PRIV_REQUIRES esp_netif
```

Functions

esp_err_t **esp_netif_sntp_init** (const *esp_sntp_config_t* *config)

Initialize SNTP with supplied config struct.

Parameters **config** -- Config struct

Returns ESP_OK on success

esp_err_t **esp_netif_sntp_start** (void)

Start SNTP service if it wasn't started during init (config.start = false) or restart it if already started.

Returns ESP_OK on success

void **esp_netif_sntp_deinit** (void)

Deinitialize esp_netif SNTP module.

esp_err_t **esp_netif_sntp_sync_wait** (TickType_t tout)

Wait for time sync event.

Parameters **tout** -- Specified timeout in RTOS ticks

Returns ESP_TIMEOUT if sync event didn't come within the timeout
ESP_ERR_NOT_FINISHED if the sync event came, but we're in smooth update mode
and still in progress (SNTP_SYNC_STATUS_IN_PROGRESS) ESP_OK if time sync'ed

esp_err_t **esp_netif_sntp_reachability** (unsigned int index, unsigned int *reachability)

Returns SNTP server's reachability shift register as described in RFC 5905.

Parameters

- **index** -- Index of the SERVER
- **reachability** -- reachability shift register

Returns ESP_OK on success, ESP_ERR_INVALID_STATE if SNTP not initialized
ESP_ERR_INVALID_ARG if invalid arguments

Structures

struct **esp_sntp_config**

SNTP configuration struct.

Public Members

bool **smooth_sync**

set to true if smooth sync required

bool **server_from_dhcp**

set to true to request NTP server config from DHCP

bool **wait_for_sync**

if true, we create a semaphore to signal time sync event

bool **start**

set to true to automatically start the SNTP service

esp_sntp_time_cb_t **sync_cb**

optionally sets callback function on time sync event

bool **renew_servers_after_new_IP**

this is used to refresh server list if NTP provided by DHCP (which cleans other pre-configured servers)

ip_event_t **ip_event_to_renew**

set the IP event id on which we refresh server list (if renew_servers_after_new_IP=true)

size_t **index_of_first_server**

refresh server list after this server (if renew_servers_after_new_IP=true)

size_t **num_of_servers**

number of preconfigured NTP servers

const char ***servers**[1]

list of servers

Macros

ESP_SNTP_SERVER_LIST (...)

Utility macro for providing multiple servers in parentheses.

ESP_NETIF_SNTP_DEFAULT_CONFIG_MULTIPLE (servers_in_list, list_of_servers)

Default configuration to init SNTP with multiple servers.

Parameters

- **servers_in_list** -- Number of servers in the list
- **list_of_servers** -- List of servers (use [ESP_SNTP_SERVER_LIST\(...\)](#))

ESP_NETIF_SNTP_DEFAULT_CONFIG (server)

Default configuration with a single server.

Type Definitions

typedef void (***esp_sntp_time_cb_t**)(struct timeval *tv)

Time sync notification function.

typedef struct [esp_sntp_config](#) **esp_sntp_config_t**

SNTP configuration struct.

Header File

- [components/esp_netif/include/esp_netif_types.h](#)
- This header file can be included with:

```
#include "esp_netif_types.h"
```

- This header file is a part of the API provided by the `esp_netif` component. To declare that your component depends on `esp_netif`, add the following to your CMakeLists.txt:

```
REQUIRES esp_netif
```

or

```
PRIV_REQUIRES esp_netif
```

Structures

struct **esp_netif_dns_info_t**

DNS server info.

Public Members

[esp_ip_addr_t](#) **ip**

IPV4 address of DNS server

struct **esp_netif_ip_info_t**

Event structure for IP_EVENT_STA_GOT_IP, IP_EVENT_ETH_GOT_IP events

Public Members

[esp_ip4_addr_t](#) **ip**

Interface IPV4 address

[esp_ip4_addr_t](#) **netmask**

Interface IPV4 netmask

esp_ip4_addr_t **gw**

Interface IPV4 gateway address

struct **esp_netif_ip6_info_t**

IPV6 IP address information.

Public Members

esp_ip6_addr_t **ip**

Interface IPV6 address

struct **ip_event_got_ip_t**

Event structure for IP_EVENT_GOT_IP event.

Public Members

esp_netif_t ***esp_netif**

Pointer to corresponding esp-netif object

esp_netif_ip_info_t **ip_info**

IP address, netmask, gateway IP address

bool **ip_changed**

Whether the assigned IP has changed or not

struct **ip_event_got_ip6_t**

Event structure for IP_EVENT_GOT_IP6 event

Public Members

esp_netif_t ***esp_netif**

Pointer to corresponding esp-netif object

esp_netif_ip6_info_t **ip6_info**

IPv6 address of the interface

int **ip_index**

IPv6 address index

struct **ip_event_add_ip6_t**

Event structure for ADD_IP6 event

Public Members

esp_ip6_addr_t **addr**

The address to be added to the interface

bool **preferred**

The default preference of the address

struct **ip_event_ap_staipassigned_t**

Event structure for IP_EVENT_AP_STAIPASSIGNED event

Public Members

esp_netif_t ***esp_netif**

Pointer to the associated netif handle

esp_ip4_addr_t **ip**

IP address which was assigned to the station

uint8_t **mac**[6]

MAC address of the connected client

struct **bridgeif_config**

LwIP bridge configuration

Public Members

uint16_t **max_fdb_dyn_entries**

maximum number of entries in dynamic forwarding database

uint16_t **max_fdb_sta_entries**

maximum number of entries in static forwarding database

uint8_t **max_ports**

maximum number of ports the bridge can consist of

struct **esp_netif_inherent_config**

ESP-netif inherent config parameters.

Public Members

esp_netif_flags_t **flags**

flags that define esp-netif behavior

uint8_t **mac**[6]

initial mac address for this interface

const *esp_netif_ip_info_t* ***ip_info**

initial ip address for this interface

uint32_t **get_ip_event**

event id to be raised when interface gets an IP

uint32_t **lost_ip_event**

event id to be raised when interface loses its IP

const char ***if_key**

string identifier of the interface

const char ***if_desc**

textual description of the interface

int **route_prio**

numeric priority of this interface to become a default routing if (if other netifs are up). A higher value of route_prio indicates a higher priority

bridgeif_config_t ***bridge_info**

LwIP bridge configuration

struct **esp_netif_driver_base_s**

ESP-netif driver base handle.

Public Members

esp_err_t (***post_attach**)(*esp_netif_t* *netif, *esp_netif_iodriver_handle* h)

post attach function pointer

esp_netif_t ***netif**

netif handle

struct **esp_netif_driver_ifconfig**

Specific IO driver configuration.

Public Members

esp_netif_iodriver_handle **handle**

io-driver handle

esp_err_t (***transmit**)(void *h, void *buffer, size_t len)

transmit function pointer

esp_err_t (***transmit_wrap**)(void *h, void *buffer, size_t len, void *netstack_buffer)

transmit wrap function pointer

void (***driver_free_rx_buffer**)(void *h, void *buffer)

free rx buffer function pointer

struct **esp_netif_config**

Generic esp_netif configuration.

Public Members

const *esp_netif_inherent_config_t* ***base**

base config

const *esp_netif_driver_ifconfig_t* ***driver**

driver config

const *esp_netif_netstack_config_t* ***stack**

stack config

struct **esp_netif_pair_mac_ip_t**

DHCP client's addr info (pair of MAC and IP address)

Public Members

uint8_t **mac**[6]

Clients MAC address

esp_ip4_addr_t **ip**

Clients IP address

Macros

ESP_ERR_ESP_NETIF_BASE

Definition of ESP-NETIF based errors.

ESP_ERR_ESP_NETIF_INVALID_PARAMS

ESP_ERR_ESP_NETIF_IF_NOT_READY

ESP_ERR_ESP_NETIF_DHCP_START_FAILED

ESP_ERR_ESP_NETIF_DHCP_ALREADY_STARTED

ESP_ERR_ESP_NETIF_DHCP_ALREADY_STOPPED

ESP_ERR_ESP_NETIF_NO_MEM

ESP_ERR_ESP_NETIF_DHCP_NOT_STOPPED

ESP_ERR_ESP_NETIF_DRIVER_ATTACH_FAILED

ESP_ERR_ESP_NETIF_INIT_FAILED

ESP_ERR_ESP_NETIF_DNS_NOT_CONFIGURED

ESP_ERR_ESP_NETIF_MLD6_FAILED

ESP_ERR_ESP_NETIF_IP6_ADDR_FAILED

ESP_ERR_ESP_NETIF_DHCP_START_FAILED

ESP_NETIF_BR_FLOOD

Definition of ESP-NETIF bridge controll.

ESP_NETIF_BR_DROP

ESP_NETIF_BR_FDW_CPU

Type Definitions

typedef struct esp_netif_obj **esp_netif_t**

typedef enum *esp_netif_flags* **esp_netif_flags_t**

typedef enum *esp_netif_ip_event_type* **esp_netif_ip_event_type_t**

typedef struct *bridgeif_config* **bridgeif_config_t**

LwIP bridge configuration

typedef struct *esp_netif_inherent_config* **esp_netif_inherent_config_t**

ESP-netif inherent config parameters.

typedef struct *esp_netif_config* **esp_netif_config_t**

typedef void ***esp_netif_iodriver_handle**

IO driver handle type.

typedef struct *esp_netif_driver_base_s* **esp_netif_driver_base_t**

ESP-netif driver base handle.

typedef struct *esp_netif_driver_ifconfig* **esp_netif_driver_ifconfig_t**

typedef struct esp_netif_netstack_config **esp_netif_netstack_config_t**

Specific L3 network stack configuration.

typedef *esp_err_t* (***esp_netif_receive_t**)(*esp_netif_t* *esp_netif, void *buffer, size_t len, void *eb)

ESP-NETIF Receive function type.

Enumerations

enum **esp_netif_dns_type_t**

Type of DNS server.

Values:

enumerator **ESP_NETIF_DNS_MAIN**

DNS main server address

enumerator **ESP_NETIF_DNS_BACKUP**

DNS backup server address (Wi-Fi STA and Ethernet only)

enumerator **ESP_NETIF_DNS_FALLBACK**

DNS fallback server address (Wi-Fi STA and Ethernet only)

enumerator **ESP_NETIF_DNS_MAX**

enum **esp_netif_dhcp_status_t**

Status of DHCP client or DHCP server.

Values:

enumerator **ESP_NETIF_DHCP_INIT**

DHCP client/server is in initial state (not yet started)

enumerator **ESP_NETIF_DHCP_STARTED**

DHCP client/server has been started

enumerator **ESP_NETIF_DHCP_STOPPED**

DHCP client/server has been stopped

enumerator **ESP_NETIF_DHCP_STATUS_MAX**

enum **esp_netif_dhcp_option_mode_t**

Mode for DHCP client or DHCP server option functions.

Values:

enumerator **ESP_NETIF_OP_START**

enumerator **ESP_NETIF_OP_SET**

Set option

enumerator **ESP_NETIF_OP_GET**

Get option

enumerator **ESP_NETIF_OP_MAX**

enum **esp_netif_dhcp_option_id_t**

Supported options for DHCP client or DHCP server.

Values:

enumerator **ESP_NETIF_SUBNET_MASK**

Network mask

enumerator **ESP_NETIF_DOMAIN_NAME_SERVER**

Domain name server

enumerator **ESP_NETIF_ROUTER_SOLICITATION_ADDRESS**

Solicitation router address

enumerator **ESP_NETIF_REQUESTED_IP_ADDRESS**

Request specific IP address

enumerator **ESP_NETIF_IP_ADDRESS_LEASE_TIME**

Request IP address lease time

enumerator **ESP_NETIF_IP_REQUEST_RETRY_TIME**

Request IP address retry counter

enumerator **ESP_NETIF_VENDOR_CLASS_IDENTIFIER**

Vendor Class Identifier of a DHCP client

enumerator **ESP_NETIF_VENDOR_SPECIFIC_INFO**

Vendor Specific Information of a DHCP server

enum **ip_event_t**

IP event declarations

Values:

enumerator **IP_EVENT_STA_GOT_IP**

station got IP from connected AP

enumerator **IP_EVENT_STA_LOST_IP**

station lost IP and the IP is reset to 0

enumerator **IP_EVENT_AP_STAIPASSIGNED**

soft-AP assign an IP to a connected station

enumerator **IP_EVENT_GOT_IP6**

station or ap or ethernet interface v6IP addr is preferred

enumerator **IP_EVENT_ETH_GOT_IP**

ethernet got IP from connected AP

enumerator **IP_EVENT_ETH_LOST_IP**

ethernet lost IP and the IP is reset to 0

enumerator **IP_EVENT_PPP_GOT_IP**

PPP interface got IP

enumerator **IP_EVENT_PPP_LOST_IP**

PPP interface lost IP

enum **esp_netif_flags**

Values:

enumerator **ESP_NETIF_DHCP_CLIENT**

enumerator **ESP_NETIF_DHCP_SERVER**

enumerator **ESP_NETIF_FLAG_AUTOUP**

enumerator **ESP_NETIF_FLAG_GARP**

enumerator **ESP_NETIF_FLAG_EVENT_IP_MODIFIED**

enumerator **ESP_NETIF_FLAG_IS_PPP**

enumerator **ESP_NETIF_FLAG_IS_BRIDGE**

enumerator **ESP_NETIF_FLAG_MLDV6_REPORT**

enum **esp_netif_ip_event_type**

Values:

enumerator **ESP_NETIF_IP_EVENT_GOT_IP**

enumerator **ESP_NETIF_IP_EVENT_LOST_IP**

Header File

- [components/esp_netif/include/esp_netif_ip_addr.h](#)
- This header file can be included with:

```
#include "esp_netif_ip_addr.h"
```

- This header file is a part of the API provided by the `esp_netif` component. To declare that your component depends on `esp_netif`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_netif
```

or

```
PRIV_REQUIRES esp_netif
```

Functions

[*esp_ip6_addr_type_t*](#) **esp_netif_ip6_get_addr_type** ([*esp_ip6_addr_t*](#) *ip6_addr)

Get the IPv6 address type.

Parameters **ip6_addr** -- [in] IPv6 type

Returns IPv6 type in form of enum `esp_ip6_addr_type_t`

static inline void **esp_netif_ip_addr_copy** (*esp_ip_addr_t* *dest, const *esp_ip_addr_t* *src)

Copy IP addresses.

Parameters

- **dest** -- [out] destination IP
- **src** -- [in] source IP

Structures

struct **esp_ip6_addr**

IPv6 address.

Public Members

uint32_t **addr**[4]

IPv6 address

uint8_t **zone**

zone ID

struct **esp_ip4_addr**

IPv4 address.

Public Members

uint32_t **addr**

IPv4 address

struct **_ip_addr**

IP address.

Public Members

esp_ip6_addr_t **ip6**

IPv6 address type

esp_ip4_addr_t **ip4**

IPv4 address type

union *_ip_addr*::[anonymous] **u_addr**

IP address union

uint8_t **type**

ipaddress type

Macros**esp_netif_htonl** (x)**esp_netif_ip4_makeu32** (a, b, c, d)**ESP_IP6_ADDR_BLOCK1** (ip6addr)**ESP_IP6_ADDR_BLOCK2** (ip6addr)**ESP_IP6_ADDR_BLOCK3** (ip6addr)**ESP_IP6_ADDR_BLOCK4** (ip6addr)**ESP_IP6_ADDR_BLOCK5** (ip6addr)**ESP_IP6_ADDR_BLOCK6** (ip6addr)**ESP_IP6_ADDR_BLOCK7** (ip6addr)**ESP_IP6_ADDR_BLOCK8** (ip6addr)**IPSTR****esp_ip4_addr_get_byte** (ipaddr, idx)**esp_ip4_addr1** (ipaddr)**esp_ip4_addr2** (ipaddr)**esp_ip4_addr3** (ipaddr)**esp_ip4_addr4** (ipaddr)**esp_ip4_addr1_16** (ipaddr)**esp_ip4_addr2_16** (ipaddr)**esp_ip4_addr3_16** (ipaddr)**esp_ip4_addr4_16** (ipaddr)**IP2STR** (ipaddr)**IPV6STR****IPV62STR** (ipaddr)**ESP_IPADDR_TYPE_V4****ESP_IPADDR_TYPE_V6****ESP_IPADDR_TYPE_ANY****ESP_IP4TOUINT32** (a, b, c, d)**ESP_IP4TOADDR** (a, b, c, d)**ESP_IP4ADDR_INIT** (a, b, c, d)**ESP_IP6ADDR_INIT** (a, b, c, d)**IP4ADDR_STRLEN_MAX****ESP_IP_IS_ANY** (addr)

Type Definitions

typedef struct *esp_ip4_addr* **esp_ip4_addr_t**

typedef struct *esp_ip6_addr* **esp_ip6_addr_t**

typedef struct *_ip_addr* **esp_ip_addr_t**

IP address.

Enumerations

enum **esp_ip6_addr_type_t**

Values:

enumerator **ESP_IP6_ADDR_IS_UNKNOWN**

enumerator **ESP_IP6_ADDR_IS_GLOBAL**

enumerator **ESP_IP6_ADDR_IS_LINK_LOCAL**

enumerator **ESP_IP6_ADDR_IS_SITE_LOCAL**

enumerator **ESP_IP6_ADDR_IS_UNIQUE_LOCAL**

enumerator **ESP_IP6_ADDR_IS_IPV4_MAPPED_IPV6**

Header File

- [components/esp_netif/include/esp_vfs_l2tap.h](#)
- This header file can be included with:

```
#include "esp_vfs_l2tap.h"
```

- This header file is a part of the API provided by the `esp_netif` component. To declare that your component depends on `esp_netif`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_netif
```

or

```
PRIV_REQUIRES esp_netif
```

Functions

esp_err_t **esp_vfs_l2tap_intf_register** (*l2tap_vfs_config_t* *config)

Add L2 TAP virtual filesystem driver.

This function must be called prior usage of ESP-NETIF L2 TAP Interface

Parameters **config** -- L2 TAP virtual filesystem driver configuration. Default base path /dev/net/tap is used when this parameter is NULL.

Returns *esp_err_t*

- **ESP_OK** on success

esp_err_t **esp_vfs_l2tap_intf_unregister** (const char *base_path)

Removes L2 TAP virtual filesystem driver.

Parameters **base_path** -- Base path to the L2 TAP virtual filesystem driver. Default path /dev/net/tap is used when this parameter is NULL.

Returns *esp_err_t*

- ESP_OK on success

esp_err_t **esp_vfs_l2tap_eth_filter** (*l2tap_iodriver_handle* driver_handle, void *buff, size_t *size)

Filters received Ethernet L2 frames into L2 TAP infrastructure.

Parameters

- **driver_handle** -- handle of driver at which the frame was received
- **buff** -- received L2 frame
- **size** -- input length of the L2 frame which is set to 0 when frame is filtered into L2 TAP

Returns *esp_err_t*

- ESP_OK is always returned

Structures

struct **l2tap_vfs_config_t**

L2Tap VFS config parameters.

Public Members

const char ***base_path**
vfs base path

Macros

L2TAP_VFS_DEFAULT_PATH

L2TAP_VFS_CONFIG_DEFAULT ()

Type Definitions

typedef void ***l2tap_iodriver_handle**

Enumerations

enum **l2tap_ioctl_opt_t**

Values:

enumerator **L2TAP_S_RCV_FILTER**

enumerator **L2TAP_G_RCV_FILTER**

enumerator **L2TAP_S_INTF_DEVICE**

enumerator **L2TAP_G_INTF_DEVICE**

enumerator **L2TAP_S_DEVICE_DRV_HNDL**

enumerator `L2TAP_G_DEVICE_DRV_HNDL`

2.4.4 IP Network Layer

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

ESP-NETIF Custom I/O Driver

This section outlines implementing a new I/O driver with ESP-NETIF connection capabilities.

By convention, the I/O driver has to register itself as an ESP-NETIF driver, and thus holds a dependency on ESP-NETIF component and is responsible for providing data path functions, post-attach callback and in most cases, also default event handlers to define network interface actions based on driver's lifecycle transitions.

Packet Input/Output According to the diagram shown in the [ESP-NETIF Architecture](#) part, the following three API functions for the packet data path must be defined for connecting with ESP-NETIF:

- `esp_netif_transmit()`
- `esp_netif_free_rx_buffer()`
- `esp_netif_receive()`

The first two functions for transmitting and freeing the rx buffer are provided as callbacks, i.e., they get called from ESP-NETIF (and its underlying TCP/IP stack) and I/O driver provides their implementation.

The receiving function on the other hand gets called from the I/O driver, so that the driver's code simply calls `esp_netif_receive()` on a new data received event.

Post Attach Callback A final part of the network interface initialization consists of attaching the ESP-NETIF instance to the I/O driver, by means of calling the following API:

```
esp_err_t esp_netif_attach(esp_netif_t *esp_netif, esp_netif_iodriver_handle_t
↳ driver_handle);
```

It is assumed that the `esp_netif_iodriver_handle` is a pointer to driver's object, a struct derived from `struct esp_netif_driver_base_s`, so that the first member of I/O driver structure must be this base structure with pointers to:

- post-attach function callback
- related ESP-NETIF instance

As a result, the I/O driver has to create an instance of the struct per below:

```
typedef struct my_netif_driver_s {
    esp_netif_driver_base_t base;           /*!< base structure reserved as_
↳ esp-netif driver */
    driver_impl_t *h;                       /*!< handle of driver_
↳ implementation */
} my_netif_driver_t;
```

with actual values of `my_netif_driver_t::base.post_attach` and the actual drivers handle `my_netif_driver_t::h`.

So when the `esp_netif_attach()` gets called from the initialization code, the post-attach callback from I/O driver's code gets executed to mutually register callbacks between ESP-NETIF and I/O driver instances. Typically the driver is started as well in the post-attach callback. An example of a simple post-attach callback is outlined below:

```
static esp_err_t my_post_attach_start(esp_netif_t * esp_netif, void * args)
{
    my_netif_driver_t *driver = args;
    const esp_netif_driver_ifconfig_t driver_ifconfig = {
        .driver_free_rx_buffer = my_free_rx_buf,
        .transmit = my_transmit,
        .handle = driver->driver_impl
    };
    driver->base.netif = esp_netif;
    ESP_ERROR_CHECK(esp_netif_set_driver_config(esp_netif, &driver_ifconfig));
    my_driver_start(driver->driver_impl);
    return ESP_OK;
}
```

Default Handlers I/O drivers also typically provide default definitions of lifecycle behavior of related network interfaces based on state transitions of I/O drivers. For example *driver start* → *network start*, etc.

An example of such a default handler is provided below:

```
esp_err_t my_driver_netif_set_default_handlers(my_netif_driver_t *driver, esp_
↪netif_t * esp_netif)
{
    driver_set_event_handler(driver->driver_impl, esp_netif_action_start, MY_DRV_
↪EVENT_START, esp_netif);
    driver_set_event_handler(driver->driver_impl, esp_netif_action_stop, MY_DRV_
↪EVENT_STOP, esp_netif);
    return ESP_OK;
}
```

Network Stack Connection The packet data path functions for transmitting and freeing the rx buffer (defined in the I/O driver) are called from the ESP-NETIF, specifically from its TCP/IP stack connecting layer.

Note that ESP-IDF provides several network stack configurations for the most common network interfaces, such as for the Wi-Fi station or Ethernet. These configurations are defined in [esp_netif/include/esp_netif_defaults.h](#) and should be sufficient for most network drivers. In rare cases, expert users might want to define custom lwIP based interface layers; it is possible, but an explicit dependency to lwIP needs to be set.

The following API reference outlines these network stack interaction with the ESP-NETIF:

Header File

- [components/esp_netif/include/esp_netif_net_stack.h](#)
- This header file can be included with:

```
#include "esp_netif_net_stack.h"
```

- This header file is a part of the API provided by the esp_netif component. To declare that your component depends on esp_netif, add the following to your CMakeLists.txt:

```
REQUIRES esp_netif
```

or

```
PRIV_REQUIRES esp_netif
```

Functions

[esp_netif_t](#) *[esp_netif_get_handle_from_netif_impl](#)(void *dev)

Returns esp-netif handle.

Parameters **dev** -- **[in]** opaque ptr to network interface of specific TCP/IP stack

Returns handle to related esp-netif instance

void ***esp_netif_get_netif_impl** (*esp_netif_t* *esp_netif)

Returns network stack specific implementation handle.

Parameters **esp_netif** -- **[in]** Handle to esp-netif instance

Returns handle to related network stack netif handle

esp_err_t **esp_netif_set_link_speed** (*esp_netif_t* *esp_netif, uint32_t speed)

Set link-speed for the specified network interface.

Parameters

- **esp_netif** -- **[in]** Handle to esp-netif instance
- **speed** -- **[in]** Link speed in bit/s

Returns ESP_OK on success

esp_err_t **esp_netif_transmit** (*esp_netif_t* *esp_netif, void *data, size_t len)

Outputs packets from the TCP/IP stack to the media to be transmitted.

This function gets called from network stack to output packets to IO driver.

Parameters

- **esp_netif** -- **[in]** Handle to esp-netif instance
- **data** -- **[in]** Data to be transmitted
- **len** -- **[in]** Length of the data frame

Returns ESP_OK on success, an error passed from the I/O driver otherwise

esp_err_t **esp_netif_transmit_wrap** (*esp_netif_t* *esp_netif, void *data, size_t len, void *netstack_buf)

Outputs packets from the TCP/IP stack to the media to be transmitted.

This function gets called from network stack to output packets to IO driver.

Parameters

- **esp_netif** -- **[in]** Handle to esp-netif instance
- **data** -- **[in]** Data to be transmitted
- **len** -- **[in]** Length of the data frame
- **netstack_buf** -- **[in]** net stack buffer

Returns ESP_OK on success, an error passed from the I/O driver otherwise

void **esp_netif_free_rx_buffer** (void *esp_netif, void *buffer)

Free the rx buffer allocated by the media driver.

This function gets called from network stack when the rx buffer to be freed in IO driver context, i.e. to deallocate a buffer owned by io driver (when data packets were passed to higher levels to avoid copying)

Parameters

- **esp_netif** -- **[in]** Handle to esp-netif instance
- **buffer** -- **[in]** Rx buffer pointer

Code examples for TCP/IP socket APIs are provided in the [protocols/sockets](#) directory of ESP-IDF examples.

2.4.5 Application Layer

Documentation for Application layer network protocols (above the IP Network layer) are provided in [Application Protocols](#).

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.5 Peripherals API

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.5.1 Analog Comparator

Introduction

Analog Comparator is a peripheral that can be used to compare a source signal with the internal reference voltage or an external reference signal.

It is a cost effective way to replace an amplifier comparator in some scenarios. But unlike the continuous comparing of the amplifier comparator, ESP Analog Comparator is driven by a source clock, which decides the sampling frequency.

Analog Comparator on ESP32-P4 has 2 unit(s), the channels in the unit(s) are:

UNIT0

- Source Channel: GPIO52
- External Reference Channel: GPIO51
- Internal Reference Channel: Range 0% ~ 70% of the VDD, the step is 10% of the VDD

UNIT1

- Source Channel: GPIO54
- External Reference Channel: GPIO53
- Internal Reference Channel: Range 0% ~ 70% of the VDD, the step is 10% of the VDD

Functional Overview

The following sections of this document cover the typical steps to install and operate an Analog Comparator unit:

- [Resource Allocation](#) - covers which parameters should be set up to get a unit handle and how to recycle the resources when it finishes working.
- [Further Configurations](#) - covers the other configurations that might need to specific and what they are used for.
- [Enable and Disable Unit](#) - covers how to enable and disable the unit.
- [Power Management](#) - describes how different source clock selections can affect power consumption.
- [IRAM Safe](#) - lists which functions are supposed to work even when the cache is disabled.
- [Thread Safety](#) - lists which APIs are guaranteed to be thread safe by the driver.
- [Kconfig Options](#) - lists the supported Kconfig options that can be used to make a different effect on driver behavior.
- [ETM Events](#) -

Resource Allocation An Analog Comparator unit channel is represented by `ana_cmpr_handle_t`. Each unit can support either an internal or an external reference.

To allocate the resource of the Analog Comparator unit, `ana_cmpr_new_unit()` need to be called to get the handle of the unit. Configurations `ana_cmpr_config_t` need to be specified while allocating the unit:

- `ana_cmpr_config_t::unit` selects the Analog Comparator unit.
- `ana_cmpr_config_t::clk_src` selects the source clock for Analog Comparator, it can affect the sampling frequency. Note that the clock source of the Analog Comparator comes from the io mux, it is shared with GPIO extension peripherals like SDM (Sigma-Delta Modulation) and Glitch

Filter. The configuration will fail if you specify different clock sources for multiple GPIO extension peripherals. The default clock sources of these peripherals are same, typically, we select `soc_periph_ana_cmpr_clk_src_t::ANA_CMPR_CLK_SRC_DEFAULT` as the clock source.

- `ana_cmpr_config_t::ref_src` selects the reference source from internal voltage or external signal.
- `ana_cmpr_config_t::cross_type` selects which kind of cross type can trigger the interrupt.

The function `ana_cmpr_new_unit()` can fail due to various errors such as insufficient memory, invalid arguments, etc. If a previously created Analog Comparator unit is no longer required, you should recycle it by calling `ana_cmpr_del_unit()`. It allows the underlying HW channel to be used for other purposes. Before deleting an Analog Comparator unit handle, you should disable it by `ana_cmpr_unit_disable()` in advance, or make sure it has not been enabled yet by `ana_cmpr_unit_enable()`.

```
#include "driver/ana_cmpr.h"

ana_cmpr_handle_t cmpr = NULL;
ana_cmpr_config_t config = {
    .unit = 0,
    .clk_src = ANA_CMPR_CLK_SRC_DEFAULT,
    .ref_src = ANA_CMPR_REF_SRC_INTERNAL,
    .cross_type = ANA_CMPR_CROSS_ANY,
};
ESP_ERROR_CHECK(ana_cmpr_new_unit(&config, &cmpr));
// ...
ESP_ERROR_CHECK(ana_cmpr_del_unit(cmpr));
```

Further Configurations

- `ana_cmpr_set_intl_reference()` - Specify the internal reference voltage when `ana_cmpr_ref_source_t::ANA_CMPR_REF_SRC_INTERNAL` is selected as reference source.

It requires `ana_cmpr_internal_ref_config_t::ref_volt` to specify the voltage. The voltage related to the VDD power supply, which can only support a certain fixed percentage of VDD. Currently on ESP32-P4, the internal reference voltage can range from 0 ~ 70% VDD with a step 10%.

```
#include "driver/ana_cmpr.h"

ana_cmpr_internal_ref_config_t ref_cfg = {
    .ref_volt = ANA_CMPR_REF_VOLT_50_PCT_VDD,
};
ESP_ERROR_CHECK(ana_cmpr_set_internal_reference(cmpr, &ref_cfg));
```

- `ana_cmpr_set_debounce()` - Set the debounce configuration.

It requires `ana_cmpr_debounce_config_t::wait_us` to set the interrupt waiting time. The interrupt is disabled temporarily for `ana_cmpr_debounce_config_t::wait_us` micro seconds, so that the frequent triggering can be avoided while the source signal crosses the reference signal. That is, the waiting time is supposed to be in inverse ratio to the relative frequency between the source and reference. If the waiting time is set too short, it can not bypass the jitter totally, but if too long, the next crossing interrupt might be missed.

```
#include "driver/ana_cmpr.h"

ana_cmpr_debounce_config_t dbc_cfg = {
    .wait_us = 1,
};
ESP_ERROR_CHECK(ana_cmpr_set_debounce(cmpr, &dbc_cfg));
```

- `ana_cmpr_set_cross_type()` - Set the source signal cross type.

The initial cross type is set in `ana_cmpr_new_unit()`, this function can update the cross type, even in ISR context.

```
#include "driver/ana_cmpr.h"

ESP_ERROR_CHECK(ana_cmpr_set_cross_type(cmpr, ANA_CMPR_CROSS_POS));
```

- [ana_cmpr_register_event_callbacks\(\)](#) - Register the callbacks.

Currently it supports [ana_cmpr_event_callbacks_t::on_cross](#), it will be called when the crossing event (specified by [ana_cmpr_config_t::cross_type](#)) occurs.

```
#include "driver/ana_cmpr.h"

static bool IRAM_ATTR example_ana_cmpr_on_cross_callback(ana_cmpr_handle_t cmpr,
                                                         const ana_cmpr_cross_event_
↳data_t *edata,
                                                         void *user_ctx)
{
    // ...
    return false;
}

ana_cmpr_event_callbacks_t cbs = {
    .on_cross = example_ana_cmpr_on_cross_callback,
};

ESP_ERROR_CHECK(ana_cmpr_register_event_callbacks(cmpr, &cbs, NULL));
```

Note: When [CONFIG_ANA_CMPR_ISR_IRAM_SAFE](#) is enabled, you should guarantee the callback context and involved data to be in internal RAM by add the attribute `IRAM_ATTR`. (See more in [IRAM Safe](#))

Enable and Disable Unit

- [ana_cmpr_enable\(\)](#) - Enable the Analog Comparator unit.
- [ana_cmpr_disable\(\)](#) - Disable the Analog Comparator unit.

After the Analog Comparator unit is enabled and the crossing event interrupt is enabled, a power management lock will be acquired if the power management is enabled (see [Power Management](#)). Under the **enable** state, only [ana_cmpr_set_int1_reference\(\)](#) and [ana_cmpr_set_debounce\(\)](#) can be called, other functions can only be called after the unit is disabled.

Calling [ana_cmpr_disable\(\)](#) does the opposite.

Power Management When power management is enabled (i.e., [CONFIG_PM_ENABLE](#) is on), the system will adjust the APB frequency before going into light sleep, thus potentially changing the resolution of the Analog Comparator.

However, the driver can prevent the system from changing APB frequency by acquiring a power management lock of type [ESP_PM_NO_LIGHT_SLEEP](#). Whenever the driver creates a Analog Comparator unit instance that has selected the clock source like [ANA_CMPR_CLK_SRC_DEFAULT](#) or [ANA_CMPR_CLK_SRC_XTAL](#) as its clock source, the driver guarantees that the power management lock is acquired when enable the channel by [ana_cmpr_enable\(\)](#). Likewise, the driver releases the lock when [ana_cmpr_disable\(\)](#) is called for that channel.

IRAM Safe By default, the Analog Comparator interrupt will be deferred when the Cache is disabled for reasons like programming/erasing Flash. Thus the alarm interrupt will not get executed in time, which is not expected in a real-time application.

There is a Kconfig option [CONFIG_ANA_CMPR_ISR_IRAM_SAFE](#) that:

1. Enables the interrupt being serviced even when cache is disabled

2. Places all functions that used by the ISR into IRAM¹
3. Places driver object into DRAM (in case it is allocated on PSRAM)

This allows the interrupt to run while the cache is disabled but comes at the cost of increased IRAM consumption.

There is a Kconfig option `CONFIG_ANA_CMPR_CTRL_FUNC_IN_IRAM` that can put commonly used IO control functions into IRAM as well. So that these functions can also be executable when the cache is disabled. These IO control functions are listed as follows:

- `ana_cmpr_set_internal_reference()`
- `ana_cmpr_set_debounce()`
- `ana_cmpr_set_cross_type()`

Thread Safety The factory function `ana_cmpr_new_unit()` is guaranteed to be thread safe by the driver, which means, user can call it from different RTOS tasks without protection by extra locks. The following functions are allowed to run under ISR context, the driver uses a critical section to prevent them being called concurrently in both task and ISR.

- `ana_cmpr_set_internal_reference()`
- `ana_cmpr_set_debounce()`
- `ana_cmpr_set_cross_type()`

Other functions that take the `ana_cmpr_handle_t` as the first positional parameter, are not treated as thread safe. Which means the user should avoid calling them from multiple tasks.

Kconfig Options

- `CONFIG_ANA_CMPR_ISR_IRAM_SAFE` controls whether the default ISR handler can work when cache is disabled, see *IRAM Safe* for more information.
- `CONFIG_ANA_CMPR_CTRL_FUNC_IN_IRAM` controls where to place the Analog Comparator control functions (IRAM or Flash), see *IRAM Safe* for more information.
- `CONFIG_ANA_CMPR_ENABLE_DEBUG_LOG` is used to enable the debug log output. Enabling this option increases the firmware binary size.

ETM Events To create an analog comparator cross event, you need to include `driver/ana_cmpr_etm.h` additionally, and allocate the event by `ana_cmpr_new_etm_event()`. You can refer to *ETM* for how to connect an event to a task.

Application Example

- `peripherals/analog_comparator` shows the basic usage of the analog comparator, and other potential usages like hysteresis comparator and SPWM generator.

API Reference

Header File

- `components/driver/analog_comparator/include/driver/ana_cmpr.h`
- This header file can be included with:

```
#include "driver/ana_cmpr.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

¹ `ana_cmpr_event_callbacks_t::on_cross` callback and the functions invoked by itself should also be placed in IRAM, you need to take care of them by themselves.

REQUIRES driver

or

PRIV_REQUIRES driver

Functions

esp_err_t **ana_cmpr_new_unit** (const *ana_cmpr_config_t* *config, *ana_cmpr_handle_t* *ret_cmpr)

Allocating a new analog comparator unit handle.

Parameters

- **config** -- [in] The config of the analog comparator unit
- **ret_cmpr** -- [out] The returned analog comparator unit handle

Returns

- ESP_OK Allocate analog comparator unit handle success
- ESP_ERR_NO_MEM No memory for the analog comparator structure
- ESP_ERR_INVALID_ARG NULL pointer of the parameters or wrong unit number
- ESP_ERR_INVALID_STATE The unit has been allocated or the clock source has been occupied

esp_err_t **ana_cmpr_del_unit** (*ana_cmpr_handle_t* cmpr)

Delete the analog comparator unit handle.

Parameters **cmpr** -- [in] The handle of analog comparator unit

Returns

- ESP_OK Delete analog comparator unit handle success
- ESP_ERR_INVALID_ARG NULL pointer of the parameters or wrong unit number
- ESP_ERR_INVALID_STATE The analog comparator is not disabled yet

esp_err_t **ana_cmpr_set_internal_reference** (*ana_cmpr_handle_t* cmpr, const *ana_cmpr_internal_ref_config_t* *ref_cfg)

Set internal reference configuration.

Note: This function only need to be called when *ana_cmpr_config_t::ref_src* is ANA_CMPR_REF_SRC_INTERNAL.

Note: This function is allowed to run within ISR context including intr callbacks

Note: This function will be placed into IRAM if CONFIG_ANA_CMPR_CTRL_FUNC_IN_IRAM is on, so that it's allowed to be executed when Cache is disabled

Parameters

- **cmpr** -- [in] The handle of analog comparator unit
- **ref_cfg** -- [in] Internal reference configuration

Returns

- ESP_OK Set denounce configuration success
- ESP_ERR_INVALID_ARG NULL pointer of the parameters
- ESP_ERR_INVALID_STATE The reference source is not ANA_CMPR_REF_SRC_INTERNAL

esp_err_t **ana_cmpr_set_debounce** (*ana_cmpr_handle_t* cmpr, const *ana_cmpr_debounce_config_t* *dbc_cfg)

Set debounce configuration to the analog comparator.

Note: This function is allowed to run within ISR context including intr callbacks

Note: This function will be placed into IRAM if `CONFIG_ANA_CMPR_CTRL_FUNC_IN_IRAM` is on, so that it's allowed to be executed when Cache is disabled

Parameters

- **cmpr** -- [in] The handle of analog comparator unit
- **dbc_cfg** -- [in] Debounce configuration

Returns

- `ESP_OK` Set denounce configuration success
- `ESP_ERR_INVALID_ARG` NULL pointer of the parameters

esp_err_t **ana_cmpr_set_cross_type** (*ana_cmpr_handle_t* cmpr, *ana_cmpr_cross_type_t* cross_type)

Set the source signal cross type.

Note: The initial cross type is configured in `ana_cmpr_new_unit`, this function can update the cross type

Note: This function is allowed to run within ISR context including intr callbacks

Note: This function will be placed into IRAM if `CONFIG_ANA_CMPR_CTRL_FUNC_IN_IRAM` is on, so that it's allowed to be executed when Cache is disabled

Parameters

- **cmpr** -- [in] The handle of analog comparator unit
- **cross_type** -- [in] The source signal cross type that can trigger the interrupt

Returns

- `ESP_OK` Set denounce configuration success
- `ESP_ERR_INVALID_ARG` NULL pointer of the parameters

esp_err_t **ana_cmpr_register_event_callbacks** (*ana_cmpr_handle_t* cmpr, const *ana_cmpr_event_callbacks_t* *cbs, void *user_data)

Register analog comparator interrupt event callbacks.

Note: This function can only be called before enabling the unit

Parameters

- **cmpr** -- [in] The handle of analog comparator unit
- **cbs** -- [in] Group of callback functions
- **user_data** -- [in] The user data that will be passed to callback functions directly

Returns

- `ESP_OK` Register callbacks success
- `ESP_ERR_INVALID_ARG` NULL pointer of the parameters
- `ESP_ERR_INVALID_STATE` The analog comparator has been enabled

esp_err_t **ana_cmpr_enable** (*ana_cmpr_handle_t* cmpr)

Enable the analog comparator unit.

Parameters **cmpr** -- [in] The handle of analog comparator unit

Returns

- ESP_OK Enable analog comparator unit success
- ESP_ERR_INVALID_ARG NULL pointer of the parameters
- ESP_ERR_INVALID_STATE The analog comparator has been enabled

esp_err_t **ana_cmpr_disable** (*ana_cmpr_handle_t* cmpr)

Disable the analog comparator unit.

Parameters **cmpr** -- [in] The handle of analog comparator unit

Returns

- ESP_OK Disable analog comparator unit success
- ESP_ERR_INVALID_ARG NULL pointer of the parameters
- ESP_ERR_INVALID_STATE The analog comparator has disabled already

esp_err_t **ana_cmpr_get_gpio** (*ana_cmpr_unit_t* unit, *ana_cmpr_channel_type_t* chan_type, int *gpio_num)

Get the specific GPIO number of the analog comparator unit.

Parameters

- **unit** -- [in] The handle of analog comparator unit
- **chan_type** -- [in] The channel type of analog comparator, like source channel or reference channel
- **gpio_num** -- [out] The output GPIO number of this channel

Returns

- ESP_OK Get GPIO success
- ESP_ERR_INVALID_ARG NULL pointer of the parameters or wrong unit number or wrong channel type

Structures

struct **ana_cmpr_config_t**

Analog comparator unit configuration.

Public Members

ana_cmpr_unit_t **unit**

Analog comparator unit

ana_cmpr_clk_src_t **clk_src**

The clock source of the analog comparator, which decide the resolution of the comparator

ana_cmpr_ref_source_t **ref_src**

Reference signal source of the comparator, select using ANA_CMPR_REF_SRC_INTERNAL or ANA_CMPR_REF_SRC_EXTERNAL. For internal reference, the reference voltage should be set to `internal_ref_volt`, for external reference, the reference signal should be connect to ANA_CMPRx_EXT_REF_GPIO

ana_cmpr_cross_type_t **cross_type**

The crossing types that can trigger interrupt

int **intr_priority**

The interrupt priority, range 1~3. If set to 0, the driver will automatically select a relative low priority (1,2,3)

uint32_t **io_loop_back**

Enable this field when the other signals that output on the comparison pins are supposed to be fed back.
Normally used for debug/test scenario

struct *ana_cmpr_config_t*::[anonymous] **flags**

Analog comparator driver flags

struct **ana_cmpr_internal_ref_config_t**

Analog comparator internal reference configuration.

Public Members

ana_cmpr_ref_voltage_t **ref_volt**

The internal reference voltage. It can be specified to a certain fixed percentage of the VDD power supply, currently supports 0%~70% VDD with a step 10%

struct **ana_cmpr_debounce_config_t**

Analog comparator debounce filter configuration.

Public Members

uint32_t **wait_us**

The wait time to prevent frequent interrupts caused by signal noise or bouncing. During the specified wait_us period, no new interrupts will be triggered. Set the value according to the signal characteristics. A rapid signal requires a small wait time, otherwise the next cross event may be missed.

struct **ana_cmpr_event_callbacks_t**

Group of Analog Comparator callbacks.

Note: The callbacks are all running under ISR environment

Note: When CONFIG_ANA_CMPR_ISR_IRAM_SAFE is enabled, the callback itself and functions called by it should be placed in IRAM. The variables used in the function should be in the SRAM as well.

Public Members

ana_cmpr_cross_cb_t **on_cross**

The callback function on cross interrupt

Header File

- [components/driver/analog_comparator/include/driver/ana_cmpr_types.h](#)
- This header file can be included with:

```
#include "driver/ana_cmpr_types.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

REQUIRES driver

or

PRIV_REQUIRES driver

Structures

struct **ana_cmpr_cross_event_data_t**

Analog comparator cross event data.

Public Members

ana_cmpr_cross_type_t **cross_type**

The cross type of the target signal to the reference signal. Will either be ANA_CMPR_CROSS_POS or ANA_CMPR_CROSS_NEG Always be ANA_CMPR_CROSS_ANY if target does not support independent interrupt (like ESP32H2)

Macros

ANA_CMPR_UNIT_0

Deprecated:

Analog comparator unit 0

Type Definitions

typedef int **ana_cmpr_unit_t**

Analog comparator unit.

typedef struct ana_cmpr_t ***ana_cmpr_handle_t**

Analog comparator unit handle.

typedef *soc_periph_ana_cmpr_clk_src_t* **ana_cmpr_clk_src_t**

Analog comparator clock source.

typedef bool (***ana_cmpr_cross_cb_t**)(*ana_cmpr_handle_t* cmpr, const *ana_cmpr_cross_event_data_t* *edata, void *user_ctx)

Prototype of Analog comparator event callback.

Param cmpr [in] Analog Comparator handle, created from `ana_cmpr_new_unit()`

Param edata [in] Point to Analog Comparator event data. The lifecycle of this pointer memory is inside this function, user should copy it into static memory if used outside this function. (Currently not use)

Param user_ctx [in] User registered context, passed from `ana_cmpr_register_event_callbacks()`

Return Whether a high priority task has been waken up by this callback function

Enumerations

enum **ana_cmpr_ref_source_t**

Analog comparator reference source.

Values:

enumerator **ANA_CMPR_REF_SRC_INTERNAL**

Analog Comparator internal reference source, related to VDD

enumerator **ANA_CMPR_REF_SRC_EXTERNAL**

Analog Comparator external reference source, from ANA_CMPR0_EXT_REF_GPIO

enum **ana_cmpr_channel_type_t**

Analog comparator channel type.

Values:

enumerator **ANA_CMPR_SOURCE_CHAN**

Analog Comparator source channel, which is used to input the signal that to be compared

enumerator **ANA_CMPR_EXT_REF_CHAN**

Analog Comparator external reference channel, which is used as the reference signal

enum **ana_cmpr_cross_type_t**

Analog comparator interrupt type.

Values:

enumerator **ANA_CMPR_CROSS_DISABLE**

Disable the cross event interrupt

enumerator **ANA_CMPR_CROSS_POS**

Positive cross can trigger event interrupt

enumerator **ANA_CMPR_CROSS_NEG**

Negative cross can trigger event interrupt

enumerator **ANA_CMPR_CROSS_ANY**

Any cross can trigger event interrupt

enum **ana_cmpr_ref_voltage_t**

Analog comparator internal reference voltage.

Values:

enumerator **ANA_CMPR_REF_VOLT_0_PCT_VDD**

Internal reference voltage equals to 0% VDD

enumerator **ANA_CMPR_REF_VOLT_10_PCT_VDD**

Internal reference voltage equals to 10% VDD

enumerator **ANA_CMPR_REF_VOLT_20_PCT_VDD**

Internal reference voltage equals to 20% VDD

enumerator **ANA_CMPR_REF_VOLT_30_PCT_VDD**

Internal reference voltage equals to 30% VDD

enumerator **ANA_CMPR_REF_VOLT_40_PCT_VDD**

Internal reference voltage equals to 40% VDD

enumerator **ANA_CMPR_REF_VOLT_50_PCT_VDD**

Internal reference voltage equals to 50% VDD

enumerator **ANA_CMPR_REF_VOLT_60_PCT_VDD**

Internal reference voltage equals to 60% VDD

enumerator **ANA_CMPR_REF_VOLT_70_PCT_VDD**

Internal reference voltage equals to 70% VDD

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.5.2 Clock Tree

The clock subsystem of ESP32-P4 is used to source and distribute system/module clocks from a range of root clocks. The clock tree driver maintains the basic functionality of the system clock and the intricate relationship among module clocks.

This document starts with the introduction to root and module clocks. Then it covers the clock tree APIs that can be called to monitor the status of the module clocks at runtime.

Introduction

This section lists definitions of ESP32-P4's supported root clocks and module clocks. These definitions are commonly used in the driver configuration, to help select a proper source clock for the peripheral.

Root Clocks Root clocks generate reliable clock signals. These clock signals then pass through various gates, muxes, dividers, or multipliers to become the clock sources for every functional module: the CPU core(s), Wi-Fi, Bluetooth, the RTC, and the peripherals.

ESP32-P4's root clocks are listed in [soc_root_clk_t](#):

- Internal 8 MHz RC Oscillator (RC_FAST)
This RC oscillator generates a about 8.5 MHz clock signal output as the RC_FAST_CLK. The exact frequency of RC_FAST_CLK cannot be computed in runtime through calibration, but it is still possible to get its frequency through an oscilloscope or a logic analyzer by routing the clock signal to a GPIO pin.
- External 40 MHz Crystal (XTAL)
- Internal 136 kHz RC Oscillator (RC_SLOW)
This RC oscillator generates a about 136kHz clock signal output as the RC_SLOW_CLK. The exact frequency of this clock can be computed in runtime through calibration.
- External 32 kHz Crystal - optional (XTAL32K)
The clock source for this XTAL32K_CLK can be either a 32 kHz crystal connecting to the XTAL_32K_P and XTAL_32K_N pins or a 32 kHz clock signal generated by an external circuit. The external signal must be connected to the XTAL_32K_P pin. XTAL32K_CLK can also be calibrated to get its exact frequency.

- External Slow Clock - optional (OSC_SLOW)

A clock signal generated by an external circuit can be connected to GPIO0 to be the clock source for the RTC_SLOW_CLK. This clock can also be calibrated to get its exact frequency.

Typically, the frequency of the signal generated from an RC oscillator circuit is less accurate and more sensitive to the environment compared to the signal generated from a crystal. ESP32-P4 provides several clock source options for the RTC_SLOW_CLK, and it is possible to make the choice based on the requirements for system time accuracy and power consumption. For more details, please refer to [RTC Timer Clock Sources](#).

Module Clocks ESP32-P4's available module clocks are listed in [soc_module_clk_t](#). Each module clock has a unique ID. You can get more information on each clock by checking the documented enum value.

API Usage

The clock tree driver provides an all-in-one API to get the frequency of the module clocks, [esp_clk_tree_src_get_freq_hz\(\)](#). This function allows you to obtain the clock frequency at any time by providing the clock name [soc_module_clk_t](#) and specifying the desired precision level for the returned frequency value [esp_clk_tree_src_freq_precision_t](#).

API Reference

Header File

- [components/soc/esp32p4/include/soc/clk_tree_defs.h](#)
- This header file can be included with:

```
#include "soc/clk_tree_defs.h"
```

Macros

SOC_CLK_RC_FAST_FREQ_APPROX

Approximate RC_FAST_CLK frequency in Hz

SOC_CLK_RC_SLOW_FREQ_APPROX

Approximate RC_SLOW_CLK frequency in Hz

SOC_CLK_RC32K_FREQ_APPROX

Approximate RC32K_CLK frequency in Hz

SOC_CLK_XTAL32K_FREQ_APPROX

Approximate XTAL32K_CLK frequency in Hz

SOC_CLK_OSC_SLOW_FREQ_APPROX

Approximate OSC_SLOW_CLK (external slow clock) frequency in Hz

SOC_GPTIMER_CLKS

Array initializer for all supported clock sources of GPTimer.

The following code can be used to iterate all possible clocks:

```
soc_periph_gptimer_clk_src_t gptimer_clks[] = (soc_periph_gptimer_clk_src_
↪t) SOC_GPTIMER_CLKS;
for (size_t i = 0; i < sizeof(gptimer_clks) / sizeof(gptimer_clks[0]); i++) {
    soc_periph_gptimer_clk_src_t clk = gptimer_clks[i];
    // Test GPTimer with the clock `clk`
}
```

SOC_RMT_CLKS

Array initializer for all supported clock sources of RMT.

SOC_UART_CLKS

Array initializer for all supported clock sources of UART.

SOC_LP_UART_CLKS

Array initializer for all supported clock sources of LP_UART.

SOC_MCPWM_TIMER_CLKS

Array initializer for all supported clock sources of MCPWM Timer.

SOC_MCPWM_CAPTURE_CLKS

Array initializer for all supported clock sources of MCPWM Capture Timer.

SOC_MCPWM_CARRIER_CLKS

Array initializer for all supported clock sources of MCPWM Carrier.

SOC_I2S_CLKS

Array initializer for all supported clock sources of I2S.

SOC_I2C_CLKS

Array initializer for all supported clock sources of I2C.

SOC_SPI_CLKS

Array initializer for all supported clock sources of SPI.

SOC_PSRAM_CLKS

Array initializer for all supported clock sources of PSRAM.

SOC_FLASH_CLKS

Array initializer for all supported clock sources of FLASH.

SOC_ANA_CMPR_CLKS

Array initializer for all supported clock sources of Analog Comparator.

SOC_MWDT_CLKS

Array initializer for all supported clock sources of MWDT.

SOC_LEDC_CLKS

Array initializer for all supported clock sources of LEDC.

SOC_PARLIO_CLKS

Array initializer for all supported clock sources of PARLIO.

SOC_SDMMC_CLKS

Array initializer for all supported clock sources of SDMMC.

Enumerationsenum **soc_root_clk_t**

Root clock.

Values:

enumerator **SOC_ROOT_CLK_INT_RC_FAST**

Internal 17.5MHz RC oscillator

enumerator **SOC_ROOT_CLK_INT_RC_SLOW**

Internal 136kHz RC oscillator

enumerator **SOC_ROOT_CLK_EXT_XTAL**

External 40MHz crystal

enumerator **SOC_ROOT_CLK_EXT_XTAL32K**

External 32kHz crystal

enumerator **SOC_ROOT_CLK_INT_RC32K**

Internal 32kHz RC oscillator

enumerator **SOC_ROOT_CLK_EXT_OSC_SLOW**

External slow clock signal at pin1

enum **soc_cpu_clk_src_t**

CPU_CLK mux inputs, which are the supported clock sources for the CPU_CLK.

Note: Enum values are matched with the register field values on purpose

Values:

enumerator **SOC_CPU_CLK_SRC_XTAL**

Select XTAL_CLK as CPU_CLK source

enumerator **SOC_CPU_CLK_SRC_PLL**

Select (C)PLL_CLK as CPU_CLK source (CPLL_CLK is the output of 40MHz crystal oscillator frequency multiplier, 400MHz)

enumerator **SOC_CPU_CLK_SRC_RC_FAST**

Select RC_FAST_CLK as CPU_CLK source

enumerator **SOC_CPU_CLK_SRC_INVALID**

Invalid CPU_CLK source

enum **soc_rtc_slow_clk_src_t**

RTC_SLOW_CLK mux inputs, which are the supported clock sources for the RTC_SLOW_CLK.

Note: Enum values are matched with the register field values on purpose

Values:

enumerator **SOC_RTC_SLOW_CLK_SRC_RC_SLOW**

Select RC_SLOW_CLK as RTC_SLOW_CLK source

enumerator **SOC_RTC_SLOW_CLK_SRC_XTAL32K**

Select XTAL32K_CLK as RTC_SLOW_CLK source

enumerator **SOC_RTC_SLOW_CLK_SRC_RC32K**

Select RC32K_CLK as RTC_SLOW_CLK source

enumerator **SOC_RTC_SLOW_CLK_SRC_OSC_SLOW**

Select OSC_SLOW_CLK (external slow clock) as RTC_SLOW_CLK source

enumerator **SOC_RTC_SLOW_CLK_SRC_INVALID**

Invalid RTC_SLOW_CLK source

enum **soc_rtc_fast_clk_src_t**

RTC_FAST_CLK mux inputs, which are the supported clock sources for the RTC_FAST_CLK.

Note: Enum values are matched with the register field values on purpose

Values:

enumerator **SOC_RTC_FAST_CLK_SRC_RC_FAST**

Select RC_FAST_CLK as RTC_FAST_CLK source

enumerator **SOC_RTC_FAST_CLK_SRC_XTAL**

Select XTAL_CLK as RTC_FAST_CLK source

enumerator **SOC_RTC_FAST_CLK_SRC_XTAL_DIV**

Alias name for SOC_RTC_FAST_CLK_SRC_XTAL

enumerator **SOC_RTC_FAST_CLK_SRC_LP_PLL**

Select LP_PLL_CLK as RTC_FAST_CLK source (LP_PLL_CLK is a 8MHz clock sourced from RC32K or XTAL32K)

enumerator **SOC_RTC_FAST_CLK_SRC_INVALID**

Invalid RTC_FAST_CLK source

enum **soc_lp_pll_clk_src_t**

LP_PLL_CLK mux inputs, which are the supported clock sources for the LP_PLL_CLK.

Note: Enum values are matched with the register field values on purpose

Values:

enumerator **SOC_LP_PLL_CLK_SRC_RC32K**

Select RC32K_CLK as LP_PLL_CLK source

enumerator **SOC_LP_PLL_CLK_SRC_XTAL32K**

Select XTAL32K_CLK as LP_PLL_CLK source

enumerator **SOC_LP_PLL_CLK_SRC_INVALID**

Invalid LP_PLL_CLK source

enum **soc_module_clk_t**

Supported clock sources for modules (CPU, peripherals, RTC, etc.)

Note: enum starts from 1, to save 0 for special purpose

Values:

enumerator **SOC_MOD_CLK_CPU**

CPU_CLK can be sourced from XTAL, CPLL, or RC_FAST by configuring soc_cpu_clk_src_t

enumerator **SOC_MOD_CLK_RTC_FAST**

RTC_FAST_CLK can be sourced from XTAL, RC_FAST, or LP_PLL by configuring soc_rtc_fast_clk_src_t

enumerator **SOC_MOD_CLK_RTC_SLOW**

RTC_SLOW_CLK can be sourced from RC_SLOW, XTAL32K, RC32K, or OSC_SLOW by configuring soc_rtc_slow_clk_src_t

enumerator **SOC_MOD_CLK_PLL_F80M**

PLL_F80M_CLK is derived from SPLL (clock gating + fixed divider of 6), it has a fixed frequency of 80MHz

enumerator **SOC_MOD_CLK_PLL_F160M**

PLL_F160M_CLK is derived from SPLL (clock gating + fixed divider of 3), it has a fixed frequency of 160MHz

enumerator **SOC_MOD_CLK_PLL_F200M**

PLL_F200M_CLK is derived from SPLL (clock gating + fixed divider of 3), it has a fixed frequency of 200MHz

enumerator **SOC_MOD_CLK_PLL_F240M**

PLL_F240M_CLK is derived from SPLL (clock gating + fixed divider of 2), it has a fixed frequency of 240MHz

enumerator **SOC_MOD_CLK_CPLL**

CPLL is from 40MHz XTAL oscillator frequency multipliers, it has a fixed frequency of 400MHz

enumerator **SOC_MOD_CLK_SPLL**

SPLL is from 40MHz XTAL oscillator frequency multipliers, it has a fixed frequency of 480MHz

enumerator **SOC_MOD_CLK_MPLL**

MPLL is from 40MHz XTAL oscillator frequency multipliers, it has a fixed frequency of 500MHz

enumerator **SOC_MOD_CLK_XTAL32K**

XTAL32K_CLK comes from the external 32kHz crystal, passing a clock gating to the peripherals

enumerator **SOC_MOD_CLK_RC_FAST**

RC_FAST_CLK comes from the internal 20MHz rc oscillator, passing a clock gating to the peripherals

enumerator **SOC_MOD_CLK_XTAL**

XTAL_CLK comes from the external 40MHz crystal

enumerator **SOC_MOD_CLK_APLL**

Audio PLL is sourced from PLL, and its frequency is configurable through APLL configuration registers

enumerator **SOC_MOD_CLK_XTAL_D2**

XTAL_D2_CLK comes from the external 40MHz crystal, passing a div of 2 to the LP peripherals

enumerator **SOC_MOD_CLK_LP_PLL**

LP_PLL is from 32kHz XTAL oscillator frequency multipliers, it has a fixed frequency of 8MHz

enumerator **SOC_MOD_CLK_INVALID**

Indication of the end of the available module clock sources

enum **soc_periph_systimer_clk_src_t**

Type of SYSTIMER clock source.

Values:

enumerator **SYSTIMER_CLK_SRC_XTAL**

SYSTIMER source clock is XTAL

enumerator **SYSTIMER_CLK_SRC_RC_FAST**

SYSTIMER source clock is RC_FAST

enumerator **SYSTIMER_CLK_SRC_DEFAULT**

SYSTIMER source clock default choice is XTAL

enum **soc_periph_gptimer_clk_src_t**

Type of GPTimer clock source.

Values:

enumerator **GPTIMER_CLK_SRC_PLL_F80M**

Select PLL_F80M as the source clock

enumerator **GPTIMER_CLK_SRC_RC_FAST**

Select RC_FAST as the source clock

enumerator **GPTIMER_CLK_SRC_XTAL**

Select XTAL as the source clock

enumerator **GPTIMER_CLK_SRC_DEFAULT**

Select XTAL as the default choice

enum **soc_periph_tg_clk_src_legacy_t**

Type of Timer Group clock source, reserved for the legacy timer group driver.

Values:

enumerator **TIMER_SRC_CLK_PLL_F80M**

Timer group clock source is PLL_F80M

enumerator **TIMER_SRC_CLK_XTAL**

Timer group clock source is XTAL

enumerator **TIMER_SRC_CLK_DEFAULT**

Timer group clock source default choice is XTAL

enum **soc_periph_rmt_clk_src_t**

Type of RMT clock source.

Values:

enumerator **RMT_CLK_SRC_PLL_F80M**

Select PLL_F80M as the source clock

enumerator **RMT_CLK_SRC_RC_FAST**

Select RC_FAST as the source clock

enumerator **RMT_CLK_SRC_XTAL**

Select XTAL as the source clock

enumerator **RMT_CLK_SRC_DEFAULT**

Select XTAL as the default choice

enum **soc_periph_rmt_clk_src_legacy_t**

Type of RMT clock source, reserved for the legacy RMT driver.

Values:

enumerator **RMT_BASECLK_PLL_F80M**

RMT source clock is PLL_F80M

enumerator **RMT_BASECLK_XTAL**

RMT source clock is XTAL

enumerator **RMT_BASECLK_DEFAULT**

RMT source clock default choice is XTAL

enum **soc_periph_uart_clk_src_legacy_t**

Type of UART clock source, reserved for the legacy UART driver.

Values:

enumerator **UART_SCLK_PLL_F80M**

UART source clock is PLL_F80M

enumerator **UART_SCLK_RTC**

UART source clock is RC_FAST

enumerator **UART_SCLK_XTAL**

UART source clock is XTAL

enumerator **UART_SCLK_DEFAULT**

UART source clock default choice is PLL_F80M

enum **soc_periph_lp_uart_clk_src_t**

Type of LP_UART clock source.

Values:

enumerator **LP_UART_SCLK_LP_FAST**

LP_UART source clock is LP(RTC)_FAST

enumerator **LP_UART_SCLK_XTAL_D2**

LP_UART source clock is XTAL_D2

enumerator **LP_UART_SCLK_DEFAULT**

LP_UART source clock default choice is LP(RTC)_FAST

enum **soc_periph_mcpwm_timer_clk_src_t**

Type of MCPWM timer clock source.

Values:

enumerator **MCPWM_TIMER_CLK_SRC_PLL160M**

Select PLL_F160M as the source clock

enumerator **MCPWM_TIMER_CLK_SRC_XTAL**

Select XTAL as the source clock

enumerator **MCPWM_TIMER_CLK_SRC_DEFAULT**

Select XTAL as the default choice

enum **soc_periph_mcpwm_capture_clk_src_t**

Type of MCPWM capture clock source.

Values:

enumerator **MCPWM_CAPTURE_CLK_SRC_PLL160M**

Select PLL_F160M as the source clock

enumerator **MCPWM_CAPTURE_CLK_SRC_XTAL**

Select XTAL as the source clock

enumerator **MCPWM_CAPTURE_CLK_SRC_DEFAULT**

Select XTAL as the default choice

enum **soc_periph_mcpwm_carrier_clk_src_t**

Type of MCPWM carrier clock source.

Values:

enumerator **MCPWM_CARRIER_CLK_SRC_PLL160M**

Select PLL_F160M as the source clock

enumerator **MCPWM_CARRIER_CLK_SRC_XTAL**

Select XTAL as the source clock

enumerator **MCPWM_CARRIER_CLK_SRC_DEFAULT**

Select XTAL as the default choice

enum **soc_periph_i2s_clk_src_t**

I2S clock source enum.

Values:

enumerator **I2S_CLK_SRC_DEFAULT**

Select XTAL as the default source clock

enumerator **I2S_CLK_SRC_XTAL**

Select XTAL as the source clock

enumerator **I2S_CLK_SRC_APLL**

Select APLL as the source clock

enumerator **I2S_CLK_SRC_EXTERNAL**

Select external clock as source clock

enum **soc_periph_i2c_clk_src_t**

Type of I2C clock source.

Values:

enumerator **I2C_CLK_SRC_XTAL**

Select XTAL as the source clock

enumerator **I2C_CLK_SRC_RC_FAST**

Select RC_FAST as the source clock

enumerator **I2C_CLK_SRC_DEFAULT**

Select XTAL as the default source clock

enum **soc_periph_spi_clk_src_t**

Type of SPI clock source.

Values:

enumerator **SPI_CLK_SRC_XTAL**

Select XTAL as SPI source clock

enumerator **SPI_CLK_SRC_DEFAULT**

Select XTAL as SPI source clock

enum **soc_periph_psram_clk_src_t**

Type of PSRAM clock source.

Values:

enumerator **PSRAM_CLK_SRC_DEFAULT**

Select SOC_MOD_CLK_SPLL as PSRAM source clock

enumerator **PSRAM_CLK_SRC_XTAL**

Select SOC_MOD_CLK_XTAL as PSRAM source clock

enumerator **PSRAM_CLK_SRC_CPLL**

Select SOC_MOD_CLK_CPLL as PSRAM source clock

enumerator **PSRAM_CLK_SRC_SPLL**

Select SOC_MOD_CLK_SPLL as PSRAM source clock

enumerator **PSRAM_CLK_SRC_MPLL**

Select SOC_MOD_CLK_MPLL as PSRAM source clock

enum **soc_periph_flash_clk_src_t**

Type of FLASH clock source.

Values:

enumerator **FLASH_CLK_SRC_DEFAULT**

Select SOC_MOD_CLK_SPLL as FLASH source clock

enumerator **FLASH_CLK_SRC_XTAL**

Select SOC_MOD_CLK_XTAL as FLASH source clock

enumerator **FLASH_CLK_SRC_CPLL**

Select SOC_MOD_CLK_CPLL as FLASH source clock

enumerator **FLASH_CLK_SRC_SPLL**

Select SOC_MOD_CLK_SPLL as FLASH source clock

enum **soc_periph_ana_cmpr_clk_src_t**

Analog Comparator clock source.

Values:

enumerator **ANA_CMPR_CLK_SRC_XTAL**

Select XTAL clock as the source clock

enumerator **ANA_CMPR_CLK_SRC_PLL_F80M**

Select PLL_F80M clock as the source clock

enumerator **ANA_CMPR_CLK_SRC_DEFAULT**

Select PLL_F80M as the default clock choice

enum **soc_periph_mwdt_clk_src_t**

MWDT clock source.

Values:

enumerator **MWDT_CLK_SRC_XTAL**

Select XTAL as the source clock

enumerator **MWDT_CLK_SRC_PLL_F80M**

Select PLL fixed 80 MHz as the source clock

enumerator **MWDT_CLK_SRC_RC_FAST**

Select RTC fast as the source clock

enumerator **MWDT_CLK_SRC_DEFAULT**

Select XTAL 40 MHz as the default clock choice

enum **soc_periph_ledc_clk_src_legacy_t**

Type of LEDC clock source, reserved for the legacy LEDC driver.

Values:

enumerator **LEDC_AUTO_CLK**

LEDC source clock will be automatically selected based on the giving resolution and duty parameter when init the timer

enumerator **LEDC_USE_XTAL_CLK**

Select XTAL as the source clock

enumerator **LEDC_USE_PLL_DIV_CLK**

Select PLL_F80M clock as the source clock

enumerator **LEDC_USE_RC_FAST_CLK**

Select RC_FAST as the source clock

enum **soc_periph_parlio_clk_src_t**

PARLIO clock source.

Values:

enumerator **PARLIO_CLK_SRC_XTAL**

Select XTAL as the source clock

enumerator **PARLIO_CLK_SRC_PLL_F160M**

Select PLL_F160M as the source clock

enumerator **PARLIO_CLK_SRC_RC_FAST**

Select RC_FAST as the source clock

enumerator **PARLIO_CLK_SRC_EXTERNAL**

Select EXTERNAL clock as the source clock

enumerator **PARLIO_CLK_SRC_DEFAULT**

Select XTAL as the default clock choice

enum **soc_periph_sdmmc_clk_src_t**

Type of SDMMC clock source.

Values:

enumerator **SDMMC_CLK_SRC_DEFAULT**

Select PLL_160M as the default choice

enumerator **SDMMC_CLK_SRC_PLL160M**

Select PLL_160M as the source clock

enumerator **SDMMC_CLK_SRC_PLL200M**

Select PLL_200M as the source clock

Header File

- [components/esp_hw_support/include/esp_clk_tree.h](#)
- This header file can be included with:

```
#include "esp_clk_tree.h"
```

Functions

[esp_err_t esp_clk_tree_src_get_freq_hz](#) ([soc_module_clk_t](#) clk_src, [esp_clk_tree_src_freq_precision_t](#) precision, [uint32_t](#) *freq_value)

Get frequency of module clock source.

Parameters

- **clk_src** -- **[in]** Clock source available to modules, in [soc_module_clk_t](#)
- **precision** -- **[in]** Degree of precision, one of [esp_clk_tree_src_freq_precision_t](#) values This arg only applies to the clock sources that their frequencies can vary: [SOC_MOD_CLK_RTC_FAST](#), [SOC_MOD_CLK_RTC_SLOW](#), [SOC_MOD_CLK_RC_FAST](#), [SOC_MOD_CLK_RC_FAST_D256](#), [SOC_MOD_CLK_XTAL32K](#) For other clock sources, this field is ignored.

- **freq_value** -- [out] Frequency of the clock source, in Hz

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_FAIL Calibration failed

Enumerations

enum **esp_clk_tree_src_freq_precision_t**

Degree of precision of frequency value to be returned by esp_clk_tree_src_get_freq_hz()

Values:

enumerator **ESP_CLK_TREE_SRC_FREQ_PRECISION_CACHED**

enumerator **ESP_CLK_TREE_SRC_FREQ_PRECISION_APPROX**

enumerator **ESP_CLK_TREE_SRC_FREQ_PRECISION_EXACT**

enumerator **ESP_CLK_TREE_SRC_FREQ_PRECISION_INVALID**

2.5.3 Elliptic Curve Digital Signature Algorithm (ECDSA)

The Elliptic Curve Digital Signature Algorithm (ECDSA) offers a variant of the Digital Signature Algorithm (DSA) which uses elliptic-curve cryptography.

ESP32-P4's ECDSA peripheral provides a secure and efficient environment for computing ECDSA signatures. It offers fast computations while ensuring the confidentiality of the signing process to prevent information leakage. ECDSA private key used in the signing process is accessible only to the hardware peripheral, and it is not readable by software.

ECDSA peripheral can help to establish **Secure Device Identity** for TLS mutual authentication and similar use-cases.

Supported Features

- ECDSA digital signature generation and verification
- Two different elliptic curves, namely P-192 and P-256 (FIPS 186-3 specification)
- Two hash algorithms for message hash in the ECDSA operation, namely SHA-224 and SHA-256 (FIPS PUB 180-4 specification)

ECDSA on ESP32-P4

On ESP32-P4, the ECDSA module works with a secret key burnt into an eFuse block. This eFuse key is made completely inaccessible (default mode) for any resources outside the cryptographic modules, thus avoiding key leakage.

ECDSA key can be programmed externally through `espefuse.py` script using:

```
espefuse.py burn_key <BLOCK_NUM> </path/to/ecdsa_private_key.pem> ECDSA_KEY
```

Note: Six physical eFuse blocks can be used as keys for the ECDSA module: block 4 ~ block 9. E.g., for block 4 (which is the first key block), the argument should be `BLOCK_KEY0`.

Alternatively the ECDSA key can also be programmed through the application running on the target.

Following code snippet uses `esp_efuse_write_key()` to set physical key block 0 in the eFuse with key purpose as `esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_ECDSA_KEY`:

```
#include "esp_efuse.h"

const uint8_t key_data[32] = { ... };

esp_err_t status = esp_efuse_write_key(EFUSE_BLK_KEY0,
                                       ESP_EFUSE_KEY_PURPOSE_ECDSA_KEY,
                                       key_data, sizeof(key_data));

if (status == ESP_OK) {
    // written key
} else {
    // writing key failed, maybe written already
}
```

Dependency on TRNG

ECDSA peripheral relies on the hardware True Random Number Generator (TRNG) for its internal entropy requirement. During ECDSA signature creation, the algorithm requires a random integer to be generated as specified in the [RFC 6090](#) section 5.3.2.

Please ensure that hardware [RNG](#) is enabled before starting ECDSA computations (primarily signing) in the application.

Application Outline

Please refer to the [ECDSA Peripheral with ESP-TLS](#) guide for details on how-to use ECDSA peripheral for establishing a mutually authenticated TLS connection.

The ECDSA peripheral in mbedTLS stack is integrated by overriding the ECDSA sign and verify APIs. Please note that, the ECDSA peripheral does not support all curves or hash algorithms and hence for cases where the requirements do not meet the hardware, implementation falls back to the software.

For a particular TLS context, additional APIs have been supplied to populate certain fields (e.g., private key ctx) to differentiate routing to hardware. ESP-TLS layer integrates these APIs internally and hence no additional work is required at the application layer. However, for custom use-cases please refer to API details below.

API Reference

Header File

- [components/mbedtls/port/include/ecdsa/ecdsa_alt.h](#)
- This header file can be included with:

```
#include "ecdsa/ecdsa_alt.h"
```

- This header file is a part of the API provided by the `mbedtls` component. To declare that your component depends on `mbedtls`, add the following to your `CMakeLists.txt`:

```
REQUIRES mbedtls
```

or

```
PRIV_REQUIRES mbedtls
```

Functions

int **esp_ecdsa_load_pubkey** (mbedtls_ecp_keypair *keypair, int efuse_blk)

Populate the public key buffer of the mbedtls_ecp_keypair context.

Parameters

- **keypair** -- The mbedtls ECP key-pair structure
- **efuse_blk** -- The efuse key block that should be used as the private key. The key purpose of this block must be ECDSA_KEY

Returns - 0 if successful

- MBEDTLS_ERR_ECP_BAD_INPUT_DATA if invalid ecp group id specified
- MBEDTLS_ERR_ECP_INVALID_KEY if efuse block with purpose ECDSA_KEY is not found
- -1 if invalid efuse block is specified

int **esp_ecdsa_privkey_load_mpi** (mbedtls_mpi *key, int efuse_blk)

Initialize MPI to notify mbedtls_ecdsa_sign to use the private key in efuse We break the MPI struct of the private key in order to differentiate between hardware key and software key.

Parameters

- **key** -- The MPI in which this functions stores the hardware context. This must be uninitialized
- **efuse_blk** -- The efuse key block that should be used as the private key. The key purpose of this block must be ECDSA_KEY

Returns - 0 if successful

- -1 otherwise

int **esp_ecdsa_privkey_load_pk_context** (mbedtls_pk_context *key_ctx, int efuse_blk)

Initialize PK context to notify mbedtls_ecdsa_sign to use the private key in efuse We break the MPI struct used to represent the private key d in ECP keypair in order to differentiate between hardware key and software key.

Parameters

- **key_ctx** -- The context in which this functions stores the hardware context. This must be uninitialized
- **efuse_blk** -- The efuse key block that should be used as the private key. The key purpose of this block must be ECDSA_KEY

Returns - 0 if successful

- -1 otherwise

int **esp_ecdsa_set_pk_context** (mbedtls_pk_context *key_ctx, [esp_ecdsa_pk_conf_t](#) *conf)

Initialize PK context and completely populate mbedtls_ecp_keypair context. We break the MPI struct used to represent the private key d in ECP keypair in order to differentiate between hardware key and software key. We also populate the ECP group field present in the mbedtls_ecp_keypair context. If the ECDSA peripheral of the chip supports exporting the public key, we can also populate the public key buffer of the mbedtls_ecp_keypair context if the load_pubkey flag is set in the [esp_ecdsa_pk_conf_t](#) config argument.

Parameters

- **key_ctx** -- The context in which this functions stores the hardware context. This must be uninitialized
- **conf** -- ESP-ECDSA private key context initialization config structure

Returns - 0 if successful

- -1 otherwise

Structures

struct **esp_ecdsa_pk_conf_t**

ECDSA private key context initialization config structure.

Note: Contains configuration information like the efuse key block that should be used as the private key, EC group ID of the private key and if the export public key operation is supported by the peripheral, a flag load_pubkey that is used specify if the public key has to be populated

Public Members

`mbedtls_ecp_group_id` **grp_id**

MbedTLS ECP group identifier

`uint8_t` **efuse_block**

EFuse block id for ECDSA private key

`bool` **load_pubkey**

Export ECDSA public key from the hardware

2.5.4 Event Task Matrix (ETM)

Introduction

Normally, if a peripheral X needs to notify peripheral Y of a particular event, this could only be done via a CPU interrupt from peripheral X, where the CPU notifies peripheral Y on behalf of peripheral X. However, in time-critical applications, the latency introduced by CPU interrupts is non-negligible.

With the help of the Event Task Matrix (ETM) module, some peripherals can directly notify other peripherals of events through pre-set connections without the intervention of CPU interrupts. This allows precise and low latency synchronization between peripherals, and lessens the CPU's workload as the CPU no longer needs to handle these events.




Fig. 2: ETM channels Overview

The ETM module has multiple programmable channels, they are used to connect a particular **Event** to a particular **Task**. When an event is activated, the ETM channel will trigger the corresponding task automatically.

Peripherals that support ETM functionality provide their or unique set of events and tasks to be connected by the ETM. An ETM channel can connect any event to any task, even looping back an event to a task on the same peripheral. However, an ETM channel can only connect one event to one task at a time (i.e., 1 to 1 relation). If you want to use different events to trigger the same task, you can set up more ETM channels.

Typically, with the help of the ETM module, you can implement features like:

- Toggle the GPIO when a timer alarm event happens
- Start an ADC conversion when a pulse edge is detected on a GPIO

Functional Overview

The following sections of this document cover the typical steps to configure and use the ETM module.

- [ETM Channel Allocation](#) - describes how to install and uninstall the ETM channel.
- [ETM Event](#) - describes how to allocate a new ETM event handle or fetch an existing handle from various peripherals.
- [ETM Task](#) - describes how to allocate a new ETM task handle or fetch an existing handle from various peripherals.
- [ETM Channel Control](#) - describes common ETM channel control functions.
- [Thread Safety](#) - lists which APIs are guaranteed to be thread-safe by the driver.
- [Kconfig Options](#) - lists the supported Kconfig options that can be used to make a different effect on driver behavior.

ETM Channel Allocation There are many identical ETM channels in ESP32-P4¹, and each channel is represented by `esp_etm_channel_handle_t` in the software. The ETM core driver manages all available hardware resources in a pool so that you do not need to care about which channel is in use and which is not. The ETM core driver will allocate a channel for you when you call `esp_etm_new_channel()` and delete it when you call `esp_etm_del_channel()`. All requirements needed for allocating a channel are provided in `esp_etm_channel_config_t`.

Before deleting an ETM channel, please disable it by `esp_etm_channel_disable()` in advance or make sure it has not been enabled yet by `esp_etm_channel_enable()`.

ETM Event ETM Event abstracts the event source, masking the details of specific event sources, and is represented by `esp_etm_event_handle_t` in the software, allowing applications to handle different types of events more easily. ETM events can be generated from a variety of peripherals, thus the way to get the event handle differs from peripherals. When an ETM event is no longer used, you should call `esp_etm_channel_connect()` with a NULL event handle to disconnect it and then call `esp_etm_del_event()` to free the event resource.

GPIO Events GPIO **edge** event is the most common event type, it can be generated by any GPIO pin. You can call `gpio_new_etm_event()` to create a GPIO event handle, with the configurations provided in `gpio_etm_event_config_t`:

- `gpio_etm_event_config_t::edge` or `gpio_etm_event_config_t::edges` decides which edge(s) to trigger the event(s), supported edge types are listed in the `gpio_etm_event_edge_t`.

You need to build a connection between the GPIO ETM event handle and the GPIO number. So you should call `gpio_etm_event_bind_gpio()` afterwards. Please note, only the ETM event handle that created by `gpio_new_etm_event()` can set a GPIO number. Calling this function with other kinds of ETM events returns `ESP_ERR_INVALID_ARG` error. Needless to say, this function does not help with the GPIO initialization, you still need to call `gpio_config()` to set the property like direction, pull up/down mode separately.

Other Peripheral Events

- Refer to [General Purpose Timer \(GPTimer\)](#) for how to get the ETM event handle from GPTimer.
- Refer to [Motor Control Pulse Width Modulator \(MCPWM\)](#) for how to get the ETM event handle from MCPWM.
- Refer to [Analog Comparator](#) for how to get the ETM event handle from analog comparator.

ETM Task ETM Task abstracts the task action and is represented by `esp_etm_task_handle_t` in the software, allowing tasks to be managed and represented in the same way. ETM tasks can be assigned to a variety of peripherals, thus the way to get the task handle differs from peripherals. When an ETM task is no longer

¹ Different ESP chip series might have different numbers of ETM channels. For more details, please refer to [ESP32-P4 Technical Reference Manual > Chapter Event Task Matrix \(ETM\) \[PDF\]](#). The driver does not forbid you from applying for more channels, but it will return an error when all available hardware resources are used up. Please always check the return value when doing channel allocation (i.e., `esp_etm_new_channel()`).

used, you should call `esp_etm_channel_connect()` with a NULL task handle to disconnect it and then call `esp_etm_del_task()` to free the task resource.

GPIO Tasks GPIO task is the most common task type. One GPIO can take one or more GPIO ETM task actions, and one GPIO ETM task action can even manage multiple GPIOs. When the task gets activated by the ETM channel, all managed GPIOs can set/clear/toggle at the same time. You can call `gpio_new_etm_task()` to create a GPIO task handle, with the configurations provided in `gpio_etm_task_config_t`:

- `gpio_etm_task_config_t::action` or `gpio_etm_task_config_t::actions` decides what GPIO action(s) would be taken by the ETM task. Supported actions are listed in the `gpio_etm_task_action_t`. If one GPIO needs to take more than one actions, the action tasks have to be created in one `gpio_new_etm_task()` call with filling the actions into the array of `gpio_etm_task_config_t::actions`.

To build a connection between the GPIO ETM task and the GPIO number, you should call `gpio_etm_task_add_gpio()`. You can call this function by several times if you want the task handle to manage more GPIOs. Please note, only the ETM task handle that created by `gpio_new_etm_task()` can manage a GPIO. Calling this function with other kinds of ETM tasks returns `ESP_ERR_INVALID_ARG` error. Needless to say, this function does not help with the GPIO initialization, you still need to call `gpio_config()` to set the property like direction, pull up/down mode separately.

Before you call `esp_etm_del_task()` to delete the GPIO ETM task, make sure that all previously added GPIOs are removed by `gpio_etm_task_rm_gpio()` in advance.

Other Peripheral Tasks

- Refer to [GPTimer](#) for how to get the ETM task handle from GPTimer.

ETM Channel Control

Connect Event and Task An ETM event has no association with an ETM task, until they are connected to the same ETM channel by calling `esp_etm_channel_connect()`. Especially, calling the function with a NULL task/event handle means disconnecting the channel from any task or event. Note that, this function can be called either before or after the channel is enabled. But calling this function at runtime to change the connection can be dangerous, because the channel may be in the middle of a cycle, and the new connection may not take effect immediately.

Enable and Disable Channel You can call `esp_etm_channel_enable()` and `esp_etm_channel_disable()` to enable and disable the ETM channel from working.

ETM Channel Profiling To check if the ETM channels are set with proper events and tasks, you can call `esp_etm_dump()` to dump all working ETM channels with their associated events and tasks. The dumping format is like:

```
=====ETM Dump Start=====
channel 0: event 48 ==> task 17
channel 1: event 48 ==> task 90
channel 2: event 48 ==> task 94
=====ETM Dump End=====
```

The digital ID printed in the dump information is defined in the `soc/soc_etm_source.h` file.

Thread Safety The factory functions like `esp_etm_new_channel()` and `gpio_new_etm_task()` are guaranteed to be thread-safe by the driver, which means, you can call them from different RTOS tasks without protection by extra locks.

No functions are allowed to run within the ISR environment.

Other functions that take `esp_etm_channel_handle_t`, `esp_etm_task_handle_t` and `esp_etm_event_handle_t` as the first positional parameter, are not treated as thread-safe, which means you should avoid calling them from multiple tasks.

Kconfig Options

- `CONFIG_ETM_ENABLE_DEBUG_LOG` is used to enable the debug log output. Enabling this option increases the firmware binary size as well.

API Reference

Header File

- `components/esp_hw_support/include/esp_etm.h`
- This header file can be included with:

```
#include "esp_etm.h"
```

Functions

`esp_err_t esp_etm_new_channel` (const `esp_etm_channel_config_t` *config, `esp_etm_channel_handle_t` *ret_chan)

Allocate an ETM channel.

Note: The channel can later be freed by `esp_etm_del_channel`

Parameters

- **config** -- [in] ETM channel configuration
- **ret_chan** -- [out] Returned ETM channel handle

Returns

- `ESP_OK`: Allocate ETM channel successfully
- `ESP_ERR_INVALID_ARG`: Allocate ETM channel failed because of invalid argument
- `ESP_ERR_NO_MEM`: Allocate ETM channel failed because of out of memory
- `ESP_ERR_NOT_FOUND`: Allocate ETM channel failed because all channels are used up and no more free one
- `ESP_FAIL`: Allocate ETM channel failed because of other reasons

`esp_err_t esp_etm_del_channel` (`esp_etm_channel_handle_t` chan)

Delete an ETM channel.

Parameters **chan** -- [in] ETM channel handle that created by `esp_etm_new_channel`

Returns

- `ESP_OK`: Delete ETM channel successfully
- `ESP_ERR_INVALID_ARG`: Delete ETM channel failed because of invalid argument
- `ESP_FAIL`: Delete ETM channel failed because of other reasons

`esp_err_t esp_etm_channel_enable` (`esp_etm_channel_handle_t` chan)

Enable ETM channel.

Note: This function will transit the channel state from init to enable.

Parameters **chan** -- [in] ETM channel handle that created by `esp_etm_new_channel`

Returns

- `ESP_OK`: Enable ETM channel successfully
- `ESP_ERR_INVALID_ARG`: Enable ETM channel failed because of invalid argument

- `ESP_ERR_INVALID_STATE`: Enable ETM channel failed because the channel has been enabled already
- `ESP_FAIL`: Enable ETM channel failed because of other reasons

esp_err_t **esp_etm_channel_disable** (*esp_etm_channel_handle_t* chan)

Disable ETM channel.

Note: This function will transit the channel state from enable to init.

Parameters **chan** -- [in] ETM channel handle that created by `esp_etm_new_channel`

Returns

- `ESP_OK`: Disable ETM channel successfully
- `ESP_ERR_INVALID_ARG`: Disable ETM channel failed because of invalid argument
- `ESP_ERR_INVALID_STATE`: Disable ETM channel failed because the channel is not enabled yet
- `ESP_FAIL`: Disable ETM channel failed because of other reasons

esp_err_t **esp_etm_channel_connect** (*esp_etm_channel_handle_t* chan, *esp_etm_event_handle_t* event, *esp_etm_task_handle_t* task)

Connect an ETM event to an ETM task via a previously allocated ETM channel.

Note: Setting the ETM event/task handle to NULL means to disconnect the channel from any event/task

Parameters

- **chan** -- [in] ETM channel handle that created by `esp_etm_new_channel`
- **event** -- [in] ETM event handle obtained from a driver/peripheral, e.g. `xxx_new_etm_event`
- **task** -- [in] ETM task handle obtained from a driver/peripheral, e.g. `xxx_new_etm_task`

Returns

- `ESP_OK`: Connect ETM event and task to the channel successfully
- `ESP_ERR_INVALID_ARG`: Connect ETM event and task to the channel failed because of invalid argument
- `ESP_FAIL`: Connect ETM event and task to the channel failed because of other reasons

esp_err_t **esp_etm_del_event** (*esp_etm_event_handle_t* event)

Delete ETM event.

Note: Although the ETM event comes from various peripherals, we provide the same user API to delete the event handle seamlessly.

Parameters **event** -- [in] ETM event handle obtained from a driver/peripheral, e.g. `xxx_new_etm_event`

Returns

- `ESP_OK`: Delete ETM event successfully
- `ESP_ERR_INVALID_ARG`: Delete ETM event failed because of invalid argument
- `ESP_FAIL`: Delete ETM event failed because of other reasons

esp_err_t **esp_etm_del_task** (*esp_etm_task_handle_t* task)

Delete ETM task.

Note: Although the ETM task comes from various peripherals, we provide the same user API to delete the task handle seamlessly.

Parameters **task** -- **[in]** ETM task handle obtained from a driver/peripheral, e.g.
`xxx_new_etm_task`

Returns

- ESP_OK: Delete ETM task successfully
- ESP_ERR_INVALID_ARG: Delete ETM task failed because of invalid argument
- ESP_FAIL: Delete ETM task failed because of other reasons

esp_err_t **esp_etm_dump** (FILE *out_stream)

Dump ETM channel usages to the given IO stream.

Parameters **out_stream** -- **[in]** IO stream (e.g. stdout)

Returns

- ESP_OK: Dump ETM channel usages successfully
- ESP_ERR_INVALID_ARG: Dump ETM channel usages failed because of invalid argument
- ESP_FAIL: Dump ETM channel usages failed because of other reasons

Structures

struct **esp_etm_channel_config_t**

ETM channel configuration.

Type Definitions

typedef struct esp_etm_channel_t ***esp_etm_channel_handle_t**

ETM channel handle.

typedef struct esp_etm_event_t ***esp_etm_event_handle_t**

ETM event handle.

typedef struct esp_etm_task_t ***esp_etm_task_handle_t**

ETM task handle.

Header File

- [components/driver/gpio/include/driver/gpio_etm.h](#)
- This header file can be included with:

```
#include "driver/gpio_etm.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t **gpio_new_etm_event** (const *gpio_etm_event_config_t* *config, *esp_etm_event_handle_t* *ret_event, ...)

Create an ETM event object for the GPIO peripheral.

Note: The created ETM event object can be deleted later by calling `esp_etm_del_event`

Note: The newly created ETM event object is not bind to any GPIO, you need to call `gpio_etm_event_bind_gpio` to bind the wanted GPIO

Note: Every success call to this function will acquire a free GPIO ETM event channel

Parameters

- **config** -- [in] GPIO ETM event configuration
- **ret_event** -- [out] Returned ETM event handle
- ... -- [out] Other returned ETM event handles if any (the order of the returned event handles is aligned with the array order in field `edges` in *gpio_etm_event_config_t*)

Returns

- ESP_OK: Create ETM event successfully
- ESP_ERR_INVALID_ARG: Create ETM event failed because of invalid argument
- ESP_ERR_NO_MEM: Create ETM event failed because of out of memory
- ESP_ERR_NOT_FOUND: Create ETM event failed because all events are used up and no more free one
- ESP_FAIL: Create ETM event failed because of other reasons

esp_err_t **gpio_etm_event_bind_gpio** (*esp_etm_event_handle_t* event, int gpio_num)

Bind the GPIO with the ETM event.

Note: Calling this function multiple times with different GPIO number can override the previous setting immediately.

Note: Only GPIO ETM object can call this function

Parameters

- **event** -- [in] ETM event handle that created by `gpio_new_etm_event`
- **gpio_num** -- [in] GPIO number that can trigger the ETM event

Returns

- ESP_OK: Set the GPIO for ETM event successfully
- ESP_ERR_INVALID_ARG: Set the GPIO for ETM event failed because of invalid argument, e.g. GPIO is not input capable, ETM event is not of GPIO type
- ESP_FAIL: Set the GPIO for ETM event failed because of other reasons

esp_err_t **gpio_new_etm_task** (const *gpio_etm_task_config_t* *config, *esp_etm_task_handle_t* *ret_task, ...)

Create an ETM task object for the GPIO peripheral.

Note: The created ETM task object can be deleted later by calling `esp_etm_del_task`

Note: The GPIO ETM task works like a container, a newly created ETM task object doesn't have GPIO members to be managed. You need to call `gpio_etm_task_add_gpio` to put one or more GPIOs to the

container.

Note: Every success call to this function will acquire a free GPIO ETM task channel

Parameters

- **config** -- **[in]** GPIO ETM task configuration
- **ret_task** -- **[out]** Returned ETM task handle
- ... -- **[out]** Other returned ETM task handles if any (the order of the returned task handles is aligned with the array order in field `actions` in `gpio_etm_task_config_t`)

Returns

- ESP_OK: Create ETM task successfully
- ESP_ERR_INVALID_ARG: Create ETM task failed because of invalid argument
- ESP_ERR_NO_MEM: Create ETM task failed because of out of memory
- ESP_ERR_NOT_FOUND: Create ETM task failed because all tasks are used up and no more free one
- ESP_FAIL: Create ETM task failed because of other reasons

esp_err_t **gpio_etm_task_add_gpio** (*esp_etm_task_handle_t* task, int gpio_num)

Add GPIO to the ETM task.

Note: You can call this function multiple times to add more GPIOs

Note: Only GPIO ETM object can call this function

Parameters

- **task** -- **[in]** ETM task handle that created by `gpio_new_etm_task`
- **gpio_num** -- **[in]** GPIO number that can be controlled by the ETM task

Returns

- ESP_OK: Add GPIO to the ETM task successfully
- ESP_ERR_INVALID_ARG: Add GPIO to the ETM task failed because of invalid argument, e.g. GPIO is not output capable, ETM task is not of GPIO type
- ESP_ERR_INVALID_STATE: Add GPIO to the ETM task failed because the GPIO is used by other ETM task already
- ESP_FAIL: Add GPIO to the ETM task failed because of other reasons

esp_err_t **gpio_etm_task_rm_gpio** (*esp_etm_task_handle_t* task, int gpio_num)

Remove the GPIO from the ETM task.

Note: Before deleting the ETM task, you need to remove all the GPIOs from the ETM task by this function

Note: Only GPIO ETM object can call this function

Parameters

- **task** -- **[in]** ETM task handle that created by `gpio_new_etm_task`
- **gpio_num** -- **[in]** GPIO number that to be remove from the ETM task

Returns

- ESP_OK: Remove the GPIO from the ETM task successfully
- ESP_ERR_INVALID_ARG: Remove the GPIO from the ETM task failed because of invalid argument

- `ESP_ERR_INVALID_STATE`: Remove the GPIO from the ETM task failed because the GPIO is not controlled by this ETM task
- `ESP_FAIL`: Remove the GPIO from the ETM task failed because of other reasons

Structures

struct **gpio_etm_event_config_t**

GPIO ETM event configuration.

If more than one kind of ETM edge event want to be triggered on the same GPIO pin, you can configure them together. It helps to save GPIO ETM event channel resources for other GPIOs.

Public Members

gpio_etm_event_edge_t **edge**

Which kind of edge can trigger the ETM event module

gpio_etm_event_edge_t **edges**[GPIO_ETM_EVENT_EDGE_TYPES]

Array of kinds of edges to trigger the ETM event module on the same GPIO

struct **gpio_etm_task_config_t**

GPIO ETM task configuration.

If multiple actions wants to be added to the same GPIO pin, you have to configure all the GPIO ETM tasks together.

Public Members

gpio_etm_task_action_t **action**

Action to take by the ETM task module

gpio_etm_task_action_t **actions**[GPIO_ETM_TASK_ACTION_TYPES]

Array of actions to take by the ETM task module on the same GPIO

Macros

GPIO_ETM_EVENT_EDGE_TYPES

GPIO ETM edge events are POS/NEG/ANY

GPIO_ETM_TASK_ACTION_TYPES

GPIO ETM action tasks are SET/CLEAR/TOGGLE

Enumerations

enum **gpio_etm_event_edge_t**

GPIO edges that can be used as ETM event.

Values:

enumerator **GPIO_ETM_EVENT_EDGE_POS**

A rising edge on the GPIO will generate an ETM event signal

enumerator **GPIO_ETM_EVENT_EDGE_NEG**

A falling edge on the GPIO will generate an ETM event signal

enumerator **GPIO_ETM_EVENT_EDGE_ANY**

Any edge on the GPIO can generate an ETM event signal

enum **gpio_etm_task_action_t**

GPIO actions that can be taken by the ETM task.

Values:

enumerator **GPIO_ETM_TASK_ACTION_SET**

Set the GPIO level to high

enumerator **GPIO_ETM_TASK_ACTION_CLR**

Clear the GPIO level to low

enumerator **GPIO_ETM_TASK_ACTION_TOG**

Toggle the GPIO level

Header File

- [components/esp_system/include/esp_systick_etm.h](#)
- This header file can be included with:

```
#include "esp_systick_etm.h"
```

Functions

esp_err_t **esp_systick_new_etm_alarm_event** (int core_id, *esp_etm_event_handle_t* *out_event)

Get the ETM event handle of systick hardware's alarm/heartbeat event.

Note: The created ETM event object can be deleted later by calling `esp_etm_del_event`

Parameters

- **core_id** -- [in] CPU core ID
- **out_event** -- [out] Returned ETM event handle

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.5.5 GPIO & RTC GPIO

GPIO Summary

The ESP32-P4 chip features 57 physical GPIO pins (GPIO0 ~ GPIO56). Each pin can be used as a general-purpose I/O, or to be connected to an internal peripheral signal. Through GPIO matrix and IO MUX, peripheral input signals can be from any IO pins, and peripheral output signals can be routed to any IO pins. Together these modules provide highly configurable I/O. For more details, see *ESP32-P4 Technical Reference Manual > IO MUX and GPIO Matrix (GPIO, IO_MUX)* [PDF].

The table below provides more information on pin usage, and please note the comments in the table for GPIOs with restrictions.

GPIO	Analog Function	LP GPIO	Comments
GPIO0		LP_GPIO0	
GPIO1		LP_GPIO1	
GPIO2	TOUCH0	LP_GPIO2	
GPIO3	TOUCH1	LP_GPIO3	
GPIO4	TOUCH2	LP_GPIO4	
GPIO5	TOUCH3	LP_GPIO5	
GPIO6	TOUCH4	LP_GPIO6	
GPIO7	TOUCH5	LP_GPIO7	
GPIO8	TOUCH6	LP_GPIO8	
GPIO9	TOUCH7	LP_GPIO9	
GPIO10	TOUCH8	LP_GPIO10	
GPIO11	TOUCH9	LP_GPIO11	
GPIO12	TOUCH10	LP_GPIO12	
GPIO13	TOUCH11	LP_GPIO13	
GPIO14	TOUCH12	LP_GPIO14	
GPIO15	TOUCH13	LP_GPIO15	
GPIO16	ADC1_CH0		
GPIO17	ADC1_CH1		
GPIO18	ADC1_CH2		
GPIO19	ADC1_CH3		
GPIO20	ADC1_CH4		
GPIO21	ADC1_CH5		
GPIO22	ADC1_CH6		
GPIO23	ADC1_CH7		
GPIO24			
GPIO25			
GPIO26			
GPIO27			
GPIO28			
GPIO29			
GPIO30			
GPIO31			
GPIO32			
GPIO33			
GPIO34			Strapping pin
GPIO35			Strapping pin
GPIO36			Strapping pin
GPIO37			Strapping pin
GPIO38			Strapping pin
GPIO39			
GPIO40			
GPIO41			
GPIO42			
GPIO43			

continues on next page

Table 2 – continued from previous page

GPIO	Analog Function	LP GPIO	Comments
GPIO44			
GPIO45			
GPIO46			
GPIO47			
GPIO48			
GPIO49	ADC2_CH0		
GPIO50	ADC2_CH1		
GPIO51	ADC2_CH2, ANA_CMPR_CH0 reference voltage		
GPIO52	ADC2_CH3, ANA_CMPR_CH0 input (non-inverting)		
GPIO53	ADC2_CH4, ANA_CMPR_CH1 reference voltage		
GPIO54	ADC2_CH5, ANA_CMPR_CH1 input (non-inverting)		
GPIO55			
GPIO56			

Note:

- **Strapping pin:** GPIO34, GPIO35, GPIO36, GPIO37, and GPIO38 are strapping pins. For more information, please refer to the [Strapping Pin](#) section.
- USB-JTAG: GPIO 24 and 25 are used by USB-JTAG by default. In order to use them as GPIOs, USB-JTAG will be disabled by the drivers.

GPIO driver offers a dump function `gpio_dump_io_configuration()` to show the configurations of the IOs at the moment, such as pull-up / pull-down, input / output enable, pin mapping etc. Below is an example dump:

```
=====IO DUMP Start=====
IO[4] -
  Pullup: 1, Pulldown: 0, DriveCap: 2
  InputEn: 1, OutputEn: 0, OpenDrain: 0
  FuncSel: 1 (GPIO)
  GPIO Matrix SigIn ID: (simple GPIO input)
  SleepSelEn: 1

IO[18] -
  Pullup: 0, Pulldown: 0, DriveCap: 2
  InputEn: 0, OutputEn: 1, OpenDrain: 0
  FuncSel: 1 (GPIO)
  GPIO Matrix SigOut ID: 256 (simple GPIO output)
  SleepSelEn: 1

IO[26] **RESERVED** -
  Pullup: 1, Pulldown: 0, DriveCap: 2
  InputEn: 1, OutputEn: 0, OpenDrain: 0
  FuncSel: 0 (IOMUX)
  SleepSelEn: 1

=====IO DUMP End=====
```

If an IO pin is routed to a peripheral signal through the GPIO matrix, the signal ID printed in the dump information is defined in the `soc/gpio_sig_map.h` file. The word ****RESERVED**** indicates the IO is occupied by either

FLASH or PSRAM. It is strongly not recommended to reconfigure them for other application purposes.

There is also separate "RTC GPIO" support, which functions when GPIOs are routed to the "RTC" low-power, analog subsystem, and Low-Power(LP) peripherals. These pin functions can be used when:

- In Deep-sleep mode
- Analog functions such as ADC/DAC/etc are in use
- LP peripherals, such as LP_UART, LP_I2C, are in use

GPIO Hysteresis Filter

ESP32-P4 support the hardware hysteresis of the input pin, which can reduce the GPIO interrupt shoot by accident due to unstable sampling when the input voltage is near the criteria of logic 0 and 1, especially when the input logic level conversion is slow or the voltage setup time is too long.

Each pin can enable hysteresis function independently. By default, the function is not enabled. You can select the hysteresis control mode by configuring `gpio_config_t::hys_ctrl_mode`. Hysteresis control mode is set along with all the other GPIO configurations in `gpio_config()`.

Application Example

- GPIO output and input interrupt example: [peripherals/gpio/generic_gpio](#).

API Reference - Normal GPIO

Header File

- [components/driver/gpio/include/driver/gpio.h](#)
- This header file can be included with:

```
#include "driver/gpio.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

`esp_err_t gpio_config` (const `gpio_config_t` *pGPIOConfig)

GPIO common configuration.

```
Configure GPIO's Mode, pull-up, PullDown, IntrType
```

Parameters `pGPIOConfig` -- Pointer to GPIO configure struct

Returns

- ESP_OK success
- ESP_ERR_INVALID_ARG Parameter error

`esp_err_t gpio_reset_pin` (`gpio_num_t` gpio_num)

Reset an gpio to default state (select gpio function, enable pullup and disable input and output).

Note: This function also configures the IOMUX for this pin to the GPIO function, and disconnects any other peripheral output configured via GPIO Matrix.

Parameters `gpio_num` -- GPIO number.

Returns Always return ESP_OK.

esp_err_t **gpio_set_intr_type** (gpio_num_t gpio_num, *gpio_int_type_t* intr_type)

GPIO set interrupt trigger type.

Parameters

- **gpio_num** -- GPIO number. If you want to set the trigger type of e.g. of GPIO16, gpio_num should be GPIO_NUM_16 (16);
- **intr_type** -- Interrupt type, select from gpio_int_type_t

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **gpio_intr_enable** (gpio_num_t gpio_num)

Enable GPIO module interrupt signal.

Note: ESP32: Please do not use the interrupt of GPIO36 and GPIO39 when using ADC or Wi-Fi and Bluetooth with sleep mode enabled. Please refer to the comments of `adc1_get_raw`. Please refer to Section 3.11 of [ESP32 ECO and Workarounds for Bugs](#) for the description of this issue.

Parameters `gpio_num` -- GPIO number. If you want to enable an interrupt on e.g. GPIO16, gpio_num should be GPIO_NUM_16 (16);

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **gpio_intr_disable** (gpio_num_t gpio_num)

Disable GPIO module interrupt signal.

Note: This function is allowed to be executed when Cache is disabled within ISR context, by enabling `CONFIG_GPIO_CTRL_FUNC_IN_IRAM`

Parameters `gpio_num` -- GPIO number. If you want to disable the interrupt of e.g. GPIO16, gpio_num should be GPIO_NUM_16 (16);

Returns

- ESP_OK success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **gpio_set_level** (gpio_num_t gpio_num, uint32_t level)

GPIO set output level.

Note: This function is allowed to be executed when Cache is disabled within ISR context, by enabling `CONFIG_GPIO_CTRL_FUNC_IN_IRAM`

Parameters

- **gpio_num** -- GPIO number. If you want to set the output level of e.g. GPIO16, gpio_num should be GPIO_NUM_16 (16);
- **level** -- Output level. 0: low ; 1: high

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG GPIO number error

int **gpio_get_level** (gpio_num_t gpio_num)

GPIO get input level.

Warning: If the pad is not configured for input (or input and output) the returned value is always 0.

Parameters **gpio_num** -- GPIO number. If you want to get the logic level of e.g. pin GPIO16, gpio_num should be GPIO_NUM_16 (16);

Returns

- 0 the GPIO input level is 0
- 1 the GPIO input level is 1

esp_err_t **gpio_set_direction** (gpio_num_t gpio_num, *gpio_mode_t* mode)

GPIO set direction.

Configure GPIO direction,such as output_only,input_only,output_and_input

Parameters

- **gpio_num** -- Configure GPIO pins number, it should be GPIO number. If you want to set direction of e.g. GPIO16, gpio_num should be GPIO_NUM_16 (16);
- **mode** -- GPIO direction

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG GPIO error

esp_err_t **gpio_set_pull_mode** (gpio_num_t gpio_num, *gpio_pull_mode_t* pull)

Configure GPIO pull-up/pull-down resistors.

Note: ESP32: Only pins that support both input & output have integrated pull-up and pull-down resistors. Input-only GPIOs 34-39 do not.

Parameters

- **gpio_num** -- GPIO number. If you want to set pull up or down mode for e.g. GPIO16, gpio_num should be GPIO_NUM_16 (16);
- **pull** -- GPIO pull up/down mode.

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG : Parameter error

esp_err_t **gpio_wakeup_enable** (gpio_num_t gpio_num, *gpio_int_type_t* intr_type)

Enable GPIO wake-up function.

Parameters

- **gpio_num** -- GPIO number.
- **intr_type** -- GPIO wake-up type. Only GPIO_INTR_LOW_LEVEL or GPIO_INTR_HIGH_LEVEL can be used.

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **gpio_wakeup_disable** (gpio_num_t gpio_num)

Disable GPIO wake-up function.

Parameters **gpio_num** -- GPIO number

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **gpio_isr_register** (void (*fn)(void*), void *arg, int intr_alloc_flags, *gpio_isr_handle_t* *handle)

Register GPIO interrupt handler, the handler is an ISR. The handler will be attached to the same CPU core that this function is running on.

This ISR function is called whenever any GPIO interrupt occurs. See the alternative `gpio_install_isr_service()` and `gpio_isr_handler_add()` API in order to have the driver support per-GPIO ISRs.

To disable or remove the ISR, pass the returned handle to the *interrupt allocation functions*.

Parameters

- **fn** -- Interrupt handler function.
- **arg** -- Parameter for handler function
- **intr_alloc_flags** -- Flags used to allocate the interrupt. One or multiple (ORred) ESP_INTR_FLAG_* values. See `esp_intr_alloc.h` for more info.
- **handle** -- Pointer to return handle. If non-NULL, a handle for the interrupt will be returned here.

Returns

- ESP_OK Success ;
- ESP_ERR_INVALID_ARG GPIO error
- ESP_ERR_NOT_FOUND No free interrupt found with the specified flags

esp_err_t **gpio_pullup_en** (gpio_num_t gpio_num)

Enable pull-up on GPIO.

Parameters **gpio_num** -- GPIO number

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **gpio_pullup_dis** (gpio_num_t gpio_num)

Disable pull-up on GPIO.

Parameters **gpio_num** -- GPIO number

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **gpio_pulldown_en** (gpio_num_t gpio_num)

Enable pull-down on GPIO.

Parameters **gpio_num** -- GPIO number

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **gpio_pulldown_dis** (gpio_num_t gpio_num)

Disable pull-down on GPIO.

Parameters **gpio_num** -- GPIO number

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **gpio_install_isr_service** (int intr_alloc_flags)

Install the GPIO driver's ETS_GPIO_INTR_SOURCE ISR handler service, which allows per-pin GPIO interrupt handlers.

This function is incompatible with `gpio_isr_register()` - if that function is used, a single global ISR is registered for all GPIO interrupts. If this function is used, the ISR service provides a global GPIO ISR and individual pin handlers are registered via the `gpio_isr_handler_add()` function.

Parameters **intr_alloc_flags** -- Flags used to allocate the interrupt. One or multiple (ORred) ESP_INTR_FLAG_* values. See esp_intr_alloc.h for more info.

Returns

- ESP_OK Success
- ESP_ERR_NO_MEM No memory to install this service
- ESP_ERR_INVALID_STATE ISR service already installed.
- ESP_ERR_NOT_FOUND No free interrupt found with the specified flags
- ESP_ERR_INVALID_ARG GPIO error

void **gpio_uninstall_isr_service** (void)

Uninstall the driver's GPIO ISR service, freeing related resources.

esp_err_t **gpio_isr_handler_add** (gpio_num_t gpio_num, *gpio_isr_t* isr_handler, void *args)

Add ISR handler for the corresponding GPIO pin.

Call this function after using gpio_install_isr_service() to install the driver's GPIO ISR handler service.

The pin ISR handlers no longer need to be declared with IRAM_ATTR, unless you pass the ESP_INTR_FLAG_IRAM flag when allocating the ISR in gpio_install_isr_service().

This ISR handler will be called from an ISR. So there is a stack size limit (configurable as "ISR stack size" in menuconfig). This limit is smaller compared to a global GPIO interrupt handler due to the additional level of indirection.

Parameters

- **gpio_num** -- GPIO number
- **isr_handler** -- ISR handler function for the corresponding GPIO number.
- **args** -- parameter for ISR handler.

Returns

- ESP_OK Success
- ESP_ERR_INVALID_STATE Wrong state, the ISR service has not been initialized.
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **gpio_isr_handler_remove** (gpio_num_t gpio_num)

Remove ISR handler for the corresponding GPIO pin.

Parameters **gpio_num** -- GPIO number

Returns

- ESP_OK Success
- ESP_ERR_INVALID_STATE Wrong state, the ISR service has not been initialized.
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **gpio_set_drive_capability** (gpio_num_t gpio_num, *gpio_drive_cap_t* strength)

Set GPIO pad drive capability.

Parameters

- **gpio_num** -- GPIO number, only support output GPIOs
- **strength** -- Drive capability of the pad

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **gpio_get_drive_capability** (gpio_num_t gpio_num, *gpio_drive_cap_t* *strength)

Get GPIO pad drive capability.

Parameters

- **gpio_num** -- GPIO number, only support output GPIOs
- **strength** -- Pointer to accept drive capability of the pad

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **gpio_hold_en** (gpio_num_t gpio_num)

Enable gpio pad hold function.

When a GPIO is set to hold, its state is latched at that moment and will not change when the internal signal or the IO MUX/GPIO configuration is modified (including input enable, output enable, output value, function, and drive strength values). This function can be used to retain the state of GPIOs when the power domain of where GPIO/IOMUX belongs to becomes off. For example, chip or system is reset (e.g. watchdog time-out, deep-sleep events are triggered), or peripheral power-down in light-sleep.

This function works in both input and output modes, and only applicable to output-capable GPIOs. If this function is enabled: in output mode: the output level of the GPIO will be locked and can not be changed. in input mode: the input read value can still reflect the changes of the input signal.

However, on ESP32/S2/C3/S3/C2, this function cannot be used to hold the state of a digital GPIO during Deep-sleep. Even if this function is enabled, the digital GPIO will be reset to its default state when the chip wakes up from Deep-sleep. If you want to hold the state of a digital GPIO during Deep-sleep, please call `gpio_deep_sleep_hold_en`.

Power down or call `gpio_hold_dis` will disable this function.

Parameters `gpio_num` -- GPIO number, only support output-capable GPIOs

Returns

- ESP_OK Success
- ESP_ERR_NOT_SUPPORTED Not support pad hold function

esp_err_t **gpio_hold_dis** (gpio_num_t gpio_num)

Disable gpio pad hold function.

When the chip is woken up from peripheral power-down sleep, the gpio will be set to the default mode, so, the gpio will output the default level if this function is called. If you don't want the level changes, the gpio should be configured to a known state before this function is called. e.g. If you hold gpio18 high during Deep-sleep, after the chip is woken up and `gpio_hold_dis` is called, gpio18 will output low level(because gpio18 is input mode by default). If you don't want this behavior, you should configure gpio18 as output mode and set it to high level before calling `gpio_hold_dis`.

Parameters `gpio_num` -- GPIO number, only support output-capable GPIOs

Returns

- ESP_OK Success
- ESP_ERR_NOT_SUPPORTED Not support pad hold function

void **gpio_iomux_in** (uint32_t gpio_num, uint32_t signal_idx)

SOC_GPIO_SUPPORT_HOLD_SINGLE_IO_IN_DSLP.

Set pad input to a peripheral signal through the IOMUX.

Parameters

- `gpio_num` -- GPIO number of the pad.
- `signal_idx` -- Peripheral signal id to input. One of the `*_IN_IDX` signals in `soc/gpio_sig_map.h`.

void **gpio_iomux_out** (uint8_t gpio_num, int func, bool oen_inv)

Set peripheral output to an GPIO pad through the IOMUX.

Parameters

- `gpio_num` -- gpio_num GPIO number of the pad.
- `func` -- The function number of the peripheral pin to output pin. One of the `FUNC_X_*` of specified pin (X) in `soc/io_mux_reg.h`.
- `oen_inv` -- True if the output enable needs to be inverted, otherwise False.

esp_err_t **gpio_force_hold_all** (void)

Force hold all digital and rtc gpio pads.

GPIO force hold, no matter the chip in active mode or sleep modes.

This function will immediately cause all pads to latch the current values of input enable, output enable, output value, function, and drive strength values.

Warning:

- This function will hold flash and UART pins as well. Therefore, this function, and all code run afterwards (till calling `gpio_force_unhold_all` to disable this feature), **MUST** be placed in internal RAM as holding the flash pins will halt SPI flash operation, and holding the UART pins will halt any UART logging.
- The hold state of all pads will be cancelled during ROM boot, so it is not recommended to use this API to hold the pads state during deepsleep and reset.

esp_err_t **gpio_force_unhold_all** (void)

Unhold all digital and rtc gpio pads.

Note: The global hold signal and the hold signal of each IO act on the PAD through 'or' logic, so if a pad has already been configured to hold by `gpio_hold_en`, this API can't release its hold state.

esp_err_t **gpio_sleep_sel_en** (gpio_num_t gpio_num)

Enable SLP_SEL to change GPIO status automatically in lightsleep.

Parameters `gpio_num` -- GPIO number of the pad.

Returns

- ESP_OK Success

esp_err_t **gpio_sleep_sel_dis** (gpio_num_t gpio_num)

Disable SLP_SEL to change GPIO status automatically in lightsleep.

Parameters `gpio_num` -- GPIO number of the pad.

Returns

- ESP_OK Success

esp_err_t **gpio_sleep_set_direction** (gpio_num_t gpio_num, *gpio_mode_t* mode)

GPIO set direction at sleep.

Configure GPIO direction,such as output_only,input_only,output_and_input

Parameters

- **gpio_num** -- Configure GPIO pins number, it should be GPIO number. If you want to set direction of e.g. GPIO16, gpio_num should be GPIO_NUM_16 (16);
- **mode** -- GPIO direction

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG GPIO error

esp_err_t **gpio_sleep_set_pull_mode** (gpio_num_t gpio_num, *gpio_pull_mode_t* pull)

Configure GPIO pull-up/pull-down resistors at sleep.

Note: ESP32: Only pins that support both input & output have integrated pull-up and pull-down resistors. Input-only GPIOs 34-39 do not.

Parameters

- **gpio_num** -- GPIO number. If you want to set pull up or down mode for e.g. GPIO16, gpio_num should be GPIO_NUM_16 (16);
- **pull** -- GPIO pull up/down mode.

Returns

- ESP_OK Success

- `ESP_ERR_INVALID_ARG` : Parameter error

esp_err_t **gpio_deep_sleep_wakeup_enable** (gpio_num_t gpio_num, *gpio_int_type_t* intr_type)

Enable GPIO deep-sleep wake-up function.

Note: Called by the SDK. User shouldn't call this directly in the APP.

Parameters

- **gpio_num** -- GPIO number.
- **intr_type** -- GPIO wake-up type. Only `GPIO_INTR_LOW_LEVEL` or `GPIO_INTR_HIGH_LEVEL` can be used.

Returns

- `ESP_OK` Success
- `ESP_ERR_INVALID_ARG` Parameter error

esp_err_t **gpio_deep_sleep_wakeup_disable** (gpio_num_t gpio_num)

Disable GPIO deep-sleep wake-up function.

Parameters **gpio_num** -- GPIO number

Returns

- `ESP_OK` Success
- `ESP_ERR_INVALID_ARG` Parameter error

esp_err_t **gpio_dump_io_configuration** (FILE *out_stream, uint64_t io_bit_mask)

Dump IO configuration information to console.

Parameters

- **out_stream** -- IO stream (e.g. stdout)
- **io_bit_mask** -- IO pin bit mask, each bit maps to an IO

Returns

- `ESP_OK` Success
- `ESP_ERR_INVALID_ARG` Parameter error

Structures

struct **gpio_config_t**

Configuration parameters of GPIO pad for `gpio_config` function.

Public Members

uint64_t **pin_bit_mask**

GPIO pin: set with bit mask, each bit maps to a GPIO

gpio_mode_t **mode**

GPIO mode: set input/output mode

gpio_pullup_t **pull_up_en**

GPIO pull-up

gpio_pulldown_t **pull_down_en**

GPIO pull-down

gpio_int_type_t **intr_type**

GPIO interrupt type

gpio_hys_ctrl_mode_t **hys_ctrl_mode**

GPIO hysteresis: hysteresis filter on slope input

Macros

GPIO_PIN_COUNT

GPIO_IS_VALID_GPIO (gpio_num)

Check whether it is a valid GPIO number.

GPIO_IS_VALID_OUTPUT_GPIO (gpio_num)

Check whether it can be a valid GPIO number of output mode.

GPIO_IS_VALID_DIGITAL_IO_PAD (gpio_num)

Check whether it can be a valid digital I/O pad.

GPIO_IS_DEEP_SLEEP_WAKEUP_VALID_GPIO (gpio_num)

Type Definitions

typedef *intr_handle_t* **gpio_isr_handle_t**

typedef void (***gpio_isr_t**)(void *arg)

GPIO interrupt handler.

Param arg User registered data

Header File

- [components/hal/include/hal/gpio_types.h](#)
- This header file can be included with:

```
#include "hal/gpio_types.h"
```

Macros

GPIO_PIN_REG_0

GPIO_PIN_REG_1

GPIO_PIN_REG_2

GPIO_PIN_REG_3

GPIO_PIN_REG_4

GPIO_PIN_REG_5

GPIO_PIN_REG_6

GPIO_PIN_REG_7

GPIO_PIN_REG_8

GPIO_PIN_REG_9

GPIO_PIN_REG_10

GPIO_PIN_REG_11

GPIO_PIN_REG_12

GPIO_PIN_REG_13

GPIO_PIN_REG_14

GPIO_PIN_REG_15

GPIO_PIN_REG_16

GPIO_PIN_REG_17

GPIO_PIN_REG_18

GPIO_PIN_REG_19

GPIO_PIN_REG_20

GPIO_PIN_REG_21

GPIO_PIN_REG_22

GPIO_PIN_REG_23

GPIO_PIN_REG_24

GPIO_PIN_REG_25

GPIO_PIN_REG_26

GPIO_PIN_REG_27

GPIO_PIN_REG_28

GPIO_PIN_REG_29

GPIO_PIN_REG_30

GPIO_PIN_REG_31

GPIO_PIN_REG_32

GPIO_PIN_REG_33

GPIO_PIN_REG_34

GPIO_PIN_REG_35

GPIO_PIN_REG_36

GPIO_PIN_REG_37

GPIO_PIN_REG_38

GPIO_PIN_REG_39

GPIO_PIN_REG_40

GPIO_PIN_REG_41

GPIO_PIN_REG_42

GPIO_PIN_REG_43

GPIO_PIN_REG_44

GPIO_PIN_REG_45

GPIO_PIN_REG_46

GPIO_PIN_REG_47

GPIO_PIN_REG_48

GPIO_PIN_REG_49

GPIO_PIN_REG_50

GPIO_PIN_REG_51

GPIO_PIN_REG_52

GPIO_PIN_REG_53

GPIO_PIN_REG_54

GPIO_PIN_REG_55

GPIO_PIN_REG_56

Enumerations

enum **gpio_port_t**

Values:

enumerator **GPIO_PORT_0**

enumerator **GPIO_PORT_MAX**

enum **gpio_int_type_t**

Values:

enumerator **GPIO_INTR_DISABLE**

Disable GPIO interrupt

enumerator **GPIO_INTR_POSEDGE**

GPIO interrupt type : rising edge

enumerator **GPIO_INTR_NEGEDGE**

GPIO interrupt type : falling edge

enumerator **GPIO_INTR_ANYEDGE**

GPIO interrupt type : both rising and falling edge

enumerator **GPIO_INTR_LOW_LEVEL**

GPIO interrupt type : input low level trigger

enumerator **GPIO_INTR_HIGH_LEVEL**

GPIO interrupt type : input high level trigger

enumerator **GPIO_INTR_MAX**

enum **gpio_mode_t**

Values:

enumerator **GPIO_MODE_DISABLE**

GPIO mode : disable input and output

enumerator **GPIO_MODE_INPUT**

GPIO mode : input only

enumerator **GPIO_MODE_OUTPUT**

GPIO mode : output only mode

enumerator **GPIO_MODE_OUTPUT_OD**

GPIO mode : output only with open-drain mode

enumerator **GPIO_MODE_INPUT_OUTPUT_OD**

GPIO mode : output and input with open-drain mode

enumerator **GPIO_MODE_INPUT_OUTPUT**

GPIO mode : output and input mode

enum **gpio_pullup_t**

Values:

enumerator **GPIO_PULLUP_DISABLE**

Disable GPIO pull-up resistor

enumerator **GPIO_PULLUP_ENABLE**

Enable GPIO pull-up resistor

enum **gpio_pulldown_t**

Values:

enumerator **GPIO_PULLDOWN_DISABLE**

Disable GPIO pull-down resistor

enumerator **GPIO_PULLDOWN_ENABLE**

Enable GPIO pull-down resistor

enum **gpio_pull_mode_t**

Values:

enumerator **GPIO_PULLUP_ONLY**

Pad pull up

enumerator **GPIO_PULLDOWN_ONLY**

Pad pull down

enumerator **GPIO_PULLUP_PULLDOWN**

Pad pull up + pull down

enumerator **GPIO_FLOATING**

Pad floating

enum **gpio_drive_cap_t**

Values:

enumerator **GPIO_DRIVE_CAP_0**

Pad drive capability: weak

enumerator **GPIO_DRIVE_CAP_1**

Pad drive capability: stronger

enumerator **GPIO_DRIVE_CAP_2**

Pad drive capability: medium

enumerator **GPIO_DRIVE_CAP_DEFAULT**

Pad drive capability: medium

enumerator **GPIO_DRIVE_CAP_3**

Pad drive capability: strongest

enumerator **GPIO_DRIVE_CAP_MAX**

enum **gpio_hys_ctrl_mode_t**

Available option for configuring hysteresis feature of GPIOs.

Values:

enumerator **GPIO_HYS_SOFT_DISABLE**

Pad input hysteresis disable by software

enumerator **GPIO_HYS_SOFT_ENABLE**

Pad input hysteresis enable by software

API Reference - RTC GPIO

Header File

- [components/driver/gpio/include/driver/rtc_io.h](#)
- This header file can be included with:

```
#include "driver/rtc_io.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

bool **rtc_gpio_is_valid_gpio** (gpio_num_t gpio_num)

Determine if the specified GPIO is a valid RTC GPIO.

Parameters `gpio_num` -- GPIO number

Returns true if GPIO is valid for RTC GPIO use. false otherwise.

int **rtc_io_number_get** (gpio_num_t gpio_num)

Get RTC IO index number by gpio number.

Parameters **gpio_num** -- GPIO number

Returns ≥ 0 : Index of rtcio. -1 : The gpio is not rtcio.

esp_err_t **rtc_gpio_init** (gpio_num_t gpio_num)

Init a GPIO as RTC GPIO.

This function must be called when initializing a pad for an analog function.

Parameters **gpio_num** -- GPIO number (e.g. GPIO_NUM_12)

Returns

- ESP_OK success
- ESP_ERR_INVALID_ARG GPIO is not an RTC IO

esp_err_t **rtc_gpio_deinit** (gpio_num_t gpio_num)

Init a GPIO as digital GPIO.

Parameters **gpio_num** -- GPIO number (e.g. GPIO_NUM_12)

Returns

- ESP_OK success
- ESP_ERR_INVALID_ARG GPIO is not an RTC IO

uint32_t **rtc_gpio_get_level** (gpio_num_t gpio_num)

Get the RTC IO input level.

Parameters **gpio_num** -- GPIO number (e.g. GPIO_NUM_12)

Returns

- 1 High level
- 0 Low level
- ESP_ERR_INVALID_ARG GPIO is not an RTC IO

esp_err_t **rtc_gpio_set_level** (gpio_num_t gpio_num, uint32_t level)

Set the RTC IO output level.

Parameters

- **gpio_num** -- GPIO number (e.g. GPIO_NUM_12)
- **level** -- output level

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG GPIO is not an RTC IO

esp_err_t **rtc_gpio_set_direction** (gpio_num_t gpio_num, *rtc_gpio_mode_t* mode)

RTC GPIO set direction.

Configure RTC GPIO direction, such as output only, input only, output and input.

Parameters

- **gpio_num** -- GPIO number (e.g. GPIO_NUM_12)
- **mode** -- GPIO direction

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG GPIO is not an RTC IO

esp_err_t **rtc_gpio_set_direction_in_sleep** (gpio_num_t gpio_num, *rtc_gpio_mode_t* mode)

RTC GPIO set direction in deep sleep mode or disable sleep status (default). In some application scenarios, IO needs to have another states during deep sleep.

NOTE: ESP32 supports INPUT_ONLY mode. The rest targets support INPUT_ONLY, OUTPUT_ONLY, INPUT_OUTPUT mode.

Parameters

- **gpio_num** -- GPIO number (e.g. GPIO_NUM_12)

- **mode** -- GPIO direction

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG GPIO is not an RTC IO

esp_err_t **rtc_gpio_pullup_en** (gpio_num_t gpio_num)

RTC GPIO pullup enable.

This function only works for RTC IOs. In general, call `gpio_pullup_en`, which will work both for normal GPIOs and RTC IOs.

Parameters **gpio_num** -- GPIO number (e.g. GPIO_NUM_12)

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG GPIO is not an RTC IO

esp_err_t **rtc_gpio_pulldown_en** (gpio_num_t gpio_num)

RTC GPIO pulldown enable.

This function only works for RTC IOs. In general, call `gpio_pulldown_en`, which will work both for normal GPIOs and RTC IOs.

Parameters **gpio_num** -- GPIO number (e.g. GPIO_NUM_12)

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG GPIO is not an RTC IO

esp_err_t **rtc_gpio_pullup_dis** (gpio_num_t gpio_num)

RTC GPIO pullup disable.

This function only works for RTC IOs. In general, call `gpio_pullup_dis`, which will work both for normal GPIOs and RTC IOs.

Parameters **gpio_num** -- GPIO number (e.g. GPIO_NUM_12)

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG GPIO is not an RTC IO

esp_err_t **rtc_gpio_pulldown_dis** (gpio_num_t gpio_num)

RTC GPIO pulldown disable.

This function only works for RTC IOs. In general, call `gpio_pulldown_dis`, which will work both for normal GPIOs and RTC IOs.

Parameters **gpio_num** -- GPIO number (e.g. GPIO_NUM_12)

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG GPIO is not an RTC IO

esp_err_t **rtc_gpio_set_drive_capability** (gpio_num_t gpio_num, *gpio_drive_cap_t* strength)

Set RTC GPIO pad drive capability.

Parameters

- **gpio_num** -- GPIO number, only support output GPIOs
- **strength** -- Drive capability of the pad

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **rtc_gpio_get_drive_capability** (gpio_num_t gpio_num, *gpio_drive_cap_t* *strength)

Get RTC GPIO pad drive capability.

Parameters

- **gpio_num** -- GPIO number, only support output GPIOs
- **strength** -- Pointer to accept drive capability of the pad

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **rtc_gpio_iomux_func_sel** (gpio_num_t gpio_num, int func)

Select a RTC IOMUX function for the RTC IO.

Parameters

- **gpio_num** -- GPIO number
- **func** -- Function to assign to the pin

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **rtc_gpio_hold_en** (gpio_num_t gpio_num)

Enable hold function on an RTC IO pad.

Enabling HOLD function will cause the pad to latch current values of input enable, output enable, output value, function, drive strength values. This function is useful when going into light or deep sleep mode to prevent the pin configuration from changing.

Parameters **gpio_num** -- GPIO number (e.g. GPIO_NUM_12)

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG GPIO is not an RTC IO

esp_err_t **rtc_gpio_hold_dis** (gpio_num_t gpio_num)

Disable hold function on an RTC IO pad.

Disabling hold function will allow the pad receive the values of input enable, output enable, output value, function, drive strength from RTC_IO peripheral.

Parameters **gpio_num** -- GPIO number (e.g. GPIO_NUM_12)

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG GPIO is not an RTC IO

esp_err_t **rtc_gpio_force_hold_en_all** (void)

Enable force hold signal for all RTC IOs.

Each RTC pad has a "force hold" input signal from the RTC controller. If this signal is set, pad latches current values of input enable, function, output enable, and other signals which come from the RTC mux. Force hold signal is enabled before going into deep sleep for pins which are used for EXT1 wakeup.

esp_err_t **rtc_gpio_force_hold_dis_all** (void)

Disable force hold signal for all RTC IOs.

esp_err_t **rtc_gpio_wakeup_enable** (gpio_num_t gpio_num, *gpio_int_type_t* intr_type)

Enable wakeup from sleep mode using specific GPIO.

Parameters

- **gpio_num** -- GPIO number
- **intr_type** -- Wakeup on high level (GPIO_INTR_HIGH_LEVEL) or low level (GPIO_INTR_LOW_LEVEL)

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if gpio_num is not an RTC IO, or intr_type is not one of GPIO_INTR_HIGH_LEVEL, GPIO_INTR_LOW_LEVEL.

esp_err_t **rtc_gpio_wakeup_disable** (gpio_num_t gpio_num)

Disable wakeup from sleep mode using specific GPIO.

Parameters **gpio_num** -- GPIO number

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if gpio_num is not an RTC IO

Macros

RTC_GPIO_IS_VALID_GPIO (gpio_num)

Header File

- `components/driver/gpio/include/driver/lp_io.h`
- This header file can be included with:

```
#include "driver/lp_io.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t **lp_gpio_connect_in_signal** (gpio_num_t gpio_num, uint32_t signal_idx, bool inv)

Connect a RTC(LP) GPIO input with a peripheral signal, which tagged as input attribute.

Note: There's no limitation on the number of signals that a RTC(LP) GPIO can connect with

Parameters

- **gpio_num** -- GPIO number, especially, `LP_GPIO_MATRIX_CONST_ZERO_INPUT` means connect logic 0 to signal `LP_GPIO_MATRIX_CONST_ONE_INPUT` means connect logic 1 to signal
- **signal_idx** -- LP peripheral signal index (tagged as input attribute)
- **inv** -- Whether the RTC(LP) GPIO input to be inverted or not

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **lp_gpio_connect_out_signal** (gpio_num_t gpio_num, uint32_t signal_idx, bool out_inv, bool oen_inv)

Connect a peripheral signal which tagged as output attribute with a RTC(LP) GPIO.

Note: There's no limitation on the number of RTC(LP) GPIOs that a signal can connect with

Parameters

- **gpio_num** -- GPIO number
- **signal_idx** -- LP peripheral signal index (tagged as input attribute), especially, `SIG_LP_GPIO_OUT_IDX` means disconnect RTC(LP) GPIO and other peripherals. Only the RTC GPIO driver can control the output level
- **out_inv** -- Whether to signal to be inverted or not
- **oen_inv** -- Whether the output enable control is inverted or not

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

Header File

- [components/hal/include/hal/rtc_io_types.h](#)
- This header file can be included with:

```
#include "hal/rtc_io_types.h"
```

Enumerations

enum **rtc_gpio_mode_t**

RTCIO output/input mode type.

Values:

enumerator **RTC_GPIO_MODE_INPUT_ONLY**

Pad input

enumerator **RTC_GPIO_MODE_OUTPUT_ONLY**

Pad output

enumerator **RTC_GPIO_MODE_INPUT_OUTPUT**

Pad input + output

enumerator **RTC_GPIO_MODE_DISABLED**

Pad (output + input) disable

enumerator **RTC_GPIO_MODE_OUTPUT_OD**

Pad open-drain output

enumerator **RTC_GPIO_MODE_INPUT_OUTPUT_OD**

Pad input + open-drain output

2.5.6 General Purpose Timer (GPTimer)

Introduction

GPTimer (General Purpose Timer) is the driver of ESP32-P4 Timer Group peripheral. The hardware timer features high resolution and flexible alarm action. The behavior when the internal counter of a timer reaches a specific target value is called a timer alarm. When a timer alarms, a user registered per-timer callback would be called.

Typically, a general purpose timer can be used in scenarios like:

- Free running as a wall clock, fetching a high-resolution timestamp at any time and any places
- Generate period alarms, trigger events periodically
- Generate one-shot alarm, respond in target time

Functional Overview

The following sections of this document cover the typical steps to install and operate a timer:

- [Resource Allocation](#) - covers which parameters should be set up to get a timer handle and how to recycle the resources when GPTimer finishes working.

- [Set and Get Count Value](#) - covers how to force the timer counting from a start point and how to get the count value at anytime.
- [Set up Alarm Action](#) - covers the parameters that should be set up to enable the alarm event.
- [Register Event Callbacks](#) - covers how to hook user specific code to the alarm event callback function.
- [Enable and Disable Timer](#) - covers how to enable and disable the timer.
- [Start and Stop Timer](#) - shows some typical use cases that start the timer with different alarm behavior.
- [ETM Event and Task](#) - describes what the events and tasks can be connected to the ETM channel.
- [Power Management](#) - describes how different source clock selections can affect power consumption.
- [IRAM Safe](#) - describes tips on how to make the timer interrupt and IO control functions work better along with a disabled cache.
- [Thread Safety](#) - lists which APIs are guaranteed to be thread safe by the driver.
- [Kconfig Options](#) - lists the supported Kconfig options that can be used to make a different effect on driver behavior.

Resource Allocation Different ESP chips might have different numbers of independent timer groups, and within each group, there could also be several independent timers.¹

A GPTimer instance is represented by `gptimer_handle_t`. The driver behind manages all available hardware resources in a pool, so that you do not need to care about which timer and which group it belongs to.

To install a timer instance, there is a configuration structure that needs to be given in advance: `gptimer_config_t`:

- `gptimer_config_t::clk_src` selects the source clock for the timer. The available clocks are listed in `gptimer_clock_source_t`, you can only pick one of them. For the effect on power consumption of different clock source, please refer to Section [Power Management](#).
- `gptimer_config_t::direction` sets the counting direction of the timer, supported directions are listed in `gptimer_count_direction_t`, you can only pick one of them.
- `gptimer_config_t::resolution_hz` sets the resolution of the internal counter. Each count step is equivalent to $1 / \text{resolution\_hz}$ seconds.
- `gptimer_config::intr_priority` sets the priority of the timer interrupt. If it is set to 0, the driver will allocate an interrupt with a default priority. Otherwise, the driver will use the given priority.
- Optional `gptimer_config_t::intr_shared` sets whether or not mark the timer interrupt source as a shared one. For the pros/cons of a shared interrupt, you can refer to [Interrupt Handling](#).

With all the above configurations set in the structure, the structure can be passed to `gptimer_new_timer()` which will instantiate the timer instance and return a handle of the timer.

The function can fail due to various errors such as insufficient memory, invalid arguments, etc. Specifically, when there are no more free timers (i.e., all hardware resources have been used up), then `ESP_ERR_NOT_FOUND` will be returned. The total number of available timers is represented by the `SOC_TIMER_GROUP_TOTAL_TIMERS` and its value depends on the ESP chip.

If a previously created GPTimer instance is no longer required, you should recycle the timer by calling `gptimer_del_timer()`. This allows the underlying HW timer to be used for other purposes. Before deleting a GPTimer handle, please disable it by `gptimer_disable()` in advance or make sure it has not enabled yet by `gptimer_enable()`.

Creating a GPTimer Handle with Resolution of 1 MHz

```
gptimer_handle_t gptimer = NULL;
gptimer_config_t timer_config = {
    .clk_src = GPTIMER_CLK_SRC_DEFAULT,
    .direction = GPTIMER_COUNT_UP,
    .resolution_hz = 1 * 1000 * 1000, // 1MHz, 1 tick = 1us
};
ESP_ERROR_CHECK(gptimer_new_timer(&timer_config, &gptimer));
```

¹ Different ESP chip series might have different numbers of GPTimer instances. For more details, please refer to **ESP32-P4 Technical Reference Manual** > Chapter **Timer Group (TIMG)** [PDF]. The driver does forbid you from applying for more timers, but it returns error when all available hardware resources are used up. Please always check the return value when doing resource allocation (e.g., `gptimer_new_timer()`).

Set and Get Count Value When the GPTimer is created, the internal counter will be reset to zero by default. The counter value can be updated asynchronously by `gptimer_set_raw_count()`. The maximum count value is dependent on the bit width of the hardware timer, which is also reflected by the SOC macro `SOC_TIMER_GROUP_COUNTER_BIT_WIDTH`. When updating the raw count of an active timer, the timer will immediately start counting from the new value.

Count value can be retrieved by `gptimer_get_raw_count()`, at any time.

Set up Alarm Action For most of the use cases of GPTimer, you should set up the alarm action before starting the timer, except for the simple wall-clock scenario, where a free running timer is enough. To set up the alarm action, you should configure several members of `gptimer_alarm_config_t` based on how you make use of the alarm event:

- `gptimer_alarm_config_t::alarm_count` sets the target count value that triggers the alarm event. You should also take the counting direction into consideration when setting the alarm value. Specially, `gptimer_alarm_config_t::alarm_count` and `gptimer_alarm_config_t::reload_count` cannot be set to the same value when `gptimer_alarm_config_t::auto_reload_on_alarm` is true, as keeping reload with a target alarm count is meaningless. Please also note, because of the interrupt latency, it's not recommended to set the alarm period smaller than 5 us.
- `gptimer_alarm_config_t::reload_count` sets the count value to be reloaded when the alarm event happens. This configuration only takes effect when `gptimer_alarm_config_t::auto_reload_on_alarm` is set to true.
- `gptimer_alarm_config_t::auto_reload_on_alarm` flag sets whether to enable the auto-reload feature. If enabled, the hardware timer will reload the value of `gptimer_alarm_config_t::reload_count` into counter immediately when an alarm event happens.

To make the alarm configurations take effect, you should call `gptimer_set_alarm_action()`. Especially, if `gptimer_alarm_config_t` is set to NULL, the alarm function will be disabled.

Note: If an alarm value is set and the timer has already exceeded this value, the alarm will be triggered immediately.

Register Event Callbacks After the timer starts up, it can generate a specific event (e.g., the "Alarm Event") dynamically. If you have some functions that should be called when the event happens, please hook your function to the interrupt service routine by calling `gptimer_register_event_callbacks()`. All supported event callbacks are listed in `gptimer_event_callbacks_t`:

- `gptimer_event_callbacks_t::on_alarm` sets a callback function for alarm events. As this function is called within the ISR context, you must ensure that the function does not attempt to block (e.g., by making sure that only FreeRTOS APIs with ISR suffix are called from within the function). The function prototype is declared in `gptimer_alarm_cb_t`.

You can save your own context to `gptimer_register_event_callbacks()` as well, via the parameter `user_data`. The user data will be directly passed to the callback function.

This function lazy installs the interrupt service for the timer but not enable it. So please call this function before `gptimer_enable()`, otherwise the `ESP_ERR_INVALID_STATE` error will be returned. See Section [Enable and Disable Timer](#) for more information.

Enable and Disable Timer Before doing IO control to the timer, you needs to enable the timer first, by calling `gptimer_enable()`. This function:

- Switches the timer driver state from **init** to **enable**.
- Enables the interrupt service if it has been lazy installed by `gptimer_register_event_callbacks()`.
- Acquires a proper power management lock if a specific clock source (e.g., APB clock) is selected. See Section [Power Management](#) for more information.

Calling `gptimer_disable()` does the opposite, that is, put the timer driver back to the **init** state, disable the interrupts service and release the power management lock.

Start and Stop Timer The basic IO operation of a timer is to start and stop. Calling `gptimer_start()` can make the internal counter work, while calling `gptimer_stop()` can make the counter stop working. The following illustrates how to start a timer with or without an alarm event.

Calling `gptimer_start()` transits the driver state from **enable** to **run**, and vice versa. You need to make sure the start and stop functions are used in pairs, otherwise, the functions may return `ESP_ERR_INVALID_STATE`. Most of the time, this error means that the timer is already stopped or in the "start protection" state (i.e., `gptimer_start()` is called but not finished).

Start Timer as a Wall Clock

```
ESP_ERROR_CHECK(gptimer_enable(gptimer));
ESP_ERROR_CHECK(gptimer_start(gptimer));
// Retrieve the timestamp at any time
uint64_t count;
ESP_ERROR_CHECK(gptimer_get_raw_count(gptimer, &count));
```

Trigger Period Events

```
typedef struct {
    uint64_t event_count;
} example_queue_element_t;

static bool example_timer_on_alarm_cb(gptimer_handle_t timer, const gptimer_alarm_
↳event_data_t *edata, void *user_ctx)
{
    BaseType_t high_task_awoken = pdFALSE;
    QueueHandle_t queue = (QueueHandle_t)user_ctx;
    // Retrieve the count value from event data
    example_queue_element_t ele = {
        .event_count = edata->count_value
    };
    // Optional: send the event data to other task by OS queue
    // Do not introduce complex logics in callbacks
    // Suggest dealing with event data in the main loop, instead of in this_
↳callback
    xQueueSendFromISR(queue, &ele, &high_task_awoken);
    // return whether we need to yield at the end of ISR
    return high_task_awoken == pdTRUE;
}

gptimer_alarm_config_t alarm_config = {
    .reload_count = 0, // counter will reload with 0 on alarm event
    .alarm_count = 1000000, // period = 1s @resolution 1MHz
    .flags.auto_reload_on_alarm = true, // enable auto-reload
};
ESP_ERROR_CHECK(gptimer_set_alarm_action(gptimer, &alarm_config));

gptimer_event_callbacks_t cbs = {
    .on_alarm = example_timer_on_alarm_cb, // register user callback
};
ESP_ERROR_CHECK(gptimer_register_event_callbacks(gptimer, &cbs, queue));
ESP_ERROR_CHECK(gptimer_enable(gptimer));
ESP_ERROR_CHECK(gptimer_start(gptimer));
```

Trigger One-Shot Event

```
typedef struct {
    uint64_t event_count;
} example_queue_element_t;
```

(continues on next page)

(continued from previous page)

```

static bool example_timer_on_alarm_cb(gptimer_handle_t timer, const gptimer_alarm_
↪event_data_t *edata, void *user_ctx)
{
    BaseType_t high_task_awoken = pdFALSE;
    QueueHandle_t queue = (QueueHandle_t)user_ctx;
    // Stop timer the sooner the better
    gptimer_stop(timer);
    // Retrieve the count value from event data
    example_queue_element_t ele = {
        .event_count = edata->count_value
    };
    // Optional: send the event data to other task by OS queue
    xQueueSendFromISR(queue, &ele, &high_task_awoken);
    // return whether we need to yield at the end of ISR
    return high_task_awoken == pdTRUE;
}

gptimer_alarm_config_t alarm_config = {
    .alarm_count = 1 * 1000 * 1000, // alarm target = 1s @resolution 1MHz
};
ESP_ERROR_CHECK(gptimer_set_alarm_action(gptimer, &alarm_config));

gptimer_event_callbacks_t cbs = {
    .on_alarm = example_timer_on_alarm_cb, // register user callback
};
ESP_ERROR_CHECK(gptimer_register_event_callbacks(gptimer, &cbs, queue));
ESP_ERROR_CHECK(gptimer_enable(gptimer));
ESP_ERROR_CHECK(gptimer_start(gptimer));

```

Dynamic Alarm Update Alarm value can be updated dynamically inside the ISR handler callback, by changing `gptimer_alarm_event_data_t::alarm_value`. Then the alarm value will be updated after the callback function returns.

```

typedef struct {
    uint64_t event_count;
} example_queue_element_t;

static bool example_timer_on_alarm_cb(gptimer_handle_t timer, const gptimer_alarm_
↪event_data_t *edata, void *user_ctx)
{
    BaseType_t high_task_awoken = pdFALSE;
    QueueHandle_t queue = (QueueHandle_t)user_data;
    // Retrieve the count value from event data
    example_queue_element_t ele = {
        .event_count = edata->count_value
    };
    // Optional: send the event data to other task by OS queue
    xQueueSendFromISR(queue, &ele, &high_task_awoken);
    // reconfigure alarm value
    gptimer_alarm_config_t alarm_config = {
        .alarm_count = edata->alarm_value + 1000000, // alarm in next 1s
    };
    gptimer_set_alarm_action(timer, &alarm_config);
    // return whether we need to yield at the end of ISR
    return high_task_awoken == pdTRUE;
}

gptimer_alarm_config_t alarm_config = {
    .alarm_count = 1000000, // initial alarm target = 1s @resolution 1MHz

```

(continues on next page)

(continued from previous page)

```

};
ESP_ERROR_CHECK(gptimer_set_alarm_action(gptimer, &alarm_config));

gptimer_event_callbacks_t cbs = {
    .on_alarm = example_timer_on_alarm_cb, // register user callback
};
ESP_ERROR_CHECK(gptimer_register_event_callbacks(gptimer, &cbs, queue));
ESP_ERROR_CHECK(gptimer_enable(gptimer));
ESP_ERROR_CHECK(gptimer_start(gptimer));

```

ETM Event and Task GPTimer is able to generate various events that can interact with the [ETM](#) module. The supported events are listed in the [gptimer_etm_event_type_t](#). You can call [gptimer_new_etm_event\(\)](#) to get the corresponding ETM event handle. Likewise, GPTimer exposes several tasks that can be triggered by other ETM events. The supported tasks are listed in the [gptimer_etm_task_type_t](#). You can call [gptimer_new_etm_task\(\)](#) to get the corresponding ETM task handle.

For how to connect the event and task to an ETM channel, please refer to the [ETM](#) documentation.

Power Management There are some power management strategies, which might turn off or change the frequency of GPTimer's source clock to save power consumption. For example, during DFS, APB clock will be scaled down. If light-sleep is also enabled, PLL and XTAL clocks will be powered off. Both of them can result in an inaccurate time keeping.

The driver can prevent the above situation from happening by creating different power management lock according to different clock source. The driver increases the reference count of that power management lock in the [gptimer_enable\(\)](#) and decrease it in the [gptimer_disable\(\)](#). So we can ensure the clock source is stable between [gptimer_enable\(\)](#) and [gptimer_disable\(\)](#).

IRAM Safe By default, the GPTimer interrupt will be deferred when the cache is disabled because of writing or erasing the flash. Thus the alarm interrupt will not get executed in time, which is not expected in a real-time application.

There is a Kconfig option [CONFIG_GPTIMER_ISR_IRAM_SAFE](#) that:

- Enables the interrupt being serviced even when the cache is disabled
- Places all functions that used by the ISR into IRAM²
- Places driver object into DRAM (in case it is mapped to PSRAM by accident)

This allows the interrupt to run while the cache is disabled, but comes at the cost of increased IRAM consumption.

There is another Kconfig option [CONFIG_GPTIMER_CTRL_FUNC_IN_IRAM](#) that can put commonly used IO control functions into IRAM as well. So, these functions can also be executable when the cache is disabled. These IO control functions are as follows:

- [gptimer_start\(\)](#)
- [gptimer_stop\(\)](#)
- [gptimer_get_raw_count\(\)](#)
- [gptimer_set_raw_count\(\)](#)
- [gptimer_set_alarm_action\(\)](#)

Thread Safety All the APIs provided by the driver are guaranteed to be thread safe, which means you can call them from different RTOS tasks without protection by extra locks. The following functions are allowed to run under ISR context.

- [gptimer_start\(\)](#)
- [gptimer_stop\(\)](#)

² [gptimer_event_callbacks_t::on_alarm](#) callback and the functions invoked by the callback should also be placed in IRAM, please take care of them by yourself.

- `gptimer_get_raw_count()`
- `gptimer_set_raw_count()`
- `gptimer_get_captured_count()`
- `gptimer_set_alarm_action()`

Kconfig Options

- `CONFIG_GPTIMER_CTRL_FUNC_IN_IRAM` controls where to place the GPTimer control functions (IRAM or flash).
- `CONFIG_GPTIMER_ISR_HANDLER_IN_IRAM` controls where to place the GPTimer ISR handler (IRAM or flash).
- `CONFIG_GPTIMER_ISR_IRAM_SAFE` controls whether the default ISR handler should be masked when the cache is disabled, see Section *IRAM Safe* for more information.
- `CONFIG_GPTIMER_ENABLE_DEBUG_LOG` is used to enable the debug log output. Enable this option will increase the firmware binary size.

Application Examples

- Typical use cases of GPTimer are listed in the example [peripherals/timer_group/gptimer](#).
- GPTimer capture external event's timestamp, with the help of ETM module: [peripherals/timer_group/gptimer_capture_hc_sr04](#).

API Reference

Header File

- `components/driver/gptimer/include/driver/gptimer.h`
- This header file can be included with:

```
#include "driver/gptimer.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

`esp_err_t gptimer_new_timer` (const `gptimer_config_t` *config, `gptimer_handle_t` *ret_timer)

Create a new General Purpose Timer, and return the handle.

Note: The newly created timer is put in the "init" state.

Parameters

- **config** -- [in] GPTimer configuration
- **ret_timer** -- [out] Returned timer handle

Returns

- `ESP_OK`: Create GPTimer successfully
- `ESP_ERR_INVALID_ARG`: Create GPTimer failed because of invalid argument
- `ESP_ERR_NO_MEM`: Create GPTimer failed because out of memory

- `ESP_ERR_NOT_FOUND`: Create GPTimer failed because all hardware timers are used up and no more free one
- `ESP_FAIL`: Create GPTimer failed because of other error

esp_err_t **gptimer_del_timer** (*gptimer_handle_t* timer)

Delete the GPTimer handle.

Note: A timer must be in the "init" state before it can be deleted.

Parameters **timer** -- [in] Timer handle created by `gptimer_new_timer`

Returns

- `ESP_OK`: Delete GPTimer successfully
- `ESP_ERR_INVALID_ARG`: Delete GPTimer failed because of invalid argument
- `ESP_ERR_INVALID_STATE`: Delete GPTimer failed because the timer is not in init state
- `ESP_FAIL`: Delete GPTimer failed because of other error

esp_err_t **gptimer_set_raw_count** (*gptimer_handle_t* timer, `uint64_t` value)

Set GPTimer raw count value.

Note: When updating the raw count of an active timer, the timer will immediately start counting from the new value.

Note: This function is allowed to run within ISR context

Note: If `CONFIG_GPTIMER_CTRL_FUNC_IN_IRAM` is enabled, this function will be placed in the IRAM by linker, makes it possible to execute even when the Flash Cache is disabled.

Parameters

- **timer** -- [in] Timer handle created by `gptimer_new_timer`
- **value** -- [in] Count value to be set

Returns

- `ESP_OK`: Set GPTimer raw count value successfully
- `ESP_ERR_INVALID_ARG`: Set GPTimer raw count value failed because of invalid argument
- `ESP_FAIL`: Set GPTimer raw count value failed because of other error

esp_err_t **gptimer_get_raw_count** (*gptimer_handle_t* timer, `uint64_t *value`)

Get GPTimer raw count value.

Note: This function will trigger a software capture event and then return the captured count value.

Note: With the raw count value and the resolution returned from `gptimer_get_resolution`, you can convert the count value into seconds.

Note: This function is allowed to run within ISR context

Note: If `CONFIG_GPTIMER_CTRL_FUNC_IN_IRAM` is enabled, this function will be placed in the IRAM by linker, makes it possible to execute even when the Flash Cache is disabled.

Parameters

- **timer** -- [in] Timer handle created by `gptimer_new_timer`
- **value** -- [out] Returned GPTimer count value

Returns

- `ESP_OK`: Get GPTimer raw count value successfully
- `ESP_ERR_INVALID_ARG`: Get GPTimer raw count value failed because of invalid argument
- `ESP_FAIL`: Get GPTimer raw count value failed because of other error

`esp_err_t gptimer_get_resolution(gptimer_handle_t timer, uint32_t *out_resolution)`

Return the real resolution of the timer.

Note: usually the timer resolution is same as what you configured in the `gptimer_config_t::resolution_hz`, but some unstable clock source (e.g. `RC_FAST`) will do a calibration, the real resolution can be different from the configured one.

Parameters

- **timer** -- [in] Timer handle created by `gptimer_new_timer`
- **out_resolution** -- [out] Returned timer resolution, in Hz

Returns

- `ESP_OK`: Get GPTimer resolution successfully
- `ESP_ERR_INVALID_ARG`: Get GPTimer resolution failed because of invalid argument
- `ESP_FAIL`: Get GPTimer resolution failed because of other error

`esp_err_t gptimer_get_captured_count(gptimer_handle_t timer, uint64_t *value)`

Get GPTimer captured count value.

Note: The capture action can be issued either by ETM event or by software (see also `gptimer_get_raw_count`).

Note: This function is allowed to run within ISR context

Note: If `CONFIG_GPTIMER_CTRL_FUNC_IN_IRAM` is enabled, this function will be placed in the IRAM by linker, makes it possible to execute even when the Flash Cache is disabled.

Parameters

- **timer** -- [in] Timer handle created by `gptimer_new_timer`
- **value** -- [out] Returned captured count value

Returns

- `ESP_OK`: Get GPTimer captured count value successfully
- `ESP_ERR_INVALID_ARG`: Get GPTimer captured count value failed because of invalid argument
- `ESP_FAIL`: Get GPTimer captured count value failed because of other error

`esp_err_t gptimer_register_event_callbacks(gptimer_handle_t timer, const gptimer_event_callbacks_t *cbs, void *user_data)`

Set callbacks for GPTimer.

Note: User registered callbacks are expected to be runnable within ISR context

Note: The first call to this function needs to be before the call to `gptimer_enable`

Note: User can deregister a previously registered callback by calling this function and setting the callback member in the `cbs` structure to NULL.

Parameters

- **timer** -- **[in]** Timer handle created by `gptimer_new_timer`
- **cbs** -- **[in]** Group of callback functions
- **user_data** -- **[in]** User data, which will be passed to callback functions directly

Returns

- ESP_OK: Set event callbacks successfully
- ESP_ERR_INVALID_ARG: Set event callbacks failed because of invalid argument
- ESP_ERR_INVALID_STATE: Set event callbacks failed because the timer is not in init state
- ESP_FAIL: Set event callbacks failed because of other error

esp_err_t **gptimer_set_alarm_action** (*gptimer_handle_t* timer, const *gptimer_alarm_config_t* *config)

Set alarm event actions for GPTimer.

Note: This function is allowed to run within ISR context, so that user can set new alarm action immediately in the ISR callback.

Note: If `CONFIG_GPTIMER_CTRL_FUNC_IN_IRAM` is enabled, this function will be placed in the IRAM by linker, makes it possible to execute even when the Flash Cache is disabled.

Parameters

- **timer** -- **[in]** Timer handle created by `gptimer_new_timer`
- **config** -- **[in]** Alarm configuration, especially, set config to NULL means disabling the alarm function

Returns

- ESP_OK: Set alarm action for GPTimer successfully
- ESP_ERR_INVALID_ARG: Set alarm action for GPTimer failed because of invalid argument
- ESP_FAIL: Set alarm action for GPTimer failed because of other error

esp_err_t **gptimer_enable** (*gptimer_handle_t* timer)

Enable GPTimer.

Note: This function will transit the timer state from "init" to "enable".

Note: This function will enable the interrupt service, if it's lazy installed in `gptimer_register_event_callbacks`.

Note: This function will acquire a PM lock, if a specific source clock (e.g. APB) is selected in the `gptimer_config_t`, while `CONFIG_PM_ENABLE` is enabled.

Note: Enable a timer doesn't mean to start it. See also `gptimer_start` for how to make the timer start counting.

Parameters `timer` -- [in] Timer handle created by `gptimer_new_timer`

Returns

- `ESP_OK`: Enable GPTimer successfully
- `ESP_ERR_INVALID_ARG`: Enable GPTimer failed because of invalid argument
- `ESP_ERR_INVALID_STATE`: Enable GPTimer failed because the timer is already enabled
- `ESP_FAIL`: Enable GPTimer failed because of other error

`esp_err_t gptimer_disable(gptimer_handle_t timer)`

Disable GPTimer.

Note: This function will transit the timer state from "enable" to "init".

Note: This function will disable the interrupt service if it's installed.

Note: This function will release the PM lock if it's acquired in the `gptimer_enable`.

Note: Disable a timer doesn't mean to stop it. See also `gptimer_stop` for how to make the timer stop counting.

Parameters `timer` -- [in] Timer handle created by `gptimer_new_timer`

Returns

- `ESP_OK`: Disable GPTimer successfully
- `ESP_ERR_INVALID_ARG`: Disable GPTimer failed because of invalid argument
- `ESP_ERR_INVALID_STATE`: Disable GPTimer failed because the timer is not enabled yet
- `ESP_FAIL`: Disable GPTimer failed because of other error

`esp_err_t gptimer_start(gptimer_handle_t timer)`

Start GPTimer (internal counter starts counting)

Note: This function will transit the timer state from "enable" to "run".

Note: This function is allowed to run within ISR context

Note: If `CONFIG_GPTIMER_CTRL_FUNC_IN_IRAM` is enabled, this function will be placed in the IRAM by linker, makes it possible to execute even when the Flash Cache is disabled.

Parameters **timer** -- [in] Timer handle created by `gptimer_new_timer`

Returns

- ESP_OK: Start GPTimer successfully
- ESP_ERR_INVALID_ARG: Start GPTimer failed because of invalid argument
- ESP_ERR_INVALID_STATE: Start GPTimer failed because the timer is not enabled or is already in running
- ESP_FAIL: Start GPTimer failed because of other error

esp_err_t **gptimer_stop** (*gptimer_handle_t* timer)

Stop GPTimer (internal counter stops counting)

Note: This function will transit the timer state from "run" to "enable".

Note: This function is allowed to run within ISR context

Note: If `CONFIG_GPTIMER_CTRL_FUNC_IN_IRAM` is enabled, this function will be placed in the IRAM by linker, makes it possible to execute even when the Flash Cache is disabled.

Parameters **timer** -- [in] Timer handle created by `gptimer_new_timer`

Returns

- ESP_OK: Stop GPTimer successfully
- ESP_ERR_INVALID_ARG: Stop GPTimer failed because of invalid argument
- ESP_ERR_INVALID_STATE: Stop GPTimer failed because the timer is not in running.
- ESP_FAIL: Stop GPTimer failed because of other error

Structures

struct **gptimer_config_t**

General Purpose Timer configuration.

Public Members

gptimer_clock_source_t **clk_src**

GPTimer clock source

gptimer_count_direction_t **direction**

Count direction

uint32_t **resolution_hz**

Counter resolution (working frequency) in Hz, hence, the step size of each count tick equals to (1 / resolution_hz) seconds

int **intr_priority**

GPTimer interrupt priority, if set to 0, the driver will try to allocate an interrupt with a relative low priority (1,2,3)

uint32_t **intr_shared**

Set true, the timer interrupt number can be shared with other peripherals

```
struct gptimer_config_t::[anonymous] flags
```

GPTimer config flags

```
struct gptimer_event_callbacks_t
```

Group of supported GPTimer callbacks.

Note: The callbacks are all running under ISR environment

Note: When CONFIG_GPTIMER_ISR_IRAM_SAFE is enabled, the callback itself and functions called by it should be placed in IRAM.

Public Members

```
gptimer_alarm_cb_t on_alarm
```

Timer alarm callback

```
struct gptimer_alarm_config_t
```

General Purpose Timer alarm configuration.

Public Members

```
uint64_t alarm_count
```

Alarm target count value

```
uint64_t reload_count
```

Alarm reload count value, effect only when `auto_reload_on_alarm` is set to true

```
uint32_t auto_reload_on_alarm
```

Reload the count value by hardware, immediately at the alarm event

```
struct gptimer_alarm_config_t::[anonymous] flags
```

Alarm config flags

Header File

- `components/driver/gptimer/include/driver/gptimer_etm.h`
- This header file can be included with:

```
#include "driver/gptimer_etm.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```


Functions

esp_err_t **gptimer_new_etm_event** (*gptimer_handle_t* timer, const *gptimer_etm_event_config_t* *config, *esp_etm_event_handle_t* *out_event)

Get the ETM event for GPTimer.

Note: The created ETM event object can be deleted later by calling `esp_etm_del_event`

Parameters

- **timer** -- [in] Timer handle created by `gptimer_new_timer`
- **config** -- [in] GPTimer ETM event configuration
- **out_event** -- [out] Returned ETM event handle

Returns

- ESP_OK: Get ETM event successfully
- ESP_ERR_INVALID_ARG: Get ETM event failed because of invalid argument
- ESP_FAIL: Get ETM event failed because of other error

esp_err_t **gptimer_new_etm_task** (*gptimer_handle_t* timer, const *gptimer_etm_task_config_t* *config, *esp_etm_task_handle_t* *out_task)

Get the ETM task for GPTimer.

Note: The created ETM task object can be deleted later by calling `esp_etm_del_task`

Parameters

- **timer** -- [in] Timer handle created by `gptimer_new_timer`
- **config** -- [in] GPTimer ETM task configuration
- **out_task** -- [out] Returned ETM task handle

Returns

- ESP_OK: Get ETM task successfully
- ESP_ERR_INVALID_ARG: Get ETM task failed because of invalid argument
- ESP_FAIL: Get ETM task failed because of other error

Structures

struct **gptimer_etm_event_config_t**
GPTimer ETM event configuration.

Public Members

gptimer_etm_event_type_t **event_type**
GPTimer ETM event type

struct **gptimer_etm_task_config_t**
GPTimer ETM task configuration.

Public Members

gptimer_etm_task_type_t **task_type**
GPTimer ETM task type

Header File

- [components/driver/gptimer/include/driver/gptimer_types.h](#)
- This header file can be included with:

```
#include "driver/gptimer_types.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Structures

struct **gptimer_alarm_event_data_t**

GPTimer alarm event data.

Public Members

uint64_t **count_value**

Current count value

uint64_t **alarm_value**

Current alarm value

Type Definitions

typedef struct gptimer_t ***gptimer_handle_t**

Type of General Purpose Timer handle.

typedef bool (***gptimer_alarm_cb_t**)(*gptimer_handle_t* timer, const *gptimer_alarm_event_data_t* *edata, void *user_ctx)

Timer alarm callback prototype.

Param timer [in] Timer handle created by `gptimer_new_timer`

Param edata [in] Alarm event data, fed by driver

Param user_ctx [in] User data, passed from `gptimer_register_event_callbacks`

Return Whether a high priority task has been waken up by this function

Header File

- [components/hal/include/hal/timer_types.h](#)
- This header file can be included with:

```
#include "hal/timer_types.h"
```

Type Definitions

typedef *soc_periph_gptimer_clk_src_t* **gptimer_clock_source_t**

GPTimer clock source.

Note: User should select the clock source based on the power and resolution requirement

Enumerations

enum **gptimer_count_direction_t**

GPTimer count direction.

Values:

enumerator **GPTIMER_COUNT_DOWN**

Decrease count value

enumerator **GPTIMER_COUNT_UP**

Increase count value

enum **gptimer_etm_task_type_t**

GPTimer specific tasks that supported by the ETM module.

Values:

enumerator **GPTIMER_ETM_TASK_START_COUNT**

Start the counter

enumerator **GPTIMER_ETM_TASK_STOP_COUNT**

Stop the counter

enumerator **GPTIMER_ETM_TASK_EN_ALARM**

Enable the alarm

enumerator **GPTIMER_ETM_TASK_RELOAD**

Reload preset value into counter

enumerator **GPTIMER_ETM_TASK_CAPTURE**

Capture current count value into specific register

enumerator **GPTIMER_ETM_TASK_MAX**

Maximum number of tasks

enum **gptimer_etm_event_type_t**

GPTimer specific events that supported by the ETM module.

Values:

enumerator **GPTIMER_ETM_EVENT_ALARM_MATCH**

Count value matches the alarm target value

enumerator **GPTIMER_ETM_EVENT_MAX**

Maximum number of events

2.5.7 Hash-Based Message Authentication Code (HMAC)

Hash-based Message Authentication Code (HMAC) is a secure authentication technique that verifies the authenticity and integrity of a message with a pre-shared key. This module provides hardware acceleration for SHA256-HMAC generation using a key burned into an eFuse block.

For more detailed information on the application workflow and the HMAC calculation process, see **ESP32-P4 Technical Reference Manual > HMAC Accelerator (HMAC)** [PDF].

Generalized Application Scheme

Let there be two parties, A and B. They want to verify the authenticity and integrity of messages sent between each other. Before they can start sending messages, they need to exchange the secret key via a secure channel.

To verify A's messages, B can do the following:

- A calculates the HMAC of the message it wants to send.
- A sends the message and the HMAC to B.
- B calculates the HMAC of the received message itself.
- B checks whether the received and calculated HMACs match.

If they do match, the message is authentic.

However, the HMAC itself is not bound to this use case. It can also be used for challenge-response protocols supporting HMAC or as a key input for further security modules (see below), etc.

HMAC on ESP32-P4

On ESP32-P4, the HMAC module works with a secret key burnt into the eFuses. This eFuse key can be made completely inaccessible for any resources outside the cryptographic modules, thus avoiding key leakage.

Furthermore, ESP32-P4 has three different application scenarios for its HMAC module:

1. HMAC is generated for software use
2. HMAC is used as a key for the Digital Signature (DS) module
3. HMAC is used for enabling the soft-disabled JTAG interface

The first mode is called **Upstream** mode, while the last two modes are called **Downstream** modes.

eFuse Keys for HMAC Six physical eFuse blocks can be used as keys for the HMAC module: block 4 ~ block 9. The enum `hmac_key_id_t` in the API maps them to `HMAC_KEY0` ~ `HMAC_KEY5`.

Each key has a corresponding eFuse parameter **key purpose** determining for which of the three HMAC application scenarios (see below) the key may be used:

Key Purpose	Application Scenario
8	HMAC generated for software use
7	HMAC used as a key for the Digital Signature (DS) module
6	HMAC used for enabling the soft-disabled JTAG interface
5	HMAC both as a key for the DS module and for enabling JTAG

This is to prevent the usage of a key for a different function than originally intended.

To calculate an HMAC, the software has to provide the ID of the key block containing the secret key as well as the **key purpose** (see **ESP32-P4 Technical Reference Manual > eFuse Controller (eFuse)** [PDF]).

Before the HMAC key calculation, the HMAC module looks up the purpose of the provided key block. The calculation only proceeds if the purpose of the provided key block matches the purpose stored in the eFuses of the key block provided by the ID.

HMAC Generation for Software Key purpose value: 8

In this case, the HMAC is given out to the software, e.g., to authenticate a message.

The API to calculate the HMAC is `esp_hmac_calculate()`. The input arguments for the function are the message, message length, and the eFuse key block ID which contains the secret and has the efuse key purpose set to Upstream mode.

HMAC for Digital Signature Key purpose values: 7, 5

The HMAC can be used as a key derivation function to decrypt private key parameters which are used by the Digital Signature module. A standard message is used by the hardware in that case. You only need to provide the eFuse key block and purpose on the HMAC side, additional parameters are required for the Digital Signature component in that case.

Neither the key nor the actual HMAC is ever exposed outside the HMAC module and DS component. The calculation of the HMAC and its handover to the DS component happen internally.

For more details, see **ESP32-P4 Technical Reference Manual > Digital Signature (DS)** [PDF].

HMAC for Enabling JTAG Key purpose values: 6, 5

The third application is using the HMAC as a key to enable JTAG if it was soft-disabled before.

Following is the procedure to re-enable the JTAG:

Stage 1: Setup

1. Generate a 256-bit HMAC secret key to use for JTAG re-enable.
2. Write the key to an eFuse block with key purpose HMAC_DOWN_ALL (5) or HMAC_DOWN_JTAG (6). This can be done using the `esp_efuse_write_key()` function in the firmware or using `espefuse.py` from the host.
3. Configure the eFuse key block to be read-protected using the `esp_efuse_set_read_protect()`, so that software cannot read back the value.
4. Burn the `soft JTAG disable` bit/bits on ESP32-P4. This will permanently disable JTAG unless the correct key value is provided by the software.

Note: The API `esp_efuse_write_field_cnt(ESP_EFUSE_SOFT_DIS_JTAG, ESP_EFUSE_SOFT_DIS_JTAG[0]->bit_count)` can be used to burn `soft JTAG disable` bits on ESP32-P4.

Note: If `DIS_PAD_JTAG` eFuse is set, then `SOFT_DIS_JTAG` functionality does not work because JTAG is permanently disabled.

JTAG enables

1. The key to re-enable JTAG is the output of the HMAC-SHA256 function using the secret key in eFuse and 32 0x00 bytes as the message.
2. Pass this key value when calling the `esp_hmac_jtag_enable()` function from the firmware.
3. To re-disable JTAG in the firmware, reset the system or call `esp_hmac_jtag_disable()`.

For more details, see **ESP32-P4 Technical Reference Manual > HMAC Accelerator (HMAC)** [PDF].

Application Outline

The following code is an outline of how to set an eFuse key and then use it to calculate an HMAC for software usage.

We use `esp_efuse_write_key` to set physical key block 4 in the eFuse for the HMAC module together with its purpose. `ESP_EFUSE_KEY_PURPOSE_HMAC_UP` (8) means that this key can only be used for HMAC generation for software usage:

```
#include "esp_efuse.h"

const uint8_t key_data[32] = { ... };

esp_err_t status = esp_efuse_write_key(EFUSE_BLK_KEY4,
                                     ESP_EFUSE_KEY_PURPOSE_HMAC_UP,
                                     key_data, sizeof(key_data));

if (status == ESP_OK) {
    // written key
} else {
    // writing key failed, maybe written already
}
```

Now we can use the saved key to calculate an HMAC for software usage.

```
#include "esp_hmac.h"

uint8_t hmac[32];

const char *message = "Hello, HMAC!";
const size_t msg_len = 12;

esp_err_t result = esp_hmac_calculate(HMAC_KEY4, message, msg_len, hmac);

if (result == ESP_OK) {
    // HMAC written to hmac now
} else {
    // failure calculating HMAC
}
```

API Reference

Header File

- [components/esp_hw_support/include/esp_hmac.h](#)
- This header file can be included with:

```
#include "esp_hmac.h"
```

Functions

***esp_err_t* esp_hmac_calculate** (*hmac_key_id_t* key_id, const void *message, size_t message_len, uint8_t *hmac)

Calculate the HMAC of a given message.

Calculate the HMAC hmac of a given message message with length message_len. SHA256 is used for the calculation.

Note: Uses the HMAC peripheral in "upstream" mode.

Parameters

- **key_id** -- Determines which of the 6 key blocks in the efuses should be used for the HMAC calculation. The corresponding purpose field of the key block in the efuse must be set to the HMAC upstream purpose value.
- **message** -- the message for which to calculate the HMAC
- **message_len** -- message length return ESP_ERR_INVALID_STATE if unsuccessful

- **hmac** -- [out] the hmac result; the buffer behind the provided pointer must be a writeable buffer of 32 bytes

Returns

- ESP_OK, if the calculation was successful,
- ESP_ERR_INVALID_ARG if message or hmac is a nullptr or if key_id out of range
- ESP_FAIL, if the hmac calculation failed

esp_err_t **esp_hmac_jtag_enable** (*hmac_key_id_t* key_id, const uint8_t *token)

Use HMAC peripheral in Downstream mode to re-enable the JTAG, if it is not permanently disabled by HW. In downstream mode, HMAC calculations performed by peripheral are used internally and not provided back to user.

Note: Return value of the API does not indicate the JTAG status.

Parameters

- **key_id** -- Determines which of the 6 key blocks in the efuses should be used for the HMAC calculation. The corresponding purpose field of the key block in the efuse must be set to HMAC downstream purpose.
- **token** -- Pre calculated HMAC value of the 32-byte 0x00 using SHA-256 and the known private HMAC key. The key is already programmed to a eFuse key block. The key block number is provided as the first parameter to this function.

Returns

- ESP_OK, if the key_purpose of the key_id matches to HMAC downstream mode, The API returns success even if calculated HMAC does not match with the provided token. However, The JTAG will be re-enabled only if the calculated HMAC value matches with provided token, otherwise JTAG will remain disabled.
- ESP_FAIL, if the key_purpose of the key_id is not set to HMAC downstream purpose or JTAG is permanently disabled by EFUSE_HARD_DIS_JTAG eFuse parameter.
- ESP_ERR_INVALID_ARG, invalid input arguments

esp_err_t **esp_hmac_jtag_disable** (void)

Disable the JTAG which might be enabled using the HMAC downstream mode. This function just clears the result generated by calling esp_hmac_jtag_enable() API.

Returns

- ESP_OK return ESP_OK after writing the HMAC_SET_INVALIDATE_JTAG_REG with value 1.

Enumerations

enum **hmac_key_id_t**

The possible efuse keys for the HMAC peripheral

Values:

enumerator **HMAC_KEY0**

enumerator **HMAC_KEY1**

enumerator **HMAC_KEY2**

enumerator **HMAC_KEY3**

enumerator **HMAC_KEY4**

enumerator **HMAC_KEY5**

enumerator **HMAC_KEY_MAX**

2.5.8 Digital Signature (DS)

The Digital Signature (DS) module provides hardware acceleration of signing messages based on RSA. It uses pre-encrypted parameters to calculate a signature. The parameters are encrypted using HMAC as a key-derivation function. In turn, the HMAC uses eFuses as input key. The whole process happens in hardware so that neither the decryption key for the RSA parameters nor the input key for the HMAC key derivation function can be seen by the software while calculating the signature.

For more detailed information on the hardware involved in signature calculation and the registers used, see *ESP32-P4 Technical Reference Manual > Digital Signature (DS)* [PDF].

Private Key Parameters

The private key parameters for the RSA signature are stored in flash. To prevent unauthorized access, they are AES-encrypted. The HMAC module is used as a key-derivation function to calculate the AES encryption key for the private key parameters. In turn, the HMAC module uses a key from the eFuses key block which can be read-protected to prevent unauthorized access as well.

Upon signature calculation invocation, the software only specifies which eFuse key to use, the corresponding eFuse key purpose, the location of the encrypted RSA parameters and the message.

Key Generation

Both the HMAC key and the RSA private key have to be created and stored before the DS peripheral can be used. This needs to be done in software on the ESP32-P4 or alternatively on a host. For this context, the IDF provides `esp_efuse_write_block()` to set the HMAC key and `esp_hmac_calculate()` to encrypt the private RSA key parameters.

You can find instructions on how to calculate and assemble the private key parameters in *ESP32-P4 Technical Reference Manual > Digital Signature (DS)* [PDF].

Signature Calculation with IDF

For more detailed information on the workflow and the registers used, see *ESP32-P4 Technical Reference Manual > Digital Signature (DS)* [PDF].

Three parameters need to be prepared to calculate the digital signature:

1. the eFuse key block ID which is used as key for the HMAC,
2. the location of the encrypted private key parameters,
3. and the message to be signed.

Since the signature calculation takes some time, there are two possible API versions to use in IDF. The first one is `esp_ds_sign()` and simply blocks until the calculation is finished. If software needs to do something else during the calculation, `esp_ds_start_sign()` can be called, followed by periodic calls to `esp_ds_is_busy()` to check when the calculation has finished. Once the calculation has finished, `esp_ds_finish_sign()` can be called to get the resulting signature.

The APIs `esp_ds_sign()` and `esp_ds_start_sign()` calculate a plain RSA signature with help of the DS peripheral. This signature needs to be converted to appropriate format for further use. For example, MbedTLS SSL stack supports PKCS#1 format. The API `esp_ds_rsa_sign()` can be used to obtain the signature directly in the PKCS#1 v1.5 format. It internally uses `esp_ds_start_sign()` and converts the signature into PKCS#1 v1.5 format.

Note: Note that this is only the basic DS building block, the message length is fixed. To create signatures of arbitrary messages, the input is normally a hash of the actual message, padded up to the required length. An API to do this is planned in the future.

Configure the DS peripheral for a TLS connection

The DS peripheral on ESP32-P4 chip must be configured before it can be used for a TLS connection. The configuration involves the following steps -

- 1) Randomly generate a 256 bit value called the *Initialization Vector* (IV).
- 2) Randomly generate a 256 bit value called the *HMAC_KEY*.
- 3) Calculate the encrypted private key parameters from the client private key (RSA) and the parameters generated in the above steps.
- 4) Then burn the 256 bit *HMAC_KEY* on the efuse, which can only be read by the DS peripheral.

For more details, see *ESP32-P4 Technical Reference Manual > Digital Signature (DS)* [PDF].

To configure the DS peripheral for development purposes, you can use the [esp-secure-cert-tool](#).

The encrypted private key parameters obtained after the DS peripheral configuration are then to be kept in flash. Furthermore, they are to be passed to the DS peripheral which makes use of those parameters for the Digital Signature operation. The application then needs to read the ds data from the flash which has been done through the API's provided by the [esp_secure_cert_mgr](#) component. Please refer the [component/README](#) for more details.

The process of initializing the DS peripheral and then performing the Digital Signature operation is done internally with help of *ESP-TLS*. Please refer to *Digital Signature with ESP-TLS* in [ESP-TLS](#) for more details. As mentioned in the *ESP-TLS* documentation, the application only needs to provide the encrypted private key parameters to the `esp_tls` context (as `ds_data`), which internally performs all necessary operations for initializing the DS peripheral and then performing the DS operation.

Example for SSL Mutual Authentication using DS

The example [ssl_ds](#) shows how to use the DS peripheral for mutual authentication. The example uses *mqtt_client* (Implemented through *ESP-MQTT*) to connect to broker test.mosquitto.org using ssl transport with mutual authentication. The ssl part is internally performed with *ESP-TLS*. See [example README](#) for more details.

API Reference

Header File

- [components/esp_hw_support/include/esp_ds.h](#)
- This header file can be included with:

```
#include "esp_ds.h"
```

Functions

`esp_err_t esp_ds_sign` (const void *message, const [esp_ds_data_t](#) *data, [hmac_key_id_t](#) key_id, void *signature)

Sign the message with a hardware key from specific key slot. The function calculates a plain RSA signature with help of the DS peripheral. The RSA encryption operation is as follows: $Z = XY \bmod M$ where, Z is the signature, X is the input message, Y and M are the RSA private key parameters.

This function is a wrapper around `esp_ds_finish_sign()` and `esp_ds_start_sign()`, so do not use them in parallel. It blocks until the signing is finished and then returns the signature.

Note: Please see note section of `esp_ds_start_sign()` for more details about the input parameters.

Parameters

- **message** -- the message to be signed; its length should be $(data->rsa\_length + 1) * 4$ bytes, and those bytes must be in little endian format. It is your responsibility to apply your hash function and padding before calling this function, if required. (e.g. `message = padding(hash(inputMsg))`)
- **data** -- the encrypted signing key data (AES encrypted RSA key + IV)
- **key_id** -- the HMAC key ID determining the HMAC key of the HMAC which will be used to decrypt the signing key data
- **signature** -- the destination of the signature, should be $(data->rsa\_length + 1) * 4$ bytes long

Returns

- ESP_OK if successful, the signature was written to the parameter `signature`.
- ESP_ERR_INVALID_ARG if one of the parameters is NULL or `data->rsa_length` is too long or 0
- ESP_ERR_HW_CRYPTODS_HMAC_FAIL if there was an HMAC failure during retrieval of the decryption key
- ESP_ERR_NO_MEM if there hasn't been enough memory to allocate the context object
- ESP_ERR_HW_CRYPTODS_INVALID_KEY if there's a problem with passing the HMAC key to the DS component
- ESP_ERR_HW_CRYPTODS_INVALID_DIGEST if the message digest didn't match; the signature is invalid.
- ESP_ERR_HW_CRYPTODS_INVALID_PADDING if the message padding is incorrect, the signature can be read though since the message digest matches.

esp_err_t **esp_ds_start_sign** (const void *message, const *esp_ds_data_t* *data, *hmac_key_id_t* key_id, *esp_ds_context_t* **esp_ds_ctx)

Start the signing process.

This function yields a context object which needs to be passed to `esp_ds_finish_sign()` to finish the signing process. The function calculates a plain RSA signature with help of the DS peripheral. The RSA encryption operation is as follows: $Z = XY \bmod M$ where, Z is the signature, X is the input message, Y and M are the RSA private key parameters.

Note: This function locks the HMAC, SHA, AES and RSA components, so the user has to ensure to call `esp_ds_finish_sign()` in a timely manner. The numbers Y, M, Rb which are a part of `esp_ds_data_t` should be provided in little endian format and should be of length equal to the RSA private key bit length. The message length in bits should also be equal to the RSA private key bit length. No padding is applied to the message automatically, Please ensure the message is appropriately padded before calling the API.

Parameters

- **message** -- the message to be signed; its length should be $(data->rsa\_length + 1) * 4$ bytes, and those bytes must be in little endian format. It is your responsibility to apply your hash function and padding before calling this function, if required. (e.g. `message = padding(hash(inputMsg))`)
- **data** -- the encrypted signing key data (AES encrypted RSA key + IV)
- **key_id** -- the HMAC key ID determining the HMAC key of the HMAC which will be used to decrypt the signing key data
- **esp_ds_ctx** -- the context object which is needed for finishing the signing process later

Returns

- ESP_OK if successful, the ds operation was started now and has to be finished with `esp_ds_finish_sign()`
- ESP_ERR_INVALID_ARG if one of the parameters is NULL or `data->rsa_length` is too long or 0

- `ESP_ERR_HW_CRYPTODS_HMAC_FAIL` if there was an HMAC failure during retrieval of the decryption key
- `ESP_ERR_NO_MEM` if there hasn't been enough memory to allocate the context object
- `ESP_ERR_HW_CRYPTODS_INVALID_KEY` if there's a problem with passing the HMAC key to the DS component

bool **esp_ds_is_busy** (void)

Return true if the DS peripheral is busy, otherwise false.

Note: Only valid if `esp_ds_start_sign()` was called before.

esp_err_t **esp_ds_finish_sign** (void *signature, *esp_ds_context_t* *esp_ds_ctx)

Finish the signing process.

Parameters

- **signature** -- the destination of the signature, should be $(data->rsa\_length + 1) * 4$ bytes long, the resultant signature bytes shall be written in little endian format.
- **esp_ds_ctx** -- the context object retrieved by `esp_ds_start_sign()`

Returns

- `ESP_OK` if successful, the ds operation has been finished and the result is written to signature.
- `ESP_ERR_INVALID_ARG` if one of the parameters is NULL
- `ESP_ERR_HW_CRYPTODS_INVALID_DIGEST` if the message digest didn't match; the signature is invalid. This means that the encrypted RSA key parameters are invalid, indicating that they may have been tampered with or indicating a flash error, etc.
- `ESP_ERR_HW_CRYPTODS_INVALID_PADDING` if the message padding is incorrect, the signature can be read though since the message digest matches (see TRM for more details).

esp_err_t **esp_ds_encrypt_params** (*esp_ds_data_t* *data, const void *iv, const *esp_ds_p_data_t* *p_data, const void *key)

Encrypt the private key parameters.

The encryption is a prerequisite step before any signature operation can be done. It is not strictly necessary to use this encryption function, the encryption could also happen on an external device.

Note: The numbers Y, M, Rb which are a part of *esp_ds_data_t* should be provided in little endian format and should be of length equal to the RSA private key bit length. The message length in bits should also be equal to the RSA private key bit length. No padding is applied to the message automatically, Please ensure the message is appropriate padded before calling the API.

Parameters

- **data** -- Output buffer to store encrypted data, suitable for later use generating signatures.
- **iv** -- Pointer to 16 byte IV buffer, will be copied into 'data'. Should be randomly generated bytes each time.
- **p_data** -- Pointer to input plaintext key data. The expectation is this data will be deleted after this process is done and 'data' is stored.
- **key** -- Pointer to 32 bytes of key data. Type determined by `key_type` parameter. The expectation is the corresponding HMAC key will be stored to efuse and then permanently erased.

Returns

- `ESP_OK` if successful, the ds operation has been finished and the result is written to signature.
- `ESP_ERR_INVALID_ARG` if one of the parameters is NULL or `p_data->rsa_length` is too long

Structures

struct **esp_digital_signature_data**

Encrypted private key data. Recommended to store in flash in this format.

Note: This struct has to match to one from the ROM code! This documentation is mostly taken from there.

Public Members

esp_digital_signature_length_t **rsa_length**

RSA LENGTH register parameters (number of words in RSA key & operands, minus one).

This value must match the length field encrypted and stored in 'c', or invalid results will be returned. (The DS peripheral will always use the value in 'c', not this value, so an attacker can't alter the DS peripheral results this way, it will just truncate or extend the message and the resulting signature in software.)

Note: In IDF, the enum type length is the same as of type unsigned, so they can be used interchangeably. See the ROM code for the original declaration of struct `ets_ds_data_t`.

uint32_t **iv**[ESP_DS_IV_BIT_LEN / 32]

IV value used to encrypt 'c'

uint8_t **c**[ESP_DS_C_LEN]

Encrypted Digital Signature parameters. Result of AES-CBC encryption of plaintext values. Includes an encrypted message digest.

struct **esp_ds_p_data_t**

Plaintext parameters used by Digital Signature.

This is only used for encrypting the RSA parameters by calling `esp_ds_encrypt_params()`. Afterwards, the result can be stored in flash or in other persistent memory. The encryption is a prerequisite step before any signature operation can be done.

Note: Y, M, Rb, & M_Prime must all be in little endian format.

Public Members

uint32_t **Y**[ESP_DS_SIGNATURE_MAX_BIT_LEN / 32]

RSA exponent.

uint32_t **M**[ESP_DS_SIGNATURE_MAX_BIT_LEN / 32]

RSA modulus.

uint32_t **Rb**[ESP_DS_SIGNATURE_MAX_BIT_LEN / 32]

RSA r inverse operand.

uint32_t **M_prime**

RSA M prime operand.

uint32_t **length**

RSA length in words (32 bit)

Macros

ESP_DS_IV_BIT_LEN

ESP_DS_IV_LEN

ESP_DS_SIGNATURE_MAX_BIT_LEN

ESP_DS_SIGNATURE_MD_BIT_LEN

ESP_DS_SIGNATURE_M_PRIME_BIT_LEN

ESP_DS_SIGNATURE_L_BIT_LEN

ESP_DS_SIGNATURE_PADDING_BIT_LEN

ESP_DS_C_LEN

Type Definitions

typedef struct esp_ds_context **esp_ds_context_t**

typedef struct *esp_digital_signature_data* **esp_ds_data_t**

Encrypted private key data. Recommended to store in flash in this format.

Note: This struct has to match to one from the ROM code! This documentation is mostly taken from there.

Enumerations

enum **esp_digital_signature_length_t**

Values:

enumerator **ESP_DS_RSA_1024**

enumerator **ESP_DS_RSA_2048**

enumerator **ESP_DS_RSA_3072**

enumerator **ESP_DS_RSA_4096**

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.5.9 Inter-Integrated Circuit (I2C)

Introduction

I2C is a serial, synchronous, multi-device, half-duplex communication protocol that allows co-existence of multiple masters and slaves on the same bus. I2C uses two bidirectional open-drain lines: serial data line (SDA) and serial clock line (SCL), pulled up by resistors.

ESP32-P4 has 2 I2C controller (also called port), responsible for handling communication on the I2C bus. A single I2C controller can be a master or a slave.

Typically, an I2C slave device has a 7-bit address or 10-bit address. ESP32-P4 supports both I2C Standard-mode (Sm) and Fast-mode (Fm) which can go up to 100KHz and 400KHz respectively.

Warning: The clock frequency of SCL in master mode should not be larger than 400 KHz

Note: The frequency of SCL is influenced by both the pull-up resistor and the wire capacitance. Therefore, users are strongly recommended to choose appropriate pull-up resistors to make the frequency accurate. The recommended value for pull-up resistors usually ranges from 1K Ohms to 10K Ohms.

Keep in mind that the higher the frequency, the smaller the pull-up resistor should be (but not less than 1 KOhms). Indeed, large resistors will decline the current, which will increase the clock switching time and reduce the frequency. We usually recommend a range of 2 KOhms to 5 KOhms, but users may also need to make some adjustments depending on their current draw requirements.

I2C Clock Configuration

- `i2c_clock_source_t::I2C_CLK_SRC_DEFAULT`: Default I2C source clock.
- `i2c_clock_source_t::I2C_CLK_SRC_XTAL`: External crystal for I2C clock source.
- `i2c_clock_source_t::I2C_CLK_RC_FAST`: Internal 20MHz rc oscillator for I2C clock source.

I2C File Structure

Public headers that need to be included in the I2C application

- `i2c.h`: The header file of legacy I2C APIs (for apps using legacy driver).
- `i2c_master.h`: The header file that provides standard communication mode specific APIs (for apps using new driver with master mode).
- `i2c_slave.h`: The header file that provides standard communication mode specific APIs (for apps using new driver with slave mode).

Note: The legacy driver can't coexist with the new driver. Include `i2c.h` to use the legacy driver or the other two headers to use the new driver. Please keep in mind that the legacy driver is now deprecated and will be removed in future.

Public headers that have been included in the headers above

- `i2c_types_legacy.h`: The legacy public types that only used in the legacy driver.
- `i2c_types.h`: The header file that provides public types.




Fig. 3: I2C file structure

Functional Overview

The I2C driver offers following services:

- **Resource Allocation** - covers how to allocate I2C bus with properly set of configurations. It also covers how to recycle the resources when they finished working.
- **I2C Master Controller** - covers behavior of I2C master controller. Introduce data transmit, data receive, and data transmit and receive.
- **I2C Slave Controller** - covers behavior of I2C slave controller. Involve data transmit and data receive.
- **Power Management** - describes how different source clock will affect power consumption.
- **IRAM Safe** - describes tips on how to make the I2C interrupt work better along with a disabled cache.
- **Thread Safety** - lists which APIs are guaranteed to be thread safe by the driver.
- **Kconfig Options** - lists the supported Kconfig options that can bring different effects to the driver.

Resource Allocation Both I2C master bus and I2C slave bus, when supported, are represented by `i2c_bus_handle_t` in the driver. The available ports are managed in a resource pool that allocates a free port on request.

Install I2C master bus and device The I2C master is designed based on bus-device model. So `i2c_master_bus_config_t` and `i2c_device_config_t` are required separately to allocate the I2C master bus instance and I2C device instance.

I2C master bus requires the configuration that specified by `i2c_master_bus_config_t`:

- `i2c_master_bus_config_t::i2c_port` sets the I2C port used by the controller.
- `i2c_master_bus_config_t::sda_io_num` sets the GPIO number for the serial data bus (SDA).
- `i2c_master_bus_config_t::scl_io_num` sets the GPIO number for the serial clock bus (SCL).
- `i2c_master_bus_config_t::clk_source` selects the source clock for I2C bus. The available clocks are listed in `i2c_clock_source_t`. For the effect on power consumption of different clock source, please refer to **Power Management** section.




Fig. 4: I2C master bus-device module

- `i2c_master_bus_config_t::glitch_ignore_cnt` sets the glitch period of master bus, if the glitch period on the line is less than this value, it can be filtered out, typically value is 7.
- `i2c_master_bus_config_t::intr_priority` Set the priority of the interrupt. If set to 0, then the driver will use a interrupt with low or medium priority (priority level may be one of 1,2 or 3), otherwise use the priority indicated by `i2c_master_bus_config_t::intr_priority`. Please use the number form (1,2,3), not the bitmask form ((1<<1),(1<<2),(1<<3)).
- `i2c_master_bus_config_t::trans_queue_depth` Depth of internal transfer queue. Only valid in asynchronous transaction.
- `i2c_master_bus_config_t::enable_internal_pullup` Enable internal pullups. Note: This is not strong enough to pullup buses under high-speed frequency. A suitable external pullup is recommended.

If the configurations in `i2c_master_bus_config_t` is specified, users can call `i2c_new_master_bus()` to allocate and initialize an I2C master bus. This function will return an I2C bus handle if it runs correctly. Specifically, when there are no more I2C port available, this function will return `ESP_ERR_NOT_FOUND` error.

I2C master device requires the configuration that specified by `i2c_device_config_t`:

- `i2c_device_config_t::dev_addr_length` configure the address bit length of the slave device. User can choose from enumerator `I2C_ADDR_BIT_LEN_7` or `I2C_ADDR_BIT_LEN_10` (if supported).
- `i2c_device_config_t::device_address` I2C device raw address. Please parse the device address to this member directly. For example, the device address is 0x28, then parse 0x28 to `i2c_device_config_t::device_address`, don't carry a write/read bit.
- `i2c_device_config_t::scl_speed_hz` set the scl line frequency of this device.
- `i2c_device_config_t::scl_wait_us`. SCL await time (in us). Usually this value should not be very small because slave stretch will happen in pretty long time. (It's possible even stretch for 12ms). Set 0 means use default reg value.

Once the `i2c_device_config_t` structure is populated with mandatory parameters, users can call `i2c_master_bus_add_device()` to allocate an I2C device instance and mounted to the master bus then. This function will return an I2C device handle if it runs correctly. Specifically, when the I2C bus is not initialized properly, calling this function will result in a `ESP_ERR_INVALID_ARG` error.

```
#include "driver/i2c_master.h"

i2c_master_bus_config_t i2c_mst_config = {
    .clk_source = I2C_CLK_SRC_DEFAULT,
    .i2c_port = TEST_I2C_PORT,
```

(continues on next page)

(continued from previous page)

```

        .scl_io_num = I2C_MASTER_SCL_IO,
        .sda_io_num = I2C_MASTER_SDA_IO,
        .glitch_ignore_cnt = 7,
        .flags.enable_internal_pullup = true,
    };

    i2c_master_bus_handle_t bus_handle;
    ESP_ERROR_CHECK(i2c_new_master_bus(&i2c_mst_config, &bus_handle));

    i2c_device_config_t dev_cfg = {
        .dev_addr_length = I2C_ADDR_BIT_LEN_7,
        .device_address = 0x58,
        .scl_speed_hz = 100000,
    };

    i2c_master_dev_handle_t dev_handle;
    ESP_ERROR_CHECK(i2c_master_bus_add_device(bus_handle, &dev_cfg, &dev_handle));

```

Get I2C master handle via port Given that i2c master handle has been initialized in some module(e.g. an audio module), an another module(e.g. a video module) is not convenient to reused the handle. We have a helper function, [i2c_master_get_bus_handle\(\)](#) for getting the initialized handle via port. However, please make sure the handle has already been initialized ahead of time. Otherwise an error would be reported.

```

// Source File 1
#include "driver/i2c_master.h"
i2c_master_bus_handle_t bus_handle;
i2c_master_bus_config_t i2c_mst_config = {
    ... // same as others
};
ESP_ERROR_CHECK(i2c_new_master_bus(&i2c_mst_config, &bus_handle));

// Source File 2
#include "driver/i2c_master.h"
i2c_master_bus_handle_t handle;
ESP_ERROR_CHECK(i2c_master_get_bus_handle(0, &handle));

```

Uninstall I2C master bus and device If a previously installed I2C bus or device is no longer needed, it's recommended to recycle the resource by calling [i2c_master_bus_rm_device\(\)](#) or [i2c_del_master_bus\(\)](#), so that to release the underlying hardware.

Install I2C slave device I2C slave requires the configuration that specified by [i2c_slave_config_t](#):

- [i2c_slave_config_t::i2c_port](#) sets the I2C port used by the controller.
- [i2c_slave_config_t::sda_io_num](#) sets the GPIO number for serial data bus (SDA).
- [i2c_slave_config_t::scl_io_num](#) sets the GPIO number for serial clock bus (SCL).
- [i2c_slave_config_t::clk_source](#) selects the source clock for I2C bus. The available clocks are listed in [i2c_clock_source_t](#). For the effect on power consumption of different clock source, please refer to [Power Management](#) section.
- [i2c_slave_config_t::send_buf_depth](#) sets the sending buffer length.
- [i2c_slave_config_t::slave_addr](#) sets the slave address
- [i2c_master_bus_config_t::intr_priority](#) Set the priority of the interrupt. If set to 0, then the driver will use a interrupt with low or medium priority (priority level may be one of 1,2 or 3), otherwise use the priority indicated by [i2c_master_bus_config_t::intr_priority](#). Please use the number form (1,2,3), not the bitmask form ((1<<1),(1<<2),(1<<3)). Please pay attention that once the interrupt priority is set, it cannot be changed until [i2c_del_master_bus\(\)](#) is called.

- `i2c_slave_config_t::addr_bit_len` sets true if you need the slave to have a 10-bit address.
- `i2c_slave_config_t::access_ram_en` Set true to enable the non-fifo mode. Thus the I2C data fifo can be used as RAM, and double addressing will be synchronised opened.
- `i2c_slave_config_t::slave_unmatch_en` Set true to enable the slave unmatched interrupt. If master send command address cannot match the slave address, and unmatched interrupt will be triggered.


Once the `i2c_slave_config_t` structure is populated with mandatory parameters, users can call `i2c_new_slave_device()` to allocate and initialize an I2C master bus. This function will return an I2C bus handle if it runs correctly. Specifically, when there are no more I2C port available, this function will return `ESP_ERR_NOT_FOUND` error.

```
i2c_slave_config_t i2c_slv_config = {
    .addr_bit_len = I2C_ADDR_BIT_LEN_7,
    .clk_source = I2C_CLK_SRC_DEFAULT,
    .i2c_port = TEST_I2C_PORT,
    .send_buf_depth = 256,
    .scl_io_num = I2C_SLAVE_SCL_IO,
    .sda_io_num = I2C_SLAVE_SDA_IO,
    .slave_addr = 0x58,
};

i2c_slave_dev_handle_t slave_handle;
ESP_ERROR_CHECK(i2c_new_slave_device(&i2c_slv_config, &slave_handle));
```

Uninstall I2C slave device If a previously installed I2C bus is no longer needed, it's recommended to recycle the resource by calling `i2c_del_slave_device()`, so that to release the underlying hardware.

I2C Master Controller After installing the i2c master driver by `i2c_new_master_bus()`, ESP32-P4 is ready to communicate with other I2C devices. I2C APIs allow the standard transactions. Like the wave as follows:



I2C Master Write After installing I2C master bus successfully, you can simply call `i2c_master_transmit()` to write data to the slave device. The principle of this function can be explained by following chart.

In order to organize the process, the driver uses a command link, that should be populated with a sequence of commands and then passed to I2C controller for execution.




Fig. 5: I2C master write to slave

Simple example for writing data to slave:

```
#define DATA_LENGTH 100
i2c_master_bus_config_t i2c_mst_config = {
```

(continues on next page)

(continued from previous page)

```

        .clk_source = I2C_CLK_SRC_DEFAULT,
        .i2c_port = I2C_PORT_NUM_0,
        .scl_io_num = I2C_MASTER_SCL_IO,
        .sda_io_num = I2C_MASTER_SDA_IO,
        .glitch_ignore_cnt = 7,
    };
    i2c_master_bus_handle_t bus_handle;

    ESP_ERROR_CHECK(i2c_new_master_bus(&i2c_mst_config, &bus_handle));

    i2c_device_config_t dev_cfg = {
        .dev_addr_length = I2C_ADDR_BIT_LEN_7,
        .device_address = 0x58,
        .scl_speed_hz = 100000,
    };

    i2c_master_dev_handle_t dev_handle;
    ESP_ERROR_CHECK(i2c_master_bus_add_device(bus_handle, &dev_cfg, &dev_handle));

    ESP_ERROR_CHECK(i2c_master_transmit(dev_handle, data_wr, DATA_LENGTH, -1));

```

I2C Master Read After installing I2C master bus successfully, you can simply call `i2c_master_receive()` to read data from the slave device. The principle of this function can be explained by following chart.



Fig. 6: I2C master read from slave

Simple example for reading data from slave:

```

#define DATA_LENGTH 100
i2c_master_bus_config_t i2c_mst_config = {
    .clk_source = I2C_CLK_SRC_DEFAULT,
    .i2c_port = I2C_PORT_NUM_0,
    .scl_io_num = I2C_MASTER_SCL_IO,
    .sda_io_num = I2C_MASTER_SDA_IO,
    .glitch_ignore_cnt = 7,
};
i2c_master_bus_handle_t bus_handle;

ESP_ERROR_CHECK(i2c_new_master_bus(&i2c_mst_config, &bus_handle));

i2c_device_config_t dev_cfg = {
    .dev_addr_length = I2C_ADDR_BIT_LEN_7,
    .device_address = 0x58,
    .scl_speed_hz = 100000,
};

i2c_master_dev_handle_t dev_handle;
ESP_ERROR_CHECK(i2c_master_bus_add_device(bus_handle, &dev_cfg, &dev_handle));

i2c_master_receive(dev_handle, data_rd, DATA_LENGTH, -1);

```

I2C Master Write and Read Some I2C device needs write configurations before reading data from it, therefore, an interface called `i2c_master_transmit_receive()` can help. The principle of this function can be explained by following chart.




Fig. 7: I2C master write to slave and read from slave

Simple example for writing and reading from slave:

```

i2c_device_config_t dev_cfg = {
    .dev_addr_length = I2C_ADDR_BIT_LEN_7,
    .device_address = 0x58,
    .scl_speed_hz = 100000,
};

i2c_master_dev_handle_t dev_handle;
ESP_ERROR_CHECK(i2c_master_bus_add_device(I2C_PORT_NUM_0, &dev_cfg, &dev_handle));
uint8_t buf[20] = {0x20};
uint8_t buffer[2];
ESP_ERROR_CHECK(i2c_master_transmit_receive(i2c_bus_handle, buf, sizeof(buf),
    <--buffer, 2, -1));
  
```

I2C Master Probe I2C driver can use `i2c_master_probe()` to detect whether the specific device has been connected on I2C bus. If this function return `ESP_OK`, that means the device has been detected.

Important: Pull-ups must be connected to the SCL and SDA pins when this function is called. If you get `ESP_ERR_TIMEOUT` while `xfer_timeout_ms` was parsed correctly, you should check the pull-up resistors. If you do not have proper resistors nearby, setting `flags.enable_internal_pullup` as true is also acceptable.




Fig. 8: I2C master probe

Simple example for probing an I2C device:

```

i2c_master_bus_config_t i2c_mst_config_1 = {
    .clk_source = I2C_CLK_SRC_DEFAULT,
    .i2c_port = TEST_I2C_PORT,
    .scl_io_num = I2C_MASTER_SCL_IO,
    .sda_io_num = I2C_MASTER_SDA_IO,
    .glitch_ignore_cnt = 7,
    .flags.enable_internal_pullup = true,
};

i2c_master_bus_handle_t bus_handle;

ESP_ERROR_CHECK(i2c_new_master_bus(&i2c_mst_config_1, &bus_handle));
ESP_ERROR_CHECK(i2c_master_probe(bus_handle, 0x22, -1));
ESP_ERROR_CHECK(i2c_del_master_bus(bus_handle));
  
```

I2C Master Execute Customized Transactions Not all I2C devices strictly adhere to the standard I2C protocol, as different manufacturers may implement custom variations. For example, some devices require the address to be shifted, while others do not. Similarly, certain devices mandate acknowledgment (ACK) checks for specific operations, whereas others might not. To accommodate these variations, `i2c_master_execute_defined_operations()` function allow developers to define and execute fully

customized I2C transactions. This flexibility ensures seamless communication with non-standard devices by tailoring the transaction sequence, addressing, and acknowledgment behavior to the device's specific requirements.

Note: If you want to define your address in `i2c_operation_job_t`, please set `i2c_device_config_t::device_address` as `I2C_DEVICE_ADDRESS_NOT_USED` to skip internal address configuration in driver.

For address configuration of using defined transactions, given that a device address is 0x20, there are two situations, see following example:

```
i2c_device_config_t i2c_device = {
    .device_address = I2C_DEVICE_ADDRESS_NOT_USED,
    .scl_speed_hz = 100 * 1000,
    .scl_wait_us = 20000,
};

i2c_master_dev_handle_t dev_handle;

i2c_master_bus_add_device(bus_handle, &i2c_device, &dev_handle);

// Situation one: The device does not allow device address shift
uint8_t address1 = 0x20;
i2c_operation_job_t i2c_ops1[] = {
    { .command = I2C_MASTER_CMD_START },
    { .command = I2C_MASTER_CMD_WRITE, .write = { .ack_check = false, .data =
↪(uint8_t *) &address1, .total_bytes = 1 } },
    { .command = I2C_MASTER_CMD_STOP },
};

// Situation one: The device should left shift one byte with carrying a write or
↪read bit (official protocol)
uint8_t address2 = (0x20 << 1 | 0); // (0x20 << 1 | 1)
i2c_operation_job_t i2c_ops2[] = {
    { .command = I2C_MASTER_CMD_START },
    { .command = I2C_MASTER_CMD_WRITE, .write = { .ack_check = false, .data =
↪(uint8_t *) &address2, .total_bytes = 1 } },
    { .command = I2C_MASTER_CMD_STOP },
};
```

There are also some devices does not need an address, you can directly do transaction with data:

```
uint8_t data[8] = {0x11, 0x22, 0x33, 0x44, 0x55, 0x66, 0x77, 0x88};

i2c_operation_job_t i2c_ops[] = {
    { .command = I2C_MASTER_CMD_START },
    { .command = I2C_MASTER_CMD_WRITE, .write = { .ack_check = false, .data =
↪(uint8_t *)data, .total_bytes = 8 } },
    { .command = I2C_MASTER_CMD_STOP },
};

i2c_master_execute_defined_operations(dev_handle, i2c_ops, sizeof(i2c_ops) /
↪sizeof(i2c_operation_job_t), -1);
```

As for read direction, the theory is same but please always be aware the last byte of read before stop should always be nack. Example is as follows:

```
uint8_t address = (0x20 << 1 | 1);
uint8_t rcv_data[10] = {};

i2c_operation_job_t i2c_ops[] = {
    { .command = I2C_MASTER_CMD_START },
```

(continues on next page)

(continued from previous page)

```

    { .command = I2C_MASTER_CMD_WRITE, .write = { .ack_check = false, .data =
↪(uint8_t *) &address, .total_bytes = 1 } },
    { .command = I2C_MASTER_CMD_READ, .read = { .ack_value = I2C_ACK_VAL, .data =
↪(uint8_t *)rcv_data, .total_bytes = 9 } },
    { .command = I2C_MASTER_CMD_READ, .read = { .ack_value = I2C_NACK_VAL, .data =
↪(uint8_t *) (rcv_data + 9), .total_bytes = 1 } }, // This must be nack.
    { .command = I2C_MASTER_CMD_STOP },
};

i2c_master_execute_defined_operations(dev_handle, i2c_ops, sizeof(i2c_ops) /
↪sizeof(i2c_operation_job_t), -1);

```

I2C Slave Controller After installing the i2c slave driver by `i2c_new_slave_device()`, ESP32-P4 is ready to communicate with other I2C master as a slave.

I2C Slave Write The send buffer of the I2C slave is used as a FIFO to store the data to be sent. The data will queue up until the master requests them. You can call `i2c_slave_transmit()` to transfer data.

Simple example for writing data to FIFO:

```

uint8_t *data_wr = (uint8_t *) malloc(DATA_LENGTH);

i2c_slave_config_t i2c_slv_config = {
    .addr_bit_len = I2C_ADDR_BIT_LEN_7,    // 7-bit address
    .clk_source = I2C_CLK_SRC_DEFAULT,     // set the clock source
    .i2c_port = 0,                         // set I2C port number
    .send_buf_depth = 256,                 // set tx buffer length
    .scl_io_num = 2,                       // SCL gpio number
    .sda_io_num = 1,                       // SDA gpio number
    .slave_addr = 0x58,                    // slave address
};

i2c_bus_handle_t i2c_bus_handle;
ESP_ERROR_CHECK(i2c_new_slave_device(&i2c_slv_config, &i2c_bus_handle));
for (int i = 0; i < DATA_LENGTH; i++) {
    data_wr[i] = i;
}

ESP_ERROR_CHECK(i2c_slave_transmit(i2c_bus_handle, data_wr, DATA_LENGTH, 10000));

```

I2C Slave Read Whenever the master writes data to the slave, the slave will automatically store data in the receive buffer. This allows the slave application to call the function `i2c_slave_receive()` as its own discretion. As `i2c_slave_receive()` is designed as a non-blocking interface. So the user needs to register callback `i2c_slave_register_event_callbacks()` to know when the receive has finished.

```

static IRAM_ATTR bool i2c_slave_rx_done_callback(i2c_slave_dev_handle_t channel,
↪const i2c_slave_rx_done_event_data_t *edata, void *user_data)
{
    BaseType_t high_task_wakeup = pdFALSE;
    QueueHandle_t receive_queue = (QueueHandle_t)user_data;
    xQueueSendFromISR(receive_queue, edata, &high_task_wakeup);
    return high_task_wakeup == pdTRUE;
}

uint8_t *data_rd = (uint8_t *) malloc(DATA_LENGTH);
uint32_t size_rd = 0;

```

(continues on next page)

(continued from previous page)

```

i2c_slave_config_t i2c_slv_config = {
    .addr_bit_len = I2C_ADDR_BIT_LEN_7,
    .clk_source = I2C_CLK_SRC_DEFAULT,
    .i2c_port = TEST_I2C_PORT,
    .send_buf_depth = 256,
    .scl_io_num = I2C_SLAVE_SCL_IO,
    .sda_io_num = I2C_SLAVE_SDA_IO,
    .slave_addr = 0x58,
};

i2c_slave_dev_handle_t slave_handle;
ESP_ERROR_CHECK(i2c_new_slave_device(&i2c_slv_config, &slave_handle));

s_receive_queue = xQueueCreate(1, sizeof(i2c_slave_rx_done_event_data_t));
i2c_slave_event_callbacks_t cbs = {
    .on_recv_done = i2c_slave_rx_done_callback,
};
ESP_ERROR_CHECK(i2c_slave_register_event_callbacks(slave_handle, &cbs, s_receive_
↪queue));

i2c_slave_rx_done_event_data_t rx_data;
ESP_ERROR_CHECK(i2c_slave_receive(slave_handle, data_rd, DATA_LENGTH));
xQueueReceive(s_receive_queue, &rx_data, pdMS_TO_TICKS(10000));
// Receive done.

```

Put Data In I2C Slave RAM I2C slave fifo mentioned above can be used as RAM, which means user can access the RAM directly via address fields. For example, writing data to the 3rd ram block with following graph. Before using this, please note that `i2c_slave_config_t::access_ram_en` needs to be set to true.




Fig. 9: Put data in I2C slave RAM

```

uint8_t data_rd[DATA_LENGTH_RAM] = {0};

i2c_slave_config_t i2c_slv_config = {
    .addr_bit_len = I2C_ADDR_BIT_LEN_7,
    .clk_source = I2C_CLK_SRC_DEFAULT,
    .i2c_port = TEST_I2C_PORT,
    .send_buf_depth = 256,
    .scl_io_num = I2C_SLAVE_SCL_IO,
    .sda_io_num = I2C_SLAVE_SDA_IO,
    .slave_addr = 0x58,
    .flags.access_ram_en = true,
};

// Master write to slave.

i2c_slave_dev_handle_t slave_handle;
ESP_ERROR_CHECK(i2c_new_slave_device(&i2c_slv_config, &slave_handle));
ESP_ERROR_CHECK(i2c_slave_read_ram(slave_handle, 0x5, data_rd, DATA_LENGTH_RAM));
ESP_ERROR_CHECK(i2c_del_slave_device(slave_handle));

```

Get Data From I2C Slave RAM Data can be stored in the RAM with a specific offset by the slave controller, and the master can read this data directly via the RAM address. For example, if the data is stored in 3rd ram block, master can read this data by following graph. Before using this, please note that `i2c_slave_config_t::access_ram_en` needs to be set to true.




Fig. 10: Get data from I2C slave RAM

```

uint8_t data_wr[DATA_LENGTH_RAM] = {0};

i2c_slave_config_t i2c_slv_config = {
    .addr_bit_len = I2C_ADDR_BIT_LEN_7,
    .clk_source = I2C_CLK_SRC_DEFAULT,
    .i2c_port = TEST_I2C_PORT,
    .send_buf_depth = 256,
    .scl_io_num = I2C_SLAVE_SCL_IO,
    .sda_io_num = I2C_SLAVE_SDA_IO,
    .slave_addr = 0x58,
    .flags.access_ram_en = true,
};

i2c_slave_dev_handle_t slave_handle;
ESP_ERROR_CHECK(i2c_new_slave_device(&i2c_slv_config, &slave_handle));
ESP_ERROR_CHECK(i2c_slave_write_ram(slave_handle, 0x2, data_wr, DATA_LENGTH_RAM));
ESP_ERROR_CHECK(i2c_del_slave_device(slave_handle));
  
```

Register Event Callbacks

I2C master callbacks When an I2C master bus triggers an interrupt, a specific event will be generated and notify the CPU. If you have some functions that need to be called when those events occurred, you can hook your functions to the ISR (Interrupt Service Routine) by calling `i2c_master_register_event_callbacks()`. Since the registered callback functions are called in the interrupt context, user should ensure the callback function doesn't attempt to block (e.g. by making sure that only FreeRTOS APIs with `ISR` suffix are called from within the function). The callback functions are required to return a boolean value, to tell the ISR whether a high priority task is woke up by it.

I2C master event callbacks are listed in the `i2c_master_event_callbacks_t`.

Although I2C is a synchronous communication protocol, we also support asynchronous behavior by registering above callback. In this way, I2C APIs will be non-blocking interface. But note that on the same bus, only one device can adopt asynchronous operation.

Important: I2C master asynchronous transaction is still an experimental feature. (The issue is when asynchronous transaction is very large, it will cause memory problem.)

- `i2c_master_event_callbacks_t::on_recv_done` sets a callback function for master "transaction-done" event. The function prototype is declared in `i2c_master_callback_t`.

I2C slave callbacks When an I2C slave bus triggers an interrupt, a specific event will be generated and notify the CPU. If you have some function that needs to be called when those events occurred, you can hook your function to the ISR (Interrupt Service Routine) by calling `i2c_slave_register_event_callbacks()`. Since the registered callback functions are called in the interrupt context, user should ensure the callback function doesn't

attempt to block (e.g. by making sure that only FreeRTOS APIs with `ISR` suffix are called from within the function). The callback function has a boolean return value, to tell the caller whether a high priority task is woke up by it.

I2C slave event callbacks are listed in the [i2c_slave_event_callbacks_t](#).

- [i2c_slave_event_callbacks_t::on_recv_done](#) sets a callback function for "receive-done" event. The function prototype is declared in [i2c_slave_received_callback_t](#).

Power Management If the controller clock source is selected to [I2C_CLK_SRC_XTAL](#), then the driver won't install power management lock for it, which is more suitable for a low power application as long as the source clock can still provide sufficient resolution.

IRAM Safe By default, the I2C interrupt will be deferred when the Cache is disabled for reasons like writing/erasing Flash. Thus the event callback functions will not get executed in time, which is not expected in a real-time application.

There's a Kconfig option [CONFIG_I2C_ISR_IRAM_SAFE](#) that will:

1. Enable the interrupt being serviced even when cache is disabled
2. Place all functions that used by the ISR into IRAM
3. Place driver object into DRAM (in case it's mapped to PSRAM by accident)

This will allow the interrupt to run while the cache is disabled but will come at the cost of increased IRAM consumption.

Thread Safety The factory function [i2c_new_master_bus\(\)](#) and [i2c_new_slave_device\(\)](#) are guaranteed to be thread safe by the driver, which means, user can call them from different RTOS tasks without protection by extra locks. Other public I2C APIs are not thread safe. which means the user should avoid calling them from multiple tasks, if user strongly needs to call them in multiple tasks, please add extra lock.

Kconfig Options

- [CONFIG_I2C_ISR_IRAM_SAFE](#) controls whether the default ISR handler can work when cache is disabled, see also [IRAM Safe](#) for more information.
- [CONFIG_I2C_ENABLE_DEBUG_LOG](#) is used to enable the debug log at the cost of increased firmware binary size.

API Reference

Header File

- [components/driver/i2c/include/driver/i2c_master.h](#)
- This header file can be included with:

```
#include "driver/i2c_master.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t **i2c_new_master_bus** (const *i2c_master_bus_config_t* *bus_config, *i2c_master_bus_handle_t* *ret_bus_handle)

Allocate an I2C master bus.

Parameters

- **bus_config** -- [in] I2C master bus configuration.
- **ret_bus_handle** -- [out] I2C bus handle

Returns

- ESP_OK: I2C master bus initialized successfully.
- ESP_ERR_INVALID_ARG: I2C bus initialization failed because of invalid argument.
- ESP_ERR_NO_MEM: Create I2C bus failed because of out of memory.
- ESP_ERR_NOT_FOUND: No more free bus.

esp_err_t **i2c_master_bus_add_device** (*i2c_master_bus_handle_t* bus_handle, const *i2c_device_config_t* *dev_config, *i2c_master_dev_handle_t* *ret_handle)

Add I2C master BUS device.

Parameters

- **bus_handle** -- [in] I2C bus handle.
- **dev_config** -- [in] device config.
- **ret_handle** -- [out] device handle.

Returns

- ESP_OK: Create I2C master device successfully.
- ESP_ERR_INVALID_ARG: I2C bus initialization failed because of invalid argument.
- ESP_ERR_NO_MEM: Create I2C bus failed because of out of memory.

esp_err_t **i2c_del_master_bus** (*i2c_master_bus_handle_t* bus_handle)

Deinitialize the I2C master bus and delete the handle.

Parameters **bus_handle** -- [in] I2C bus handle.

Returns

- ESP_OK: Delete I2C bus success, otherwise, failed.
- Otherwise: Some module delete failed.

esp_err_t **i2c_master_bus_rm_device** (*i2c_master_dev_handle_t* handle)

I2C master bus delete device.

Parameters **handle** -- i2c device handle

Returns

- ESP_OK: If device is successfully deleted.

esp_err_t **i2c_master_transmit** (*i2c_master_dev_handle_t* i2c_dev, const uint8_t *write_buffer, size_t write_size, int xfer_timeout_ms)

Perform a write transaction on the I2C bus. The transaction will be undergoing until it finishes or it reaches the timeout provided.

Note: If a callback was registered with `i2c_master_register_event_callbacks`, the transaction will be asynchronous, and thus, this function will return directly, without blocking. You will get finish information from callback. Besides, data buffer should always be completely prepared when callback is registered, otherwise, the data will get corrupt.

Parameters

- **i2c_dev** -- [in] I2C master device handle that created by `i2c_master_bus_add_device`.
- **write_buffer** -- [in] Data bytes to send on the I2C bus.
- **write_size** -- [in] Size, in bytes, of the write buffer.
- **xfer_timeout_ms** -- [in] Wait timeout, in ms. Note: -1 means wait forever.

Returns

- ESP_OK: I2C master transmit success

- `ESP_ERR_INVALID_ARG`: I2C master transmit parameter invalid.
- `ESP_ERR_TIMEOUT`: Operation timeout(larger than `xfer_timeout_ms`) because the bus is busy or hardware crash.

esp_err_t **i2c_master_transmit_receive** (*i2c_master_dev_handle_t* i2c_dev, const uint8_t *write_buffer, size_t write_size, uint8_t *read_buffer, size_t read_size, int xfer_timeout_ms)

Perform a write-read transaction on the I2C bus. The transaction will be undergoing until it finishes or it reaches the timeout provided.

Note: If a callback was registered with `i2c_master_register_event_callbacks`, the transaction will be asynchronous, and thus, this function will return directly, without blocking. You will get finish information from callback. Besides, data buffer should always be completely prepared when callback is registered, otherwise, the data will get corrupt.

Parameters

- **i2c_dev** -- [in] I2C master device handle that created by `i2c_master_bus_add_device`.
- **write_buffer** -- [in] Data bytes to send on the I2C bus.
- **write_size** -- [in] Size, in bytes, of the write buffer.
- **read_buffer** -- [out] Data bytes received from i2c bus.
- **read_size** -- [in] Size, in bytes, of the read buffer.
- **xfer_timeout_ms** -- [in] Wait timeout, in ms. Note: -1 means wait forever.

Returns

- `ESP_OK`: I2C master transmit-receive success
- `ESP_ERR_INVALID_ARG`: I2C master transmit parameter invalid.
- `ESP_ERR_TIMEOUT`: Operation timeout(larger than `xfer_timeout_ms`) because the bus is busy or hardware crash.

esp_err_t **i2c_master_receive** (*i2c_master_dev_handle_t* i2c_dev, uint8_t *read_buffer, size_t read_size, int xfer_timeout_ms)

Perform a read transaction on the I2C bus. The transaction will be undergoing until it finishes or it reaches the timeout provided.

Note: If a callback was registered with `i2c_master_register_event_callbacks`, the transaction will be asynchronous, and thus, this function will return directly, without blocking. You will get finish information from callback. Besides, data buffer should always be completely prepared when callback is registered, otherwise, the data will get corrupt.

Parameters

- **i2c_dev** -- [in] I2C master device handle that created by `i2c_master_bus_add_device`.
- **read_buffer** -- [out] Data bytes received from i2c bus.
- **read_size** -- [in] Size, in bytes, of the read buffer.
- **xfer_timeout_ms** -- [in] Wait timeout, in ms. Note: -1 means wait forever.

Returns

- `ESP_OK`: I2C master receive success
- `ESP_ERR_INVALID_ARG`: I2C master receive parameter invalid.
- `ESP_ERR_TIMEOUT`: Operation timeout(larger than `xfer_timeout_ms`) because the bus is busy or hardware crash.

esp_err_t **i2c_master_probe** (*i2c_master_bus_handle_t* bus_handle, uint16_t address, int xfer_timeout_ms)

Probe I2C address, if address is correct and ACK is received, this function will return `ESP_OK`.

Attention Pull-ups must be connected to the SCL and SDA pins when this function is called. If you get `ESP_ERR_TIMEOUT` while `xfer_timeout_ms` was parsed correctly, you should check the pull-up resistors. If you do not have proper resistors nearby, `flags.enable_internal_pullup`` is also acceptable.

Note: The principle of this function is to send device address with a write command. If the device on your I2C bus, there would be an ACK signal and function returns `ESP_OK`. If the device is not on your I2C bus, there would be a NACK signal and function returns `ESP_ERR_NOT_FOUND`. `ESP_ERR_TIMEOUT` is not an expected failure, which indicated that the i2c probe not works properly, usually caused by pull-up resistors not be connected properly. Suggestion check data on SDA/SCL line to see whether there is ACK/NACK signal is on line when i2c probe function fails.

Note: There are lots of I2C devices all over the world, we assume that not all I2C device support the behavior like `device_address+nack/ack`. So, if the on line data is strange and no ack/nack got respond. Please check the device datasheet.

Parameters

- **bus_handle** -- [in] I2C master device handle that created by `i2c_master_bus_add_device`.
- **address** -- [in] I2C device address that you want to probe.
- **xfer_timeout_ms** -- [in] Wait timeout, in ms. Note: -1 means wait forever (Not recommended in this function).

Returns

- `ESP_OK`: I2C device probe successfully
- `ESP_ERR_NOT_FOUND`: I2C probe failed, doesn't find the device with specific address you gave.
- `ESP_ERR_TIMEOUT`: Operation timeout (larger than `xfer_timeout_ms`) because the bus is busy or hardware crash.

esp_err_t **i2c_master_execute_defined_operations** (*i2c_master_dev_handle_t* i2c_dev, *i2c_operation_job_t* *i2c_operation, size_t operation_list_num, int xfer_timeout_ms)

Execute a series of pre-defined I2C operations.

This function processes a list of I2C operations, such as start, write, read, and stop, according to the user-defined *i2c_operation_job_t* array. It performs these operations sequentially on the specified I2C master device.

Note: The `ack_value` field in the READ operation must be set to `I2C_NACK_VAL` if the next operation is a STOP command.

Parameters

- **i2c_dev** -- [in] Handle to the I2C master device.
- **i2c_operation** -- [in] Pointer to an array of user-defined I2C operation jobs. Each job specifies a command and associated parameters.
- **operation_list_num** -- [in] The number of operations in the `i2c_operation` array.
- **xfer_timeout_ms** -- [in] Timeout for the transaction, in milliseconds.

Returns

- `ESP_OK`: Transaction completed successfully.
- `ESP_ERR_INVALID_ARG`: One or more arguments are invalid.
- `ESP_ERR_TIMEOUT`: Transaction timed out.
- `ESP_FAIL`: Other error during transaction.

esp_err_t **i2c_master_register_event_callbacks** (*i2c_master_dev_handle_t* i2c_dev, const *i2c_master_event_callbacks_t* *cbs, void *user_data)

Register I2C transaction callbacks for a master device.

Note: User can deregister a previously registered callback by calling this function and setting the callback member in the `cbs` structure to NULL.

Note: When `CONFIG_I2C_ISR_IRAM_SAFE` is enabled, the callback itself and functions called by it should be placed in IRAM. The variables used in the function should be in the SRAM as well. The `user_data` should also reside in SRAM.

Note: If the callback is used for helping asynchronous transaction. On the same bus, only one device can be used for performing asynchronous operation.

Parameters

- **i2c_dev** -- [in] I2C master device handle that created by `i2c_master_bus_add_device`.
- **cbs** -- [in] Group of callback functions
- **user_data** -- [in] User data, which will be passed to callback functions directly

Returns

- `ESP_OK`: Set I2C transaction callbacks successfully
- `ESP_ERR_INVALID_ARG`: Set I2C transaction callbacks failed because of invalid argument
- `ESP_FAIL`: Set I2C transaction callbacks failed because of other error

esp_err_t **i2c_master_bus_reset** (*i2c_master_bus_handle_t* bus_handle)

Reset the I2C master bus.

Parameters **bus_handle** -- I2C bus handle.

Returns

- `ESP_OK`: Reset succeed.
- `ESP_ERR_INVALID_ARG`: I2C master bus handle is not initialized.
- Otherwise: Reset failed.

esp_err_t **i2c_master_bus_wait_all_done** (*i2c_master_bus_handle_t* bus_handle, int timeout_ms)

Wait for all pending I2C transactions done.

Parameters

- **bus_handle** -- [in] I2C bus handle
- **timeout_ms** -- [in] Wait timeout, in ms. Specially, -1 means to wait forever.

Returns

- `ESP_OK`: Flush transactions successfully
- `ESP_ERR_INVALID_ARG`: Flush transactions failed because of invalid argument
- `ESP_ERR_TIMEOUT`: Flush transactions failed because of timeout
- `ESP_FAIL`: Flush transactions failed because of other error

esp_err_t **i2c_master_get_bus_handle** (*i2c_port_num_t* port_num, *i2c_master_bus_handle_t* *ret_handle)

Retrieves the I2C master bus handle for a specified I2C port number.

This function retrieves the I2C master bus handle for the given I2C port number. Please make sure the handle has already been initialized, and this function would simply returns the existing handle. Note that the returned handle still can't be used concurrently

Parameters

- **port_num** -- I2C port number for which the handle is to be retrieved.
- **ret_handle** -- Pointer to a variable where the retrieved handle will be stored.

Returns

- ESP_OK: Success. The handle is retrieved successfully.
- ESP_ERR_INVALID_ARG: Invalid argument, such as invalid port number
- ESP_ERR_INVALID_STATE: Invalid state, such as the I2C port is not initialized.

Structures

struct **i2c_master_bus_config_t**

I2C master bus specific configurations.

Public Members

i2c_port_num_t **i2c_port**

I2C port number, -1 for auto selecting

gpio_num_t **sda_io_num**

GPIO number of I2C SDA signal, pulled-up internally

gpio_num_t **scl_io_num**

GPIO number of I2C SCL signal, pulled-up internally

i2c_clock_source_t **clk_source**

Clock source of I2C master bus, channels in the same group must use the same clock source

uint8_t **glitch_ignore_cnt**

If the glitch period on the line is less than this value, it can be filtered out, typically value is 7 (unit: I2C module clock cycle)

int **intr_priority**

I2C interrupt priority, if set to 0, driver will select the default priority (1,2,3).

size_t **trans_queue_depth**

Depth of internal transfer queue, increase this value can support more transfers pending in the background, only valid in asynchronous transaction. (Typically max_device_num * per_transaction)

uint32_t **enable_internal_pullup**

Enable internal pullups. Note: This is not strong enough to pullup buses under high-speed frequency. Recommend proper external pull-up if possible

struct *i2c_master_bus_config_t*::[anonymous] **flags**

I2C master config flags

struct **i2c_device_config_t**

I2C device configuration.

Public Members

***i2c_addr_bit_len_t* dev_addr_length**

Select the address length of the slave device.

uint16_t device_address

I2C device raw address. (The 7/10 bit address without read/write bit). Macro I2C_DEVICE_ADDRESS_NOT_USED (0xFFFF) stands for skip the address config inside driver.

uint32_t scl_speed_hz

I2C SCL line frequency.

uint32_t scl_wait_us

Timeout value. (unit: us). Please note this value should not be so small that it can handle stretch/disturbance properly. If 0 is set, that means use the default reg value

uint32_t disable_ack_check

Disable ACK check. If this is set false, that means ack check is enabled, the transaction will be stoped and API returns error when nack is detected.

struct *i2c_device_config_t*::[anonymous] flags

I2C device config flags

struct i2c_operation_job_t

Structure representing an I2C operation job.

This structure is used to define individual I2C operations (write or read) within a sequence of I2C master transactions.

Public Members***i2c_master_command_t* command**

I2C command indicating the type of operation (START, WRITE, READ, or STOP)

bool ack_check

Whether to enable ACK check during WRITE operation

uint8_t *data

Pointer to the data to be written

Pointer to the buffer for storing the data read from the bus

size_t total_bytes

Total number of bytes to write

Total number of bytes to read

struct *i2c_operation_job_t*::[anonymous]::[anonymous] write

Structure for WRITE command.

Used when the command is set to I2C_MASTER_CMD_WRITE.

***i2c_ack_value_t* ack_value**

ACK value to send after the read (ACK or NACK)

struct *i2c_operation_job_t*::[anonymous]::[anonymous] **read**

Structure for READ command.

Used when the `command` is set to `I2C_MASTER_CMD_READ`.

struct **i2c_master_event_callbacks_t**

Group of I2C master callbacks, can be used to get status during transaction or doing other small things. But take care potential concurrency issues.

Note: The callbacks are all running under ISR context

Note: When `CONFIG_I2C_ISR_IRAM_SAFE` is enabled, the callback itself and functions called by it should be placed in IRAM. The variables used in the function should be in the SRAM as well.

Public Members

***i2c_master_callback_t* on_trans_done**

I2C master transaction finish callback

Macros

I2C_DEVICE_ADDRESS_NOT_USED

Skip carry address bit in driver transmit and receive

Header File

- `components/driver/i2c/include/driver/i2c_slave.h`
- This header file can be included with:

```
#include "driver/i2c_slave.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

***esp_err_t* i2c_new_slave_device** (const *i2c_slave_config_t* *slave_config, *i2c_slave_dev_handle_t* *ret_handle)

Initialize an I2C slave device.

Parameters

- **slave_config** -- [in] I2C slave device configurations
- **ret_handle** -- [out] Return a generic I2C device handle

Returns

- ESP_OK: I2C slave device initialized successfully
- ESP_ERR_INVALID_ARG: I2C device initialization failed because of invalid argument.
- ESP_ERR_NO_MEM: Create I2C device failed because of out of memory.

esp_err_t **i2c_del_slave_device** (*i2c_slave_dev_handle_t* i2c_slave)

Deinitialize the I2C slave device.

Parameters **i2c_slave** -- **[in]** I2C slave device handle that created by `i2c_new_slave_device`.

Returns

- ESP_OK: Delete I2C device successfully.
- ESP_ERR_INVALID_ARG: I2C device initialization failed because of invalid argument.

esp_err_t **i2c_slave_receive** (*i2c_slave_dev_handle_t* i2c_slave, `uint8_t *data`, `size_t buffer_size`)

Read bytes from I2C internal buffer. Start a job to receive I2C data.

Note: This function is non-blocking, it initiates a new receive job and then returns. User should check the received data from the `on_recv_done` callback that registered by `i2c_slave_register_event_callbacks()`.

Parameters

- **i2c_slave** -- **[in]** I2C slave device handle that created by `i2c_new_slave_device`.
- **data** -- **[out]** Buffer to store data from I2C fifo. Should be valid until `on_recv_done` is triggered.
- **buffer_size** -- **[in]** Buffer size of data that provided by users.

Returns

- ESP_OK: I2C slave receive success.
- ESP_ERR_INVALID_ARG: I2C slave receive parameter invalid.
- ESP_ERR_NOT_SUPPORTED: This function should be work in fifo mode, but I2C_SLAVE_NONFIFO mode is configured

esp_err_t **i2c_slave_transmit** (*i2c_slave_dev_handle_t* i2c_slave, `const uint8_t *data`, `int size`, `int xfer_timeout_ms`)

Write bytes to internal ringbuffer of the I2C slave data. When the TX fifo empty, the ISR will fill the hardware FIFO with the internal ringbuffer's data.

Note: If you connect this slave device to some master device, the data transaction direction is from slave device to master device.

Parameters

- **i2c_slave** -- **[in]** I2C slave device handle that created by `i2c_new_slave_device`.
- **data** -- **[in]** Buffer to write to slave fifo, can pickup by master. Can be freed after this function returns. Equal or larger than `size`.
- **size** -- **[in]** In bytes, of `data` buffer.
- **xfer_timeout_ms** -- **[in]** Wait timeout, in ms. Note: -1 means wait forever.

Returns

- ESP_OK: I2C slave transmit success.
- ESP_ERR_INVALID_ARG: I2C slave transmit parameter invalid.
- ESP_ERR_TIMEOUT: Operation timeout (larger than `xfer_timeout_ms`) because the device is busy or hardware crash.
- ESP_ERR_NOT_SUPPORTED: This function should be work in fifo mode, but I2C_SLAVE_NONFIFO mode is configured

esp_err_t **i2c_slave_register_event_callbacks** (*i2c_slave_dev_handle_t* i2c_slave, const *i2c_slave_event_callbacks_t* *cbs, void *user_data)

Set I2C slave event callbacks for I2C slave channel.

Note: User can deregister a previously registered callback by calling this function and setting the callback member in the `cbs` structure to NULL.

Note: When `CONFIG_I2C_ISR_IRAM_SAFE` is enabled, the callback itself and functions called by it should be placed in IRAM. The variables used in the function should be in the SRAM as well. The `user_data` should also reside in SRAM.

Parameters

- **i2c_slave** -- [in] I2C slave device handle that created by `i2c_new_slave_device`.
- **cbs** -- [in] Group of callback functions
- **user_data** -- [in] User data, which will be passed to callback functions directly

Returns

- `ESP_OK`: Set I2C transaction callbacks successfully
- `ESP_ERR_INVALID_ARG`: Set I2C transaction callbacks failed because of invalid argument
- `ESP_FAIL`: Set I2C transaction callbacks failed because of other error

esp_err_t **i2c_slave_read_ram** (*i2c_slave_dev_handle_t* i2c_slave, uint8_t ram_address, uint8_t *data, size_t receive_size)

Read bytes from I2C internal ram. This can be only used when `access_ram_en` in configuration structure set to true.

Parameters

- **i2c_slave** -- [in] I2C slave device handle that created by `i2c_new_slave_device`.
- **ram_address** -- [in] The offset of RAM (Cannot larger than I2C RAM memory)
- **data** -- [out] Buffer to store data read from I2C ram.
- **receive_size** -- [in] Received size from RAM.

Returns

- `ESP_OK`: I2C slave transmit success.
- `ESP_ERR_INVALID_ARG`: I2C slave transmit parameter invalid.
- `ESP_ERR_NOT_SUPPORTED`: This function should be work in non-fifo mode, but `I2C_SLAVE_FIFO` mode is configured

esp_err_t **i2c_slave_write_ram** (*i2c_slave_dev_handle_t* i2c_slave, uint8_t ram_address, const uint8_t *data, size_t size)

Write bytes to I2C internal ram. This can be only used when `access_ram_en` in configuration structure set to true.

Parameters

- **i2c_slave** -- [in] I2C slave device handle that created by `i2c_new_slave_device`.
- **ram_address** -- [in] The offset of RAM (Cannot larger than I2C RAM memory)
- **data** -- [in] Buffer to fill.
- **size** -- [in] Received size from RAM.

Returns

- `ESP_OK`: I2C slave transmit success.
- `ESP_ERR_INVALID_ARG`: I2C slave transmit parameter invalid.
- `ESP_ERR_INVALID_SIZE`: Write size is larger than
- `ESP_ERR_NOT_SUPPORTED`: This function should be work in non-fifo mode, but `I2C_SLAVE_FIFO` mode is configured

Structures

struct **i2c_slave_config_t**

I2C slave specific configurations.

Public Members

i2c_port_num_t **i2c_port**

I2C port number, -1 for auto selecting

gpio_num_t **sda_io_num**

SDA IO number used by I2C bus

gpio_num_t **scl_io_num**

SCL IO number used by I2C bus

i2c_clock_source_t **clk_source**

Clock source of I2C bus.

uint32_t **send_buf_depth**

Depth of internal transfer ringbuffer, increase this value can support more transfers pending in the background

uint16_t **slave_addr**

I2C slave address

i2c_addr_bit_len_t **addr_bit_len**

I2C slave address in bit length

int **intr_priority**

I2C interrupt priority, if set to 0, driver will select the default priority (1,2,3).

uint32_t **broadcast_en**

I2C slave enable broadcast

uint32_t **access_ram_en**

Can get access to I2C RAM directly

uint32_t **slave_unmatch_en**

Can trigger unmatch interrupt when slave address does not match what master sends

struct *i2c_slave_config_t*::[anonymous] **flags**

I2C slave config flags

struct **i2c_slave_event_callbacks_t**

Group of I2C slave callbacks (e.g. get i2c slave stretch cause). But take care of potential concurrency issues.

Note: The callbacks are all running under ISR context

Note: When CONFIG_I2C_ISR_IRAM_SAFE is enabled, the callback itself and functions called by it should be placed in IRAM. The variables used in the function should be in the SRAM as well.

Public Members

i2c_slave_received_callback_t **on_recv_done**

I2C slave receive done callback

Header File

- `components/driver/i2c/include/driver/i2c_types.h`
- This header file can be included with:

```
#include "driver/i2c_types.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Structures

struct **i2c_master_event_data_t**

Data type used in I2C event callback.

Public Members

i2c_master_event_t **event**

The I2C hardware event that I2C callback is called.

struct **i2c_slave_rx_done_event_data_t**

Event structure used in I2C slave.

Public Members

uint8_t ***buffer**

Pointer for buffer received in callback.

Type Definitions

typedef int **i2c_port_num_t**

I2C port number.

typedef struct i2c_master_bus_t ***i2c_master_bus_handle_t**

Type of I2C master bus handle.

```
typedef struct i2c_master_dev_t *i2c_master_dev_handle_t
```

Type of I2C master bus device handle.

```
typedef struct i2c_slave_dev_t *i2c_slave_dev_handle_t
```

Type of I2C slave device handle.

```
typedef bool (*i2c_master_callback_t)(i2c_master_dev_handle_t i2c_dev, const i2c_master_event_data_t *evt_data, void *arg)
```

An callback for I2C transaction.

Param i2c_dev [in] Handle for I2C device.

Param evt_data [out] I2C capture event data, fed by driver

Param arg [in] User data, set in `i2c_master_register_event_callbacks()`

Return Whether a high priority task has been waken up by this function

```
typedef bool (*i2c_slave_received_callback_t)(i2c_slave_dev_handle_t i2c_slave, const i2c_slave_rx_done_event_data_t *evt_data, void *arg)
```

Callback signature for I2C slave.

Param i2c_slave [in] Handle for I2C slave.

Param evt_data [out] I2C capture event data, fed by driver

Param arg [in] User data, set in `i2c_slave_register_event_callbacks()`

Return Whether a high priority task has been waken up by this function

Enumerations

```
enum i2c_master_status_t
```

Enumeration for I2C fsm status.

Values:

enumerator **I2C_STATUS_READ**

read status for current master command

enumerator **I2C_STATUS_WRITE**

write status for current master command

enumerator **I2C_STATUS_START**

Start status for current master command

enumerator **I2C_STATUS_STOP**

stop status for current master command

enumerator **I2C_STATUS_IDLE**

idle status for current master command

enumerator **I2C_STATUS_ACK_ERROR**

ack error status for current master command

enumerator **I2C_STATUS_DONE**

I2C command done

enumerator **I2C_STATUS_TIMEOUT**

I2C bus status error, and operation timeout

enum **i2c_master_event_t**

Enumeration for I2C event.

Values:

enumerator **I2C_EVENT_ALIVE**

i2c bus in alive status.

enumerator **I2C_EVENT_DONE**

i2c bus transaction done

enumerator **I2C_EVENT_NACK**

i2c bus nack

enumerator **I2C_EVENT_TIMEOUT**

i2c bus timeout

enum **i2c_master_command_t**

Enum for I2C master commands.

These commands are used to define the I2C master operations. They correspond to hardware-level commands supported by the I2C peripheral.

Values:

enumerator **I2C_MASTER_CMD_START**

Start or Restart condition

enumerator **I2C_MASTER_CMD_WRITE**

Write operation

enumerator **I2C_MASTER_CMD_READ**

Read operation

enumerator **I2C_MASTER_CMD_STOP**

Stop condition

enum **i2c_ack_value_t**

Enum for I2C master ACK values.

These values define the acknowledgment (ACK) behavior during read operations.

Values:

enumerator **I2C_ACK_VAL**

Acknowledge (ACK) signal

enumerator **I2C_NACK_VAL**

Not Acknowledge (NACK) signal

Header File

- [components/hal/include/hal/i2c_types.h](#)
- This header file can be included with:

```
#include "hal/i2c_types.h"
```

Structures

struct **i2c_hal_clk_config_t**

Data structure for calculating I2C bus timing.

Public Members

uint16_t **clk_div**

I2C core clock divider

uint16_t **scl_low**

I2C scl low period

uint16_t **scl_high**

I2C scl high period

uint16_t **scl_wait_high**

I2C scl wait_high period

uint16_t **sda_hold**

I2C scl low period

uint16_t **sda_sample**

I2C sda sample time

uint16_t **setup**

I2C start and stop condition setup period

uint16_t **hold**

I2C start and stop condition hold period

uint16_t **tout**

I2C bus timeout period

Type Definitions

typedef [soc_periph_i2c_clk_src_t](#) **i2c_clock_source_t**

I2C group clock source.

Enumerations

enum **i2c_port_t**

I2C port number, can be I2C_NUM_0 ~ (I2C_NUM_MAX-1).

Values:

enumerator **I2C_NUM_0**

I2C port 0

enumerator **I2C_NUM_1**

I2C port 1

enumerator **I2C_NUM_MAX**

I2C port max

enum **i2c_addr_bit_len_t**

Enumeration for I2C device address bit length.

Values:

enumerator **I2C_ADDR_BIT_LEN_7**

i2c address bit length 7

enumerator **I2C_ADDR_BIT_LEN_10**

i2c address bit length 10

enum **i2c_mode_t**

Values:

enumerator **I2C_MODE_SLAVE**

I2C slave mode

enumerator **I2C_MODE_MASTER**

I2C master mode

enumerator **I2C_MODE_MAX**

enum **i2c_rw_t**

Values:

enumerator **I2C_MASTER_WRITE**

I2C write data

enumerator **I2C_MASTER_READ**

I2C read data

enum **i2c_trans_mode_t**

Values:

enumerator **I2C_DATA_MODE_MSB_FIRST**

I2C data msb first

enumerator **I2C_DATA_MODE_LSB_FIRST**

I2C data lsb first

enumerator **I2C_DATA_MODE_MAX**

enum **i2c_addr_mode_t**

Values:

enumerator **I2C_ADDR_BIT_7**

I2C 7bit address for slave mode

enumerator **I2C_ADDR_BIT_10**

I2C 10bit address for slave mode

enumerator **I2C_ADDR_BIT_MAX**

enum **i2c_ack_type_t**

Values:

enumerator **I2C_MASTER_ACK**

I2C ack for each byte read

enumerator **I2C_MASTER_NACK**

I2C nack for each byte read

enumerator **I2C_MASTER_LAST_NACK**

I2C nack for the last byte

enumerator **I2C_MASTER_ACK_MAX**

enum **i2c_slave_stretch_cause_t**

Enum for I2C slave stretch causes.

Values:

enumerator **I2C_SLAVE_STRETCH_CAUSE_ADDRESS_MATCH**

Stretching SCL low when the slave is read by the master and the address just matched

enumerator **I2C_SLAVE_STRETCH_CAUSE_TX_EMPTY**

Stretching SCL low when TX FIFO is empty in slave mode

enumerator **I2C_SLAVE_STRETCH_CAUSE_RX_FULL**

Stretching SCL low when RX FIFO is full in slave mode

enumerator **I2C_SLAVE_STRETCH_CAUSE_SENDING_ACK**

Stretching SCL low when slave sending ACK

enum **i2c_slave_read_write_status_t**

Values:

enumerator **I2C_SLAVE_WRITE_BY_MASTER**

enumerator **I2C_SLAVE_READ_BY_MASTER**

enum **i2c_bus_mode_t**

Enum for i2c working modes.

Values:

enumerator **I2C_BUS_MODE_MASTER**

I2C works under master mode

enumerator **I2C_BUS_MODE_SLAVE**

I2C works under slave mode

2.5.10 Inter-IC Sound (I2S)

Introduction

I2S (Inter-IC Sound) is a synchronous serial communication protocol usually used for transmitting audio data between two digital audio devices.

ESP32-P4 contains one I2S peripheral(s). These peripherals can be configured to input and output sample data via the I2S driver.

An I2S bus that communicates in standard or TDM mode consists of the following lines:

- **MCLK:** Master clock line. It is an optional signal depending on the slave side, mainly used for offering a reference clock to the I2S slave device.
- **BCLK:** Bit clock line. The bit clock for data line.
- **WS:** Word (Slot) select line. It is usually used to identify the vocal tract except PDM mode.
- **DIN/DOUT:** Serial data input/output line. Data will loopback internally if DIN and DOUT are set to a same GPIO.

An I2S bus that communicates in PDM mode consists of the following lines:

- **CLK:** PDM clock line.
- **DIN/DOUT:** Serial data input/output line.

Each I2S controller has the following features that can be configured by the I2S driver:

- Operation as system master or slave
- Capable of acting as transmitter or receiver
- DMA controller that allows stream sampling of data without requiring the CPU to copy each data sample

Each controller has separate RX and TX channels. That means they are able to work under different clocks and slot configurations with separate GPIO pins. Note that although the internal MCLKs of TX channel and RX channel are separate on a controller, the output MCLK signal can only be attached to one channel. If independent MCLK output is required for each channel, they must be allocated on different I2S controllers.

I2S File Structure

Public headers that need to be included in the I2S application are as follows:

- `i2s.h`: The header file that provides legacy I2S APIs (for apps using legacy driver).




Fig. 11: I2S File Structure

- `i2s_std.h`: The header file that provides standard communication mode specific APIs (for apps using new driver with standard mode).
- `i2s_pdm.h`: The header file that provides PDM communication mode specific APIs (for apps using new driver with PDM mode).
- `i2s_tdm.h`: The header file that provides TDM communication mode specific APIs (for apps using new driver with TDM mode).

Note: The legacy driver cannot coexist with the new driver. Include `i2s.h` to use the legacy driver, or include the other three headers to use the new driver. The legacy driver might be removed in future.

Public headers that have been included in the headers above are as follows:

- `i2s_types_legacy.h`: The header file that provides legacy public types that are only used in the legacy driver.
- `i2s_types.h`: The header file that provides public types.
- `i2s_common.h`: The header file that provides common APIs for all communication modes.

I2S Clock

Clock Source

- `i2s_clock_src_t::I2S_CLK_SRC_DEFAULT`: Default PLL clock.
- `i2s_clock_src_t::I2S_CLK_SRC_PLL_160M`: 160 MHz PLL clock.
- `i2s_clock_src_t::I2S_CLK_SRC_APLL`: Audio PLL clock, which is more precise than `I2S_CLK_SRC_PLL_160M` in high sample rate applications. Its frequency is configurable according to the sample rate. However, if APLL has been occupied by EMAC or other channels, the APLL frequency cannot be changed, and the driver will try to work under this APLL frequency. If this frequency cannot meet the requirements of I2S, the clock configuration will fail.

Clock Terminology

- **Sample rate**: The number of sampled data in one second per slot.

- **SCLK**: Source clock frequency. It is the frequency of the clock source.
- **MCLK**: Master clock frequency. BCLK is generated from this clock. The MCLK signal usually serves as a reference clock and is mostly needed to synchronize BCLK and WS between I2S master and slave roles.
- **BCLK**: Bit clock frequency. Every tick of this clock stands for one data bit on data pin. The slot bit width configured in `i2s_std_slot_config_t::slot_bit_width` is equal to the number of BCLK ticks, which means there will be 8/16/24/32 BCLK ticks in one slot.
- **LRCK / WS**: Left/right clock or word select clock. For non-PDM mode, its frequency is equal to the sample rate.

Note: Normally, MCLK should be the multiple of sample rate and BCLK at the same time. The field `i2s_std_clk_config_t::mclk_multiple` indicates the multiple of MCLK to the sample rate. In most cases, `I2S_MCLK_MULTIPLE_256` should be enough. However, if `slot_bit_width` is set to `I2S_SLOT_BIT_WIDTH_24BIT`, to keep MCLK a multiple to the BCLK, `i2s_std_clk_config_t::mclk_multiple` should be set to multiples that are divisible by 3 such as `I2S_MCLK_MULTIPLE_384`. Otherwise, WS will be inaccurate.


I2S Communication Mode

Overview of All Modes


Target	Standard	PDM TX	PDM RX	TDM	ADC/DAC	LCD/Camera
ESP32	I2S 0/1	I2S 0	I2S 0	none	I2S 0	I2S 0
ESP32-S2	I2S 0	none	none	none	none	I2S 0
ESP32-C3	I2S 0	I2S 0	none	I2S 0	none	none
ESP32-C6	I2S 0	I2S 0	none	I2S 0	none	none
ESP32-S3	I2S 0/1	I2S 0	I2S 0	I2S 0/1	none	none
ESP32-H2	I2S 0	I2S 0	none	I2S 0	none	none
ESP32-P4	I2S 0~2	I2S 0	I2S 0	I2S 0~2	none	none

Standard Mode In standard mode, there are always two sound channels, i.e., the left and right channels, which are called "slots". These slots support 8/16/24/32-bit width sample data. The communication format for the slots mainly includes the followings:


- **Philips Format**: Data signal has one-bit shift comparing to the WS signal, and the duty of WS signal is 50%.



- **MSB Format**: Basically the same as Philips format, but without data shift.




- **PCM Short Format:** Data has one-bit shift and meanwhile the WS signal becomes a pulse lasting for one BCLK cycle.



PDM Mode (TX) PDM (Pulse-density Modulation) mode for the TX channel can convert PCM data into PDM format which always has left and right slots. PDM TX is only supported on I2S0 and it only supports 16-bit width sample data. It needs at least a CLK pin for clock signal and a DOUT pin for data signal (i.e., the WS and SD signal in the following figure; the BCK signal is an internal bit sampling clock, which is not needed between PDM devices). This mode allows users to configure the up-sampling parameters `i2s_pdm_tx_clk_config_t::up_sample_fp` and `i2s_pdm_tx_clk_config_t::up_sample_fs`. The up-sampling rate can be calculated by $\text{up\_sample\_rate} = \text{i2s\_pdm\_tx\_clk\_config\_t::up\_sample\_fp} / \text{i2s\_pdm\_tx\_clk\_config\_t::up\_sample\_fs}$. There are two up-sampling modes in PDM TX:

- **Fixed Clock Frequency:** In this mode, the up-sampling rate changes according to the sample rate. Setting $\text{fp} = 960$ and $\text{fs} = \text{sample\_rate} / 100$, then the clock frequency (F_{pdm}) on CLK pin will be fixed to $128 * 48 \text{ KHz} = 6.144 \text{ MHz}$. Note that this frequency is not equal to the sample rate (F_{pcm}).
- **Fixed Up-sampling Rate:** In this mode, the up-sampling rate is fixed to 2. Setting $\text{fp} = 960$ and $\text{fs} = 480$, then the clock frequency (F_{pdm}) on CLK pin will be $128 * \text{sample\_rate}$.




PDM Mode (RX) PDM (Pulse-density Modulation) mode for RX channel can receive PDM-format data and convert the data into PCM format. PDM RX is only supported on I2S0, and it only supports 16-bit width sample data. PDM RX needs at least a CLK pin for clock signal and a DIN pin for data signal. This mode allows users to configure the down-sampling parameter `i2s_pdm_rx_clk_config_t::dn_sample_mode`. There are two down-sampling modes in PDM RX:

- `i2s_pdm_dsr_t::I2S_PDM_DSR_8S`: In this mode, the clock frequency (F_{pdm}) on the WS pin is $sample\_rate (F_{pcm}) * 64$.
- `i2s_pdm_dsr_t::I2S_PDM_DSR_16S`: In this mode, the clock frequency (F_{pdm}) on the WS pin is $sample\_rate (F_{pcm}) * 128$.


TDM Mode TDM (Time Division Multiplexing) mode supports up to 16 slots. These slots can be enabled by `i2s_tdm_slot_config_t::slot_mask`.

Any data bit-width is supported no matter how many slots are enabled, which means there can be up to $32 \text{ bit-width} * 16 \text{ slots} = 512 \text{ bit data}$ in one frame.


- **Philips Format**: Data signal has one-bit shift comparing to the WS signal. And no matter how many slots are contained in one frame, the duty of WS signal always keeps 50%.




- **MSB Format**: Basically the same as the Philips format, but without data shift.



- **PCM Short Format**: Data has one-bit shift and the WS signal becomes a pulse lasting one BCLK cycle for every frame.



- **PCM Long Format**: Data has one-bit shift and the WS signal lasts one-slot bit width for every frame. For example, the duty of WS will be 25% if there are four slots enabled, and 20% if there are five slots.



Functional Overview

The I2S driver offers the following services:

Resource Management There are three levels of resources in the I2S driver:

- platform level: Resources of all I2S controllers in the current target.
- controller level: Resources in one I2S controller.
- channel level: Resources of TX or RX channel in one I2S controller.

The public APIs are all channel-level APIs. The channel handle `i2s_chan_handle_t` can help users to manage the resources under a specific channel without considering the other two levels. The other two upper levels' resources are private and are managed by the driver automatically. Users can call `i2s_new_channel()` to allocate a channel handle and call `i2s_del_channel()` to delete it.

Power Management When the power management is enabled (i.e., `CONFIG_PM_ENABLE` is on), the system will adjust or stop the source clock of I2S before entering Light-sleep, thus potentially changing the I2S signals and leading to transmitting or receiving invalid data.

The I2S driver can prevent the system from changing or stopping the source clock by acquiring a power management lock. When the source clock is generated from APB, the lock type will be set to `esp_pm_lock_type_t::ESP_PM_APB_FREQ_MAX` and when the source clock is APLL (if supported), it will be set to `esp_pm_lock_type_t::ESP_PM_NO_LIGHT_SLEEP`. Whenever the user is reading or writing via I2S (i.e., calling `i2s_channel_read()` or `i2s_channel_write()`), the driver guarantees that the power management lock is acquired. Likewise, the driver releases the lock after the reading or writing finishes.

Finite State Machine There are three states for an I2S channel, namely, registered, ready, and running. Their relationship is shown in the following diagram:

The `<mode>` in the diagram can be replaced by corresponding I2S communication modes, e.g., `std` for standard two-slot mode. For more information about communication modes, please refer to the [I2S Communication Mode](#) section.

Data Transport The data transport of the I2S peripheral, including sending and receiving, is realized by DMA. Before transporting data, please call `i2s_channel_enable()` to enable the specific channel. When the sent or received data reaches the size of one DMA buffer, the `I2S_OUT_EOF` or `I2S_IN_SUC_EOF` interrupt will be triggered. Note that the DMA buffer size is not equal to `i2s_chan_config_t::dma_frame_num`. One frame here refers to all the sampled data in one WS circle. Therefore, `dma_buffer_size = dma_frame_num * slot_num * slot_bit_width / 8`. For the data transmitting, users can input the data by calling `i2s_channel_write()`. This function helps users to copy the data from the source buffer to the DMA TX buffer and wait for the transmission to finish. Then it will repeat until the sent bytes reach the given size. For the data receiving, the function `i2s_channel_read()` waits to receive the message queue which contains the DMA buffer address. It helps users copy the data from the DMA RX buffer to the destination buffer.




Fig. 12: I2S Finite State Machine

Both `i2s_channel_write()` and `i2s_channel_read()` are blocking functions. They keep waiting until the whole source buffer is sent or the whole destination buffer is loaded, unless they exceed the max blocking time, where the error code `ESP_ERR_TIMEOUT` returns. To send or receive data asynchronously, callbacks can be registered by `i2s_channel_register_event_callback()`. Users are able to access the DMA buffer directly in the callback function instead of transmitting or receiving by the two blocking functions. However, please be aware that it is an interrupt callback, so do not add complex logic, run floating operation, or call non-reentrant functions in the callback.

Configuration Users can initialize a channel by calling corresponding functions (i.e., `i2s_channel_init_std_mode()`, `i2s_channel_init_pdm_rx_mode()`, `i2s_channel_init_pdm_tx_mode()`, or `i2s_channel_init_tdm_mode()`) to a specific mode. If the configurations need to be updated after initialization, users have to first call `i2s_channel_disable()` to ensure that the channel has stopped, and then call corresponding reconfig functions, like `i2s_channel_reconfig_std_slot()`, `i2s_channel_reconfig_std_clock()`, and `i2s_channel_reconfig_std_gpio()`.

IRAM Safe By default, the I2S interrupt will be deferred when the cache is disabled for reasons like writing/erasing flash. Thus the EOF interrupt will not get executed in time.

To avoid such case in real-time applications, you can enable the Kconfig option `CONFIG_I2S_ISR_IRAM_SAFE` that:

1. Keeps the interrupt being serviced even when the cache is disabled.
2. Places driver object into DRAM (in case it is linked to PSRAM by accident).

This allows the interrupt to run while the cache is disabled, but comes at the cost of increased IRAM consumption.

Thread Safety All the public I2S APIs are guaranteed to be thread safe by the driver, which means users can call them from different RTOS tasks without protection by extra locks. Notice that the I2S driver uses mutex lock to ensure the thread safety, thus these APIs are not allowed to be used in ISR.

Kconfig Options

- `CONFIG_I2S_ISR_IRAM_SAFE` controls whether the default ISR handler can work when the cache is disabled. See *IRAM Safe* for more information.
- `CONFIG_I2S_SUPPRESS_DEPRECATED_WARN` controls whether to suppress the compiling warning message while using the legacy I2S driver.
- `CONFIG_I2S_ENABLE_DEBUG_LOG` is used to enable the debug log output. Enable this option increases the firmware binary size.

Application Example

The examples of the I2S driver can be found in the directory `peripherals/i2s`. Here are some simple usages of each mode:

Standard TX/RX Usage Different slot communication formats can be generated by the following helper macros for standard mode. As described above, there are three formats in standard mode, and their helper macros are:

- `I2S_STD_PHILIPS_SLOT_DEFAULT_CONFIG`
- `I2S_STD_PCM_SLOT_DEFAULT_CONFIG`
- `I2S_STD_MSB_SLOT_DEFAULT_CONFIG`

The clock config helper macro is:

- `I2S_STD_CLK_DEFAULT_CONFIG`

Please refer to *Standard Mode* for information about STD API. And for more details, please refer to `driver/i2s/include/driver/i2s_std.h`.

STD TX Mode Take 16-bit data width for example. When the data in a `uint16_t` writing buffer are:

data 0	data 1	data 2	data 3	data 4	data 5	data 6	data 7	...
0x0001	0x0002	0x0003	0x0004	0x0005	0x0006	0x0007	0x0008	...

Here is the table of the real data on the line with different `i2s_std_slot_config_t::slot_mode` and `i2s_std_slot_config_t::slot_mask`.

data width	bit	slot mode	slot mask	WS low	WS high	WS low	WS high	WS low	WS high	WS low	WS high
16 bit	mono	left		0x0001	0x0000	0x0002	0x0000	0x0003	0x0000	0x0004	0x0000
		right		0x0000	0x0001	0x0000	0x0002	0x0000	0x0003	0x0000	0x0004
		both		0x0001	0x0001	0x0002	0x0002	0x0003	0x0003	0x0004	0x0004
	stereo	left		0x0001	0x0000	0x0003	0x0000	0x0005	0x0000	0x0007	0x0000
		right		0x0000	0x0002	0x0000	0x0004	0x0000	0x0006	0x0000	0x0008
		both		0x0001	0x0002	0x0003	0x0004	0x0005	0x0006	0x0007	0x0008

Note: Similar for 8-bit and 32-bit data widths, the type of the buffer is better to be `uint8_t` and `uint32_t`. But specially, when the data width is 24-bit, the data buffer should be aligned with 3-byte (i.e., every 3 bytes stands for a 24-bit data in one slot). Additionally, `i2s_chan_config_t::dma_frame_num`, `i2s_std_clk_config_t::mclk_multiple`, and the writing buffer size should be the multiple of 3, otherwise the data on the line or the sample rate will be incorrect.

```
#include "driver/i2s_std.h"
#include "driver/gpio.h"

i2s_chan_handle_t tx_handle;
/* Get the default channel configuration by the helper macro.
 * This helper macro is defined in `i2s_common.h` and shared by all the I2S_
↳ communication modes.
 * It can help to specify the I2S role and port ID */
i2s_chan_config_t chan_cfg = I2S_CHANNEL_DEFAULT_CONFIG(I2S_NUM_AUTO, I2S_ROLE_
↳ MASTER);
/* Allocate a new TX channel and get the handle of this channel */
i2s_new_channel(&chan_cfg, &tx_handle, NULL);

/* Setting the configurations, the slot configuration and clock configuration can_
↳ be generated by the macros
 * These two helper macros are defined in `i2s_std.h` which can only be used in_
↳ STD mode.
 * They can help to specify the slot and clock configurations for initialization_
↳ or updating */
i2s_std_config_t std_cfg = {
    .clk_cfg = I2S_STD_CLK_DEFAULT_CONFIG(48000),
    .slot_cfg = I2S_STD_MSB_SLOT_DEFAULT_CONFIG(I2S_DATA_BIT_WIDTH_32BIT, I2S_SLOT_
↳ MODE_STEREO),
    .gpio_cfg = {
        .mclk = I2S_GPIO_UNUSED,
        .bclk = GPIO_NUM_4,
        .ws = GPIO_NUM_5,
        .dout = GPIO_NUM_18,
        .din = I2S_GPIO_UNUSED,
        .invert_flags = {
            .mclk_inv = false,
            .bclk_inv = false,
            .ws_inv = false,
        },
    },
};
```

(continues on next page)

(continued from previous page)

```

    },
};
/* Initialize the channel */
i2s_channel_init_std_mode(tx_handle, &std_cfg);

/* Before writing data, start the TX channel first */
i2s_channel_enable(tx_handle);
i2s_channel_write(tx_handle, src_buf, bytes_to_write, bytes_written, ticks_to_
↳wait);

/* If the configurations of slot or clock need to be updated,
 * stop the channel first and then update it */
// i2s_channel_disable(tx_handle);
// std_cfg.slot_cfg.slot_mode = I2S_SLOT_MODE_MONO; // Default is stereo
// i2s_channel_reconfig_std_slot(tx_handle, &std_cfg.slot_cfg);
// std_cfg.clk_cfg.sample_rate_hz = 96000;
// i2s_channel_reconfig_std_clock(tx_handle, &std_cfg.clk_cfg);

/* Have to stop the channel before deleting it */
i2s_channel_disable(tx_handle);
/* If the handle is not needed any more, delete it to release the channel_
↳resources */
i2s_del_channel(tx_handle);

```

STD RX Mode Taking 16-bit data width for example, when the data on the line are:

WS low	WS high	WS low	WS high	WS low	WS high	WS low	WS high	...
0x0001	0x0002	0x0003	0x0004	0x0005	0x0006	0x0007	0x0008	...

Here is the table of the data received in the buffer with different `i2s_std_slot_config_t::slot_mode` and `i2s_std_slot_config_t::slot_mask`.

data width	bit	slot mode	slot mask	data 0	data 1	data 2	data 3	data 4	data 5	data 6	data 7
16 bit	mono	left		0x0001	0x0003	0x0005	0x0007	0x0009	0x000b	0x000d	0x000f
		right		0x0002	0x0004	0x0006	0x0008	0x000a	0x000c	0x000e	0x0010
	stereo	any		0x0001	0x0002	0x0003	0x0004	0x0005	0x0006	0x0007	0x0008

Note: 8-bit, 24-bit, and 32-bit are similar as 16-bit, the data bit-width in the receiving buffer is equal to the data bit-width on the line. Additionally, when using 24-bit data width, `i2s_chan_config_t::dma_frame_num`, `i2s_std_clk_config_t::mclk_multiple`, and the receiving buffer size should be the multiple of 3, otherwise the data on the line or the sample rate will be incorrect.

```

#include "driver/i2s_std.h"
#include "driver/gpio.h"

i2s_chan_handle_t rx_handle;
/* Get the default channel configuration by helper macro.
 * This helper macro is defined in `i2s_common.h` and shared by all the I2S_
↳communication modes.
 * It can help to specify the I2S role and port ID */
i2s_chan_config_t chan_cfg = I2S_CHANNEL_DEFAULT_CONFIG(I2S_NUM_AUTO, I2S_ROLE_
↳MASTER);
/* Allocate a new RX channel and get the handle of this channel */
i2s_new_channel(&chan_cfg, NULL, &rx_handle);

```

(continues on next page)

(continued from previous page)

```

/* Setting the configurations, the slot configuration and clock configuration can
↳ be generated by the macros
 * These two helper macros are defined in `i2s_std.h` which can only be used in
↳ STD mode.
 * They can help to specify the slot and clock configurations for initialization
↳ or updating */
i2s_std_config_t std_cfg = {
    .clk_cfg = I2S_STD_CLK_DEFAULT_CONFIG(48000),
    .slot_cfg = I2S_STD_MSB_SLOT_DEFAULT_CONFIG(I2S_DATA_BIT_WIDTH_32BIT, I2S_SLOT_
↳ MODE_STEREO),
    .gpio_cfg = {
        .mclk = I2S_GPIO_UNUSED,
        .bclk = GPIO_NUM_4,
        .ws = GPIO_NUM_5,
        .dout = I2S_GPIO_UNUSED,
        .din = GPIO_NUM_19,
        .invert_flags = {
            .mclk_inv = false,
            .bclk_inv = false,
            .ws_inv = false,
        },
    },
};
/* Initialize the channel */
i2s_channel_init_std_mode(rx_handle, &std_cfg);

/* Before reading data, start the RX channel first */
i2s_channel_enable(rx_handle);
i2s_channel_read(rx_handle, desc_buf, bytes_to_read, bytes_read, ticks_to_wait);

/* Have to stop the channel before deleting it */
i2s_channel_disable(rx_handle);
/* If the handle is not needed any more, delete it to release the channel
↳ resources */
i2s_del_channel(rx_handle);

```

PDM TX Usage For PDM mode in TX channel, the slot configuration helper macro is:

- `I2S_PDM_TX_SLOT_DEFAULT_CONFIG`

The clock configuration helper macro is:

- `I2S_PDM_TX_CLK_DEFAULT_CONFIG`

Please refer to [PDM Mode](#) for information about PDM TX API. And for more details, please refer to [driver/i2s/include/driver/i2s_pdm.h](#).

The PDM data width is fixed to 16-bit. When the data in an `int16_t` writing buffer is:

data 0	data 1	data 2	data 3	data 4	data 5	data 6	data 7	...
0x0001	0x0002	0x0003	0x0004	0x0005	0x0006	0x0007	0x0008	...

Here is the table of the real data on the line with different `i2s_pdm_tx_slot_config_t::slot_mode` and `i2s_pdm_tx_slot_config_t::line_mode` (The PDM format on the line is transferred to PCM format for easier comprehension).

line mode	slot mode	line	left	right	left	right	left	right	left	right
one-line Codec	mono	dout	0x0001	0x0000	0x0002	0x0000	0x0003	0x0000	0x0004	0x0000
	stereo	dout	0x0001	0x0002	0x0003	0x0004	0x0005	0x0006	0x0007	0x0008
one-line DAC	mono	dout	0x0001	0x0001	0x0002	0x0002	0x0003	0x0003	0x0004	0x0004
two-line DAC	mono	dout	0x0002	0x0002	0x0004	0x0004	0x0006	0x0006	0x0008	0x0008
		dout2	0x0000	0x0000	0x0000	0x0000	0x0000	0x0000	0x0000	0x0000
	stereo	dout	0x0002	0x0002	0x0004	0x0004	0x0006	0x0006	0x0008	0x0008
		dout2	0x0001	0x0001	0x0003	0x0003	0x0005	0x0005	0x0007	0x0007

Note: There are three line modes for PDM TX mode, i.e., `I2S_PDM_TX_ONE_LINE_CODEC`, `I2S_PDM_TX_ONE_LINE_DAC`, and `I2S_PDM_TX_TWO_LINE_DAC`. One-line codec is for the PDM codecs that require clock signal. The PDM codec can differentiate the left and right slots by the clock level. The other two modes are used to drive power amplifiers directly with a low-pass filter. They do not need the clock signal, so there are two lines to differentiate the left and right slots. Additionally, for the mono mode of one-line codec, users can force change the slot to the right by setting the clock invert flag in GPIO configuration.

```
#include "driver/i2s_pdm.h"
#include "driver/gpio.h"

/* Allocate an I2S TX channel */
i2s_chan_config_t chan_cfg = I2S_CHANNEL_DEFAULT_CONFIG(I2S_NUM_0, I2S_ROLE_
↳MASTER);
i2s_new_channel(&chan_cfg, &tx_handle, NULL);

/* Init the channel into PDM TX mode */
i2s_pdm_tx_config_t pdm_tx_cfg = {
    .clk_cfg = I2S_PDM_TX_CLK_DEFAULT_CONFIG(36000),
    .slot_cfg = I2S_PDM_TX_SLOT_DEFAULT_CONFIG(I2S_DATA_BIT_WIDTH_16BIT, I2S_SLOT_
↳MODE_MONO),
    .gpio_cfg = {
        .clk = GPIO_NUM_5,
        .dout = GPIO_NUM_18,
        .invert_flags = {
            .clk_inv = false,
        },
    },
};
i2s_channel_init_pdm_tx_mode(tx_handle, &pdm_tx_cfg);

...
```

PDM RX Usage For PDM mode in RX channel, the slot configuration helper macro is:

- `I2S_PDM_RX_SLOT_DEFAULT_CONFIG`

The clock configuration helper macro is:

- `I2S_PDM_RX_CLK_DEFAULT_CONFIG`

Please refer to [PDM Mode](#) for information about PDM RX API. And for more details, please refer to [driver/i2s/include/driver/i2s_pdm.h](#).

The PDM data width is fixed to 16-bit. When the data on the line (The PDM format on the line is transferred to PCM format for easier comprehension) is:

left	right	left	right	left	right	left	right	...
0x0001	0x0002	0x0003	0x0004	0x0005	0x0006	0x0007	0x0008	...

Here is the table of the data received in a `int16_t` buffer with different `i2s_pdm_rx_slot_config_t::slot_mode` and `i2s_pdm_rx_slot_config_t::slot_mask`.

```
#include "driver/i2s_pdm.h"
#include "driver/gpio.h"

i2s_chan_handle_t rx_handle;

/* Allocate an I2S RX channel */
i2s_chan_config_t chan_cfg = I2S_CHANNEL_DEFAULT_CONFIG(I2S_NUM_0, I2S_ROLE_
↪MASTER);
i2s_new_channel(&chan_cfg, NULL, &rx_handle);

/* Init the channel into PDM RX mode */
i2s_pdm_rx_config_t pdm_rx_cfg = {
    .clk_cfg = I2S_PDM_RX_CLK_DEFAULT_CONFIG(36000),
    .slot_cfg = I2S_PDM_RX_SLOT_DEFAULT_CONFIG(I2S_DATA_BIT_WIDTH_16BIT, I2S_SLOT_
↪MODE_MONO),
    .gpio_cfg = {
        .clk = GPIO_NUM_5,
        .din = GPIO_NUM_19,
        .invert_flags = {
            .clk_inv = false,
        },
    },
};
i2s_channel_init_pdm_rx_mode(rx_handle, &pdm_rx_cfg);

...
```

TDM TX/RX Usage Different slot communication formats can be generated by the following helper macros for TDM mode. As described above, there are four formats in TDM mode, and their helper macros are:

- `I2S_TDM_PHILIPS_SLOT_DEFAULT_CONFIG`
- `I2S_TDM_MSB_SLOT_DEFAULT_CONFIG`
- `I2S_TDM_PCM_SHORT_SLOT_DEFAULT_CONFIG`
- `I2S_TDM_PCM_LONG_SLOT_DEFAULT_CONFIG`

The clock config helper macro is:

- `I2S_TDM_CLK_DEFAULT_CONFIG`

Please refer to [TDM Mode](#) for information about TDM API. And for more details, please refer to [driver/i2s/include/driver/i2s_tdm.h](#).

Note: Due to hardware limitation, when setting the clock configuration for a slave role, please be aware that `i2s_tdm_clk_config_t::bclk_div` should not be smaller than 8. Increasing this field can reduce the lagging of the data sent from the slave. In the high sample rate case, the data might lag behind for more than one BCLK which leads to data malposition. Users may gradually increase `i2s_tdm_clk_config_t::bclk_div` to correct it.

As `i2s_tdm_clk_config_t::bclk_div` is the division of MCLK to BCLK, increasing it also increases the MCLK frequency. Therefore, the clock calculation may fail if MCLK is too high to divide from the source clock. This means that a larger value for `i2s_tdm_clk_config_t::bclk_div` is not necessarily better.

TDM TX Mode

```
#include "driver/i2s_tdm.h"
#include "driver/gpio.h"
```

(continues on next page)

(continued from previous page)

```

/* Allocate an I2S TX channel */
i2s_chan_config_t chan_cfg = I2S_CHANNEL_DEFAULT_CONFIG(I2S_NUM_AUTO, I2S_ROLE_
↪MASTER);
i2s_new_channel(&chan_cfg, &tx_handle, NULL);

/* Init the channel into TDM mode */
i2s_tdm_config_t tdm_cfg = {
    .clk_cfg = I2S_TDM_CLK_DEFAULT_CONFIG(44100),
    .slot_cfg = I2S_TDM_MSB_SLOT_DEFAULT_CONFIG(I2S_DATA_BIT_WIDTH_16BIT, I2S_SLOT_
↪MODE_STEREO,
        I2S_TDM_SLOT0 | I2S_TDM_SLOT1 | I2S_TDM_SLOT2 | I2S_TDM_SLOT3),
    .gpio_cfg = {
        .mclk = I2S_GPIO_UNUSED,
        .bclk = GPIO_NUM_4,
        .ws = GPIO_NUM_5,
        .dout = GPIO_NUM_18,
        .din = I2S_GPIO_UNUSED,
        .invert_flags = {
            .mclk_inv = false,
            .bclk_inv = false,
            .ws_inv = false,
        },
    },
};
i2s_channel_init_tdm_mode(tx_handle, &tdm_cfg);

...

```

TDM RX Mode

```

#include "driver/i2s_tdm.h"
#include "driver/gpio.h"

/* Set the channel mode to TDM */
i2s_chan_config_t chan_cfg = I2S_CHANNEL_CONFIG(I2S_ROLE_MASTER, I2S_COMM_MODE_TDM,
↪ &i2s_pin);
i2s_new_channel(&chan_cfg, NULL, &rx_handle);

/* Init the channel into TDM mode */
i2s_tdm_config_t tdm_cfg = {
    .clk_cfg = I2S_TDM_CLK_DEFAULT_CONFIG(44100),
    .slot_cfg = I2S_TDM_MSB_SLOT_DEFAULT_CONFIG(I2S_DATA_BIT_WIDTH_16BIT, I2S_SLOT_
↪MODE_STEREO,
        I2S_TDM_SLOT0 | I2S_TDM_SLOT1 | I2S_TDM_SLOT2 | I2S_TDM_SLOT3),
    .gpio_cfg = {
        .mclk = I2S_GPIO_UNUSED,
        .bclk = GPIO_NUM_4,
        .ws = GPIO_NUM_5,
        .dout = I2S_GPIO_UNUSED,
        .din = GPIO_NUM_18,
        .invert_flags = {
            .mclk_inv = false,
            .bclk_inv = false,
            .ws_inv = false,
        },
    },
};
i2s_channel_init_tdm_mode(rx_handle, &tdm_cfg);

...

```

Full-duplex Full-duplex mode registers TX and RX channel in an I2S port at the same time, and the channels share the BCLK and WS signals. Currently, STD and TDM communication modes supports full-duplex mode in the following way, but PDM full-duplex is not supported because due to different PDM TX and RX clocks.

Note that one handle can only stand for one channel. Therefore, it is still necessary to configure the slot and clock for both TX and RX channels one by one.

There are two methods to allocate a pair of full-duplex channels:

1. Allocate both TX and RX handles in a single call of `i2s_new_channel()`.

```
#include "driver/i2s_std.h"
#include "driver/gpio.h"

i2s_chan_handle_t tx_handle;
i2s_chan_handle_t rx_handle;

/* Allocate a pair of I2S channel */
i2s_chan_config_t chan_cfg = I2S_CHANNEL_DEFAULT_CONFIG(I2S_NUM_AUTO, I2S_ROLE_
↪MASTER);
/* Allocate for TX and RX channel at the same time, then they will work in full-
↪duplex mode */
i2s_new_channel(&chan_cfg, &tx_handle, &rx_handle);

/* Set the configurations for BOTH TWO channels, since TX and RX channel have to
↪be same in full-duplex mode */
i2s_std_config_t std_cfg = {
    .clk_cfg = I2S_STD_CLK_DEFAULT_CONFIG(32000),
    .slot_cfg = I2S_STD_PHILIPS_SLOT_DEFAULT_CONFIG(I2S_DATA_BIT_WIDTH_16BIT, I2S_
↪SLOT_MODE_STEREO),
    .gpio_cfg = {
        .mclk = I2S_GPIO_UNUSED,
        .bclk = GPIO_NUM_4,
        .ws = GPIO_NUM_5,
        .dout = GPIO_NUM_18,
        .din = GPIO_NUM_19,
        .invert_flags = {
            .mclk_inv = false,
            .bclk_inv = false,
            .ws_inv = false,
        },
    },
};
i2s_channel_init_std_mode(tx_handle, &std_cfg);
i2s_channel_init_std_mode(rx_handle, &std_cfg);

i2s_channel_enable(tx_handle);
i2s_channel_enable(rx_handle);

...
```

2. Allocate TX and RX handles separately, and initialize them with the same configuration.

```
#include "driver/i2s_std.h"
#include "driver/gpio.h"

i2s_chan_handle_t tx_handle;
i2s_chan_handle_t rx_handle;

/* Allocate a pair of I2S channels on a same port */
i2s_chan_config_t chan_cfg = I2S_CHANNEL_DEFAULT_CONFIG(I2S_NUM_0, I2S_ROLE_
↪MASTER);
/* Allocate for TX and RX channel separately, they are not full-duplex yet */
```

(continues on next page)

(continued from previous page)

```

ESP_ERROR_CHECK(i2s_new_channel(&chan_cfg, &tx_handle, NULL));

/* Set the configurations for BOTH TWO channels, they will constitute in full-
↳duplex mode automatically */
i2s_std_config_t std_cfg = {
    .clk_cfg = I2S_STD_CLK_DEFAULT_CONFIG(32000),
    .slot_cfg = I2S_STD_PHILIPS_SLOT_DEFAULT_CONFIG(I2S_DATA_BIT_WIDTH_16BIT, I2S_
↳SLOT_MODE_STEREO),
    .gpio_cfg = {
        .mclk = I2S_GPIO_UNUSED,
        .bclk = GPIO_NUM_4,
        .ws = GPIO_NUM_5,
        .dout = GPIO_NUM_18,
        .din = GPIO_NUM_19,
        .invert_flags = {
            .mclk_inv = false,
            .bclk_inv = false,
            .ws_inv = false,
        },
    },
};
ESP_ERROR_CHECK(i2s_channel_init_std_mode(tx_handle, &std_cfg));
ESP_ERROR_CHECK(i2s_channel_enable(tx_handle));
// ...
ESP_ERROR_CHECK(i2s_new_channel(&chan_cfg, NULL, &rx_handle));
ESP_ERROR_CHECK(i2s_channel_init_std_mode(rx_handle, &std_cfg));
ESP_ERROR_CHECK(i2s_channel_enable(rx_handle));

...

```

Simplex Mode To allocate a channel in simplex mode, `i2s_new_channel()` should be called for each channel. The clock and GPIO pins of TX/RX channel on ESP32-P4 are independent, so they can be configured with different modes and clocks, and are able to coexist on the same I2S port in simplex mode. PDM duplex can be realized by registering PDM TX simplex and PDM RX simplex on the same I2S port. But in this way, PDM TX/RX might work with different clocks, so take care when configuring the GPIO pins and clocks.

The following example offers a use case for the simplex mode, but note that although the internal MCLK signals for TX and RX channel are separate, the output MCLK can only be bound to one of them if they are from the same controller. If MCLK has been initialized by both channels, it will be bound to the channel that initializes later.

```

#include "driver/i2s_std.h"
#include "driver/gpio.h"

i2s_chan_handle_t tx_handle;
i2s_chan_handle_t rx_handle;
i2s_chan_config_t chan_cfg = I2S_CHANNEL_DEFAULT_CONFIG(I2S_NUM_0, I2S_ROLE_
↳MASTER);
ESP_ERROR_CHECK(i2s_new_channel(&chan_cfg, &tx_handle, NULL));
i2s_std_config_t std_tx_cfg = {
    .clk_cfg = I2S_STD_CLK_DEFAULT_CONFIG(48000),
    .slot_cfg = I2S_STD_PHILIPS_SLOT_DEFAULT_CONFIG(I2S_DATA_BIT_WIDTH_16BIT, I2S_
↳SLOT_MODE_STEREO),
    .gpio_cfg = {
        .mclk = GPIO_NUM_0,
        .bclk = GPIO_NUM_4,
        .ws = GPIO_NUM_5,
        .dout = GPIO_NUM_18,
        .din = I2S_GPIO_UNUSED,
        .invert_flags = {
            .mclk_inv = false,

```

(continues on next page)

(continued from previous page)

```

        .bclk_inv = false,
        .ws_inv = false,
    },
},
};
/* Initialize the channel */
ESP_ERROR_CHECK(i2s_channel_init_std_mode(tx_handle, &std_tx_cfg));
ESP_ERROR_CHECK(i2s_channel_enable(tx_handle));

/* RX channel will be registered on another I2S, if no other available I2S unit
↳ found
* it will return ESP_ERR_NOT_FOUND */
ESP_ERROR_CHECK(i2s_new_channel(&chan_cfg, NULL, &rx_handle)); // Both RX and TX
↳ channel will be registered on I2S0, but they can work with different
↳ configurations.
i2s_std_config_t std_rx_cfg = {
    .clk_cfg = I2S_STD_CLK_DEFAULT_CONFIG(16000),
    .slot_cfg = I2S_STD_MSB_SLOT_DEFAULT_CONFIG(I2S_DATA_BIT_WIDTH_32BIT, I2S_SLOT
↳ MODE_STEREO),
    .gpio_cfg = {
        .mclk = I2S_GPIO_UNUSED,
        .bclk = GPIO_NUM_6,
        .ws = GPIO_NUM_7,
        .dout = I2S_GPIO_UNUSED,
        .din = GPIO_NUM_19,
        .invert_flags = {
            .mclk_inv = false,
            .bclk_inv = false,
            .ws_inv = false,
        },
    },
},
};
ESP_ERROR_CHECK(i2s_channel_init_std_mode(rx_handle, &std_rx_cfg));
ESP_ERROR_CHECK(i2s_channel_enable(rx_handle));

```

Application Notes

How to Prevent Data Lost For applications that need a high frequency sample rate, the massive data throughput may cause data lost. Users can receive data lost event by registering the ISR callback function to receive the event queue:

```

static IRAM_ATTR bool i2s_rx_queue_overflow_callback(i2s_chan_handle_t
↳ handle, i2s_event_data_t *event, void *user_ctx)
{
    // handle RX queue overflow event ...
    return false;
}

i2s_event_callbacks_t cbs = {
    .on_recv = NULL,
    .on_recv_q_ovf = i2s_rx_queue_overflow_callback,
    .on_sent = NULL,
    .on_send_q_ovf = NULL,
};
TEST_ESP_OK(i2s_channel_register_event_callback(rx_handle, &cbs, NULL));

```

Please follow these steps to prevent data lost:

1. Determine the interrupt interval. Generally, when data lost happens, the bigger the interval, the better, which helps to reduce the interrupt times. This means `dma_frame_num` should be as big as possible while the

DMA buffer size is below the maximum value of 4092. The relationships are:

```
interrupt_interval(unit: sec) = dma_frame_num / sample_rate
dma_buffer_size = dma_frame_num * slot_num * data_bit_width / 8 <= 4092
```

2. Determine `dma_desc_num`. `dma_desc_num` is decided by the maximum time of `i2s_channel_read` polling cycle. All the received data is supposed to be stored between two `i2s_channel_read`. This cycle can be measured by a timer or an outputting GPIO signal. The relationship is:

```
dma_desc_num > polling_cycle / interrupt_interval
```

3. Determine the receiving buffer size. The receiving buffer offered by users in `i2s_channel_read` should be able to take all the data in all DMA buffers, which means that it should be larger than the total size of all the DMA buffers:

```
recv_buffer_size > dma_desc_num * dma_buffer_size
```

For example, if there is an I2S application, and the known values are:

```
sample_rate = 144000 Hz
data_bit_width = 32 bits
slot_num = 2
polling_cycle = 10 ms
```

Then the parameters `dma_frame_num`, `dma_desc_num`, and `recv_buf_size` can be calculated as follows:

```
dma_frame_num * slot_num * data_bit_width / 8 = dma_buffer_size <= 4092
dma_frame_num <= 511
interrupt_interval = dma_frame_num / sample_rate = 511 / 144000 = 0.003549 s = 3.
↪ 549 ms
dma_desc_num > polling_cycle / interrupt_interval = cell(10 / 3.549) = cell(2.818)
↪ = 3
recv_buffer_size > dma_desc_num * dma_buffer_size = 3 * 4092 = 12276 bytes
```

API Reference

Standard Mode

Header File

- [components/driver/i2s/include/driver/i2s_std.h](#)
- This header file can be included with:

```
#include "driver/i2s_std.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

[esp_err_t i2s_channel_init_std_mode](#) ([i2s_chan_handle_t](#) handle, const [i2s_std_config_t](#) *std_cfg)

Initialize I2S channel to standard mode.

Note: Only allowed to be called when the channel state is REGISTERED, (i.e., channel has been allocated, but not initialized) and the state will be updated to READY if initialization success, otherwise the state will return to REGISTERED.

Note: When initialize the STD mode with a same configuration as another channel on a same port, these two channels can constitute as full-duplex mode automatically

Parameters

- **handle** -- [in] I2S channel handler
- **std_cfg** -- [in] Configurations for standard mode, including clock, slot and GPIO The clock configuration can be generated by the helper macro `I2S_STD_CLK_DEFAULT_CONFIG` The slot configuration can be generated by the helper macro `I2S_STD_PHILIPS_SLOT_DEFAULT_CONFIG`, `I2S_STD_PCM_SLOT_DEFAULT_CONFIG` or `I2S_STD_MSB_SLOT_DEFAULT_CONFIG`

Returns

- `ESP_OK` Initialize successfully
- `ESP_ERR_NO_MEM` No memory for storing the channel information
- `ESP_ERR_INVALID_ARG` NULL pointer or invalid configuration
- `ESP_ERR_INVALID_STATE` This channel is not registered

esp_err_t **i2s_channel_reconfig_std_clock** (*i2s_chan_handle_t* handle, const *i2s_std_clk_config_t* *clk_cfg)

Reconfigure the I2S clock for standard mode.

Note: Only allowed to be called when the channel state is READY, i.e., channel has been initialized, but not started this function won't change the state. `i2s_channel_disable` should be called before calling this function if I2S has started.

Note: The input channel handle has to be initialized to standard mode, i.e., `i2s_channel_init_std_mode` has been called before reconfiguring

Parameters

- **handle** -- [in] I2S channel handler
- **clk_cfg** -- [in] Standard mode clock configuration, can be generated by `I2S_STD_CLK_DEFAULT_CONFIG`

Returns

- `ESP_OK` Set clock successfully
- `ESP_ERR_INVALID_ARG` NULL pointer, invalid configuration or not standard mode
- `ESP_ERR_INVALID_STATE` This channel is not initialized or not stopped

esp_err_t **i2s_channel_reconfig_std_slot** (*i2s_chan_handle_t* handle, const *i2s_std_slot_config_t* *slot_cfg)

Reconfigure the I2S slot for standard mode.

Note: Only allowed to be called when the channel state is READY, i.e., channel has been initialized, but not started this function won't change the state. `i2s_channel_disable` should be called before calling this function if I2S has started.

Note: The input channel handle has to be initialized to standard mode, i.e.,

`i2s_channel_init_std_mode` has been called before reconfiguring

Parameters

- **handle** -- [in] I2S channel handler
- **slot_cfg** -- [in] Standard mode slot configuration, can be generated by `I2S_STD_PHILIPS_SLOT_DEFAULT_CONFIG`, `I2S_STD_PCM_SLOT_DEFAULT_CONFIG` and `I2S_STD_MSB_SLOT_DEFAULT_CONFIG`.

Returns

- `ESP_OK` Set clock successfully
- `ESP_ERR_NO_MEM` No memory for DMA buffer
- `ESP_ERR_INVALID_ARG` NULL pointer, invalid configuration or not standard mode
- `ESP_ERR_INVALID_STATE` This channel is not initialized or not stopped

esp_err_t **i2s_channel_reconfig_std_gpio** (*i2s_chan_handle_t* handle, const *i2s_std_gpio_config_t* *gpio_cfg)

Reconfigure the I2S GPIO for standard mode.

Note: Only allowed to be called when the channel state is `READY`, i.e., channel has been initialized, but not started this function won't change the state. `i2s_channel_disable` should be called before calling this function if I2S has started.

Note: The input channel handle has to be initialized to standard mode, i.e., `i2s_channel_init_std_mode` has been called before reconfiguring

Parameters

- **handle** -- [in] I2S channel handler
- **gpio_cfg** -- [in] Standard mode GPIO configuration, specified by user

Returns

- `ESP_OK` Set clock successfully
- `ESP_ERR_INVALID_ARG` NULL pointer, invalid configuration or not standard mode
- `ESP_ERR_INVALID_STATE` This channel is not initialized or not stopped

Structures

struct **i2s_std_slot_config_t**

I2S slot configuration for standard mode.

Public Members

i2s_data_bit_width_t **data_bit_width**

I2S sample data bit width (valid data bits per sample)

i2s_slot_bit_width_t **slot_bit_width**

I2S slot bit width (total bits per slot)

i2s_slot_mode_t **slot_mode**

Set mono or stereo mode with `I2S_SLOT_MODE_MONO` or `I2S_SLOT_MODE_STEREO` In TX direction, mono means the written buffer contains only one slot data and stereo means the written buffer contains both left and right data

i2s_std_slot_mask_t **slot_mask**

Select the left, right or both slot

uint32_t **ws_width**

WS signal width (i.e. the number of BCLK ticks that WS signal is high)

bool **ws_pol**

WS signal polarity, set true to enable high lever first

bool **bit_shift**

Set to enable bit shift in Philips mode

bool **left_align**

Set to enable left alignment

bool **big_endian**

Set to enable big endian

bool **bit_order_lsb**

Set to enable lsb first

struct **i2s_std_clk_config_t**

I2S clock configuration for standard mode.

Public Members

uint32_t **sample_rate_hz**

I2S sample rate

i2s_clock_src_t **clk_src**

Choose clock source, see `soc_periph_i2s_clk_src_t` for the supported clock sources. selected `I2S_CLK_SRC_EXTERNAL` (if supports) to enable the external source clock input via MCLK pin,

uint32_t **ext_clk_freq_hz**

External clock source frequency in Hz, only take effect when `clk_src = I2S_CLK_SRC_EXTERNAL`, otherwise this field will be ignored, Please make sure the frequency input is equal or greater than BCLK, i.e. `sample_rate_hz * slot_bits * 2`

i2s_mclk_multiple_t **mclk_multiple**

The multiple of MCLK to the sample rate Default is 256 in the helper macro, it can satisfy most of cases, but please set this field a multiple of 3 (like 384) when using 24-bit data width, otherwise the sample rate might be inaccurate

struct **i2s_std_gpio_config_t**

I2S standard mode GPIO pins configuration.

Public Members

`gpio_num_t mclk`

MCK pin, output by default, input if the clock source is selected to `I2S_CLK_SRC_EXTERNAL`

`gpio_num_t bclk`

BCK pin, input in slave role, output in master role

`gpio_num_t ws`

WS pin, input in slave role, output in master role

`gpio_num_t dout`

DATA pin, output

`gpio_num_t din`

DATA pin, input

`uint32_t mclk_inv`

Set 1 to invert the MCLK input/output

`uint32_t bclk_inv`

Set 1 to invert the BCLK input/output

`uint32_t ws_inv`

Set 1 to invert the WS input/output

struct *i2s_std_gpio_config_t*::[anonymous] **invert_flags**

GPIO pin invert flags

struct **i2s_std_config_t**

I2S standard mode major configuration that including clock/slot/GPIO configuration.

Public Members

i2s_std_clk_config_t **clk_cfg**

Standard mode clock configuration, can be generated by macro `I2S_STD_CLK_DEFAULT_CONFIG`

i2s_std_slot_config_t **slot_cfg**

Standard mode slot configuration, can be generated by macros `I2S_STD_[mode]_SLOT_DEFAULT_CONFIG`, [mode] can be replaced with `PHILIPS/MSB/PCM`

i2s_std_gpio_config_t **gpio_cfg**

Standard mode GPIO configuration, specified by user

Macros

I2S_STD_PHILIPS_SLOT_DEFAULT_CONFIG (bits_per_sample, mono_or_stereo)

Philips format in 2 slots.

This file is specified for I2S standard communication mode Features:

- Philips/MSB/PCM are supported in standard mode
- Fixed to 2 slots

Parameters

- **bits_per_sample** -- I2S data bit width
- **mono_or_stereo** -- I2S_SLOT_MODE_MONO or I2S_SLOT_MODE_STEREO

I2S_STD_PCM_SLOT_DEFAULT_CONFIG (bits_per_sample, mono_or_stereo)

PCM(short) format in 2 slots.

Note: PCM(long) is same as Philips in 2 slots

Parameters

- **bits_per_sample** -- I2S data bit width
- **mono_or_stereo** -- I2S_SLOT_MODE_MONO or I2S_SLOT_MODE_STEREO

I2S_STD_MSB_SLOT_DEFAULT_CONFIG (bits_per_sample, mono_or_stereo)

MSB format in 2 slots.

Parameters

- **bits_per_sample** -- I2S data bit width
- **mono_or_stereo** -- I2S_SLOT_MODE_MONO or I2S_SLOT_MODE_STEREO

I2S_STD_CLK_DEFAULT_CONFIG (rate)

I2S default standard clock configuration.

Note: Please set the mclk_multiple to I2S_MCLK_MULTIPLE_384 while using 24 bits data width. Otherwise the sample rate might be imprecise since the BCLK division is not an integer.

Parameters

- **rate** -- sample rate

PDM Mode

Header File

- [components/driver/i2s/include/driver/i2s_pdm.h](#)
- This header file can be included with:

```
#include "driver/i2s_pdm.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

[esp_err_t i2s_channel_init_pdm_rx_mode](#) ([i2s_chan_handle_t](#) handle, const [i2s_pdm_rx_config_t](#) *pdm_rx_cfg)

Initialize I2S channel to PDM RX mode.

Note: Only allowed to be called when the channel state is REGISTERED, (i.e., channel has been allocated, but not initialized) and the state will be updated to READY if initialization success, otherwise the state will return to REGISTERED.

Parameters

- **handle** -- [in] I2S RX channel handler
- **pdm_rx_cfg** -- [in] Configurations for PDM RX mode, including clock, slot and GPIO The clock configuration can be generated by the helper macro `I2S_PDM_RX_CLK_DEFAULT_CONFIG` The slot configuration can be generated by the helper macro `I2S_PDM_RX_SLOT_DEFAULT_CONFIG`

Returns

- `ESP_OK` Initialize successfully
- `ESP_ERR_NO_MEM` No memory for storing the channel information
- `ESP_ERR_INVALID_ARG` NULL pointer or invalid configuration
- `ESP_ERR_INVALID_STATE` This channel is not registered

esp_err_t **i2s_channel_reconfig_pdm_rx_clock** (*i2s_chan_handle_t* handle, const *i2s_pdm_rx_clk_config_t* *clk_cfg)

Reconfigure the I2S clock for PDM RX mode.

Note: Only allowed to be called when the channel state is READY, i.e., channel has been initialized, but not started this function won't change the state. `i2s_channel_disable` should be called before calling this function if I2S has started.

Note: The input channel handle has to be initialized to PDM RX mode, i.e., `i2s_channel_init_pdm_rx_mode` has been called before reconfiguring

Parameters

- **handle** -- [in] I2S RX channel handler
- **clk_cfg** -- [in] PDM RX mode clock configuration, can be generated by `I2S_PDM_RX_CLK_DEFAULT_CONFIG`

Returns

- `ESP_OK` Set clock successfully
- `ESP_ERR_INVALID_ARG` NULL pointer, invalid configuration or not PDM mode
- `ESP_ERR_INVALID_STATE` This channel is not initialized or not stopped

esp_err_t **i2s_channel_reconfig_pdm_rx_slot** (*i2s_chan_handle_t* handle, const *i2s_pdm_rx_slot_config_t* *slot_cfg)

Reconfigure the I2S slot for PDM RX mode.

Note: Only allowed to be called when the channel state is READY, i.e., channel has been initialized, but not started this function won't change the state. `i2s_channel_disable` should be called before calling this function if I2S has started.

Note: The input channel handle has to be initialized to PDM RX mode, i.e., `i2s_channel_init_pdm_rx_mode` has been called before reconfiguring

Parameters

- **handle** -- [in] I2S RX channel handler

- **slot_cfg** -- [in] PDM RX mode slot configuration, can be generated by `I2S_PDM_RX_SLOT_DEFAULT_CONFIG`

Returns

- `ESP_OK` Set clock successfully
- `ESP_ERR_NO_MEM` No memory for DMA buffer
- `ESP_ERR_INVALID_ARG` NULL pointer, invalid configuration or not PDM mode
- `ESP_ERR_INVALID_STATE` This channel is not initialized or not stopped

esp_err_t **i2s_channel_reconfig_pdm_rx_gpio** (*i2s_chan_handle_t* handle, const *i2s_pdm_rx_gpio_config_t* *gpio_cfg)

Reconfigure the I2S GPIO for PDM RX mode.

Note: Only allowed to be called when the channel state is `READY`, i.e., channel has been initialized, but not started this function won't change the state. `i2s_channel_disable` should be called before calling this function if I2S has started.

Note: The input channel handle has to be initialized to PDM RX mode, i.e., `i2s_channel_init_pdm_rx_mode` has been called before reconfiguring

Parameters

- **handle** -- [in] I2S RX channel handler
- **gpio_cfg** -- [in] PDM RX mode GPIO configuration, specified by user

Returns

- `ESP_OK` Set clock successfully
- `ESP_ERR_INVALID_ARG` NULL pointer, invalid configuration or not PDM mode
- `ESP_ERR_INVALID_STATE` This channel is not initialized or not stopped

esp_err_t **i2s_channel_init_pdm_tx_mode** (*i2s_chan_handle_t* handle, const *i2s_pdm_tx_config_t* *pdm_tx_cfg)

Initialize I2S channel to PDM TX mode.

Note: Only allowed to be called when the channel state is `REGISTERED`, (i.e., channel has been allocated, but not initialized) and the state will be updated to `READY` if initialization success, otherwise the state will return to `REGISTERED`.

Parameters

- **handle** -- [in] I2S TX channel handler
- **pdm_tx_cfg** -- [in] Configurations for PDM TX mode, including clock, slot and GPIO The clock configuration can be generated by the helper macro `I2S_PDM_TX_CLK_DEFAULT_CONFIG` The slot configuration can be generated by the helper macro `I2S_PDM_TX_SLOT_DEFAULT_CONFIG`

Returns

- `ESP_OK` Initialize successfully
- `ESP_ERR_NO_MEM` No memory for storing the channel information
- `ESP_ERR_INVALID_ARG` NULL pointer or invalid configuration
- `ESP_ERR_INVALID_STATE` This channel is not registered

esp_err_t **i2s_channel_reconfig_pdm_tx_clock** (*i2s_chan_handle_t* handle, const *i2s_pdm_tx_clk_config_t* *clk_cfg)

Reconfigure the I2S clock for PDM TX mode.

Note: Only allowed to be called when the channel state is `READY`, i.e., channel has been initialized, but not

started this function won't change the state. `i2s_channel_disable` should be called before calling this function if I2S has started.

Note: The input channel handle has to be initialized to PDM TX mode, i.e., `i2s_channel_init_pdm_tx_mode` has been called before reconfiguring

Parameters

- **handle** -- [in] I2S TX channel handler
- **clk_cfg** -- [in] PDM TX mode clock configuration, can be generated by `I2S_PDM_TX_CLK_DEFAULT_CONFIG`

Returns

- `ESP_OK` Set clock successfully
- `ESP_ERR_INVALID_ARG` NULL pointer, invalid configuration or not PDM mode
- `ESP_ERR_INVALID_STATE` This channel is not initialized or not stopped

esp_err_t **i2s_channel_reconfig_pdm_tx_slot** (*i2s_chan_handle_t* handle, const *i2s_pdm_tx_slot_config_t* *slot_cfg)

Reconfigure the I2S slot for PDM TX mode.

Note: Only allowed to be called when the channel state is READY, i.e., channel has been initialized, but not started this function won't change the state. `i2s_channel_disable` should be called before calling this function if I2S has started.

Note: The input channel handle has to be initialized to PDM TX mode, i.e., `i2s_channel_init_pdm_tx_mode` has been called before reconfiguring

Parameters

- **handle** -- [in] I2S TX channel handler
- **slot_cfg** -- [in] PDM TX mode slot configuration, can be generated by `I2S_PDM_TX_SLOT_DEFAULT_CONFIG`

Returns

- `ESP_OK` Set clock successfully
- `ESP_ERR_NO_MEM` No memory for DMA buffer
- `ESP_ERR_INVALID_ARG` NULL pointer, invalid configuration or not PDM mode
- `ESP_ERR_INVALID_STATE` This channel is not initialized or not stopped

esp_err_t **i2s_channel_reconfig_pdm_tx_gpio** (*i2s_chan_handle_t* handle, const *i2s_pdm_tx_gpio_config_t* *gpio_cfg)

Reconfigure the I2S GPIO for PDM TX mode.

Note: Only allowed to be called when the channel state is READY, i.e., channel has been initialized, but not started this function won't change the state. `i2s_channel_disable` should be called before calling this function if I2S has started.

Note: The input channel handle has to be initialized to PDM TX mode, i.e., `i2s_channel_init_pdm_tx_mode` has been called before reconfiguring

Parameters

- **handle** -- [in] I2S TX channel handler

- **gpio_cfg** -- [in] PDM TX mode GPIO configuration, specified by user

Returns

- ESP_OK Set clock successfully
- ESP_ERR_INVALID_ARG NULL pointer, invalid configuration or not PDM mode
- ESP_ERR_INVALID_STATE This channel is not initialized or not stopped

Structures

struct **i2s_pdm_rx_slot_config_t**

I2S slot configuration for PDM RX mode.

Public Members

i2s_data_bit_width_t **data_bit_width**

I2S sample data bit width (valid data bits per sample), only support 16 bits for PDM mode

i2s_slot_bit_width_t **slot_bit_width**

I2S slot bit width (total bits per slot) , only support 16 bits for PDM mode

i2s_slot_mode_t **slot_mode**

Set mono or stereo mode with I2S_SLOT_MODE_MONO or I2S_SLOT_MODE_STEREO

i2s_pdm_slot_mask_t **slot_mask**

Choose the slots to activate

bool **hp_en**

High pass filter enable

float **hp_cut_off_freq_hz**

High pass filter cut-off frequency, range 23.3Hz ~ 185Hz, see cut-off frequency sheet above

uint32_t **amplify_num**

The amplification number of the final conversion result. The data that have converted from PDM to PCM module, will time `amplify_num` additionally to amplify the final result. Note that it's only a multiplier of the digital PCM data, not the gain of the analog signal range 1~15, default 1

struct **i2s_pdm_rx_clk_config_t**

I2S clock configuration for PDM RX mode.

Public Members

uint32_t **sample_rate_hz**

I2S sample rate

i2s_clock_src_t **clk_src**

Choose clock source

i2s_mclk_multiple_t **mclk_multiple**

The multiple of MCLK to the sample rate

i2s_pdm_dsr_t **dn_sample_mode**

Down-sampling rate mode

uint32_t **bclk_div**

The division from MCLK to BCLK. The typical and minimum value is I2S_PDM_RX_BCLK_DIV_MIN. It will be set to I2S_PDM_RX_BCLK_DIV_MIN by default if it is smaller than I2S_PDM_RX_BCLK_DIV_MIN

struct **i2s_pdm_rx_gpio_config_t**

I2S PDM TX mode GPIO pins configuration.

Public Members

gpio_num_t **clk**

PDM clk pin, output

gpio_num_t **din**

DATA pin 0, input

gpio_num_t **dins**[(4)]

DATA pins, input, only take effect when corresponding I2S_PDM_RX_LINE_x_SLOT_{xxx} is enabled in *i2s_pdm_rx_slot_config_t::slot_mask*

uint32_t **clk_inv**

Set 1 to invert the clk output

struct *i2s_pdm_rx_gpio_config_t*::[anonymous] **invert_flags**

GPIO pin invert flags

struct **i2s_pdm_rx_config_t**

I2S PDM RX mode major configuration that including clock/slot/GPIO configuration.

Public Members

i2s_pdm_rx_clk_config_t **clk_cfg**

PDM RX clock configurations, can be generated by macro I2S_PDM_RX_CLK_DEFAULT_CONFIG

i2s_pdm_rx_slot_config_t **slot_cfg**

PDM RX slot configurations, can be generated by macro I2S_PDM_RX_SLOT_DEFAULT_CONFIG

i2s_pdm_rx_gpio_config_t **gpio_cfg**

PDM RX slot configurations, specified by user

struct **i2s_pdm_tx_slot_config_t**

I2S slot configuration for PDM TX mode.

Public Members

i2s_data_bit_width_t **data_bit_width**

I2S sample data bit width (valid data bits per sample), only support 16 bits for PDM mode

i2s_slot_bit_width_t **slot_bit_width**

I2S slot bit width (total bits per slot), only support 16 bits for PDM mode

i2s_slot_mode_t **slot_mode**

Set mono or stereo mode with I2S_SLOT_MODE_MONO or I2S_SLOT_MODE_STEREO For PDM TX mode, mono means the data buffer only contains one slot data, Stereo means the data buffer contains two slots data

uint32_t **sd_prescale**

Sigma-delta filter prescale

i2s_pdm_sig_scale_t **sd_scale**

Sigma-delta filter scaling value

i2s_pdm_sig_scale_t **hp_scale**

High pass filter scaling value

i2s_pdm_sig_scale_t **lp_scale**

Low pass filter scaling value

i2s_pdm_sig_scale_t **sinc_scale**

Sinc filter scaling value

i2s_pdm_tx_line_mode_t **line_mode**

PDM TX line mode, one-line codec, one-line dac, two-line dac mode can be selected

bool **hp_en**

High pass filter enable

float **hp_cut_off_freq_hz**

High pass filter cut-off frequency, range 23.3Hz ~ 185Hz, see cut-off frequency sheet above

uint32_t **sd_dither**

Sigma-delta filter dither

uint32_t **sd_dither2**

Sigma-delta filter dither2

struct **i2s_pdm_tx_clk_config_t**

I2S clock configuration for PDM TX mode.

Public Members

uint32_t **sample_rate_hz**

I2S sample rate, not suggest to exceed 48000 Hz, otherwise more glitches and noise may appear

i2s_clock_src_t **clk_src**

Choose clock source

i2s_mclk_multiple_t **mclk_multiple**

The multiple of MCLK to the sample rate

uint32_t **up_sample_fp**

Up-sampling param fp

uint32_t **up_sample_fs**

Up-sampling param fs, not allowed to be greater than 480

uint32_t **bclk_div**

The division from MCLK to BCLK. The minimum value is I2S_PDM_TX_BCLK_DIV_MIN. It will be set to I2S_PDM_TX_BCLK_DIV_MIN by default if it is smaller than I2S_PDM_TX_BCLK_DIV_MIN

struct **i2s_pdm_tx_gpio_config_t**

I2S PDM TX mode GPIO pins configuration.

Public Members

gpio_num_t **clk**

PDM clk pin, output

gpio_num_t **dout**

DATA pin, output

gpio_num_t **dout2**

The second data pin for the DAC dual-line mode, only take effect when the line mode is I2S_PDM_TX_TWO_LINE_DAC

uint32_t **clk_inv**

Set 1 to invert the clk output

struct *i2s_pdm_tx_gpio_config_t*::[anonymous] **invert_flags**

GPIO pin invert flags

struct **i2s_pdm_tx_config_t**

I2S PDM TX mode major configuration that including clock/slot/GPIO configuration.

Public Members

i2s_pdm_tx_clk_config_t **clk_cfg**

PDM TX clock configurations, can be generated by macro I2S_PDM_TX_CLK_DEFAULT_CONFIG

i2s_pdm_tx_slot_config_t slot_cfg

PDM TX slot configurations, can be generated by macro I2S_PDM_TX_SLOT_DEFAULT_CONFIG

i2s_pdm_tx_gpio_config_t gpio_cfg

PDM TX GPIO configurations, specified by user

Macros

I2S_PDM_RX_SLOT_DEFAULT_CONFIG (bits_per_sample, mono_or_stereo)

PDM format in 2 slots(RX)

This file is specified for I2S PDM communication mode Features:

- Only support PDM TX/RX mode
- Fixed to 2 slots
- Data bit width only support 16 bits

Parameters

- **bits_per_sample** -- I2S data bit width, only support 16 bits for PDM mode
- **mono_or_stereo** -- I2S_SLOT_MODE_MONO or I2S_SLOT_MODE_STEREO

I2S_PDM_RX_CLK_DEFAULT_CONFIG (rate)

I2S default PDM RX clock configuration.

Parameters

- **rate** -- sample rate

I2S_PDM_TX_SLOT_DEFAULT_CONFIG (bits_per_sample, mono_or_stereo)

PDM style in 2 slots(TX) for codec line mode.

Parameters

- **bits_per_sample** -- I2S data bit width, only support 16 bits for PDM mode
- **mono_or_stereo** -- I2S_SLOT_MODE_MONO or I2S_SLOT_MODE_STEREO

I2S_PDM_TX_SLOT_DAC_DEFAULT_CONFIG (bits_per_sample, mono_or_stereo)

PDM style in 1 slots(TX) for DAC line mode.

Note: The noise might be different with different configurations, this macro provides a set of configurations that have relatively high SNR (Signal Noise Ratio), you can also adjust them to fit your case.

Parameters

- **bits_per_sample** -- I2S data bit width, only support 16 bits for PDM mode
- **mono_or_stereo** -- I2S_SLOT_MODE_MONO or I2S_SLOT_MODE_STEREO

I2S_PDM_TX_CLK_DEFAULT_CONFIG (rate)

I2S default PDM TX clock configuration for codec line mode.

Note: TX PDM can only be set to the following two up-sampling rate configurations: 1: fp = 960, fs = sample_rate_hz / 100, in this case, Fpdm = 128*48000 2: fp = 960, fs = 480, in this case, Fpdm = 128*Fpcm = 128*sample_rate_hz If the PDM receiver do not care the PDM serial clock, it's recommended set Fpdm = 128*48000. Otherwise, the second configuration should be adopted.

Parameters

- **rate** -- sample rate (not suggest to exceed 48000 Hz, otherwise more glitches and noise may appear)

I2S_PDM_TX_CLK_DAC_DEFAULT_CONFIG (rate)

I2S default PDM TX clock configuration for DAC line mode.

Note: TX PDM can only be set to the following two up-sampling rate configurations: 1: $fp = 960$, $fs = \text{sample\_rate\_hz} / 100$, in this case, $Fpdm = 128 * 48000$ 2: $fp = 960$, $fs = 480$, in this case, $Fpdm = 128 * Fpcm = 128 * \text{sample\_rate\_hz}$ If the PDM receiver do not care the PDM serial clock, it's recommended set $Fpdm = 128 * 48000$. Otherwise, the second configuration should be adopted.

Note: The noise might be different with different configurations, this macro provides a set of configurations that have relatively high SNR (Signal Noise Ratio), you can also adjust them to fit your case.

Parameters

- **rate** -- sample rate (not suggest to exceed 48000 Hz, otherwise more glitches and noise may appear)

TDM Mode**Header File**

- [components/driver/i2s/include/driver/i2s_tdm.h](#)
- This header file can be included with:

```
#include "driver/i2s_tdm.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

[esp_err_t](#) **i2s_channel_init_tdm_mode** ([i2s_chan_handle_t](#) handle, const [i2s_tdm_config_t](#) *tdm_cfg)

Initialize I2S channel to TDM mode.

Note: Only allowed to be called when the channel state is REGISTERED, (i.e., channel has been allocated, but not initialized) and the state will be updated to READY if initialization success, otherwise the state will return to REGISTERED.

Note: When initialize the TDM mode with a same configuration as another channel on a same port, these two channels can constitute as full-duplex mode automatically

Parameters

- **handle** -- [in] I2S channel handler
- **tdm_cfg** -- [in] Configurations for TDM mode, including clock, slot and GPIO The clock configuration can be generated by the helper macro `I2S_TDM_CLK_DEFAULT_CONFIG` The slot configuration can be generated by the helper macro `I2S_TDM_PHILIPS_SLOT_DEFAULT_CONFIG`, `I2S_TDM_PCM_SHORT_SLOT_DEFAULT_CONFIG`,

I2S_TDM_PCM_LONG_SLOT_DEFAULT_CONFIG or
 I2S_TDM_MSB_SLOT_DEFAULT_CONFIG

Returns

- ESP_OK Initialize successfully
- ESP_ERR_NO_MEM No memory for storing the channel information
- ESP_ERR_INVALID_ARG NULL pointer or invalid configuration
- ESP_ERR_INVALID_STATE This channel is not registered

esp_err_t **i2s_channel_reconfig_tdm_clock** (*i2s_chan_handle_t* handle, const *i2s_tdm_clk_config_t* *clk_cfg)

Reconfigure the I2S clock for TDM mode.

Note: Only allowed to be called when the channel state is READY, i.e., channel has been initialized, but not started this function won't change the state. `i2s_channel_disable` should be called before calling this function if I2S has started.

Note: The input channel handle has to be initialized to TDM mode, i.e., `i2s_channel_init_tdm_mode` has been called before reconfiguring

Parameters

- **handle** -- [in] I2S channel handler
- **clk_cfg** -- [in] Standard mode clock configuration, can be generated by I2S_TDM_CLK_DEFAULT_CONFIG

Returns

- ESP_OK Set clock successfully
- ESP_ERR_INVALID_ARG NULL pointer, invalid configuration or not TDM mode
- ESP_ERR_INVALID_STATE This channel is not initialized or not stopped

esp_err_t **i2s_channel_reconfig_tdm_slot** (*i2s_chan_handle_t* handle, const *i2s_tdm_slot_config_t* *slot_cfg)

Reconfigure the I2S slot for TDM mode.

Note: Only allowed to be called when the channel state is READY, i.e., channel has been initialized, but not started this function won't change the state. `i2s_channel_disable` should be called before calling this function if I2S has started.

Note: The input channel handle has to be initialized to TDM mode, i.e., `i2s_channel_init_tdm_mode` has been called before reconfiguring

Parameters

- **handle** -- [in] I2S channel handler
- **slot_cfg** -- [in] Standard mode slot configuration, can be generated by I2S_TDM_PHILIPS_SLOT_DEFAULT_CONFIG, I2S_TDM_PCM_SHORT_SLOT_DEFAULT_CONFIG, I2S_TDM_PCM_LONG_SLOT_DEFAULT_CONFIG or
I2S_TDM_MSB_SLOT_DEFAULT_CONFIG.

Returns

- ESP_OK Set clock successfully
- ESP_ERR_NO_MEM No memory for DMA buffer
- ESP_ERR_INVALID_ARG NULL pointer, invalid configuration or not TDM mode
- ESP_ERR_INVALID_STATE This channel is not initialized or not stopped

esp_err_t **i2s_channel_reconfig_tdm_gpio** (*i2s_chan_handle_t* handle, const *i2s_tdm_gpio_config_t* *gpio_cfg)

Reconfigure the I2S GPIO for TDM mode.

Note: Only allowed to be called when the channel state is READY, i.e., channel has been initialized, but not started this function won't change the state. `i2s_channel_disable` should be called before calling this function if I2S has started.

Note: The input channel handle has to be initialized to TDM mode, i.e., `i2s_channel_init_tdm_mode` has been called before reconfiguring

Parameters

- **handle** -- [in] I2S channel handler
- **gpio_cfg** -- [in] Standard mode GPIO configuration, specified by user

Returns

- ESP_OK Set clock successfully
- ESP_ERR_INVALID_ARG NULL pointer, invalid configuration or not TDM mode
- ESP_ERR_INVALID_STATE This channel is not initialized or not stopped

Structures

struct **i2s_tdm_slot_config_t**

I2S slot configuration for TDM mode.

Public Members

i2s_data_bit_width_t **data_bit_width**

I2S sample data bit width (valid data bits per sample)

i2s_slot_bit_width_t **slot_bit_width**

I2S slot bit width (total bits per slot)

i2s_slot_mode_t **slot_mode**

Set mono or stereo mode with I2S_SLOT_MODE_MONO or I2S_SLOT_MODE_STEREO

i2s_tdm_slot_mask_t **slot_mask**

Slot mask. Activating slots by setting 1 to corresponding bits. When the activated slots is not consecutive, those data in inactivated slots will be ignored

uint32_t **ws_width**

WS signal width (i.e. the number of BCLK ticks that WS signal is high)

bool **ws_pol**

WS signal polarity, set true to enable high level first

bool **bit_shift**

Set true to enable bit shift in Philips mode

bool **left_align**

Set true to enable left alignment

bool **big_endian**

Set true to enable big endian

bool **bit_order_lsb**

Set true to enable lsb first

bool **skip_mask**

Set true to enable skip mask. If it is enabled, only the data of the enabled channels will be sent, otherwise all data stored in DMA TX buffer will be sent

uint32_t **total_slot**

I2S total number of slots. If it is smaller than the biggest activated channel number, it will be set to this number automatically.

struct **i2s_tdm_clk_config_t**

I2S clock configuration for TDM mode.

Public Members

uint32_t **sample_rate_hz**

I2S sample rate

i2s_clock_src_t **clk_src**

Choose clock source, see `soc_periph_i2s_clk_src_t` for the supported clock sources. selected `I2S_CLK_SRC_EXTERNAL` (if supports) to enable the external source clock inputted via MCLK pin, please make sure the frequency inputted is equal or greater than `sample_rate_hz * mclk_multiple`

uint32_t **ext_clk_freq_hz**

External clock source frequency in Hz, only take effect when `clk_src = I2S_CLK_SRC_EXTERNAL`, otherwise this field will be ignored Please make sure the frequency inputted is equal or greater than BCLK, i.e. `sample_rate_hz * slot_bits * slot_num`

i2s_mclk_multiple_t **mclk_multiple**

The multiple of MCLK to the sample rate, only take effect for master role

uint32_t **bclk_div**

The division from MCLK to BCLK, only take effect for slave role, it shouldn't be smaller than 8. Increase this field when data sent by slave lag behind

struct **i2s_tdm_gpio_config_t**

I2S TDM mode GPIO pins configuration.

Public Members

`gpio_num_t mclk`

MCK pin, output by default, input if the clock source is selected to `I2S_CLK_SRC_EXTERNAL`

`gpio_num_t bclk`

BCK pin, input in slave role, output in master role

`gpio_num_t ws`

WS pin, input in slave role, output in master role

`gpio_num_t dout`

DATA pin, output

`gpio_num_t din`

DATA pin, input

`uint32_t mclk_inv`

Set 1 to invert the MCLK input/output

`uint32_t bclk_inv`

Set 1 to invert the BCLK input/output

`uint32_t ws_inv`

Set 1 to invert the WS input/output

struct *`i2s_tdm_gpio_config_t`*::[anonymous] **`invert_flags`**

GPIO pin invert flags

struct **`i2s_tdm_config_t`**

I2S TDM mode major configuration that including clock/slot/GPIO configuration.

Public Members

`i2s_tdm_clk_config_t` **`clk_cfg`**

TDM mode clock configuration, can be generated by macro `I2S_TDM_CLK_DEFAULT_CONFIG`

`i2s_tdm_slot_config_t` **`slot_cfg`**

TDM mode slot configuration, can be generated by macros `I2S_TDM_[mode]_SLOT_DEFAULT_CONFIG`, [mode] can be replaced with `PHILIPS/MSB/PCM_SHORT/PCM_LONG`

`i2s_tdm_gpio_config_t` **`gpio_cfg`**

TDM mode GPIO configuration, specified by user

Macros

`I2S_TDM_AUTO_SLOT_NUM`

This file is specified for I2S TDM communication mode Features:

- More than 2 slots

I2S_TDM_AUTO_WS_WIDTH**I2S_TDM_PHILIPS_SLOT_DEFAULT_CONFIG** (bits_per_sample, mono_or_stereo, mask)

Philips format in active slot that enabled by mask.

Parameters

- **bits_per_sample** -- I2S data bit width
- **mono_or_stereo** -- I2S_SLOT_MODE_MONO or I2S_SLOT_MODE_STEREO
- **mask** -- active slot mask

I2S_TDM_MSB_SLOT_DEFAULT_CONFIG (bits_per_sample, mono_or_stereo, mask)

MSB format in active slot enabled that by mask.

Parameters

- **bits_per_sample** -- I2S data bit width
- **mono_or_stereo** -- I2S_SLOT_MODE_MONO or I2S_SLOT_MODE_STEREO
- **mask** -- active slot mask

I2S_TDM_PCM_SHORT_SLOT_DEFAULT_CONFIG (bits_per_sample, mono_or_stereo, mask)

PCM(short) format in active slot that enabled by mask.

Parameters

- **bits_per_sample** -- I2S data bit width
- **mono_or_stereo** -- I2S_SLOT_MODE_MONO or I2S_SLOT_MODE_STEREO
- **mask** -- active slot mask

I2S_TDM_PCM_LONG_SLOT_DEFAULT_CONFIG (bits_per_sample, mono_or_stereo, mask)

PCM(long) format in active slot that enabled by mask.

Parameters

- **bits_per_sample** -- I2S data bit width
- **mono_or_stereo** -- I2S_SLOT_MODE_MONO or I2S_SLOT_MODE_STEREO
- **mask** -- active slot mask

I2S_TDM_CLK_DEFAULT_CONFIG (rate)

I2S default TDM clock configuration.

Note: Please set the mclk_multiple to I2S_MCLK_MULTIPLE_384 while the data width in slot configuration is set to 24 bits Otherwise the sample rate might be imprecise since the BCLK division is not a integer

Parameters

- **rate** -- sample rate

I2S Driver

Header File

- [components/driver/i2s/include/driver/i2s_common.h](#)
- This header file can be included with:

```
#include "driver/i2s_common.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

PRIV_REQUIRES driver

Functions

esp_err_t **i2s_new_channel** (const *i2s_chan_config_t* *chan_cfg, *i2s_chan_handle_t* *ret_tx_handle, *i2s_chan_handle_t* *ret_rx_handle)

Allocate new I2S channel(s)

Note: The new created I2S channel handle will be REGISTERED state after it is allocated successfully.

Note: When the port id in channel configuration is I2S_NUM_AUTO, driver will allocate I2S port automatically on one of the I2S controller, otherwise driver will try to allocate the new channel on the selected port.

Note: If both tx_handle and rx_handle are not NULL, it means this I2S controller will work at full-duplex mode, the RX and TX channels will be allocated on a same I2S port in this case. Note that some configurations of TX/RX channel are shared on ESP32 and ESP32S2, so please make sure they are working at same condition and under same status(start/stop). Currently, full-duplex mode can't guarantee TX/RX channels write/read synchronously, they can only share the clock signals for now.

Note: If tx_handle OR rx_handle is NULL, it means this I2S controller will work at simplex mode. For ESP32 and ESP32S2, the whole I2S controller (i.e. both RX and TX channel) will be occupied, even if only one of RX or TX channel is registered. For the other targets, another channel on this controller will still available.

Parameters

- **chan_cfg** -- [in] I2S controller channel configurations
- **ret_tx_handle** -- [out] I2S channel handler used for managing the sending channel(optional)
- **ret_rx_handle** -- [out] I2S channel handler used for managing the receiving channel(optional)

Returns

- ESP_OK Allocate new channel(s) success
- ESP_ERR_NOT_SUPPORTED The communication mode is not supported on the current chip
- ESP_ERR_INVALID_ARG NULL pointer or illegal parameter in *i2s_chan_config_t*
- ESP_ERR_NOT_FOUND No available I2S channel found

esp_err_t **i2s_del_channel** (*i2s_chan_handle_t* handle)

Delete the I2S channel.

Note: Only allowed to be called when the I2S channel is at REGISTERED or READY state (i.e., it should stop before deleting it).

Note: Resource will be free automatically if all channels in one port are deleted

Parameters **handle** -- [in] I2S channel handler

- ESP_OK Delete successfully
- ESP_ERR_INVALID_ARG NULL pointer

esp_err_t **i2s_channel_get_info** (*i2s_chan_handle_t* handle, *i2s_chan_info_t* *chan_info)

Get I2S channel information.

Parameters

- **handle** -- [in] I2S channel handler
- **chan_info** -- [out] I2S channel basic information

Returns

- ESP_OK Get I2S channel information success
- ESP_ERR_NOT_FOUND The input handle doesn't match any registered I2S channels, it may not an I2S channel handle or not available any more
- ESP_ERR_INVALID_ARG The input handle or chan_info pointer is NULL

esp_err_t **i2s_channel_enable** (*i2s_chan_handle_t* handle)

Enable the I2S channel.

Note: Only allowed to be called when the channel state is READY, (i.e., channel has been initialized, but not started) the channel will enter RUNNING state once it is enabled successfully.

Note: Enable the channel can start the I2S communication on hardware. It will start outputting BCLK and WS signal. For MCLK signal, it will start to output when initialization is finished

Parameters **handle** -- [in] I2S channel handler

- ESP_OK Start successfully
- ESP_ERR_INVALID_ARG NULL pointer
- ESP_ERR_INVALID_STATE This channel has not initialized or already started

esp_err_t **i2s_channel_disable** (*i2s_chan_handle_t* handle)

Disable the I2S channel.

Note: Only allowed to be called when the channel state is RUNNING, (i.e., channel has been started) the channel will enter READY state once it is disabled successfully.

Note: Disable the channel can stop the I2S communication on hardware. It will stop BCLK and WS signal but not MCLK signal

Parameters **handle** -- [in] I2S channel handler

Returns

- ESP_OK Stop successfully
- ESP_ERR_INVALID_ARG NULL pointer
- ESP_ERR_INVALID_STATE This channel has not stated

esp_err_t **i2s_channel_preload_data** (*i2s_chan_handle_t* tx_handle, const void *src, size_t size, size_t *bytes_loaded)

Preload the data into TX DMA buffer.

Note: Only allowed to be called when the channel state is READY, (i.e., channel has been initialized, but not started)

Note: As the initial DMA buffer has no data inside, it will transmit the empty buffer after enabled the channel, this function is used to preload the data into the DMA buffer, so that the valid data can be transmitted

immediately after the channel is enabled.

Note: This function can be called multiple times before enabling the channel, the buffer that loaded later will be concatenated behind the former loaded buffer. But when all the DMA buffers have been loaded, no more data can be preload then, please check the `bytes_loaded` parameter to see how many bytes are loaded successfully, when the `bytes_loaded` is smaller than the `size`, it means the DMA buffers are full.

Parameters

- **tx_handle** -- [in] I2S TX channel handler
- **src** -- [in] The pointer of the source buffer to be loaded
- **size** -- [in] The source buffer size
- **bytes_loaded** -- [out] The bytes that successfully been loaded into the TX DMA buffer

Returns

- ESP_OK Load data successful
- ESP_FAIL Failed to push the message queue
- ESP_ERR_INVALID_ARG NULL pointer or not TX direction
- ESP_ERR_INVALID_STATE This channel has not started

esp_err_t **i2s_channel_write** (*i2s_chan_handle_t* handle, const void *src, size_t size, size_t *bytes_written, uint32_t timeout_ms)

I2S write data.

Note: Only allowed to be called when the channel state is RUNNING, (i.e., TX channel has been started and is not writing now) but the RUNNING only stands for the software state, it doesn't mean there is no the signal transporting on line.

Parameters

- **handle** -- [in] I2S channel handler
- **src** -- [in] The pointer of sent data buffer
- **size** -- [in] Max data buffer length
- **bytes_written** -- [out] Byte number that actually be sent, can be NULL if not needed
- **timeout_ms** -- [in] Max block time

Returns

- ESP_OK Write successfully
- ESP_ERR_INVALID_ARG NULL pointer or this handle is not TX handle
- ESP_ERR_TIMEOUT Writing timeout, no writing event received from ISR within ticks_to_wait
- ESP_ERR_INVALID_STATE I2S is not ready to write

esp_err_t **i2s_channel_read** (*i2s_chan_handle_t* handle, void *dest, size_t size, size_t *bytes_read, uint32_t timeout_ms)

I2S read data.

Note: Only allowed to be called when the channel state is RUNNING but the RUNNING only stands for the software state, it doesn't mean there is no the signal transporting on line.

Parameters

- **handle** -- [in] I2S channel handler
- **dest** -- [in] The pointer of receiving data buffer
- **size** -- [in] Max data buffer length
- **bytes_read** -- [out] Byte number that actually be read, can be NULL if not needed

- **timeout_ms** -- [in] Max block time

Returns

- ESP_OK Read successfully
- ESP_ERR_INVALID_ARG NULL pointer or this handle is not RX handle
- ESP_ERR_TIMEOUT Reading timeout, no reading event received from ISR within ticks_to_wait
- ESP_ERR_INVALID_STATE I2S is not ready to read

esp_err_t **i2s_channel_register_event_callback** (*i2s_chan_handle_t* handle, const *i2s_event_callbacks_t* *callbacks, void *user_data)

Set event callbacks for I2S channel.

Note: Only allowed to be called when the channel state is REGISTERED / READY, (i.e., before channel starts)

Note: User can deregister a previously registered callback by calling this function and setting the callback member in the `callbacks` structure to NULL.

Note: When CONFIG_I2S_ISR_IRAM_SAFE is enabled, the callback itself and functions called by it should be placed in IRAM. The variables used in the function should be in the SRAM as well. The `user_data` should also reside in SRAM or internal RAM as well.

Parameters

- **handle** -- [in] I2S channel handler
- **callbacks** -- [in] Group of callback functions
- **user_data** -- [in] User data, which will be passed to callback functions directly

Returns

- ESP_OK Set event callbacks successfully
- ESP_ERR_INVALID_ARG Set event callbacks failed because of invalid argument
- ESP_ERR_INVALID_STATE Set event callbacks failed because the current channel state is not REGISTERED or READY

Structures

struct **i2s_event_callbacks_t**

Group of I2S callbacks.

Note: The callbacks are all running under ISR environment

Note: When CONFIG_I2S_ISR_IRAM_SAFE is enabled, the callback itself and functions called by it should be placed in IRAM. The variables used in the function should be in the SRAM as well.

Public Members

i2s_isr_callback_t **on_recv**

Callback of data received event, only for RX channel The event data includes DMA buffer address and size that just finished receiving data

***i2s_isr_callback_t* on_recv_q_ovf**

Callback of receiving queue overflowed event, only for RX channel The event data includes buffer size that has been overwritten

***i2s_isr_callback_t* on_sent**

Callback of data sent event, only for TX channel The event data includes DMA buffer address and size that just finished sending data

***i2s_isr_callback_t* on_send_q_ovf**

Callback of sending queue overflowed event, only for TX channel The event data includes buffer size that has been overwritten

struct **i2s_chan_config_t**

I2S controller channel configuration.

Public Members***i2s_port_t* id**

I2S port id

***i2s_role_t* role**

I2S role, I2S_ROLE_MASTER or I2S_ROLE_SLAVE

uint32_t dma_desc_num

I2S DMA buffer number, it is also the number of DMA descriptor

uint32_t dma_frame_num

I2S frame number in one DMA buffer. One frame means one-time sample data in all slots, it should be the multiple of 3 when the data bit width is 24.

bool auto_clear

Set to auto clear DMA TX buffer, I2S will always send zero automatically if no data to send

int intr_priority

I2S interrupt priority, range [0, 7], if set to 0, the driver will try to allocate an interrupt with a relative low priority (1,2,3)

struct **i2s_chan_info_t**

I2S channel information.

Public Members***i2s_port_t* id**

I2S port id

***i2s_role_t* role**

I2S role, I2S_ROLE_MASTER or I2S_ROLE_SLAVE

i2s_dir_t **dir**

I2S channel direction

i2s_comm_mode_t **mode**

I2S channel communication mode

i2s_chan_handle_t **pair_chan**

I2S pair channel handle in duplex mode, always NULL in simplex mode

Macros

I2S_CHANNEL_DEFAULT_CONFIG (i2s_num, i2s_role)

get default I2S property

I2S_GPIO_UNUSED

Used in i2s_gpio_config_t for signals which are not used

I2S Types

Header File

- [components/driver/i2s/include/driver/i2s_types.h](#)
- This header file can be included with:

```
#include "driver/i2s_types.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Structures

struct **i2s_event_data_t**

Event structure used in I2S event queue.

Public Members

void ***data**

The pointer of DMA buffer that just finished sending or receiving for `on_recv` and `on_sent` callback
NULL for `on_recv_q_ovf` and `on_send_q_ovf` callback

size_t **size**

The buffer size of DMA buffer when success to send or receive, also the buffer size that dropped when queue overflow. It is related to the `dma_frame_num` and `data_bit_width`, typically it is fixed when `data_bit_width` is not changed.

Type Definitions

typedef struct i2s_channel_obj_t ***i2s_chan_handle_t**

I2S channel object handle, the control unit of the I2S driver

typedef bool (***i2s_isr_callback_t**)(*i2s_chan_handle_t* handle, *i2s_event_data_t* *event, void *user_ctx)

I2S event callback.

Param handle [in] I2S channel handle, created from `i2s_new_channel()`

Param event [in] I2S event data

Param user_ctx [in] User registered context, passed from `i2s_channel_register_event_callback()`

Return Whether a high priority task has been waken up by this callback function

Enumerations

enum **i2s_port_t**

I2S controller port number, the max port number is (SOC_I2S_NUM -1).

Values:

enumerator **I2S_NUM_0**

I2S controller port 0

enumerator **I2S_NUM_1**

I2S controller port 1

enumerator **I2S_NUM_AUTO**

Select whichever port is available

enum **i2s_comm_mode_t**

I2S controller communication mode.

Values:

enumerator **I2S_COMM_MODE_STD**

I2S controller using standard communication mode, support Philips/MSB/PCM format

enumerator **I2S_COMM_MODE_PDM**

I2S controller using PDM communication mode, support PDM output or input

enumerator **I2S_COMM_MODE_TDM**

I2S controller using TDM communication mode, support up to 16 slots per frame

enumerator **I2S_COMM_MODE_NONE**

Unspecified I2S controller mode

enum **i2s_mclk_multiple_t**

The multiple of MCLK to sample rate.

Values:

enumerator **I2S_MCLK_MULTIPLE_128**

MCLK = sample_rate * 128

enumerator **I2S_MCLK_MULTIPLE_256**

MCLK = sample_rate * 256

enumerator **I2S_MCLK_MULTIPLE_384**

MCLK = sample_rate * 384

enumerator **I2S_MCLK_MULTIPLE_512**

MCLK = sample_rate * 512

Header File

- [components/hal/include/hal/i2s_types.h](#)
- This header file can be included with:

```
#include "hal/i2s_types.h"
```

Type Definitions

typedef [soc_periph_i2s_clk_src_t](#) **i2s_clock_src_t**

I2S clock source

Enumerations

enum **i2s_slot_mode_t**

I2S channel slot mode.

Values:

enumerator **I2S_SLOT_MODE_MONO**

I2S channel slot format mono, transmit same data in all slots for tx mode, only receive the data in the first slots for rx mode.

enumerator **I2S_SLOT_MODE_STEREO**

I2S channel slot format stereo, transmit different data in different slots for tx mode, receive the data in all slots for rx mode.

enum **i2s_dir_t**

I2S channel direction.

Values:

enumerator **I2S_DIR_RX**

I2S channel direction RX

enumerator **I2S_DIR_TX**

I2S channel direction TX

enum **i2s_role_t**

I2S controller role.

Values:

enumerator **I2S_ROLE_MASTER**

I2S controller master role, bclk and ws signal will be set to output

enumerator **I2S_ROLE_SLAVE**

I2S controller slave role, bclk and ws signal will be set to input

enum **i2s_data_bit_width_t**

Available data bit width in one slot.

Values:

enumerator **I2S_DATA_BIT_WIDTH_8BIT**

I2S channel data bit-width: 8

enumerator **I2S_DATA_BIT_WIDTH_16BIT**

I2S channel data bit-width: 16

enumerator **I2S_DATA_BIT_WIDTH_24BIT**

I2S channel data bit-width: 24

enumerator **I2S_DATA_BIT_WIDTH_32BIT**

I2S channel data bit-width: 32

enum **i2s_slot_bit_width_t**

Total slot bit width in one slot.

Values:

enumerator **I2S_SLOT_BIT_WIDTH_AUTO**

I2S channel slot bit-width equals to data bit-width

enumerator **I2S_SLOT_BIT_WIDTH_8BIT**

I2S channel slot bit-width: 8

enumerator **I2S_SLOT_BIT_WIDTH_16BIT**

I2S channel slot bit-width: 16

enumerator **I2S_SLOT_BIT_WIDTH_24BIT**

I2S channel slot bit-width: 24

enumerator **I2S_SLOT_BIT_WIDTH_32BIT**

I2S channel slot bit-width: 32

enum **i2s_pcm_compress_t**

A/U-law decompress or compress configuration.

Values:

enumerator **I2S_PCM_DISABLE**

Disable A/U law decompress or compress

enumerator **I2S_PCM_A_DECOMPRESS**

A-law decompress

enumerator **I2S_PCM_A_COMPRESS**

A-law compress

enumerator **I2S_PCM_U_DECOMPRESS**

U-law decompress

enumerator **I2S_PCM_U_COMPRESS**

U-law compress

enum **i2s_pdm_dsr_t**

I2S PDM RX down-sampling mode.

Values:

enumerator **I2S_PDM_DSR_8S**

downsampling number is 8 for PDM RX mode

enumerator **I2S_PDM_DSR_16S**

downsampling number is 16 for PDM RX mode

enumerator **I2S_PDM_DSR_MAX**

enum **i2s_pdm_sig_scale_t**

pdm tx signal scaling mode

Values:

enumerator **I2S_PDM_SIG_SCALING_DIV_2**

I2S TX PDM signal scaling: /2

enumerator **I2S_PDM_SIG_SCALING_MUL_1**

I2S TX PDM signal scaling: x1

enumerator **I2S_PDM_SIG_SCALING_MUL_2**

I2S TX PDM signal scaling: x2

enumerator **I2S_PDM_SIG_SCALING_MUL_4**

I2S TX PDM signal scaling: x4

enum **i2s_pdm_tx_line_mode_t**

PDM TX line mode.

Note: For the standard codec mode, PDM pins are connect to a codec which requires both clock signal and data signal For the DAC output mode, PDM data signal can be connected to a power amplifier directly with a low-pass filter, normally, DAC output mode doesn't need the clock signal.

Values:

enumerator **I2S_PDM_TX_ONE_LINE_CODEC**

Standard PDM format output, left and right slot data on a single line

enumerator **I2S_PDM_TX_ONE_LINE_DAC**

PDM DAC format output, left or right slot data on a single line

enumerator **I2S_PDM_TX_TWO_LINE_DAC**

PDM DAC format output, left and right slot data on separated lines

enum **i2s_std_slot_mask_t**

I2S slot select in standard mode.

Note: It has different meanings in tx/rx/mono/stereo mode, and it may have different behaviors on different targets. For the details, please refer to the I2S API reference

Values:

enumerator **I2S_STD_SLOT_LEFT**

I2S transmits or receives left slot

enumerator **I2S_STD_SLOT_RIGHT**

I2S transmits or receives right slot

enumerator **I2S_STD_SLOT_BOTH**

I2S transmits or receives both left and right slot

enum **i2s_pdm_slot_mask_t**

I2S slot select in PDM mode.

Values:

enumerator **I2S_PDM_SLOT_RIGHT**

I2S PDM only transmits or receives the PDM device whose 'select' pin is pulled up

enumerator **I2S_PDM_SLOT_LEFT**

I2S PDM only transmits or receives the PDM device whose 'select' pin is pulled down

enumerator **I2S_PDM_SLOT_BOTH**

I2S PDM transmits or receives both two slots

enumerator **I2S_PDM_RX_LINE0_SLOT_RIGHT**

I2S PDM receives the right slot on line 0

enumerator **I2S_PDM_RX_LINE0_SLOT_LEFT**

I2S PDM receives the left slot on line 0

enumerator **I2S_PDM_RX_LINE1_SLOT_RIGHT**

I2S PDM receives the right slot on line 1

enumerator **I2S_PDM_RX_LINE1_SLOT_LEFT**

I2S PDM receives the left slot on line 1

enumerator **I2S_PDM_RX_LINE2_SLOT_RIGHT**

I2S PDM receives the right slot on line 2

enumerator **I2S_PDM_RX_LINE2_SLOT_LEFT**

I2S PDM receives the left slot on line 2

enumerator **I2S_PDM_RX_LINE3_SLOT_RIGHT**

I2S PDM receives the right slot on line 3

enumerator **I2S_PDM_RX_LINE3_SLOT_LEFT**

I2S PDM receives the left slot on line 3

enumerator **I2S_PDM_LINE_SLOT_ALL**

I2S PDM receives all slots

enum **i2s_tdm_slot_mask_t**

tdm slot number

Note: Multiple slots in TDM mode. For TX module, only the active slot send the audio data, the inactive slot send a constant or will be skipped if 'skip_msk' is set. For RX module, only receive the audio data in active slots, the data in inactive slots will be ignored. the bit map of active slot can not exceed (0x1<<total_slot_num). e.g: slot_mask = (I2S_TDM_SLOT0 | I2S_TDM_SLOT3), here the active slot number is 2 and total_slot is not supposed to be smaller than 4.

Values:

enumerator **I2S_TDM_SLOT0**

I2S slot 0 enabled

enumerator **I2S_TDM_SLOT1**

I2S slot 1 enabled

enumerator **I2S_TDM_SLOT2**

I2S slot 2 enabled

enumerator **I2S_TDM_SLOT3**

I2S slot 3 enabled

enumerator **I2S_TDM_SLOT4**

I2S slot 4 enabled

enumerator **I2S_TDM_SLOT5**

I2S slot 5 enabled

enumerator **I2S_TDM_SLOT6**

I2S slot 6 enabled

enumerator **I2S_TDM_SLOT7**

I2S slot 7 enabled

enumerator **I2S_TDM_SLOT8**

I2S slot 8 enabled

enumerator **I2S_TDM_SLOT9**

I2S slot 9 enabled

enumerator **I2S_TDM_SLOT10**

I2S slot 10 enabled

enumerator **I2S_TDM_SLOT11**

I2S slot 11 enabled

enumerator **I2S_TDM_SLOT12**

I2S slot 12 enabled

enumerator **I2S_TDM_SLOT13**

I2S slot 13 enabled

enumerator **I2S_TDM_SLOT14**

I2S slot 14 enabled

enumerator **I2S_TDM_SLOT15**

I2S slot 15 enabled

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.5.11 LCD

Introduction

ESP chips can generate various kinds of timings that needed by common LCDs on the market, like SPI LCD, I80 LCD (a.k.a Intel 8080 parallel LCD), RGB/SRGB LCD, I2C LCD, etc. The `esp_lcd` component is officially to support those LCDs with a group of universal APIs across chips.

Functional Overview

In `esp_lcd`, an LCD panel is represented by `esp_lcd_panel_handle_t`, which plays the role of an **abstract frame buffer**, regardless of the frame memory is allocated inside ESP chip or in external LCD controller. Based on the location of the frame buffer and the hardware connection interface, the LCD panel drivers are mainly grouped into the following categories:

- Controller based LCD driver involves multiple steps to get a panel handle, like bus allocation, IO device registration and controller driver install. The frame buffer is located in the controller's internal GRAM (Graphical RAM). ESP-IDF provides only a limited number of LCD controller drivers out of the box (e.g., ST7789, SSD1306), *More Controller Based LCD Drivers* are maintained in the [Espressif Component Registry](#).

- [SPI Interfaced LCD](#) describes the steps to install the SPI LCD IO driver and then get the panel handle.
- [I2C Interfaced LCD](#) describes the steps to install the I2C LCD IO driver and then get the panel handle.
- [LCD Panel IO Operations](#) - provides a set of APIs to operate the LCD panel, like turning on/off the display, setting the orientation, etc. These operations are common for either controller-based LCD panel driver or RGB LCD panel driver.

SPI Interfaced LCD

1. Create an SPI bus. Please refer to [SPI Master API doc](#) for more details.

Currently the driver supports SPI, Quad SPI and Octal SPI (simulate Intel 8080 timing) modes.

```
spi_bus_config_t buscfg = {
    .sclk_io_num = EXAMPLE_PIN_NUM_SCLK,
    .mosi_io_num = EXAMPLE_PIN_NUM_MOSI,
    .miso_io_num = EXAMPLE_PIN_NUM_MISO,
    .quadwp_io_num = -1,
    .quadhd_io_num = -1,
    .max_transfer_sz = EXAMPLE_LCD_H_RES * 80 * sizeof(uint16_t), //_
    ↪transfer 80 lines of pixels (assume pixel is RGB565) at most in one_
    ↪SPI transaction
};
ESP_ERROR_CHECK(spi_bus_initialize(LCD_HOST, &buscfg, SPI_DMA_CH_
    ↪AUTO)); // Enable the DMA feature
```

2. Allocate an LCD IO device handle from the SPI bus. In this step, you need to provide the following information:

- [esp_lcd_panel_io_spi_config_t::dc_gpio_num](#): Sets the gpio number for the DC signal line (some LCD calls this RS line). The LCD driver uses this GPIO to switch between sending command and sending data.
- [esp_lcd_panel_io_spi_config_t::cs_gpio_num](#): Sets the gpio number for the CS signal line. The LCD driver uses this GPIO to select the LCD chip. If the SPI bus only has one device attached (i.e., this LCD), you can set the gpio number to -1 to occupy the bus exclusively.
- [esp_lcd_panel_io_spi_config_t::pclk_hz](#) sets the frequency of the pixel clock, in Hz. The value should not exceed the range recommended in the LCD spec.
- [esp_lcd_panel_io_spi_config_t::spi_mode](#) sets the SPI mode. The LCD driver uses this mode to communicate with the LCD. For the meaning of the SPI mode, please refer to the [SPI Master API doc](#).
- [esp_lcd_panel_io_spi_config_t::lcd_cmd_bits](#) and [esp_lcd_panel_io_spi_config_t::lcd_param_bits](#) set the bit width of the command and parameter that recognized by the LCD controller chip. This is chip specific, you should refer to your LCD spec in advance.
- [esp_lcd_panel_io_spi_config_t::trans_queue_depth](#) sets the depth of the SPI transaction queue. A bigger value means more transactions can be queued up, but it also consumes more memory.

```
esp_lcd_panel_io_handle_t io_handle = NULL;
esp_lcd_panel_io_spi_config_t io_config = {
    .dc_gpio_num = EXAMPLE_PIN_NUM_LCD_DC,
    .cs_gpio_num = EXAMPLE_PIN_NUM_LCD_CS,
    .pclk_hz = EXAMPLE_LCD_PIXEL_CLOCK_HZ,
    .lcd_cmd_bits = EXAMPLE_LCD_CMD_BITS,
    .lcd_param_bits = EXAMPLE_LCD_PARAM_BITS,
    .spi_mode = 0,
    .trans_queue_depth = 10,
};
// Attach the LCD to the SPI bus
ESP_ERROR_CHECK(esp_lcd_new_panel_io_spi((esp_lcd_spi_bus_handle_t)LCD_
    ↪HOST, &io_config, &io_handle));
```

3. Install the LCD controller driver. The LCD controller driver is responsible for sending the commands and parameters to the LCD controller chip. In this step, you need to specify the SPI IO device handle that allocated

in the last step, and some panel specific configurations:

- `esp_lcd_panel_dev_config_t::reset_gpio_num` sets the LCD's hardware reset GPIO number. If the LCD does not have a hardware reset pin, set this to -1.
- `esp_lcd_panel_dev_config_t::rgb_ele_order` sets the R-G-B element order of each color data.
- `esp_lcd_panel_dev_config_t::bits_per_pixel` sets the bit width of the pixel color data. The LCD driver uses this value to calculate the number of bytes to send to the LCD controller chip.
- `esp_lcd_panel_dev_config_t::data_endian` specifies the data endian to be transmitted to the screen. No need to specify for color data within 1 byte, like RGB232. For drivers that do not support specifying data endian, this field would be ignored.

```
esp_lcd_panel_handle_t panel_handle = NULL;
esp_lcd_panel_dev_config_t panel_config = {
    .reset_gpio_num = EXAMPLE_PIN_NUM_RST,
    .rgb_ele_order = LCD_RGB_ELEMENT_ORDER_BGR,
    .bits_per_pixel = 16,
};
// Create LCD panel handle for ST7789, with the SPI IO device handle
ESP_ERROR_CHECK(esp_lcd_new_panel_st7789(io_handle, &panel_config, &
↪panel_handle));
```

I2C Interfaced LCD

1. Create I2C bus. Please refer to [I2C API doc](#) for more details.

```
i2c_master_bus_handle_t i2c_bus = NULL;
i2c_master_bus_config_t bus_config = {
    .clk_source = I2C_CLK_SRC_DEFAULT,
    .glitch_ignore_cnt = 7,
    .i2c_port = I2C_BUS_PORT,
    .sda_io_num = EXAMPLE_PIN_NUM_SDA,
    .scl_io_num = EXAMPLE_PIN_NUM_SCL,
    .flags.enable_internal_pullup = true,
};
ESP_ERROR_CHECK(i2c_new_master_bus(&bus_config, &i2c_bus));
```

2. Allocate an LCD IO device handle from the I2C bus. In this step, you need to provide the following information:

- `esp_lcd_panel_io_i2c_config_t::dev_addr` sets the I2C device address of the LCD controller chip. The LCD driver uses this address to communicate with the LCD controller chip.
- `esp_lcd_panel_io_i2c_config_t::scl_speed_hz` sets the I2C clock frequency in Hz. The value should not exceed the range recommended in the LCD spec.
- `esp_lcd_panel_io_i2c_config_t::lcd_cmd_bits` and `esp_lcd_panel_io_i2c_config_t::lcd_param_bits` set the bit width of the command and parameter that recognized by the LCD controller chip. This is chip specific, you should refer to your LCD spec in advance.

```
esp_lcd_panel_io_handle_t io_handle = NULL;
esp_lcd_panel_io_i2c_config_t io_config = {
    .dev_addr = EXAMPLE_I2C_HW_ADDR,
    .scl_speed_hz = EXAMPLE_LCD_PIXEL_CLOCK_HZ,
    .control_phase_bytes = 1, // refer to LCD spec
    .dc_bit_offset = 6, // refer to LCD spec
    .lcd_cmd_bits = EXAMPLE_LCD_CMD_BITS,
    .lcd_param_bits = EXAMPLE_LCD_CMD_BITS,
};
ESP_ERROR_CHECK(esp_lcd_new_panel_io_i2c(i2c_bus, &io_config, &io_
↪handle));
```

3. Install the LCD controller driver. The LCD controller driver is responsible for sending the commands and parameters to the LCD controller chip. In this step, you need to specify the I2C IO device handle that allocated in the last step, and some panel specific configurations:

- `esp_lcd_panel_dev_config_t::reset_gpio_num` sets the LCD's hardware reset GPIO number. If the LCD does not have a hardware reset pin, set this to `-1`.
- `esp_lcd_panel_dev_config_t::bits_per_pixel` sets the bit width of the pixel color data. The LCD driver uses this value to calculate the number of bytes to send to the LCD controller chip.

```
esp_lcd_panel_handle_t panel_handle = NULL;
esp_lcd_panel_dev_config_t panel_config = {
    .bits_per_pixel = 1,
    .reset_gpio_num = EXAMPLE_PIN_NUM_RST,
};
ESP_ERROR_CHECK(esp_lcd_new_panel_ssd1306(io_handle, &panel_config, &
    ↪panel_handle));
```

More Controller Based LCD Drivers

More LCD panel drivers and touch drivers are available in [ESP-IDF Component Registry](#). The list of available and planned drivers with links is in this [table](#).

LCD Panel IO Operations

- `esp_lcd_panel_reset()` can reset the LCD panel.
- `esp_lcd_panel_init()` performs a basic initialization of the panel. To perform more manufacture specific initialization, please go to [Steps to Add Manufacture Specific Initialization](#).
- Through combined use of `esp_lcd_panel_swap_xy()` and `esp_lcd_panel_mirror()`, you can rotate the LCD screen.
- `esp_lcd_panel_disp_on_off()` can turn on or off the LCD screen by cutting down the output path from the frame buffer to the LCD screen.
- `esp_lcd_panel_disp_sleep()` can reduce the power consumption of the LCD screen by entering the sleep mode. The internal frame buffer is still retained.
- `esp_lcd_panel_draw_bitmap()` is the most significant function, which does the magic to draw the user provided color buffer to the LCD screen, where the draw window is also configurable.

Steps to Add Manufacture Specific Initialization

The LCD controller drivers (e.g., `st7789`) in `esp-idf` only provide basic initialization in the `esp_lcd_panel_init()`, leaving the vast majority of settings to the default values. Some LCD modules need to set a bunch of manufacture specific configurations before it can display normally. These configurations usually include gamma, power voltage and so on. If you want to add manufacture specific initialization, please follow the steps below:

```
esp_lcd_panel_reset(panel_handle);
esp_lcd_panel_init(panel_handle);
// set extra configurations e.g., gamma control
// with the underlying IO handle
// please consult your manufacture for special commands and corresponding values
esp_lcd_panel_io_tx_param(io_handle, GAMMA_CMD, (uint8_t[]) {
    GAMMA_ARRAY
}, N);
// turn on the display
esp_lcd_panel_disp_on_off(panel_handle, true);
```

Application Example

LCD examples are located under: [peripherals/lcd](#):

- Universal SPI LCD example with SPI touch - [peripherals/lcd/spi_lcd_touch](#)
- Jpeg decoding and LCD display - [peripherals/lcd/tjpgd](#)
- I2C interfaced OLED display scrolling text - [peripherals/lcd/i2c_oled](#)

API Reference

Header File

- [components/hal/include/hal/lcd_types.h](#)
- This header file can be included with:

```
#include "hal/lcd_types.h"
```

Enumerations

enum **lcd_rgb_element_order_t**

RGB color endian.

Values:

enumerator **LCD_RGB_ELEMENT_ORDER_RGB**

RGB element order: RGB

enumerator **LCD_RGB_ELEMENT_ORDER_BGR**

RGB element order: BGR

enum **lcd_rgb_data_endian_t**

RGB data endian.

Values:

enumerator **LCD_RGB_DATA_ENDIAN_BIG**

RGB data endian: MSB first

enumerator **LCD_RGB_DATA_ENDIAN_LITTLE**

RGB data endian: LSB first

enum **lcd_color_space_t**

LCD color space.

Values:

enumerator **LCD_COLOR_SPACE_RGB**

Color space: RGB

enumerator **LCD_COLOR_SPACE_YUV**

Color space: YUV

enum **lcd_color_range_t**

LCD color range.

Values:

enumerator **LCD_COLOR_RANGE_LIMIT**

Limited color range

enumerator **LCD_COLOR_RANGE_FULL**

Full color range

enum **lcd_yuv_sample_t**

YUV sampling method.

Values:

enumerator **LCD_YUV_SAMPLE_422**

YUV 4:2:2 sampling

enumerator **LCD_YUV_SAMPLE_420**

YUV 4:2:0 sampling

enumerator **LCD_YUV_SAMPLE_411**

YUV 4:1:1 sampling

enum **lcd_yuv_conv_std_t**

The standard used for conversion between RGB and YUV.

Values:

enumerator **LCD_YUV_CONV_STD_BT601**

YUV<->RGB conversion standard: BT.601

enumerator **LCD_YUV_CONV_STD_BT709**

YUV<->RGB conversion standard: BT.709

Header File

- [components/esp_lcd/include/esp_lcd_types.h](#)
- This header file can be included with:

```
#include "esp_lcd_types.h"
```

- This header file is a part of the API provided by the `esp_lcd` component. To declare that your component depends on `esp_lcd`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_lcd
```

or

```
PRIV_REQUIRES esp_lcd
```


Type Definitions

typedef struct esp_lcd_panel_io_t ***esp_lcd_panel_io_handle_t**

Type of LCD panel IO handle

typedef struct esp_lcd_panel_t ***esp_lcd_panel_handle_t**

Type of LCD panel handle

Header File

- [components/esp_lcd/include/esp_lcd_panel_io.h](#)
- This header file can be included with:

```
#include "esp_lcd_panel_io.h"
```

- This header file is a part of the API provided by the `esp_lcd` component. To declare that your component depends on `esp_lcd`, add the following to your CMakeLists.txt:

```
REQUIRES esp_lcd
```

or

```
PRIV_REQUIRES esp_lcd
```

Functions

esp_err_t esp_lcd_panel_io_rx_param (*esp_lcd_panel_io_handle_t* io, int lcd_cmd, void *param, size_t param_size)

Transmit LCD command and receive corresponding parameters.

Note: Commands sent by this function are short, so they are sent using polling transactions. The function does not return before the command transfer is completed. If any queued transactions sent by `esp_lcd_panel_io_tx_color()` are still pending when this function is called, this function will wait until they are finished and the queue is empty before sending the command(s).

Parameters

- **io** -- [in] LCD panel IO handle, which is created by other factory API like `esp_lcd_new_panel_io_spi()`
- **lcd_cmd** -- [in] The specific LCD command, set to -1 if no command needed
- **param** -- [out] Buffer for the command data
- **param_size** -- [in] Size of param buffer

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_ERR_NOT_SUPPORTED if read is not supported by transport
- ESP_OK on success

esp_err_t esp_lcd_panel_io_tx_param (*esp_lcd_panel_io_handle_t* io, int lcd_cmd, const void *param, size_t param_size)

Transmit LCD command and corresponding parameters.

Note: Commands sent by this function are short, so they are sent using polling transactions. The function does not return before the command transfer is completed. If any queued transactions sent by `esp_lcd_panel_io_tx_color()` are still pending when this function is called, this function will wait until they are finished and the queue is empty before sending the command(s).

Parameters

- **io** -- [in] LCD panel IO handle, which is created by other factory API like `esp_lcd_new_panel_io_spi()`
- **lcd_cmd** -- [in] The specific LCD command, set to -1 if no command needed
- **param** -- [in] Buffer that holds the command specific parameters, set to NULL if no parameter is needed for the command
- **param_size** -- [in] Size of `param` in memory, in bytes, set to zero if no parameter is needed for the command

Returns

- `ESP_ERR_INVALID_ARG` if parameter is invalid
- `ESP_OK` on success

esp_err_t **esp_lcd_panel_io_tx_color** (*esp_lcd_panel_io_handle_t* io, int lcd_cmd, const void *color, size_t color_size)

Transmit LCD RGB data.

Note: This function will package the command and RGB data into a transaction, and push into a queue. The real transmission is performed in the background (DMA+interrupt). The caller should take care of the lifecycle of the `color` buffer. Recycling of color buffer should be done in the callback `on_color_trans_done()`.

Parameters

- **io** -- [in] LCD panel IO handle, which is created by factory API like `esp_lcd_new_panel_io_spi()`
- **lcd_cmd** -- [in] The specific LCD command, set to -1 if no command needed
- **color** -- [in] Buffer that holds the RGB color data
- **color_size** -- [in] Size of `color` in memory, in bytes

Returns

- `ESP_ERR_INVALID_ARG` if parameter is invalid
- `ESP_OK` on success

esp_err_t **esp_lcd_panel_io_del** (*esp_lcd_panel_io_handle_t* io)

Destroy LCD panel IO handle (deinitialize panel and free all corresponding resource)

Parameters **io** -- [in] LCD panel IO handle, which is created by factory API like `esp_lcd_new_panel_io_spi()`

Returns

- `ESP_ERR_INVALID_ARG` if parameter is invalid
- `ESP_OK` on success

esp_err_t **esp_lcd_panel_io_register_event_callbacks** (*esp_lcd_panel_io_handle_t* io, const *esp_lcd_panel_io_callbacks_t* *cbs, void *user_ctx)

Register LCD panel IO callbacks.

Parameters

- **io** -- [in] LCD panel IO handle, which is created by factory API like `esp_lcd_new_panel_io_spi()`
- **cbs** -- [in] structure with all LCD panel IO callbacks
- **user_ctx** -- [in] User private data, passed directly to callback's `user_ctx`

Returns

- `ESP_ERR_INVALID_ARG` if parameter is invalid
- `ESP_OK` on success

esp_err_t **esp_lcd_new_panel_io_spi** (*esp_lcd_spi_bus_handle_t* bus, const *esp_lcd_panel_io_spi_config_t* *io_config, *esp_lcd_panel_io_handle_t* *ret_io)

Create LCD panel IO handle, for SPI interface.

Parameters

- **bus** -- [in] SPI bus handle

- **io_config** -- **[in]** IO configuration, for SPI interface
- **ret_io** -- **[out]** Returned IO handle

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_ERR_NO_MEM if out of memory
- ESP_OK on success

esp_err_t **esp_lcd_new_panel_io_i2c_v1** (uint32_t bus, const *esp_lcd_panel_io_i2c_config_t* *io_config, *esp_lcd_panel_io_handle_t* *ret_io)

Create LCD panel IO handle, for I2C interface in legacy implementation.

Note: Please don't call this function in your project directly. Please call `esp_lcd_new_panel_to_i2c` instead.

Parameters

- **bus** -- **[in]** I2C bus handle, (in uint32_t)
- **io_config** -- **[in]** IO configuration, for I2C interface
- **ret_io** -- **[out]** Returned IO handle

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_ERR_NO_MEM if out of memory
- ESP_OK on success

esp_err_t **esp_lcd_new_panel_io_i2c_v2** (*i2c_master_bus_handle_t* bus, const *esp_lcd_panel_io_i2c_config_t* *io_config, *esp_lcd_panel_io_handle_t* *ret_io)

Create LCD panel IO handle, for I2C interface in new implementation.

Note: Please don't call this function in your project directly. Please call `esp_lcd_new_panel_to_i2c` instead.

Parameters

- **bus** -- **[in]** I2C bus handle, (in *i2c_master_dev_handle_t*)
- **io_config** -- **[in]** IO configuration, for I2C interface
- **ret_io** -- **[out]** Returned IO handle

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_ERR_NO_MEM if out of memory
- ESP_OK on success

Structures

struct **esp_lcd_panel_io_event_data_t**

Type of LCD panel IO event data.

struct **esp_lcd_panel_io_callbacks_t**

Type of LCD panel IO callbacks.

Public Members

esp_lcd_panel_io_color_trans_done_cb_t **on_color_trans_done**

Callback invoked when color data transfer has finished

struct **esp_lcd_panel_io_spi_config_t**

Panel IO configuration structure, for SPI interface.

Public Members

int **cs_gpio_num**

GPIO used for CS line

int **dc_gpio_num**

GPIO used to select the D/C line, set this to -1 if the D/C line is not used

int **spi_mode**

Traditional SPI mode (0~3)

unsigned int **pclk_hz**

Frequency of pixel clock

size_t **trans_queue_depth**

Size of internal transaction queue

[*esp_lcd_panel_io_color_trans_done_cb_t*](#) **on_color_trans_done**

Callback invoked when color data transfer has finished

void ***user_ctx**

User private data, passed directly to on_color_trans_done's user_ctx

int **lcd_cmd_bits**

Bit-width of LCD command

int **lcd_param_bits**

Bit-width of LCD parameter

unsigned int **dc_high_on_cmd**

If enabled, DC level = 1 indicates command transfer

unsigned int **dc_low_on_data**

If enabled, DC level = 0 indicates color data transfer

unsigned int **dc_low_on_param**

If enabled, DC level = 0 indicates parameter transfer

unsigned int **octal_mode**

transmit with octal mode (8 data lines), this mode is used to simulate Intel 8080 timing

unsigned int **quad_mode**

transmit with quad mode (4 data lines), this mode is useful when transmitting LCD parameters (Only use one line for command)

unsigned int **sio_mode**

Read and write through a single data line (MOSI)

unsigned int **lsb_first**

transmit LSB bit first

unsigned int **cs_high_active**

CS line is high active

struct *esp_lcd_panel_io_spi_config_t*::[anonymous] **flags**

Extra flags to fine-tune the SPI device

struct **esp_lcd_panel_io_i2c_config_t**

Panel IO configuration structure, for I2C interface.

Public Members

uint32_t **dev_addr**

I2C device address

esp_lcd_panel_io_color_trans_done_cb_t **on_color_trans_done**

Callback invoked when color data transfer has finished

void ***user_ctx**

User private data, passed directly to on_color_trans_done's user_ctx

size_t **control_phase_bytes**

I2C LCD panel will encode control information (e.g. D/C selection) into control phase, in several bytes

unsigned int **dc_bit_offset**

Offset of the D/C selection bit in control phase

int **lcd_cmd_bits**

Bit-width of LCD command

int **lcd_param_bits**

Bit-width of LCD parameter

unsigned int **dc_low_on_data**

If this flag is enabled, DC line = 0 means transfer data, DC line = 1 means transfer command; vice versa

unsigned int **disable_control_phase**

If this flag is enabled, the control phase isn't used

struct *esp_lcd_panel_io_i2c_config_t*::[anonymous] **flags**

Extra flags to fine-tune the I2C device

uint32_t **scl_speed_hz**

I2C LCD SCL frequency (hz)

Macros

esp_lcd_new_panel_io_i2c (bus, io_config, ret_io)

Create LCD panel IO handle.

Parameters

- **bus** -- [in] I2C bus handle
- **io_config** -- [in] IO configuration, for I2C interface
- **ret_io** -- [out] Returned IO handle

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_ERR_NO_MEM if out of memory
- ESP_OK on success

Type Definitions

typedef void ***esp_lcd_spi_bus_handle_t**

Type of LCD SPI bus handle

typedef uint32_t **esp_lcd_i2c_bus_handle_t**

Type of LCD I2C bus handle

typedef struct esp_lcd_i80_bus_t ***esp_lcd_i80_bus_handle_t**

Type of LCD intel 8080 bus handle

typedef bool (***esp_lcd_panel_io_color_trans_done_cb_t**)([esp_lcd_panel_io_handle_t](#) panel_io, [esp_lcd_panel_io_event_data_t](#) *edata, void *user_ctx)

Declare the prototype of the function that will be invoked when panel IO finishes transferring color data.

Param panel_io [in] LCD panel IO handle, which is created by factory API like `esp_lcd_new_panel_io_spi()`

Param edata [in] Panel IO event data, fed by driver

Param user_ctx [in] User data, passed from `esp_lcd_panel_io_xxx_config_t`

Return Whether a high priority task has been waken up by this function

Header File

- [components/esp_lcd/include/esp_lcd_panel_ops.h](#)
- This header file can be included with:

```
#include "esp_lcd_panel_ops.h"
```

- This header file is a part of the API provided by the `esp_lcd` component. To declare that your component depends on `esp_lcd`, add the following to your CMakeLists.txt:

```
REQUIRES esp_lcd
```

or

```
PRIV_REQUIRES esp_lcd
```

Functions

[esp_err_t](#) **esp_lcd_panel_reset** ([esp_lcd_panel_handle_t](#) panel)

Reset LCD panel.

Note: Panel reset must be called before attempting to initialize the panel using `esp_lcd_panel_init()`.

Parameters **panel** -- [in] LCD panel handle, which is created by other factory API like `esp_lcd_new_panel_st7789()`

Returns

- ESP_OK on success

esp_err_t **esp_lcd_panel_init** (*esp_lcd_panel_handle_t* panel)

Initialize LCD panel.

Note: Before calling this function, make sure the LCD panel has finished the `reset` stage by `esp_lcd_panel_reset()`.

Parameters **panel** -- [in] LCD panel handle, which is created by other factory API like `esp_lcd_new_panel_st7789()`

Returns

- ESP_OK on success

esp_err_t **esp_lcd_panel_del** (*esp_lcd_panel_handle_t* panel)

Deinitialize the LCD panel.

Parameters **panel** -- [in] LCD panel handle, which is created by other factory API like `esp_lcd_new_panel_st7789()`

Returns

- ESP_OK on success

esp_err_t **esp_lcd_panel_draw_bitmap** (*esp_lcd_panel_handle_t* panel, int x_start, int y_start, int x_end, int y_end, const void *color_data)

Draw bitmap on LCD panel.

Parameters

- **panel** -- [in] LCD panel handle, which is created by other factory API like `esp_lcd_new_panel_st7789()`
- **x_start** -- [in] Start index on x-axis (x_start included)
- **y_start** -- [in] Start index on y-axis (y_start included)
- **x_end** -- [in] End index on x-axis (x_end not included)
- **y_end** -- [in] End index on y-axis (y_end not included)
- **color_data** -- [in] RGB color data that will be dumped to the specific window range

Returns

- ESP_OK on success

esp_err_t **esp_lcd_panel_mirror** (*esp_lcd_panel_handle_t* panel, bool mirror_x, bool mirror_y)

Mirror the LCD panel on specific axis.

Note: Combined with `esp_lcd_panel_swap_xy()`, one can realize screen rotation

Parameters

- **panel** -- [in] LCD panel handle, which is created by other factory API like `esp_lcd_new_panel_st7789()`
- **mirror_x** -- [in] Whether the panel will be mirrored about the x axis
- **mirror_y** -- [in] Whether the panel will be mirrored about the y axis

Returns

- ESP_OK on success
- ESP_ERR_NOT_SUPPORTED if this function is not supported by the panel

esp_err_t **esp_lcd_panel_swap_xy** (*esp_lcd_panel_handle_t* panel, bool swap_axes)

Swap/Exchange x and y axis.

Note: Combined with `esp_lcd_panel_mirror()`, one can realize screen rotation

Parameters

- **panel** -- [in] LCD panel handle, which is created by other factory API like `esp_lcd_new_panel_st7789()`
- **swap_axes** -- [in] Whether to swap the x and y axis

Returns

- ESP_OK on success
- ESP_ERR_NOT_SUPPORTED if this function is not supported by the panel

esp_err_t **esp_lcd_panel_set_gap** (*esp_lcd_panel_handle_t* panel, int x_gap, int y_gap)

Set extra gap in x and y axis.

The gap is the space (in pixels) between the left/top sides of the LCD panel and the first row/column respectively of the actual contents displayed.

Note: Setting a gap is useful when positioning or centering a frame that is smaller than the LCD.

Parameters

- **panel** -- [in] LCD panel handle, which is created by other factory API like `esp_lcd_new_panel_st7789()`
- **x_gap** -- [in] Extra gap on x axis, in pixels
- **y_gap** -- [in] Extra gap on y axis, in pixels

Returns

- ESP_OK on success

esp_err_t **esp_lcd_panel_invert_color** (*esp_lcd_panel_handle_t* panel, bool invert_color_data)

Invert the color (bit-wise invert the color data line)

Parameters

- **panel** -- [in] LCD panel handle, which is created by other factory API like `esp_lcd_new_panel_st7789()`
- **invert_color_data** -- [in] Whether to invert the color data

Returns

- ESP_OK on success

esp_err_t **esp_lcd_panel_disp_on_off** (*esp_lcd_panel_handle_t* panel, bool on_off)

Turn on or off the display.

Parameters

- **panel** -- [in] LCD panel handle, which is created by other factory API like `esp_lcd_new_panel_st7789()`
- **on_off** -- [in] True to turns on display, False to turns off display

Returns

- ESP_OK on success
- ESP_ERR_NOT_SUPPORTED if this function is not supported by the panel

esp_err_t **esp_lcd_panel_disp_off** (*esp_lcd_panel_handle_t* panel, bool off)

Turn off the display.

Parameters

- **panel** -- [in] LCD panel handle, which is created by other factory API like `esp_lcd_new_panel_st7789()`
- **off** -- [in] Whether to turn off the screen

Returns

- ESP_OK on success
- ESP_ERR_NOT_SUPPORTED if this function is not supported by the panel

esp_err_t **esp_lcd_panel_disp_sleep** (*esp_lcd_panel_handle_t* panel, bool sleep)

Enter or exit sleep mode.

Parameters

- **panel** -- [in] LCD panel handle, which is created by other factory API like `esp_lcd_new_panel_st7789()`
- **sleep** -- [in] True to enter sleep mode, False to wake up

Returns

- ESP_OK on success
- ESP_ERR_NOT_SUPPORTED if this function is not supported by the panel

Header File

- [components/esp_lcd/include/esp_lcd_panel_rgb.h](#)
- This header file can be included with:

```
#include "esp_lcd_panel_rgb.h"
```

- This header file is a part of the API provided by the `esp_lcd` component. To declare that your component depends on `esp_lcd`, add the following to your CMakeLists.txt:

```
REQUIRES esp_lcd
```

or

```
PRIV_REQUIRES esp_lcd
```

Header File

- [components/esp_lcd/include/esp_lcd_panel_vendor.h](#)
- This header file can be included with:

```
#include "esp_lcd_panel_vendor.h"
```

- This header file is a part of the API provided by the `esp_lcd` component. To declare that your component depends on `esp_lcd`, add the following to your CMakeLists.txt:

```
REQUIRES esp_lcd
```

or

```
PRIV_REQUIRES esp_lcd
```

Functions

esp_err_t **esp_lcd_new_panel_st7789** (const *esp_lcd_panel_io_handle_t* io, const *esp_lcd_panel_dev_config_t* *panel_dev_config, *esp_lcd_panel_handle_t* *ret_panel)

Create LCD panel for model ST7789.

Parameters

- **io** -- [in] LCD panel IO handle
- **panel_dev_config** -- [in] general panel device configuration
- **ret_panel** -- [out] Returned LCD panel handle

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_ERR_NO_MEM if out of memory
- ESP_OK on success

esp_err_t **esp_lcd_new_panel_nt35510** (const *esp_lcd_panel_io_handle_t* io, const *esp_lcd_panel_dev_config_t* *panel_dev_config, *esp_lcd_panel_handle_t* *ret_panel)

Create LCD panel for model NT35510.

Parameters

- **io** -- **[in]** LCD panel IO handle
- **panel_dev_config** -- **[in]** general panel device configuration
- **ret_panel** -- **[out]** Returned LCD panel handle

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_ERR_NO_MEM if out of memory
- ESP_OK on success

```
esp_err_t esp_lcd_new_panel_ssd1306 (const esp_lcd_panel_io_handle_t io, const  
                                     esp_lcd_panel_dev_config_t *panel_dev_config,  
                                     esp_lcd_panel_handle_t *ret_panel)
```

Create LCD panel for model SSD1306.

Parameters

- **io** -- **[in]** LCD panel IO handle
- **panel_dev_config** -- **[in]** general panel device configuration
- **ret_panel** -- **[out]** Returned LCD panel handle

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_ERR_NO_MEM if out of memory
- ESP_OK on success

Structures

struct **esp_lcd_panel_dev_config_t**

Configuration structure for panel device.

Public Members

int **reset_gpio_num**

GPIO used to reset the LCD panel, set to -1 if it's not used

esp_lcd_color_space_t **color_space**

Deprecated:

Set RGB color space, please use *rgb_ele_order* instead

lcd_color_rgb_endian_t **rgb_endian**

Deprecated:

Set RGB data endian, please use *rgb_ele_order* instead

lcd_rgb_element_order_t **rgb_ele_order**

Set RGB element order, RGB or BGR

lcd_rgb_data_endian_t **data_endian**

Set the data endian for color data larger than 1 byte

unsigned int **bits_per_pixel**

Color depth, in bpp

unsigned int **reset_active_high**

Setting this if the panel reset is high level active

struct *esp_lcd_panel_dev_config_t*::[anonymous] **flags**

LCD panel config flags

void ***vendor_config**

vendor specific configuration, optional, left as NULL if not used

2.5.12 LED Control (LEDC)

Introduction

The LED control (LEDC) peripheral is primarily designed to control the intensity of LEDs, although it can also be used to generate PWM signals for other purposes. It has 8 channels which can generate independent waveforms that can be used, for example, to drive RGB LED devices.

The PWM controller can automatically increase or decrease the duty cycle gradually, allowing for fades without any processor interference.

Functionality Overview

Setting up a channel of the LEDC is done in three steps. Note that unlike ESP32, ESP32-P4 only supports configuring channels in "low speed" mode.

1. *Timer Configuration* by specifying the PWM signal's frequency and duty cycle resolution.
2. *Channel Configuration* by associating it with the timer and GPIO to output the PWM signal.
3. *Change PWM Signal* that drives the output in order to change LED's intensity. This can be done under the full control of software or with hardware fading functions.

As an optional step, it is also possible to set up an interrupt on fade end.




Fig. 13: Key Settings of LED PWM Controller's API

Note: For an initial setup, it is recommended to configure for the timers first (by calling `ledc_timer_config()`), and then for the channels (by calling `ledc_channel_config()`). This ensures the PWM frequency is at the desired value since the appearance of the PWM signal from the IO pad.

Timer Configuration Setting the timer is done by calling the function `ledc_timer_config()` and passing the data structure `ledc_timer_config_t` that contains the following configuration settings:

- Speed mode (value must be `LEDC_LOW_SPEED_MODE`)
- Timer number `ledc_timer_t`
- PWM signal frequency in Hz
- Resolution of PWM duty
- Source clock `ledc_clk_cfg_t`

The frequency and the duty resolution are interdependent. The higher the PWM frequency, the lower the duty resolution which is available, and vice versa. This relationship might be important if you are planning to use this API for purposes other than changing the intensity of LEDs. For more details, see Section [Supported Range of Frequency and Duty Resolutions](#).

The source clock can also limit the PWM frequency. The higher the source clock frequency, the higher the maximum PWM frequency can be configured.

Table 3: Characteristics of ESP32-P4 LEDC source clocks

Clock name	Clock freq	Clock capabilities
PLL_80M_CLK	80 MHz	/
RC_FAST_CLK	~ 20 MHz	Dynamic Frequency Scaling compatible, Light sleep compatible
XTAL_CLK	40 MHz	Dynamic Frequency Scaling compatible

Note:

1. On ESP32-P4, if `RC_FAST_CLK` is chosen as the LEDC clock source, you may see the frequency of output PWM signal is not very accurate. This is because no internal calibration is performed to get the exact frequency of the clock due to hardware limitation, a theoretic frequency value is used.
 2. For ESP32-P4, all timers share one clock source. In other words, it is impossible to use different clock sources for different timers.
-

The LEDC driver offers a helper function `ledc_find_suitable_duty_resolution()` to find the maximum possible resolution for the timer, given the source clock frequency and the desired PWM signal frequency.

When a timer is no longer needed by any channel, it can be deconfigured by calling the same function `ledc_timer_config()`. The configuration structure `ledc_timer_config_t` passes in should be:

- `ledc_timer_config_t::speed_mode` The speed mode of the timer which wants to be deconfigured belongs to (`ledc_mode_t`)
- `ledc_timer_config_t::timer_num` The ID of the timers which wants to be deconfigured (`ledc_timer_t`)
- `ledc_timer_config_t::deconfigure` Set this to true so that the timer specified can be deconfigured

Channel Configuration When the timer is set up, configure the desired channel (one out of `ledc_channel_t`). This is done by calling the function `ledc_channel_config()`.

Similar to the timer configuration, the channel setup function should be passed a structure `ledc_channel_config_t` that contains the channel's configuration parameters.

At this point, the channel should start operating and generating the PWM signal on the selected GPIO, as configured in `ledc_channel_config_t`, with the frequency specified in the timer settings and the given duty cycle. The channel operation (signal generation) can be suspended at any time by calling the function `ledc_stop()`.

Change PWM Signal Once the channel starts operating and generating the PWM signal with the constant duty cycle and frequency, there are a couple of ways to change this signal. When driving LEDs, primarily the duty cycle is changed to vary the light intensity.

The following two sections describe how to change the duty cycle using software and hardware fading. If required, the signal's frequency can also be changed; it is covered in Section [Change PWM Frequency](#).

Note: All the timers and channels in the ESP32-P4's LED PWM Controller only support low speed mode. Any change of PWM settings must be explicitly triggered by software (see below).

Change PWM Duty Cycle Using Software To set the duty cycle, use the dedicated function `ledc_set_duty()`. After that, call `ledc_update_duty()` to activate the changes. To check the currently set value, use the corresponding `_get_` function `ledc_get_duty()`.

Another way to set the duty cycle, as well as some other channel parameters, is by calling `ledc_channel_config()` covered in Section [Channel Configuration](#).

The range of the duty cycle values passed to functions depends on selected `duty_resolution` and should be from 0 to $(2 ** \text{duty\_resolution})$. For example, if the selected duty resolution is 10, then the duty cycle values can range from 0 to 1024. This provides the resolution of ~ 0.1%.

Warning: On ESP32-P4, when channel's binded timer selects its maximum duty resolution, the duty cycle value cannot be set to $(2 ** \text{duty\_resolution})$. Otherwise, the internal duty counter in the hardware will overflow and be messed up.

Change PWM Duty Cycle Using Hardware The LEDC hardware provides the means to gradually transition from one duty cycle value to another. To use this functionality, enable fading with `ledc_fade_func_install()` and then configure it by calling one of the available fading functions:

- `ledc_set_fade_with_time()`
- `ledc_set_fade_with_step()`
- `ledc_set_fade()`

On ESP32-P4, the hardware additionally allows to perform up to 16 consecutive linear fades without CPU intervention. This feature can be useful if you want to do a fade with gamma correction.

The luminance perceived by human eyes does not have a linear relationship with the PWM duty cycle. In order to make human feel the LED is dimming or lightening linearly, the change in duty cycle should be non-linear, which is the so-called gamma correction. The LED controller can simulate a gamma curve fading by piecewise linear approximation. `ledc_fill_multi_fade_param_list()` is a function that can help to construct the parameters for the piecewise linear fades. First, you need to allocate a memory block for saving the fade parameters, then by providing start/end PWM duty cycle values, gamma correction function, and the total number of desired linear segments to the helper function, it will fill the calculation results into the allocated space. You can also construct the array of `ledc_fade_param_config_t` manually. Once the fade parameter structs are prepared, a consecutive fading can be configured by passing the pointer to the prepared `ledc_fade_param_config_t` list and the total number of fade ranges to `ledc_set_multi_fade()`.

Start fading with `ledc_fade_start()`. A fade can be operated in blocking or non-blocking mode, please check `ledc_fade_mode_t` for the difference between the two available fade modes. Note that with either fade mode, the next fade or fixed-duty update will not take effect until the last fade finishes or is stopped. `ledc_fade_stop()` has to be called to stop a fade that is in progress.

To get a notification about the completion of a fade operation, a fade end callback function can be registered for each channel by calling `ledc_cb_register()` after the fade service being installed. The fade end callback prototype is defined in `ledc_cb_t`, where you should return a boolean value from the callback function, indicating whether a high priority task is woken up by this callback function. It is worth mentioning, the callback and the function invoked by itself should be placed in IRAM, as the interrupt service routine is in IRAM. `ledc_cb_register()` will print a warning message if it finds the addresses of callback and user context are incorrect.

If not required anymore, fading and an associated interrupt can be disabled with `ledc_fade_func_uninstall()`.

Change PWM Frequency The LEDC API provides several ways to change the PWM frequency "on the fly":

- Set the frequency by calling `ledc_set_freq()`. There is a corresponding function `ledc_get_freq()` to check the current frequency.
- Change the frequency and the duty resolution by calling `ledc_bind_channel_timer()` to bind some other timer to the channel.
- Change the channel's timer by calling `ledc_channel_config()`.

More Control Over PWM There are several individual timer-specific functions that can be used to change PWM output:

- `ledc_timer_set()`
- `ledc_timer_rst()`
- `ledc_timer_pause()`
- `ledc_timer_resume()`

The first two functions are called "behind the scenes" by `ledc_timer_config()` to provide a startup of a timer after it is configured.

Use Interrupts When configuring an LEDC channel, one of the parameters selected within `ledc_channel_config_t` is `ledc_intr_type_t` which triggers an interrupt on fade completion.

For registration of a handler to address this interrupt, call `ledc_isr_register()`.

Supported Range of Frequency and Duty Resolutions

The LED PWM Controller is designed primarily to drive LEDs. It provides a large flexibility of PWM duty cycle settings. For instance, the PWM frequency of 5 kHz can have the maximum duty resolution of 13 bits. This means that the duty can be set anywhere from 0 to 100% with a resolution of $\sim 0.012\%$ ($2^{13} = 8192$ discrete levels of the LED intensity). Note, however, that these parameters depend on the clock signal clocking the LED PWM Controller timer which in turn clocks the channel (see [timer configuration](#) and the [ESP32-P4 Technical Reference Manual > LED PWM Controller \(LEDC\) \[PDF\]](#)).

The LEDC can be used for generating signals at much higher frequencies that are sufficient enough to clock other devices, e.g., a digital camera module. In this case, the maximum available frequency is 40 MHz with duty resolution of 1 bit. This means that the duty cycle is fixed at 50% and cannot be adjusted.

The LEDC API is designed to report an error when trying to set a frequency and a duty resolution that exceed the range of LEDC's hardware. For example, an attempt to set the frequency to 20 MHz and the duty resolution to 3 bits results in the following error reported on a serial monitor:

```
E (196) ledc: requested frequency and duty resolution cannot be achieved, try_
↪reducing freq_hz or duty_resolution. div_param=128
```

In such a situation, either the duty resolution or the frequency must be reduced. For example, setting the duty resolution to 2 resolves this issue and makes it possible to set the duty cycle at 25% steps, i.e., at 25%, 50% or 75%.

The LEDC driver also captures and reports attempts to configure frequency/duty resolution combinations that are below the supported minimum, e.g.,:

```
E (196) ledc: requested frequency and duty resolution cannot be achieved, try_
↪increasing freq_hz or duty_resolution. div_param=128000000
```

The duty resolution is normally set using `ledc_timer_bit_t`. This enumeration covers the range from 10 to 15 bits. If a smaller duty resolution is required (from 10 down to 1), enter the equivalent numeric values directly.

Application Example

The LEDC basic example: [peripherals/ledc/ledc_basic](#).

The LEDC change duty cycle and fading control example: [peripherals/ledc/ledc_fade](#).

The LEDC color control with Gamma correction on RGB LED example: [peripherals/ledc/ledc_gamma_curve_fade](#).

API Reference

Header File

- [components/driver/ledc/include/driver/ledc.h](#)
- This header file can be included with:

```
#include "driver/ledc.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

`esp_err_t ledc_channel_config` (const `ledc_channel_config_t` *ledc_conf)

LEDC channel configuration Configure LEDC channel with the given channel/output gpio_num/interrupt/source timer/frequency(Hz)/LEDC duty.

Parameters `ledc_conf` -- Pointer of LEDC channel configure struct

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

`uint32_t ledc_find_suitable_duty_resolution` (uint32_t src_clk_freq, uint32_t timer_freq)

Helper function to find the maximum possible duty resolution in bits for `ledc_timer_config()`

Parameters

- `src_clk_freq` -- LEDC timer source clock frequency (Hz) (See doxygen comments of `ledc_clk_cfg_t` or get from `esp_clk_tree_src_get_freq_hz`)
- `timer_freq` -- Desired LEDC timer frequency (Hz)

Returns

- 0 The timer frequency cannot be achieved
- Others The largest duty resolution value to be set

`esp_err_t ledc_timer_config` (const `ledc_timer_config_t` *timer_conf)

LEDC timer configuration Configure LEDC timer with the given source timer/frequency(Hz)/duty_resolution.

Parameters `timer_conf` -- Pointer of LEDC timer configure struct

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

- **ESP_FAIL** Can not find a proper pre-divider number base on the given frequency and the current `duty_resolution`.
- **ESP_ERR_INVALID_STATE** Timer cannot be de-configured because timer is not configured or is not paused

esp_err_t **ledc_update_duty** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel)

LEDC update channel parameters.

Note: Call this function to activate the LEDC updated parameters. After `ledc_set_duty`, we need to call this function to update the settings. And the new LEDC parameters don't take effect until the next PWM cycle.

Note: `ledc_set_duty`, `ledc_set_duty_with_hpoint` and `ledc_update_duty` are not thread-safe, do not call these functions to control one LEDC channel in different tasks at the same time. A thread-safe version of API is `ledc_set_duty_and_update`

Note: If `CONFIG_LEDC_CTRL_FUNC_IN_IRAM` is enabled, this function will be placed in the IRAM by linker, makes it possible to execute even when the Cache is disabled.

Note: This function is allowed to run within ISR context.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel (0 - `LEDC_CHANNEL_MAX-1`), select from `ledc_channel_t`

Returns

- **ESP_OK** Success
- **ESP_ERR_INVALID_ARG** Parameter error

esp_err_t **ledc_set_pin** (int gpio_num, *ledc_mode_t* speed_mode, *ledc_channel_t* ledc_channel)

Set LEDC output gpio.

Note: This function only routes the LEDC signal to GPIO through matrix, other LEDC resources initialization are not involved. Please use `ledc_channel_config()` instead to fully configure a LEDC channel.

Parameters

- **gpio_num** -- The LEDC output gpio
- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **ledc_channel** -- LEDC channel (0 - `LEDC_CHANNEL_MAX-1`), select from `ledc_channel_t`

Returns

- **ESP_OK** Success
- **ESP_ERR_INVALID_ARG** Parameter error

esp_err_t **ledc_stop** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel, uint32_t idle_level)

LEDC stop. Disable LEDC output, and set idle level.

Note: If `CONFIG_LEDC_CTRL_FUNC_IN_IRAM` is enabled, this function will be placed in the IRAM by

linker, makes it possible to execute even when the Cache is disabled.

Note: This function is allowed to run within ISR context.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel (0 - LEDC_CHANNEL_MAX-1), select from `ledc_channel_t`
- **idle_level** -- Set output idle level after LEDC stops.

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **ledc_set_freq**(*ledc_mode_t* speed_mode, *ledc_timer_t* timer_num, uint32_t freq_hz)

LEDC set channel frequency (Hz)

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **timer_num** -- LEDC timer index (0-3), select from `ledc_timer_t`
- **freq_hz** -- Set the LEDC frequency

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_FAIL Can not find a proper pre-divider number base on the given frequency and the current duty_resolution.

uint32_t **ledc_get_freq**(*ledc_mode_t* speed_mode, *ledc_timer_t* timer_num)

LEDC get channel frequency (Hz)

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **timer_num** -- LEDC timer index (0-3), select from `ledc_timer_t`

Returns

- 0 error
- Others Current LEDC frequency

esp_err_t **ledc_set_duty_with_hpoint**(*ledc_mode_t* speed_mode, *ledc_channel_t* channel, uint32_t duty, uint32_t hpoint)

LEDC set duty and hpoint value Only after calling `ledc_update_duty` will the duty update.

Note: `ledc_set_duty`, `ledc_set_duty_with_hpoint` and `ledc_update_duty` are not thread-safe, do not call these functions to control one LEDC channel in different tasks at the same time. A thread-safe version of API is `ledc_set_duty_and_update`

Note: For ESP32, hardware does not support any duty change while a fade operation is running in progress on that channel. Other duty operations will have to wait until the fade operation has finished.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.

- **channel** -- LEDC channel (0 - LEDC_CHANNEL_MAX-1), select from ledc_channel_t
- **duty** -- Set the LEDC duty, the range of duty setting is [0, (2**duty_resolution)]
- **hpoint** -- Set the LEDC hpoint value, the range is [0, (2**duty_resolution)-1]

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

int **ledc_get_hpoint** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel)

LEDC get hpoint value, the counter value when the output is set high level.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel (0 - LEDC_CHANNEL_MAX-1), select from ledc_channel_t

Returns

- LEDC_ERR_VAL if parameter error
- Others Current hpoint value of LEDC channel

esp_err_t **ledc_set_duty** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel, uint32_t duty)

LEDC set duty This function do not change the hpoint value of this channel. if needed, please call ledc_set_duty_with_hpoint. only after calling ledc_update_duty will the duty update.

Note: ledc_set_duty, ledc_set_duty_with_hpoint and ledc_update_duty are not thread-safe, do not call these functions to control one LEDC channel in different tasks at the same time. A thread-safe version of API is ledc_set_duty_and_update.

Note: For ESP32, hardware does not support any duty change while a fade operation is running in progress on that channel. Other duty operations will have to wait until the fade operation has finished.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel (0 - LEDC_CHANNEL_MAX-1), select from ledc_channel_t
- **duty** -- Set the LEDC duty, the range of duty setting is [0, (2**duty_resolution)]

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

uint32_t **ledc_get_duty** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel)

LEDC get duty This function returns the duty at the present PWM cycle. You shouldn't expect the function to return the new duty in the same cycle of calling ledc_update_duty, because duty update doesn't take effect until the next cycle.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel (0 - LEDC_CHANNEL_MAX-1), select from ledc_channel_t

Returns

- LEDC_ERR_DUTY if parameter error
- Others Current LEDC duty

esp_err_t **ledc_set_fade** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel, uint32_t duty, *ledc_duty_direction_t* fade_direction, uint32_t step_num, uint32_t duty_cycle_num, uint32_t duty_scale)

LEDC set gradient Set LEDC gradient, After the function calls the ledc_update_duty function, the function can take effect.

Note: For ESP32, hardware does not support any duty change while a fade operation is running in progress on that channel. Other duty operations will have to wait until the fade operation has finished.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel (0 - LEDC_CHANNEL_MAX-1), select from ledc_channel_t
- **duty** -- Set the start of the gradient duty, the range of duty setting is [0, (2*duty_resolution)]
- **fade_direction** -- Set the direction of the gradient
- **step_num** -- Set the number of the gradient
- **duty_cycle_num** -- Set how many LEDC tick each time the gradient lasts
- **duty_scale** -- Set gradient change amplitude

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **ledc_isr_register** (void (*fn)(void*), void *arg, int intr_alloc_flags, *ledc_isr_handle_t* *handle)

Register LEDC interrupt handler, the handler is an ISR. The handler will be attached to the same CPU core that this function is running on.

Parameters

- **fn** -- Interrupt handler function.
- **arg** -- User-supplied argument passed to the handler function.
- **intr_alloc_flags** -- Flags used to allocate the interrupt. One or multiple (ORred) ESP_INTR_FLAG_* values. See esp_intr_alloc.h for more info.
- **handle** -- Pointer to return handle. If non-NULL, a handle for the interrupt will be returned here.

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_ERR_NOT_FOUND Failed to find available interrupt source

esp_err_t **ledc_timer_set** (*ledc_mode_t* speed_mode, *ledc_timer_t* timer_sel, uint32_t clock_divider, uint32_t duty_resolution, *ledc_clk_src_t* clk_src)

Configure LEDC settings.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **timer_sel** -- Timer index (0-3), there are 4 timers in LEDC module
- **clock_divider** -- Timer clock divide value, the timer clock is divided from the selected clock source
- **duty_resolution** -- Resolution of duty setting in number of bits. The range is [1, SOC_LEDC_TIMER_BIT_WIDTH]
- **clk_src** -- Select LEDC source clock.

Returns

- (-1) Parameter error
- Other Current LEDC duty

esp_err_t **ledc_timer_rst** (*ledc_mode_t* speed_mode, *ledc_timer_t* timer_sel)

Reset LEDC timer.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **timer_sel** -- LEDC timer index (0-3), select from *ledc_timer_t*

Returns

- ESP_ERR_INVALID_ARG Parameter error
- ESP_OK Success

esp_err_t **ledc_timer_pause** (*ledc_mode_t* speed_mode, *ledc_timer_t* timer_sel)

Pause LEDC timer counter.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **timer_sel** -- LEDC timer index (0-3), select from *ledc_timer_t*

Returns

- ESP_ERR_INVALID_ARG Parameter error
- ESP_OK Success

esp_err_t **ledc_timer_resume** (*ledc_mode_t* speed_mode, *ledc_timer_t* timer_sel)

Resume LEDC timer.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **timer_sel** -- LEDC timer index (0-3), select from *ledc_timer_t*

Returns

- ESP_ERR_INVALID_ARG Parameter error
- ESP_OK Success

esp_err_t **ledc_bind_channel_timer** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel, *ledc_timer_t* timer_sel)

Bind LEDC channel with the selected timer.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel index (0 - LEDC_CHANNEL_MAX-1), select from *ledc_channel_t*
- **timer_sel** -- LEDC timer index (0-3), select from *ledc_timer_t*

Returns

- ESP_ERR_INVALID_ARG Parameter error
- ESP_OK Success

esp_err_t **ledc_set_fade_with_step** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel, uint32_t target_duty, uint32_t scale, uint32_t cycle_num)

Set LEDC fade function.

Note: Call `ledc_fade_func_install()` once before calling this function. Call `ledc_fade_start()` after this to start fading.

Note: `ledc_set_fade_with_step`, `ledc_set_fade_with_time` and `ledc_fade_start` are not thread-safe, do not call these functions to control one LEDC channel in different tasks at the same time. A thread-safe version of API is `ledc_set_fade_step_and_start`

Note: For ESP32, hardware does not support any duty change while a fade operation is running in progress on that channel. Other duty operations will have to wait until the fade operation has finished.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel index (0 - LEDC_CHANNEL_MAX-1), select from `ledc_channel_t`
- **target_duty** -- Target duty of fading [0, (2**duty_resolution)]
- **scale** -- Controls the increase or decrease step scale.
- **cycle_num** -- increase or decrease the duty every cycle_num cycles

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_ERR_INVALID_STATE Channel not initialized
- ESP_FAIL Fade function init error

esp_err_t **ledc_set_fade_with_time** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel, uint32_t target_duty, int desired_fade_time_ms)

Set LEDC fade function, with a limited time.

Note: Call `ledc_fade_func_install()` once before calling this function. Call `ledc_fade_start()` after this to start fading.

Note: `ledc_set_fade_with_step`, `ledc_set_fade_with_time` and `ledc_fade_start` are not thread-safe, do not call these functions to control one LEDC channel in different tasks at the same time. A thread-safe version of API is `ledc_set_fade_step_and_start`

Note: For ESP32, hardware does not support any duty change while a fade operation is running in progress on that channel. Other duty operations will have to wait until the fade operation has finished.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel index (0 - LEDC_CHANNEL_MAX-1), select from `ledc_channel_t`
- **target_duty** -- Target duty of fading [0, (2**duty_resolution)]
- **desired_fade_time_ms** -- The intended time of the fading (ms). Note that the actual time it takes to complete the fade could vary by a factor of up to 2x shorter or longer than the expected time due to internal rounding errors in calculations. Specifically:
 - The total number of cycles ($\text{total\_cycle\_num} = \text{desired\_fade\_time\_ms} * \text{freq} / 1000$)
 - The difference in duty cycle ($\text{duty\_delta} = |\text{target\_duty} - \text{current\_duty}|$) The fade may complete faster than expected if `total_cycle_num` larger than `duty_delta`. Conversely, it may take longer than expected if `total_cycle_num` is less than `duty_delta`. The closer the ratio of `total_cycle_num/duty_delta` (or its inverse) is to a whole number (the floor value), the more accurately the actual fade duration will match the intended time. If an exact fade time is expected, please consider to split the entire fade into several smaller linear fades. The split should make each fade step has a divisible `total_cycle_num/duty_delta` (or its inverse) ratio.

Returns

- ESP_OK Success

- ESP_ERR_INVALID_ARG Parameter error
- ESP_ERR_INVALID_STATE Channel not initialized
- ESP_FAIL Fade function init error

esp_err_t **ledc_fade_func_install** (int intr_alloc_flags)

Install LEDC fade function. This function will occupy interrupt of LEDC module.

Parameters **intr_alloc_flags** -- Flags used to allocate the interrupt. One or multiple (ORred) ESP_INTR_FLAG_* values. See esp_intr_alloc.h for more info.

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Intr flag error
- ESP_ERR_NOT_FOUND Failed to find available interrupt source
- ESP_ERR_INVALID_STATE Fade function already installed

void **ledc_fade_func_uninstall** (void)

Uninstall LEDC fade function.

esp_err_t **ledc_fade_start** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel, *ledc_fade_mode_t* fade_mode)

Start LEDC fading.

Note: Call ledc_fade_func_install() once before calling this function. Call this API right after ledc_set_fade_with_time or ledc_set_fade_with_step before to start fading.

Note: Starting fade operation with this API is not thread-safe, use with care.

Note: For ESP32, hardware does not support any duty change while a fade operation is running in progress on that channel. Other duty operations will have to wait until the fade operation has finished.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel number
- **fade_mode** -- Whether to block until fading done. See ledc_types.h ledc_fade_mode_t for more info. Note that this function will not return until fading to the target duty if LEDC_FADE_WAIT_DONE mode is selected.

Returns

- ESP_OK Success
- ESP_ERR_INVALID_STATE Channel not initialized or fade function not installed.
- ESP_ERR_INVALID_ARG Parameter error.

esp_err_t **ledc_fade_stop** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel)

Stop LEDC fading. The duty of the channel is guaranteed to be fixed at most one PWM cycle after the function returns.

Note: This API can be called if a new fixed duty or a new fade want to be set while the last fade operation is still running in progress.

Note: Call this API will abort the fading operation only if it was started by calling ledc_fade_start with LEDC_FADE_NO_WAIT mode.

Note: If a fade was started with LEDC_FADE_WAIT_DONE mode, calling this API afterwards has no use in stopping the fade. Fade will continue until it reaches the target duty.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel number

Returns

- ESP_OK Success
- ESP_ERR_INVALID_STATE Channel not initialized
- ESP_ERR_INVALID_ARG Parameter error
- ESP_FAIL Fade function init error

esp_err_t **ledc_set_duty_and_update** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel, uint32_t duty, uint32_t hpoint)

A thread-safe API to set duty for LEDC channel and return when duty updated.

Note: For ESP32, hardware does not support any duty change while a fade operation is running in progress on that channel. Other duty operations will have to wait until the fade operation has finished.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel (0 - LEDC_CHANNEL_MAX-1), select from ledc_channel_t
- **duty** -- Set the LEDC duty, the range of duty setting is [0, (2**duty_resolution)]
- **hpoint** -- Set the LEDC hpoint value, the range is [0, (2**duty_resolution)-1]

Returns

- ESP_OK Success
- ESP_ERR_INVALID_STATE Channel not initialized
- ESP_ERR_INVALID_ARG Parameter error
- ESP_FAIL Fade function init error

esp_err_t **ledc_set_fade_time_and_start** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel, uint32_t target_duty, uint32_t desired_fade_time_ms, *ledc_fade_mode_t* fade_mode)

A thread-safe API to set and start LEDC fade function, with a limited time.

Note: Call ledc_fade_func_install() once, before calling this function.

Note: For ESP32, hardware does not support any duty change while a fade operation is running in progress on that channel. Other duty operations will have to wait until the fade operation has finished.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel index (0 - LEDC_CHANNEL_MAX-1), select from ledc_channel_t
- **target_duty** -- Target duty of fading [0, (2**duty_resolution)]

- **desired_fade_time_ms** -- The intended time of the fading (ms). Note that the actual time it takes to complete the fade could vary by a factor of up to 2x shorter or longer than the expected time due to internal rounding errors in calculations. Specifically:
 - The total number of cycles ($\text{total\_cycle\_num} = \text{desired\_fade\_time\_ms} * \text{freq} / 1000$)
 - The difference in duty cycle ($\text{duty\_delta} = |\text{target\_duty} - \text{current\_duty}|$) The fade may complete faster than expected if total_cycle_num larger than duty_delta . Conversely, it may take longer than expected if total_cycle_num is less than duty_delta . The closer the ratio of $\text{total\_cycle\_num}/\text{duty\_delta}$ (or its inverse) is to a whole number (the floor value), the more accurately the actual fade duration will match the intended time. If an exact fade time is expected, please consider to split the entire fade into several smaller linear fades. The split should make each fade step has a divisible $\text{total\_cycle\_num}/\text{duty\_delta}$ (or its inverse) ratio.
- **fade_mode** -- choose blocking or non-blocking mode

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_ERR_INVALID_STATE Channel not initialized
- ESP_FAIL Fade function init error

esp_err_t **ledc_set_fade_step_and_start** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel, uint32_t target_duty, uint32_t scale, uint32_t cycle_num, *ledc_fade_mode_t* fade_mode)

A thread-safe API to set and start LEDC fade function.

Note: Call `ledc_fade_func_install()` once before calling this function.

Note: For ESP32, hardware does not support any duty change while a fade operation is running in progress on that channel. Other duty operations will have to wait until the fade operation has finished.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel index (0 - LEDC_CHANNEL_MAX-1), select from `ledc_channel_t`
- **target_duty** -- Target duty of fading [0, (2**duty_resolution)]
- **scale** -- Controls the increase or decrease step scale.
- **cycle_num** -- increase or decrease the duty every cycle_num cycles
- **fade_mode** -- choose blocking or non-blocking mode

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_ERR_INVALID_STATE Channel not initialized
- ESP_FAIL Fade function init error

esp_err_t **ledc_cb_register** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel, *ledc_cbs_t* *cbs, void *user_arg)

LEDC callback registration function.

Note: The callback is called from an ISR, it must never attempt to block, and any FreeRTOS API called must be ISR capable.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel index (0 - LEDC_CHANNEL_MAX-1), select from `ledc_channel_t`
- **cbs** -- Group of LEDC callback functions
- **user_arg** -- user registered data for the callback function

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_ERR_INVALID_STATE Channel not initialized
- ESP_FAIL Fade function init error

esp_err_t **ledc_set_multi_fade** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel, uint32_t start_duty, const *ledc_fade_param_config_t* *fade_params_list, uint32_t list_len)

Set a LEDC multi-fade.

Note: Call `ledc_fade_func_install()` once before calling this function. Call `ledc_fade_start()` after this to start fading.

Note: This function is not thread-safe, do not call it to control one LEDC channel in different tasks at the same time. A thread-safe version of API is `ledc_set_multi_fade_and_start`

Note: This function does not prohibit from duty overflow. User should take care of this by themselves. If duty overflow happens, the PWM signal will suddenly change from 100% duty cycle to 0%, or the other way around.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel index (0 - LEDC_CHANNEL_MAX-1), select from `ledc_channel_t`
- **start_duty** -- Set the start of the gradient duty, the range of duty setting is [0, (2*duty_resolution)]
- **fade_params_list** -- Pointer to the array of fade parameters for a multi-fade
- **list_len** -- Length of the `fade_params_list`, i.e. number of fade ranges for a multi-fade (1 - SOC_LEDC_GAMMA_CURVE_FADE_RANGE_MAX)

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_ERR_INVALID_STATE Channel not initialized
- ESP_FAIL Fade function init error

esp_err_t **ledc_set_multi_fade_and_start** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel, uint32_t start_duty, const *ledc_fade_param_config_t* *fade_params_list, uint32_t list_len, *ledc_fade_mode_t* fade_mode)

A thread-safe API to set and start LEDC multi-fade function.

Note: Call `ledc_fade_func_install()` once before calling this function.

Note: Fade will always begin from the current duty cycle. Make sure it is stable and synchronized to the

desired initial value before calling this function. Otherwise, you may see unexpected duty change.

Note: This function does not prohibit from duty overflow. User should take care of this by themselves. If duty overflow happens, the PWM signal will suddenly change from 100% duty cycle to 0%, or the other way around.

Parameters

- **speed_mode** -- Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- LEDC channel index (0 - LEDC_CHANNEL_MAX-1), select from `ledc_channel_t`
- **start_duty** -- Set the start of the gradient duty, the range of duty setting is [0, (2**duty_resolution)]
- **fade_params_list** -- Pointer to the array of fade parameters for a multi-fade
- **list_len** -- Length of the `fade_params_list`, i.e. number of fade ranges for a multi-fade (1 - SOC_LEDC_GAMMA_CURVE_FADE_RANGE_MAX)
- **fade_mode** -- Choose blocking or non-blocking mode

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_ERR_INVALID_STATE Channel not initialized
- ESP_FAIL Fade function init error

```
esp_err_t ledc_fill_multi_fade_param_list (ledc_mode_t speed_mode, ledc_channel_t channel,  
                                           uint32_t start_duty, uint32_t end_duty, uint32_t  
                                           linear_phase_num, uint32_t max_fade_time_ms,  
                                           uint32_t (*gamma_correction_operator)(uint32_t),  
                                           uint32_t fade_params_list_size,  
                                           ledc_fade_param_config_t *fade_params_list, uint32_t  
                                           *hw_fade_range_num)
```

Helper function to fill the fade params for a multi-fade. Useful if desires a gamma curve fading.

Note: The fade params are calculated based on the given `start_duty` and `end_duty`. If the duty is not at the start duty (gamma-corrected) when the fade begins, you may see undesired brightness change. Therefore, please always remember that when passing the `fade_params` to either `ledc_set_multi_fade` or `ledc_set_multi_fade_and_start`, the `start_duty` argument has to be the gamma-corrected `start_duty`.

Parameters

- **speed_mode** -- **[in]** Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- **[in]** LEDC channel index (0 - LEDC_CHANNEL_MAX-1), select from `ledc_channel_t`
- **start_duty** -- **[in]** Duty cycle [0, (2**duty_resolution)] where the multi-fade begins with. This value should be a non-gamma-corrected duty cycle.
- **end_duty** -- **[in]** Duty cycle [0, (2**duty_resolution)] where the multi-fade ends with. This value should be a non-gamma-corrected duty cycle.
- **linear_phase_num** -- **[in]** Number of linear fades to simulate a gamma curved fade (1 - SOC_LEDC_GAMMA_CURVE_FADE_RANGE_MAX)
- **max_fade_time_ms** -- **[in]** The maximum time of the fading (ms).
- **gamma_correction_operator** -- **[in]** User provided gamma correction function. The function argument should be able to take any value within [0, (2**duty_resolution)]. And returns the gamma-corrected duty cycle.

- **fade_params_list_size** -- [in] The size of the fade_params_list user allocated (1 - SOC_LEDC_GAMMA_CURVE_FADE_RANGE_MAX)
- **fade_params_list** -- [out] Pointer to the array of *ledc_fade_param_config_t* structure
- **hw_fade_range_num** -- [out] Number of fade ranges for this multi-fade

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_ERR_INVALID_STATE Channel not initialized
- ESP_FAIL Required number of hardware ranges exceeds the size of the *ledc_fade_param_config_t* array user allocated

esp_err_t **ledc_read_fade_param** (*ledc_mode_t* speed_mode, *ledc_channel_t* channel, uint32_t range, uint32_t *dir, uint32_t *cycle, uint32_t *scale, uint32_t *step)

Get the fade parameters that are stored in gamma ram for a certain fade range.

Gamma ram is where saves the fade parameters for each fade range. The fade parameters are written in during fade configuration. When fade begins, the duty will change according to the parameters in gamma ram.

Parameters

- **speed_mode** -- [in] Select the LEDC channel group with specified speed mode. Note that not all targets support high speed mode.
- **channel** -- [in] LEDC channel index (0 - LEDC_CHANNEL_MAX-1), select from *ledc_channel_t*
- **range** -- [in] Range index (0 - (SOC_LEDC_GAMMA_CURVE_FADE_RANGE_MAX-1)), it specifies to which range in gamma ram to read
- **dir** -- [out] Pointer to accept fade direction value
- **cycle** -- [out] Pointer to accept fade cycle value
- **scale** -- [out] Pointer to accept fade scale value
- **step** -- [out] Pointer to accept fade step value

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_ERR_INVALID_STATE Channel not initialized

Structures

struct **ledc_channel_config_t**

Configuration parameters of LEDC channel for *ledc_channel_config* function.

Public Members

int **gpio_num**

the LEDC output gpio_num, if you want to use gpio16, gpio_num = 16

ledc_mode_t **speed_mode**

LEDC speed speed_mode, high-speed mode (only exists on esp32) or low-speed mode

ledc_channel_t **channel**

LEDC channel (0 - LEDC_CHANNEL_MAX-1)

ledc_intr_type_t **intr_type**

configure interrupt, Fade interrupt enable or Fade interrupt disable

ledc_timer_t **timer_sel**

Select the timer source of channel (0 - LEDC_TIMER_MAX-1)

uint32_t **duty**

LEDC channel duty, the range of duty setting is [0, (2**duty_resolution)]

int **hpoint**

LEDC channel hpoint value, the range is [0, (2**duty_resolution)-1]

unsigned int **output_invert**

Enable (1) or disable (0) gpio output invert

struct *ledc_channel_config_t*::[anonymous] **flags**

LEDC flags

struct **ledc_timer_config_t**

Configuration parameters of LEDC timer for ledc_timer_config function.

Public Members

ledc_mode_t **speed_mode**

LEDC speed speed_mode, high-speed mode (only exists on esp32) or low-speed mode

ledc_timer_bit_t **duty_resolution**

LEDC channel duty resolution

ledc_timer_t **timer_num**

The timer source of channel (0 - LEDC_TIMER_MAX-1)

uint32_t **freq_hz**

LEDC timer frequency (Hz)

ledc_clk_cfg_t **clk_cfg**

Configure LEDC source clock from ledc_clk_cfg_t. Note that LEDC_USE_RC_FAST_CLK and LEDC_USE_XTAL_CLK are non-timer-specific clock sources. You can not have one LEDC timer uses RC_FAST_CLK as the clock source and have another LEDC timer uses XTAL_CLK as its clock source. All chips except esp32 and esp32s2 do not have timer-specific clock sources, which means clock source for all timers must be the same one.

bool **deconfigure**

Set this field to de-configure a LEDC timer which has been configured before. Note that it will not check whether the timer wants to be de-configured is binded to any channel. Also, the timer has to be paused first before it can be de-configured. When this field is set, duty_resolution, freq_hz, clk_cfg fields are ignored.

struct **ledc_cb_param_t**

LEDC callback parameter.

Public Members

ledc_cb_event_t **event**

Event name

uint32_t **speed_mode**

Speed mode of the LEDC channel group

uint32_t **channel**

LEDC channel (0 - LEDC_CHANNEL_MAX-1)

uint32_t **duty**

LEDC current duty of the channel, the range of duty is [0, (2**duty_resolution)]

struct **ledc_cbs_t**

Group of supported LEDC callbacks.

Note: The callbacks are all running under ISR environment


Public Members

ledc_cb_t **fade_cb**

LEDC fade_end callback function


struct **ledc_fade_param_config_t**

Structure for the fade parameters for one hardware fade to be written to gamma wr register.



(continues on next page)

(continued from previous page)



Note: Be aware of the maximum value available on each element

Public Members

uint32_t dir

Duty change direction. Set 1 as increase, 0 as decrease

uint32_t cycle_num

Number of PWM cycles of each step [0, 2**SOC_LEDC_FADE_PARAMS_BIT_WIDTH-1]

uint32_t scale

Duty change of each step [0, 2**SOC_LEDC_FADE_PARAMS_BIT_WIDTH-1]

uint32_t **step_num**

Total number of steps in one hardware fade [0, 2**SOC_LEDC_FADE_PARAMS_BIT_WIDTH-1]

Macros

LEDC_ERR_DUTY

LEDC_ERR_VAL

Type Definitions

typedef *intr_handle_t* **ledc_isr_handle_t**

typedef bool (***ledc_cb_t**)(const *ledc_cb_param_t* *param, void *user_arg)

Type of LEDC event callback.

Param param LEDC callback parameter

Param user_arg User registered data

Return Whether a high priority task has been waken up by this function

Enumerations

enum **ledc_cb_event_t**

LEDC callback event type.

Values:

enumerator **LEDC_FADE_END_EVT**

LEDC fade end event

Header File

- [components/hal/include/hal/ledc_types.h](#)
- This header file can be included with:

```
#include "hal/ledc_types.h"
```

Type Definitions

typedef *soc_periph_ledc_clk_src_legacy_t* **ledc_clk_cfg_t**

LEDC clock source configuration struct.

In theory, the following enumeration shall be placed in LEDC driver's header. However, as the next enumeration, *ledc_clk_src_t*, makes the use of some of these values and to avoid mutual inclusion of the headers, we must define it here.

Enumerations

enum **ledc_mode_t**

Values:

enumerator **LEDC_LOW_SPEED_MODE**

LEDC low speed speed_mode

enumerator **LEDC_SPEED_MODE_MAX**

LEDC speed limit

enum **ledc_intr_type_t**

Values:

enumerator **LEDC_INTR_DISABLE**

Disable LEDC interrupt

enumerator **LEDC_INTR_FADE_END**

Enable LEDC interrupt

enumerator **LEDC_INTR_MAX**

enum **ledc_duty_direction_t**

Values:

enumerator **LEDC_DUTY_DIR_DECREASE**

LEDC duty decrease direction

enumerator **LEDC_DUTY_DIR_INCREASE**

LEDC duty increase direction

enumerator **LEDC_DUTY_DIR_MAX**

enum **ledc_slow_clk_sel_t**

LEDC global clock sources.

Values:

enumerator **LEDC_SLOW_CLK_RC_FAST**

LEDC low speed timer clock source is RC_FAST clock

enumerator **LEDC_SLOW_CLK_PLL_DIV**

LEDC low speed timer clock source is a PLL_DIV clock

enumerator **LEDC_SLOW_CLK_XTAL**

LEDC low speed timer clock source XTAL clock

enumerator **LEDC_SLOW_CLK_RTC8M**

Alias of 'LEDC_SLOW_CLK_RC_FAST'

enum **ledc_clk_src_t**

LEDC timer-specific clock sources.

Note: Setting numeric values to match `ledc_clk_cfg_t` values are a hack to avoid collision with `LEDC_AUTO_CLK` in the driver, as these enums have very similar names and user may pass one of these by mistake.

Values:

enumerator **LEDC_SCLK**

Selecting this value for `LEDC_TICK_SEL_TIMER` let the hardware take its source clock from `LEDC_CLK_SEL`

enum **ledc_timer_t**

Values:

enumerator **LEDC_TIMER_0**

LEDC timer 0

enumerator **LEDC_TIMER_1**

LEDC timer 1

enumerator **LEDC_TIMER_2**

LEDC timer 2

enumerator **LEDC_TIMER_3**

LEDC timer 3

enumerator **LEDC_TIMER_MAX**

enum **ledc_channel_t**

Values:

enumerator **LEDC_CHANNEL_0**

LEDC channel 0

enumerator **LEDC_CHANNEL_1**

LEDC channel 1

enumerator **LEDC_CHANNEL_2**

LEDC channel 2

enumerator **LEDC_CHANNEL_3**

LEDC channel 3

enumerator **LEDC_CHANNEL_4**

LEDC channel 4

enumerator **LEDC_CHANNEL_5**

LEDC channel 5

enumerator **LEDC_CHANNEL_6**

LEDC channel 6

enumerator **LEDC_CHANNEL_7**

LEDC channel 7

enumerator **LEDC_CHANNEL_MAX**

enum **ledc_timer_bit_t**

Values:

enumerator **LEDC_TIMER_1_BIT**

LEDC PWM duty resolution of 1 bits

enumerator **LEDC_TIMER_2_BIT**

LEDC PWM duty resolution of 2 bits

enumerator **LEDC_TIMER_3_BIT**

LEDC PWM duty resolution of 3 bits

enumerator **LEDC_TIMER_4_BIT**

LEDC PWM duty resolution of 4 bits

enumerator **LEDC_TIMER_5_BIT**

LEDC PWM duty resolution of 5 bits

enumerator **LEDC_TIMER_6_BIT**

LEDC PWM duty resolution of 6 bits

enumerator **LEDC_TIMER_7_BIT**

LEDC PWM duty resolution of 7 bits

enumerator **LEDC_TIMER_8_BIT**

LEDC PWM duty resolution of 8 bits

enumerator **LEDC_TIMER_9_BIT**

LEDC PWM duty resolution of 9 bits

enumerator **LEDC_TIMER_10_BIT**

LEDC PWM duty resolution of 10 bits

enumerator **LEDC_TIMER_11_BIT**

LEDC PWM duty resolution of 11 bits

enumerator **LEDC_TIMER_12_BIT**

LEDC PWM duty resolution of 12 bits

enumerator **LEDC_TIMER_13_BIT**

LEDC PWM duty resolution of 13 bits

enumerator **LEDC_TIMER_14_BIT**

LEDC PWM duty resolution of 14 bits

enumerator **LEDC_TIMER_15_BIT**

LEDC PWM duty resolution of 15 bits

enumerator **LEDC_TIMER_16_BIT**

LEDC PWM duty resolution of 16 bits

enumerator **LEDC_TIMER_17_BIT**

LEDC PWM duty resolution of 17 bits

enumerator **LEDC_TIMER_18_BIT**

LEDC PWM duty resolution of 18 bits

enumerator **LEDC_TIMER_19_BIT**

LEDC PWM duty resolution of 19 bits

enumerator **LEDC_TIMER_20_BIT**

LEDC PWM duty resolution of 20 bits

enumerator **LEDC_TIMER_BIT_MAX**

enum **ledc_fade_mode_t**

Values:

enumerator **LEDC_FADE_NO_WAIT**

LEDC fade function will return immediately

enumerator **LEDC_FADE_WAIT_DONE**

LEDC fade function will block until fading to the target duty

enumerator **LEDC_FADE_MAX**

2.5.13 Motor Control Pulse Width Modulator (MCPWM)

The MCPWM peripheral is a versatile PWM generator, which contains various submodules to make it a key element in power electronic applications like motor control, digital power, and so on. Typically, the MCPWM peripheral can be used in the following scenarios:

- Digital motor control, e.g., brushed/brushless DC motor, RC servo motor
- Switch mode-based digital power conversion
- Power DAC, where the duty cycle is equivalent to a DAC analog value
- Calculate external pulse width, and convert it into other analog values like speed, distance
- Generate Space Vector PWM (SVPWM) signals for Field Oriented Control (FOC)

The main submodules are listed in the following diagram:

- **MCPWM Timer:** The time base of the final PWM signal. It also determines the event timing of other submodules.
- **MCPWM Operator:** The key module that is responsible for generating the PWM waveforms. It consists of other submodules, like comparator, PWM generator, dead time, and carrier modulator.
- **MCPWM Comparator:** The compare module takes the time-base count value as input, and continuously compares it to the threshold value configured. When the timer is equal to any of the threshold values, a compare event will be generated and the MCPWM generator can update its level accordingly.
- **MCPWM Generator:** One MCPWM generator can generate a pair of PWM waves, complementarily or independently, based on various events triggered by other submodules like MCPWM Timer and MCPWM Comparator.
- **MCPWM Fault:** The fault module is used to detect the fault condition from outside, mainly via the GPIO matrix. Once the fault signal is active, MCPWM Operator will force all the generators into a predefined state to protect the system from damage.




Fig. 14: MCPWM Overview

- **MCPWM Sync:** The sync module is used to synchronize the MCPWM timers, so that the final PWM signals generated by different MCPWM generators can have a fixed phase difference. The sync signal can be routed from the GPIO matrix or from an MCPWM Timer event.
- **Dead Time:** This submodule is used to insert extra delay to the existing PWM edges generated in the previous steps.
- **Carrier Modulation:** The carrier submodule can modulate a high-frequency carrier signal into PWM waveforms by the generator and dead time submodules. This capability is mandatory for controlling the power-switching elements.
- **Brake:** MCPWM operator can set how to brake the generators when a particular fault is detected. You can shut down the PWM output immediately or regulate the PWM output cycle by cycle, depending on how critical the fault is.
- **MCPWM Capture:** This is a standalone submodule that can work even without the above MCPWM operators. The capture consists one dedicated timer and several independent channels, with each channel connected to the GPIO. A pulse on the GPIO triggers the capture timer to store the time-base count value and then notify you by an interrupt. Using this feature, you can measure a pulse width precisely. What is more, the capture timer can also be synchronized by the MCPWM Sync submodule.

Functional Overview

Description of the MCPWM functionality is divided into the following sections:

- [Resource Allocation and Initialization](#) - covers how to allocate various MCPWM objects, like timers, operators, comparators, generators and so on. These objects are the basis of the following IO setting and control functions.
- [Timer Operations and Events](#) - describes control functions and event callbacks supported by the MCPWM timer.
- [Comparator Operations and Events](#) - describes control functions and event callbacks supported by the MCPWM comparator.
- [Generator Actions on Events](#) - describes how to set actions for MCPWM generators on particular events that are generated by the MCPWM timer and comparators.
- [Generator Configurations for Classical PWM Waveforms](#) - demonstrates some classical PWM waveforms that can be achieved by configuring generator actions.
- [Dead Time](#) - describes how to set dead time for MCPWM generators.
- [Dead Time Configurations for Classical PWM Waveforms](#) - demonstrates some classical PWM waveforms that can be achieved by configuring dead time.
- [Carrier Modulation](#) - describes how to set and modulate a high frequency onto the final PWM waveforms.
- [Faults and Brake Actions](#) - describes how to set brake actions for MCPWM operators on particular fault events.
- [Generator Force Actions](#) - describes how to control the generator output level asynchronously in a forceful way.
- [Synchronization](#) - describes how to synchronize the MCPWM timers and get a fixed phase difference between the generated PWM signals.

- *Capture* - describes how to use the MCPWM capture module to measure the pulse width of a signal.
- *ETM Event and Task* - describes what the events and tasks can be connected to the ETM channel.
- *Power Management* - describes how different source clocks affects power consumption.
- *IRAM Safe* - describes tips on how to make the RMT interrupt work better along with a disabled cache.
- *Thread Safety* - lists which APIs are guaranteed to be thread-safe by the driver.
- *Kconfig Options* - lists the supported Kconfig options that can bring different effects to the driver.

Resource Allocation and Initialization As displayed in the diagram above, the MCPWM peripheral consists of several submodules. Each submodule has its own resource allocation, which is described in the following sections.

MCPWM Timers You can allocate a MCPWM timer object by calling `mcpwm_new_timer()` function, with a configuration structure `mcpwm_timer_config_t` as the parameter. The configuration structure is defined as:

- `mcpwm_timer_config_t::group_id` specifies the MCPWM group ID. The ID should belong to `[0, SOC_MCPWM_GROUPS - 1]` range. Please note, timers located in different groups are totally independent.
- `mcpwm_timer_config_t::intr_priority` sets the priority of the interrupt. If it is set to 0, the driver will allocate an interrupt with a default priority. Otherwise, the driver will use the given priority.
- `mcpwm_timer_config_t::clk_src` sets the clock source of the timer.
- `mcpwm_timer_config_t::resolution_hz` sets the expected resolution of the timer. The driver internally sets a proper divider based on the clock source and the resolution.
- `mcpwm_timer_config_t::count_mode` sets the count mode of the timer.
- `mcpwm_timer_config_t::period_ticks` sets the period of the timer, in ticks (the tick resolution is set in the `mcpwm_timer_config_t::resolution_hz`).
- `mcpwm_timer_config_t::update_period_on_empty` sets whether to update the period value when the timer counts to zero.
- `mcpwm_timer_config_t::update_period_on_sync` sets whether to update the period value when the timer takes a sync signal.

The `mcpwm_new_timer()` will return a pointer to the allocated timer object if the allocation succeeds. Otherwise, it will return an error code. Specifically, when there are no more free timers in the MCPWM group, this function will return the `ESP_ERR_NOT_FOUND` error.¹

On the contrary, calling the `mcpwm_del_timer()` function will free the allocated timer object.

MCPWM Operators You can allocate a MCPWM operator object by calling `mcpwm_new_operator()` function, with a configuration structure `mcpwm_operator_config_t` as the parameter. The configuration structure is defined as:

- `mcpwm_operator_config_t::group_id` specifies the MCPWM group ID. The ID should belong to `[0, SOC_MCPWM_GROUPS - 1]` range. Please note, operators located in different groups are totally independent.
- `mcpwm_operator_config_t::intr_priority` sets the priority of the interrupt. If it is set to 0, the driver will allocate an interrupt with a default priority. Otherwise, the driver will use the given priority.
- `mcpwm_operator_config_t::update_gen_action_on_tez` sets whether to update the generator action when the timer counts to zero. Here and below, the timer refers to the one that is connected to the operator by `mcpwm_operator_connect_timer()`.
- `mcpwm_operator_config_t::update_gen_action_on_tep` sets whether to update the generator action when the timer counts to peak.
- `mcpwm_operator_config_t::update_gen_action_on_sync` sets whether to update the generator action when the timer takes a sync signal.
- `mcpwm_operator_config_t::update_dead_time_on_tez` sets whether to update the dead time when the timer counts to zero.
- `mcpwm_operator_config_t::update_dead_time_on_tep` sets whether to update the dead time when the timer counts to the peak.

¹ Different ESP chip series might have a different number of MCPWM resources (e.g., groups, timers, comparators, operators, generators, triggers and so on). Please refer to the [TRM] for details. The driver does not forbid you from applying for more MCPWM resources, but it returns an error when there are no hardware resources available. Please always check the return value when doing *Resource Allocation and Initialization*.

- `mcpwm_operator_config_t::update_dead_time_on_sync` sets whether to update the dead time when the timer takes a sync signal.

The `mcpwm_new_operator()` will return a pointer to the allocated operator object if the allocation succeeds. Otherwise, it will return an error code. Specifically, when there are no more free operators in the MCPWM group, this function will return the `ESP_ERR_NOT_FOUND` error. [Page 456, 1](#)

On the contrary, calling `mcpwm_del_operator()` function will free the allocated operator object.

MCPWM Comparators You can allocate a MCPWM comparator object by calling the `mcpwm_new_comparator()` function, with a MCPWM operator handle and configuration structure `mcpwm_comparator_config_t` as the parameter. The operator handle is created by `mcpwm_new_operator()`. The configuration structure is defined as:

- `mcpwm_comparator_config_t::intr_priority` sets the priority of the interrupt. If it is set to 0, the driver will allocate an interrupt with a default priority. Otherwise, the driver will use the given priority.
- `mcpwm_comparator_config_t::update_cmp_on_tez` sets whether to update the compare threshold when the timer counts to zero.
- `mcpwm_comparator_config_t::update_cmp_on_tep` sets whether to update the compare threshold when the timer counts to the peak.
- `mcpwm_comparator_config_t::update_cmp_on_sync` sets whether to update the compare threshold when the timer takes a sync signal.

The `mcpwm_new_comparator()` will return a pointer to the allocated comparator object if the allocation succeeds. Otherwise, it will return an error code. Specifically, when there are no more free comparators in the MCPWM operator, this function will return the `ESP_ERR_NOT_FOUND` error. [Page 456, 1](#)

On the contrary, calling the `mcpwm_del_comparator()` function will free the allocated comparator object.

There's another kind of comparator called "Event Comparator", which **can not** control the final PWM directly but only generates the ETM events at a configurable time stamp. You can allocate an event comparator by calling the `mcpwm_new_event_comparator()` function. This function will return the same handle type as `mcpwm_new_comparator()`, but with a different configuration structure `mcpwm_event_comparator_config_t`. For more information, please refer to [ETM Event and Task](#).

MCPWM Generators You can allocate a MCPWM generator object by calling the `mcpwm_new_generator()` function, with a MCPWM operator handle and configuration structure `mcpwm_generator_config_t` as the parameter. The operator handle is created by `mcpwm_new_operator()`. The configuration structure is defined as:

- `mcpwm_generator_config_t::gen_gpio_num` sets the GPIO number used by the generator.
- `mcpwm_generator_config_t::invert_pwm` sets whether to invert the PWM signal.
- `mcpwm_generator_config_t::io_loop_back` sets whether to enable the Loop-back mode. It is for debugging purposes only. It enables both the GPIO's input and output ability through the GPIO matrix peripheral.
- `mcpwm_generator_config_t::io_od_mode` configures the PWM GPIO as open-drain output.
- `mcpwm_generator_config_t::pull_up` and `mcpwm_generator_config_t::pull_down` controls whether to enable the internal pull-up and pull-down resistors accordingly.

The `mcpwm_new_generator()` will return a pointer to the allocated generator object if the allocation succeeds. Otherwise, it will return an error code. Specifically, when there are no more free generators in the MCPWM operator, this function will return the `ESP_ERR_NOT_FOUND` error. [Page 456, 1](#)

On the contrary, calling the `mcpwm_del_generator()` function will free the allocated generator object.

MCPWM Faults There are two types of faults: A fault signal reflected from the GPIO and a fault generated by software.

To allocate a GPIO fault object, you can call the `mcpwm_new_gpio_fault()` function, with the configuration structure `mcpwm_gpio_fault_config_t` as the parameter. The configuration structure is defined as:

- `mcpwm_gpio_fault_config_t::group_id` sets the MCPWM group ID. The ID should belong to [0, `SOC_MCPWM_GROUPS` - 1] range. Please note, GPIO faults located in different groups are totally independent, i.e., GPIO faults in group 0 can not be detected by the operator in group 1.
- `mcpwm_gpio_fault_config_t::intr_priority` sets the priority of the interrupt. If it is set to 0, the driver will allocate an interrupt with a default priority. Otherwise, the driver will use the given priority.
- `mcpwm_gpio_fault_config_t::gpio_num` sets the GPIO number used by the fault.
- `mcpwm_gpio_fault_config_t::active_level` sets the active level of the fault signal.
- `mcpwm_gpio_fault_config_t::pull_up` and `mcpwm_gpio_fault_config_t::pull_down` set whether to pull up and/or pull down the GPIO internally.
- `mcpwm_gpio_fault_config_t::io_loop_back` sets whether to enable the loopback mode. It is for debugging purposes only. It enables both the GPIO's input and output ability through the GPIO matrix peripheral.

The `mcpwm_new_gpio_fault()` will return a pointer to the allocated fault object if the allocation succeeds. Otherwise, it will return an error code. Specifically, when there are no more free GPIO faults in the MCPWM group, this function will return the `ESP_ERR_NOT_FOUND` error. ^{Page 456, 1}

Software fault object can be used to trigger a fault by calling the function `mcpwm_soft_fault_activate()` instead of waiting for a real fault signal on the GPIO. A software fault object can be allocated by calling the `mcpwm_new_soft_fault()` function, with configuration structure `mcpwm_soft_fault_config_t` as the parameter. Currently, this configuration structure is left for future purposes.

The `mcpwm_new_soft_fault()` function will return a pointer to the allocated fault object if the allocation succeeds. Otherwise, it will return an error code. Specifically, when there is no memory left for the fault object, this function will return the `ESP_ERR_NO_MEM` error. Although the software fault and GPIO fault are of different types, the returned fault handle is of the same type.

On the contrary, calling the `mcpwm_del_fault()` function will free the allocated fault object, this function works for both software and GPIO fault.

MCPWM Sync Sources The sync source is what can be used to synchronize the MCPWM timer and MCPWM capture timer. There are three types of sync sources: a sync source reflected from the GPIO, a sync source generated by software, and a sync source generated by an MCPWM timer event.

To allocate a GPIO sync source, you can call the `mcpwm_new_gpio_sync_src()` function, with configuration structure `mcpwm_gpio_sync_src_config_t` as the parameter. The configuration structure is defined as:

- `mcpwm_gpio_sync_src_config_t::group_id` sets the MCPWM group ID. The ID should belong to [0, `SOC_MCPWM_GROUPS` - 1] range. Please note, the GPIO sync sources located in different groups are totally independent, i.e., GPIO sync source in group 0 can not be detected by the timers in group 1.
- `mcpwm_gpio_sync_src_config_t::gpio_num` sets the GPIO number used by the sync source.
- `mcpwm_gpio_sync_src_config_t::active_neg` sets whether the sync signal is active on falling edges.
- `mcpwm_gpio_sync_src_config_t::pull_up` and `mcpwm_gpio_sync_src_config_t::pull_down` set whether to pull up and/or pull down the GPIO internally.
- `mcpwm_gpio_sync_src_config_t::io_loop_back` sets whether to enable the Loop-back mode. It is for debugging purposes only. It enables both the GPIO's input and output ability through the GPIO matrix peripheral.

The `mcpwm_new_gpio_sync_src()` will return a pointer to the allocated sync source object if the allocation succeeds. Otherwise, it will return an error code. Specifically, when there are no more free GPIO sync sources in the MCPWM group, this function will return the `ESP_ERR_NOT_FOUND` error. ^{Page 456, 1}

To allocate a timer event sync source, you can call the `mcpwm_new_timer_sync_src()` function, with configuration structure `mcpwm_timer_sync_src_config_t` as the parameter. The configuration structure is defined as:

- `mcpwm_timer_sync_src_config_t::timer_event` specifies on what timer event to generate the sync signal.
- `mcpwm_timer_sync_src_config_t::propagate_input_sync` sets whether to propagate the input sync signal (i.e., the input sync signal will be routed to its sync output).

The `mcpwm_new_timer_sync_src()` will return a pointer to the allocated sync source object if the allocation succeeds. Otherwise, it will return an error code. Specifically, if a sync source has been allocated from the same timer before, this function will return the `ESP_ERR_INVALID_STATE` error.

Last but not least, to allocate a software sync source, you can call the `mcpwm_new_soft_sync_src()` function, with configuration structure `mcpwm_soft_sync_config_t` as the parameter. Currently, this configuration structure is left for future purposes.

`mcpwm_new_soft_sync_src()` will return a pointer to the allocated sync source object if the allocation succeeds. Otherwise, it will return an error code. Specifically, when there is no memory left for the sync source object, this function will return the `ESP_ERR_NO_MEM` error. Please note, to make a software sync source take effect, do not forget to call `mcpwm_soft_sync_activate()`.

On the contrary, calling the `mcpwm_del_sync_src()` function will free the allocated sync source object. This function works for all types of sync sources.

MCPWM Capture Timer and Channels The MCPWM group has a dedicated timer which is used to capture the timestamp when a specific event occurred. The capture timer is connected to several independent channels, each channel is assigned a GPIO.

To allocate a capture timer, you can call the `mcpwm_new_capture_timer()` function, with configuration structure `mcpwm_capture_timer_config_t` as the parameter. The configuration structure is defined as:

- `mcpwm_capture_timer_config_t::group_id` sets the MCPWM group ID. The ID should belong to `[0, SOC_MCPWM_GROUPS - 1]` range.
- `mcpwm_capture_timer_config_t::clk_src` sets the clock source of the capture timer.
- `mcpwm_capture_timer_config_t::resolution_hz` The driver internally will set a proper divider based on the clock source and the resolution. If it is set to 0, the driver will pick an appropriate resolution on its own, and you can subsequently view the current timer resolution via `mcpwm_capture_timer_get_resolution()`.

The `mcpwm_new_capture_timer()` will return a pointer to the allocated capture timer object if the allocation succeeds. Otherwise, it will return an error code. Specifically, when there is no free capture timer left in the MCPWM group, this function will return the `ESP_ERR_NOT_FOUND` error. ^{Page 456, 1}

Next, to allocate a capture channel, you can call the `mcpwm_new_capture_channel()` function, with a capture timer handle and configuration structure `mcpwm_capture_channel_config_t` as the parameter. The configuration structure is defined as:

- `mcpwm_capture_channel_config_t::intr_priority` sets the priority of the interrupt. If it is set to 0, the driver will allocate an interrupt with a default priority. Otherwise, the driver will use the given priority.
- `mcpwm_capture_channel_config_t::gpio_num` sets the GPIO number used by the capture channel.
- `mcpwm_capture_channel_config_t::prescale` sets the prescaler of the input signal.
- `mcpwm_capture_channel_config_t::pos_edge` and `mcpwm_capture_channel_config_t::neg_edge` set whether to capture on the positive and/or falling edge of the input signal.
- `mcpwm_capture_channel_config_t::pull_up` and `mcpwm_capture_channel_config_t::pull_down` set whether to pull up and/or pull down the GPIO internally.
- `mcpwm_capture_channel_config_t::invert_cap_signal` sets whether to invert the capture signal.
- `mcpwm_capture_channel_config_t::io_loop_back` sets whether to enable the Loop-back mode. It is for debugging purposes only. It enables both the GPIO's input and output ability through the GPIO matrix peripheral.

The `mcpwm_new_capture_channel()` will return a pointer to the allocated capture channel object if the allocation succeeds. Otherwise, it will return an error code. Specifically, when there is no free capture channel left in the capture timer, this function will return the `ESP_ERR_NOT_FOUND` error.

On the contrary, calling `mcpwm_del_capture_channel()` and `mcpwm_del_capture_timer()` will free the allocated capture channel and timer object accordingly.

MCPWM Interrupt Priority MCPWM allows configuring interrupts separately for timer, operator, comparator, fault, and capture events. The interrupt priority is determined by the respective `config_t::intr_priority`. Additionally, events within the same MCPWM group share a common interrupt source. When registering multiple interrupt events, the interrupt priorities need to remain consistent.

Note: When registering multiple interrupt events within an MCPWM group, the driver will use the interrupt priority of the first registered event as the MCPWM group's interrupt priority.

Timer Operations and Events

Update Period The timer period is initialized by the `mcpwm_timer_config_t::period_ticks` parameter in `mcpwm_timer_config_t`. You can update the period at runtime by calling `mcpwm_timer_set_period()` function. The new period will take effect based on how you set the `mcpwm_timer_config_t::update_period_on_empty` and `mcpwm_timer_config_t::update_period_on_sync` parameters in `mcpwm_timer_config_t`. If none of them are set, the timer period will take effect immediately.

Register Timer Event Callbacks The MCPWM timer can generate different events at runtime. If you have some function that should be called when a particular event happens, you should hook your function to the interrupt service routine by calling `mcpwm_timer_register_event_callbacks()`. The callback function prototype is declared in `mcpwm_timer_event_cb_t`. All supported event callbacks are listed in the `mcpwm_timer_event_callbacks_t`:

- `mcpwm_timer_event_callbacks_t::on_full` sets the callback function for the timer when it counts to peak value.
- `mcpwm_timer_event_callbacks_t::on_empty` sets the callback function for the timer when it counts to zero.
- `mcpwm_timer_event_callbacks_t::on_stop` sets the callback function for the timer when it is stopped.

The callback functions above are called within the ISR context, so they should **not** attempt to block. For example, you may make sure that only FreeRTOS APIs with the `ISR` suffix are called within the function.

The parameter `user_data` of the `mcpwm_timer_register_event_callbacks()` function is used to save your own context. It is passed to each callback function directly.

This function will lazy the install interrupt service for the MCPWM timer without enabling it. It is only allowed to be called before `mcpwm_timer_enable()`, otherwise the `ESP_ERR_INVALID_STATE` error will be returned. See also [Enable and Disable timer](#) for more information.

Enable and Disable Timer Before doing IO control to the timer, you need to enable the timer first, by calling `mcpwm_timer_enable()`. This function:

- switches the timer state from **init** to **enable**.
- enables the interrupt service if it has been lazy installed by `mcpwm_timer_register_event_callbacks()`.
- acquire a proper power management lock if a specific clock source (e.g., PLL_160M clock) is selected. See also [Power management](#) for more information.

On the contrary, calling `mcpwm_timer_disable()` will put the timer driver back to the **init** state, disable the interrupt service and release the power management lock.

Start and Stop Timer The basic IO operation of a timer is to start and stop. Calling `mcpwm_timer_start_stop()` with different `mcpwm_timer_start_stop_cmd_t` commands can start the timer immediately or stop the timer at a specific event. What is more, you can even start the timer for only one round, which means, the timer will count to peak value or zero, and then stop itself.

Connect Timer with Operator The allocated MCPWM timer should be connected with an MCPWM operator by calling `mcpwm_operator_connect_timer()`, so that the operator can take that timer as its time base, and generate the required PWM waves. Please make sure the MCPWM timer and operator are in the same group. Otherwise, this function will return the `ESP_ERR_INVALID_ARG` error.

Comparator Operations and Events

Register Comparator Event Callbacks The MCPWM comparator can inform you when the timer counter equals the compare value. If you have some function that should be called when this event happens, you should hook your function to the interrupt service routine by calling `mcpwm_comparator_register_event_callbacks()`. The callback function prototype is declared in `mcpwm_compare_event_cb_t`. All supported event callbacks are listed in the `mcpwm_comparator_event_callbacks_t`:

- `mcpwm_comparator_event_callbacks_t::on_reach` sets the callback function for the comparator when the timer counter equals the compare value.

The callback function provides event-specific data of type `mcpwm_compare_event_data_t` to you. The callback function is called within the ISR context, so it should **not** attempt to block. For example, you may make sure that only FreeRTOS APIs with the `_ISR` suffix are called within the function.

The parameter `user_data` of `mcpwm_comparator_register_event_callbacks()` function is used to save your own context. It is passed to the callback function directly.

This function will lazy the installation of interrupt service for the MCPWM comparator, whereas the service can only be removed in `mcpwm_del_comparator`.

Note: It is not supported to register event callbacks for an **Event Comparator** because it can not generate any interrupt.

Set Compare Value You can set the compare value for the MCPWM comparator at runtime by calling `mcpwm_comparator_set_compare_value()`. There are a few points to note:

- A new compare value might not take effect immediately. The update time for the compare value is set by `mcpwm_comparator_config_t::update_cmp_on_tez` or `mcpwm_comparator_config_t::update_cmp_on_tep` or `mcpwm_comparator_config_t::update_cmp_on_tez_or_tep`.
- Make sure the operator has connected to one MCPWM timer already by `mcpwm_operator_connect_timer()`. Otherwise, it will return the error code `ESP_ERR_INVALID_STATE`.
- The compare value should not exceed the timer's count peak, otherwise, the compare event will never get triggered.

Generator Actions on Events

Set Generator Action on Timer Event One generator can set multiple actions on different timer events, by calling `mcpwm_generator_set_actions_on_timer_event()` with a variable number of action configurations. The action configuration is defined in `mcpwm_gen_timer_event_action_t`:

- `mcpwm_gen_timer_event_action_t::direction` specifies the timer direction. The supported directions are listed in `mcpwm_timer_direction_t`.
- `mcpwm_gen_timer_event_action_t::event` specifies the timer event. The supported timer events are listed in `mcpwm_timer_event_t`.
- `mcpwm_gen_timer_event_action_t::action` specifies the generator action to be taken. The supported actions are listed in `mcpwm_generator_action_t`.

There is a helper macro `MCPWM_GEN_TIMER_EVENT_ACTION` to simplify the construction of a timer event action entry.

Please note, the argument list of `mcpwm_generator_set_actions_on_timer_event()` **must** be terminated by `MCPWM_GEN_TIMER_EVENT_ACTION_END`.

You can also set the timer action one by one by calling `mcpwm_generator_set_action_on_timer_event()` without varargs.

Set Generator Action on Compare Event One generator can set multiple actions on different compare events, by calling `mcpwm_generator_set_actions_on_compare_event()` with a variable number of action configurations. The action configuration is defined in `mcpwm_gen_compare_event_action_t`:

- `mcpwm_gen_compare_event_action_t::direction` specifies the timer direction. The supported directions are listed in `mcpwm_timer_direction_t`.
- `mcpwm_gen_compare_event_action_t::comparator` specifies the comparator handle. See [MCPWM Comparators](#) for how to allocate a comparator.
- `mcpwm_gen_compare_event_action_t::action` specifies the generator action to be taken. The supported actions are listed in `mcpwm_generator_action_t`.

There is a helper macro `MCPWM_GEN_COMPARE_EVENT_ACTION` to simplify the construction of a compare event action entry.

Please note, the argument list of `mcpwm_generator_set_actions_on_compare_event()` **must** be terminated by `MCPWM_GEN_COMPARE_EVENT_ACTION_END`.

You can also set the compare action one by one by calling `mcpwm_generator_set_action_on_compare_event()` without varargs.

Set Generator Action on Fault Event One generator can set action on fault based trigger events, by calling `mcpwm_generator_set_action_on_fault_event()` with an action configurations. The action configuration is defined in `mcpwm_gen_fault_event_action_t`:

- `mcpwm_gen_fault_event_action_t::direction` specifies the timer direction. The supported directions are listed in `mcpwm_timer_direction_t`.
- `mcpwm_gen_fault_event_action_t::fault` specifies the fault used for the trigger. See [MCPWM Faults](#) for how to allocate a fault.
- `mcpwm_gen_fault_event_action_t::action` specifies the generator action to be taken. The supported actions are listed in `mcpwm_generator_action_t`.

When no free trigger slot is left in the operator to which the generator belongs, this function will return the `ESP_ERR_NOT_FOUND` error. ^{Page 456, 1}

The trigger only support GPIO fault. when the input is not a GPIO fault, this function will return the `ESP_ERR_NOT_SUPPORTED` error.

There is a helper macro `MCPWM_GEN_FAULT_EVENT_ACTION` to simplify the construction of a trigger event action entry.

Please note, fault event does not have variadic function like `mcpwm_generator_set_actions_on_fault_event()`.

Set Generator Action on Sync Event One generator can set action on sync based trigger events, by calling `mcpwm_generator_set_action_on_sync_event()` with an action configurations. The action configuration is defined in `mcpwm_gen_sync_event_action_t`:

- `mcpwm_gen_sync_event_action_t::direction` specifies the timer direction. The supported directions are listed in `mcpwm_timer_direction_t`.
- `mcpwm_gen_sync_event_action_t::sync` specifies the sync source used for the trigger. See [MCPWM Sync Sources](#) for how to allocate a sync source.
- `mcpwm_gen_sync_event_action_t::action` specifies the generator action to be taken. The supported actions are listed in `mcpwm_generator_action_t`.

When no free trigger slot is left in the operator to which the generator belongs, this function will return the `ESP_ERR_NOT_FOUND` error. ^{Page 456, 1}


The trigger only support one sync action, regardless of the kinds. When set sync actions more than once, this function will return the `ESP_ERR_INVALID_STATE` error.

There is a helper macro `MCPWM_GEN_SYNC_EVENT_ACTION` to simplify the construction of a trigger event action entry.


Please note, sync event does not have variadic function like `mcpwm_generator_set_actions_on_sync_event()`.

Generator Configurations for Classical PWM Waveforms This section will demonstrate the classical PWM waveforms that can be generated by the pair of generators. The code snippet that is used to generate the waveforms is also provided below the diagram. Some general summary:

- The **Symmetric** or **Asymmetric** of the waveforms is determined by the count mode of the MCPWM timer.
- The **active level** of the waveform pair is determined by the level of the PWM with a smaller duty cycle.
- The period of the PWM waveform is determined by the timer's period and count mode.
- The duty cycle of the PWM waveform is determined by the generator's various action combinations.



```
static void gen_action_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb,
    ↪mcpwm_cmpr_handle_t cmpa, mcpwm_cmpr_handle_t cmpb)
{
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_timer_event(gena,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
    ↪TIMER_EVENT_EMPTY, MCPWM_GEN_ACTION_HIGH)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_compare_event(gena,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpa,
    ↪MCPWM_GEN_ACTION_LOW)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_timer_event(genb,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
    ↪TIMER_EVENT_EMPTY, MCPWM_GEN_ACTION_HIGH)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_compare_event(genb,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpb,
    ↪MCPWM_GEN_ACTION_LOW)));
}
```



```
static void gen_action_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb,
    ↪mcpwm_cmpr_handle_t cmpa, mcpwm_cmpr_handle_t cmpb)
{
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_timer_event(gena,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
    ↪TIMER_EVENT_FULL, MCPWM_GEN_ACTION_LOW)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_compare_event(gena,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpa,
    ↪MCPWM_GEN_ACTION_HIGH)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_timer_event(genb,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
    ↪TIMER_EVENT_FULL, MCPWM_GEN_ACTION_LOW)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_compare_event(genb,
```


(continues on next page)

(continued from previous page)

```

        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpb,
↪MCPWM_GEN_ACTION_HIGH));
    }


```

Pulse Placement Asymmetric Waveform

```

static void gen_action_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb,
↪mcpwm_cmpr_handle_t cmpa, mcpwm_cmpr_handle_t cmpb)
{
    ESP_ERROR_CHECK(mcpwm_generator_set_actions_on_compare_event(gena,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpa,
↪MCPWM_GEN_ACTION_HIGH),
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpb,
↪MCPWM_GEN_ACTION_LOW),
        MCPWM_GEN_COMPARE_EVENT_ACTION_END()));
    ESP_ERROR_CHECK(mcpwm_generator_set_actions_on_timer_event(genb,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
↪TIMER_EVENT_EMPTY, MCPWM_GEN_ACTION_TOGGLE),
        MCPWM_GEN_TIMER_EVENT_ACTION_END()));
}


```

Dual Edge Asymmetric Waveform - Active Low

```

static void gen_action_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb,
↪mcpwm_cmpr_handle_t cmpa, mcpwm_cmpr_handle_t cmpb)
{
    ESP_ERROR_CHECK(mcpwm_generator_set_actions_on_compare_event(gena,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpa,
↪MCPWM_GEN_ACTION_HIGH),
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_DOWN,
↪cmpb, MCPWM_GEN_ACTION_LOW),
        MCPWM_GEN_COMPARE_EVENT_ACTION_END()));
    ESP_ERROR_CHECK(mcpwm_generator_set_actions_on_timer_event(genb,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
↪TIMER_EVENT_EMPTY, MCPWM_GEN_ACTION_LOW),
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_DOWN, MCPWM_
↪TIMER_EVENT_FULL, MCPWM_GEN_ACTION_HIGH),
        MCPWM_GEN_TIMER_EVENT_ACTION_END()));
}

```

Dual Edge Symmetric Waveform - Active Low

```
static void gen_action_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb,
↪mcpwm_cmpr_handle_t cmpa, mcpwm_cmpr_handle_t cmpb)
{
    ESP_ERROR_CHECK(mcpwm_generator_set_actions_on_compare_event(gena,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpa,
↪MCPWM_GEN_ACTION_HIGH),
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_DOWN,
↪cmpa, MCPWM_GEN_ACTION_LOW),
        MCPWM_GEN_COMPARE_EVENT_ACTION_END()));
    ESP_ERROR_CHECK(mcpwm_generator_set_actions_on_compare_event(genb,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpb,
↪MCPWM_GEN_ACTION_HIGH),
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_DOWN,
↪cmpb, MCPWM_GEN_ACTION_LOW),
        MCPWM_GEN_COMPARE_EVENT_ACTION_END()));
}
```

Dual Edge Symmetric Waveform, Co

Dual Edge Symmetric Waveform - Complementary



```
static void gen_action_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb,
↪mcpwm_cmpr_handle_t cmpa, mcpwm_cmpr_handle_t cmpb)
{
    ESP_ERROR_CHECK(mcpwm_generator_set_actions_on_compare_event(gena,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpa,
↪MCPWM_GEN_ACTION_HIGH),
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_DOWN,
↪cmpa, MCPWM_GEN_ACTION_LOW),
        MCPWM_GEN_COMPARE_EVENT_ACTION_END()));
    ESP_ERROR_CHECK(mcpwm_generator_set_actions_on_compare_event(genb,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpb,
↪MCPWM_GEN_ACTION_LOW),
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_DOWN,
↪cmpb, MCPWM_GEN_ACTION_HIGH),
        MCPWM_GEN_COMPARE_EVENT_ACTION_END()));
}
```

Dead Time In power electronics, the rectifier and inverter are commonly used. This requires the use of a rectifier bridge and an inverter bridge. Each bridge arm has two power electronic devices, such as MOSFET, IGBT, etc. The two MOSFETs on the same arm can not conduct at the same time, otherwise there will be a short circuit. The fact is that, although the PWM wave shows it is turning off the switch, the MOSFET still needs a small time window to make that happen. This requires an extra delay to be added to the existing PWM wave generated by setting *Generator Actions on Events*.

The dead time driver works like a **decorator**. This is also reflected in the function parameters of *mcpwm_generator_set_dead_time()*, where it takes the primary generator handle (*in_generator*), and returns a new generator (*out_generator*) after applying the dead time. Please note, if the *out_generator* and *in_generator* are the same, it means you are adding the time delay to the PWM waveform in an "in-place" fashion. In turn, if the *out_generator* and *in_generator* are different, it means you are deriving a new PWM waveform from the existing *in_generator*.

Dead time specific configuration is listed in the *mcpwm_dead_time_config_t* structure:

- *mcpwm_dead_time_config_t::posedge_delay_ticks* and *mcpwm_dead_time_config_t::negedge_delay_ticks* set the number of ticks to delay the PWM waveform on the rising and falling edge. Specifically, setting both

of them to zero means bypassing the dead time module. The resolution of the dead time tick is the same as the timer that is connected with the operator by `mcpwm_operator_connect_timer()`.

- `mcpwm_dead_time_config_t::invert_output` sets whether to invert the signal after applying the dead time, which can be used to control the delay edge polarity.


Warning: Due to the hardware limitation, one delay module (either posedge delay or negedge delay) can not be applied to multiple MCPWM generators at the same time. e.g., the following configuration is **invalid**:

```
mcpwm_dead_time_config_t dt_config = {
    .posedge_delay_ticks = 10,
};
// Set posedge delay to generator A
mcpwm_generator_set_dead_time(mcpwm_gen_a, mcpwm_gen_a, &dt_config);
// NOTE: This is invalid, you can not apply the posedge delay to another
// generator
mcpwm_generator_set_dead_time(mcpwm_gen_b, mcpwm_gen_b, &dt_config);
```

However, you can apply posedge delay to generator A and negedge delay to generator B. You can also set both posedge delay and negedge delay for generator B, while letting generator A bypass the dead time module. Note that if negedge delay and posedge delay are both set for generator A, generator B will not be available. Where generator A is the first generator requested through the operator handle and generator B is the second generator requested through an operator handle.

Note: It is also possible to generate the required dead time by setting *Generator Actions on Events*, especially by controlling edge placement using different comparators. However, if the more classical edge delay-based dead time with polarity control is required, then the dead time submodule should be used.

Dead Time Configurations for Classical PWM Waveforms This section demonstrates the classical PWM waveforms that can be generated by the dead time submodule. The code snippet that is used to generate the waveforms is also provided below the diagram.



```
static void gen_action_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb,
    ↪ mcpwm_cmpr_handle_t cmpa, mcpwm_cmpr_handle_t cmpb)
{
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_timer_event(gena,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
    ↪ TIMER_EVENT_EMPTY, MCPWM_GEN_ACTION_HIGH)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_compare_event(gena,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpa,
    ↪ MCPWM_GEN_ACTION_LOW)));
}

static void dead_time_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb)
{
    mcpwm_dead_time_config_t dead_time_config = {
        .posedge_delay_ticks = 50,
```


(continues on next page)

(continued from previous page)

```

        .negedge_delay_ticks = 0
    };
    ESP_ERROR_CHECK(mcpwm_generator_set_dead_time(gena, gena, &dead_time_config));
    dead_time_config.posedge_delay_ticks = 0;
    dead_time_config.negedge_delay_ticks = 100;
    dead_time_config.flags.invert_output = true;
    ESP_ERROR_CHECK(mcpwm_generator_set_dead_time(gena, genb, &dead_time_config));
}

```




```

static void gen_action_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb,
    ↪mcpwm_cmpr_handle_t cmpa, mcpwm_cmpr_handle_t cmpb)
{
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_timer_event(gena,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
    ↪TIMER_EVENT_EMPTY, MCPWM_GEN_ACTION_HIGH)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_compare_event(gena,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpa,
    ↪MCPWM_GEN_ACTION_LOW)));
}

static void dead_time_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb)
{
    mcpwm_dead_time_config_t dead_time_config = {
        .posedge_delay_ticks = 50,
        .negedge_delay_ticks = 0,
        .flags.invert_output = true
    };
    ESP_ERROR_CHECK(mcpwm_generator_set_dead_time(gena, gena, &dead_time_config));
    dead_time_config.posedge_delay_ticks = 0;
    dead_time_config.negedge_delay_ticks = 100;
    dead_time_config.flags.invert_output = false;
    ESP_ERROR_CHECK(mcpwm_generator_set_dead_time(gena, genb, &dead_time_config));
}

```



```

static void gen_action_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb,
    ↪mcpwm_cmpr_handle_t cmpa, mcpwm_cmpr_handle_t cmpb)
{
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_timer_event(gena,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
    ↪TIMER_EVENT_EMPTY, MCPWM_GEN_ACTION_HIGH)));
}

```

(continues on next page)


(continued from previous page)

```

    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_compare_event(gena,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpa,
        ↪MCPWM_GEN_ACTION_LOW)));
}

static void dead_time_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb)
{
    mcpwm_dead_time_config_t dead_time_config = {
        .posedge_delay_ticks = 50,
        .negedge_delay_ticks = 0,
    };
    ESP_ERROR_CHECK(mcpwm_generator_set_dead_time(gena, gena, &dead_time_config));
    dead_time_config.posedge_delay_ticks = 0;
    dead_time_config.negedge_delay_ticks = 100;
    ESP_ERROR_CHECK(mcpwm_generator_set_dead_time(gena, genb, &dead_time_config));
}

```




```

static void gen_action_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb,
    ↪mcpwm_cmpr_handle_t cmpa, mcpwm_cmpr_handle_t cmpb)
{
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_timer_event(gena,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
        ↪TIMER_EVENT_EMPTY, MCPWM_GEN_ACTION_HIGH)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_compare_event(gena,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpa,
        ↪MCPWM_GEN_ACTION_LOW)));
}

static void dead_time_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb)
{
    mcpwm_dead_time_config_t dead_time_config = {
        .posedge_delay_ticks = 50,
        .negedge_delay_ticks = 0,
        .flags.invert_output = true
    };
    ESP_ERROR_CHECK(mcpwm_generator_set_dead_time(gena, gena, &dead_time_config));
    dead_time_config.posedge_delay_ticks = 0;
    dead_time_config.negedge_delay_ticks = 100;
    ESP_ERROR_CHECK(mcpwm_generator_set_dead_time(gena, genb, &dead_time_config));
}

```


RED on A, Byp

**Rising Delay on PWMA and Bypass Dead Time for PWMB**

```
static void gen_action_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb,
↪mcpwm_cmpr_handle_t cmpa, mcpwm_cmpr_handle_t cmpb)
{
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_timer_event(gena,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
↪TIMER_EVENT_EMPTY, MCPWM_GEN_ACTION_HIGH)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_compare_event(gena,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpa,
↪MCPWM_GEN_ACTION_LOW)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_timer_event(genb,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
↪TIMER_EVENT_EMPTY, MCPWM_GEN_ACTION_HIGH)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_compare_event(genb,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpb,
↪MCPWM_GEN_ACTION_LOW)));
}

static void dead_time_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb)
{
    mcpwm_dead_time_config_t dead_time_config = {
        .posedge_delay_ticks = 50,
        .negedge_delay_ticks = 0,
    };
    // apply deadtime to generator_a
    ESP_ERROR_CHECK(mcpwm_generator_set_dead_time(gena, gena, &dead_time_config));
    // bypass deadtime module for generator_b
    dead_time_config.posedge_delay_ticks = 0;
    ESP_ERROR_CHECK(mcpwm_generator_set_dead_time(genb, genb, &dead_time_config));
}
```

FED on B, Byp

**Falling Delay on PWMB and Bypass Dead Time for PWMA**

```
static void gen_action_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb,
↪mcpwm_cmpr_handle_t cmpa, mcpwm_cmpr_handle_t cmpb)
{
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_timer_event(gena,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
↪TIMER_EVENT_EMPTY, MCPWM_GEN_ACTION_HIGH)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_compare_event(gena,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpa,
↪MCPWM_GEN_ACTION_LOW)));
}
```

(continues on next page)

(continued from previous page)


```

    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_timer_event(genb,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
    ↪TIMER_EVENT_EMPTY, MCPWM_GEN_ACTION_HIGH)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_compare_event(genb,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpb, ↪
    ↪MCPWM_GEN_ACTION_LOW)));
}

static void dead_time_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb)
{
    mcpwm_dead_time_config_t dead_time_config = {
        .posedge_delay_ticks = 0,
        .negedge_delay_ticks = 0,
    };
    // generator_a bypass the deadtime module (no delay)
    ESP_ERROR_CHECK(mcpwm_generator_set_dead_time(gena, gena, &dead_time_config));
    // apply dead time to generator_b
    dead_time_config.negedge_delay_ticks = 50;
    ESP_ERROR_CHECK(mcpwm_generator_set_dead_time(genb, genb, &dead_time_config));
}

```

Rising and Falling Delay on PWMB and Bypass Dead Time for PWMA



```

static void gen_action_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb, ↪
    ↪mcpwm_cmpr_handle_t cmpa, mcpwm_cmpr_handle_t cmpb)
{
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_timer_event(gena,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
    ↪TIMER_EVENT_EMPTY, MCPWM_GEN_ACTION_HIGH)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_compare_event(gena,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpa, ↪
    ↪MCPWM_GEN_ACTION_LOW)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_timer_event(genb,
        MCPWM_GEN_TIMER_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, MCPWM_
    ↪TIMER_EVENT_EMPTY, MCPWM_GEN_ACTION_HIGH)));
    ESP_ERROR_CHECK(mcpwm_generator_set_action_on_compare_event(genb,
        MCPWM_GEN_COMPARE_EVENT_ACTION(MCPWM_TIMER_DIRECTION_UP, cmpb, ↪
    ↪MCPWM_GEN_ACTION_LOW)));
}

static void dead_time_config(mcpwm_gen_handle_t gena, mcpwm_gen_handle_t genb)
{
    mcpwm_dead_time_config_t dead_time_config = {
        .posedge_delay_ticks = 0,
        .negedge_delay_ticks = 0,
    };
    // generator_a bypass the deadtime module (no delay)
    ESP_ERROR_CHECK(mcpwm_generator_set_dead_time(gena, gena, &dead_time_config));
}

```

(continues on next page)

(continued from previous page)

```
// apply dead time on both edge for generator_b
dead_time_config.negedge_delay_ticks = 50;
dead_time_config.posedge_delay_ticks = 50;
ESP_ERROR_CHECK(mcpwm_generator_set_dead_time(genb, genb, &dead_time_config));
}
```

Carrier Modulation The MCPWM operator has a carrier submodule that can be used if galvanic isolation from the motor driver is required (e.g., isolated digital power application) by passing the PWM output signals through transformers. Any of the PWM output signals may be at 100% duty and not changing whenever a motor is required to run steadily at the full load. Coupling with non-alternating signals with a transformer is problematic, so the signals are modulated by the carrier submodule to create an AC waveform, to make the coupling possible.

To configure the carrier submodule, you can call `mcpwm_operator_apply_carrier()`, and provide configuration structure `mcpwm_carrier_config_t`:

- `mcpwm_carrier_config_t::clk_src` sets the clock source of the carrier.
- `mcpwm_carrier_config_t::frequency_hz` indicates carrier frequency in Hz.
- `mcpwm_carrier_config_t::duty_cycle` indicates the duty cycle of the carrier. Note that, the supported choices of the duty cycle are discrete, the driver searches for the nearest one based on your configuration.
- `mcpwm_carrier_config_t::first_pulse_duration_us` indicates the duration of the first pulse in microseconds. The resolution of the first pulse duration is determined by the carrier frequency you set in the `mcpwm_carrier_config_t::frequency_hz`. The first pulse duration can not be zero, and it has to be at least one period of the carrier. A longer pulse width can help conduct the inductance quicker.
- `mcpwm_carrier_config_t::invert_before_modulate` and `mcpwm_carrier_config_t::invert_after_modulate` set whether to invert the carrier output before and after modulation.

Specifically, the carrier submodule can be disabled by calling `mcpwm_operator_apply_carrier()` with a NULL configuration.

Faults and Brake Actions The MCPWM operator is able to sense external signals with information about the failure of the motor, the power driver or any other device connected. These failure signals are encapsulated into MCPWM fault objects.

You should determine possible failure modes of the motor and what action should be performed on detection of a particular fault, e.g., drive all outputs low for a brushed motor, lock current state for a stepper motor, etc. Because of this action, the motor should be put into a safe state to reduce the likelihood of damage caused by the fault.

Set Operator Brake Mode on Fault The way that MCPWM operator reacts to the fault is called **Brake**. The MCPWM operator can be configured to perform different brake modes for each fault object by calling `mcpwm_operator_set_brake_on_fault()`. Specific brake configuration is passed as a structure `mcpwm_brake_config_t`:

- `mcpwm_brake_config_t::fault` sets which fault the operator should react to.
- `mcpwm_brake_config_t::brake_mode` sets the brake mode that should be used for the fault. The supported brake modes are listed in the `mcpwm_operator_brake_mode_t`. For `MCPWM_OPER_BRAKE_MODE_CBC` mode, the operator recovers itself automatically as long as the fault disappears. You can specify the recovery time in `mcpwm_brake_config_t::cbc_recover_on_tez` and `mcpwm_brake_config_t::cbc_recover_on_tep`. For `MCPWM_OPER_BRAKE_MODE_OST` mode, the operator can not recover even though the fault disappears. You have to call `mcpwm_operator_recover_from_fault()` to manually recover it.

Set Generator Action on Brake Event One generator can set multiple actions on different brake events, by calling `mcpwm_generator_set_actions_on_brake_event()` with a variable number of action configurations. The action configuration is defined in `mcpwm_gen_brake_event_action_t`:

- `mcpwm_gen_brake_event_action_t::direction` specifies the timer direction. The supported directions are listed in `mcpwm_timer_direction_t`.
- `mcpwm_gen_brake_event_action_t::brake_mode` specifies the brake mode. The supported brake modes are listed in the `mcpwm_operator_brake_mode_t`.
- `mcpwm_gen_brake_event_action_t::action` specifies the generator action to be taken. The supported actions are listed in `mcpwm_generator_action_t`.

There is a helper macro `MCPWM_GEN_BRAKE_EVENT_ACTION` to simplify the construction of a brake event action entry.

Please note, the argument list of `mcpwm_generator_set_actions_on_brake_event()` **must** be terminated by `MCPWM_GEN_BRAKE_EVENT_ACTION_END`.

You can also set the brake action one by one by calling `mcpwm_generator_set_action_on_brake_event()` without varargs.

Register Fault Event Callbacks The MCPWM fault detector can inform you when it detects a valid fault or a fault signal disappears. If you have some function that should be called when such an event happens, you should hook your function to the interrupt service routine by calling `mcpwm_fault_register_event_callbacks()`. The callback function prototype is declared in `mcpwm_fault_event_cb_t`. All supported event callbacks are listed in the `mcpwm_fault_event_callbacks_t`:

- `mcpwm_fault_event_callbacks_t::on_fault_enter` sets the callback function that will be called when a fault is detected.
- `mcpwm_fault_event_callbacks_t::on_fault_exit` sets the callback function that will be called when a fault is cleared.

The callback function is called within the ISR context, so it should **not** attempt to block. For example, you may make sure that only FreeRTOS APIs with the `_ISR` suffix are called within the function.

The parameter `user_data` of `mcpwm_fault_register_event_callbacks()` function is used to save your own context. It is passed to the callback function directly.

This function will lazy the install interrupt service for the MCPWM fault, whereas the service can only be removed in `mcpwm_del_fault`.

Register Brake Event Callbacks The MCPWM operator can inform you when it is going to take a brake action. If you have some function that should be called when this event happens, you should hook your function to the interrupt service routine by calling `mcpwm_operator_register_event_callbacks()`. The callback function prototype is declared in `mcpwm_brake_event_cb_t`. All supported event callbacks are listed in the `mcpwm_operator_event_callbacks_t`:

- `mcpwm_operator_event_callbacks_t::on_brake_cbc` sets the callback function that will be called when the operator is going to take a **CBC** action.
- `mcpwm_operator_event_callbacks_t::on_brake_ost` sets the callback function that will be called when the operator is going to take an **OST** action.

The callback function is called within the ISR context, so it should **not** attempt to block. For example, you may make sure that only FreeRTOS APIs with the `_ISR` suffix are called within the function.

The parameter `user_data` of the `mcpwm_operator_register_event_callbacks()` function is used to save your own context. It will be passed to the callback function directly.

This function will lazy the install interrupt service for the MCPWM operator, whereas the service can only be removed in `mcpwm_del_operator`.

Generator Force Actions Software can override generator output level at runtime, by calling `mcpwm_generator_set_force_level()`. The software force level always has a higher priority than other event actions set in e.g., `mcpwm_generator_set_actions_on_timer_event()`.

- Set the `level` to -1 means to disable the force action, and the generator's output level will be controlled by the event actions again.

- Set the `hold_on` to true, and the force output level will keep alive until it is removed by assigning `level` to -1.
- Set the `hole_on` to false, the force output level will only be active for a short time, and any upcoming event can override it.

Synchronization When a sync signal is taken by the MCPWM timer, the timer will be forced into a pre-defined **phase**, where the phase is determined by count value and count direction. You can set the sync phase by calling `mcpwm_timer_set_phase_on_sync()`. The sync phase configuration is defined in `mcpwm_timer_sync_phase_config_t` structure:

- `mcpwm_timer_sync_phase_config_t::sync_src` sets the sync signal source. See [MCPWM Sync Sources](#) for how to create a sync source object. Specifically, if this is set to `NULL`, the driver will disable the sync feature for the MCPWM timer.
- `mcpwm_timer_sync_phase_config_t::count_value` sets the count value to load when the sync signal is taken.
- `mcpwm_timer_sync_phase_config_t::direction` sets the count direction when the sync signal is taken.

Likewise, the [MCPWM Capture Timer](#) can be synced as well. You can set the sync phase for the capture timer by calling `mcpwm_capture_timer_set_phase_on_sync()`. The sync phase configuration is defined in `mcpwm_capture_timer_sync_phase_config_t` structure:

- `mcpwm_capture_timer_sync_phase_config_t::sync_src` sets the sync signal source. See [MCPWM Sync Sources](#) for how to create a sync source object. Specifically, if this is set to `NULL`, the driver will disable the sync feature for the MCPWM capture timer.
- `mcpwm_capture_timer_sync_phase_config_t::count_value` sets the count value to load when the sync signal is taken.
- `mcpwm_capture_timer_sync_phase_config_t::direction` sets the count direction when the sync signal is taken. Note that, different from MCPWM Timer, the capture timer can only support one count direction: `MCPWM_TIMER_DIRECTION_UP`.




Fig. 15: GPIO Sync All MCPWM Timers

Sync Timers by GPIO

```

static void example_setup_sync_strategy(mcpwm_timer_handle_t timers[])
{
    mcpwm_sync_handle_t gpio_sync_source = NULL;
    mcpwm_gpio_sync_src_config_t gpio_sync_config = {
        .group_id = 0,           // GPIO fault should be in the same group of
        ↪the above timers
        .gpio_num = EXAMPLE_SYNC_GPIO,
        .flags.pull_down = true,
        .flags.active_neg = false, // By default, a posedge pulse can trigger a
        ↪sync event
    };
    ESP_ERROR_CHECK(mcpwm_new_gpio_sync_src(&gpio_sync_config, &gpio_sync_source));

    mcpwm_timer_sync_phase_config_t sync_phase_config = {
        .count_value = 0,           // sync phase: target count value
        .direction = MCPWM_TIMER_DIRECTION_UP, // sync phase: count direction
        .sync_src = gpio_sync_source, // sync source
    };
    for (int i = 0; i < 3; i++) {
        ESP_ERROR_CHECK(mcpwm_timer_set_phase_on_sync(timers[i], &sync_phase_
        ↪config));
    }
}

```

Capture The basic functionality of MCPWM capture is to record the time when any pulse edge of the capture signal turns active. Then you can get the pulse width and convert it into other physical quantities like distance or speed in the capture callback function. For example, in the BLDC (Brushless DC, see figure below) scenario, you can use the capture submodule to sense the rotor position from the Hall sensor.




Fig. 16: MCPWM BLDC with Hall Sensor

The capture timer is usually connected to several capture channels. Please refer to [MCPWM Capture Timer and Channels](#) for more information about resource allocation.

Register Capture Event Callbacks The MCPWM capture channel can inform you when there is a valid edge detected on the signal. You have to register a callback function to get the timer count value of the captured moment, by calling `mcpwm_capture_channel_register_event_callbacks()`. The callback function prototype is declared in `mcpwm_capture_event_cb_t`. All supported capture callbacks are listed in the `mcpwm_capture_event_callbacks_t`:

- `mcpwm_capture_event_callbacks_t::on_cap` sets the callback function for the capture channel when a valid edge is detected.

The callback function provides event-specific data of type `mcpwm_capture_event_data_t`, so that you can get the edge of the capture signal in `mcpwm_capture_event_data_t::cap_edge` and the count value of that moment in `mcpwm_capture_event_data_t::cap_value`. To convert the capture count into a timestamp, you need to know the resolution of the capture timer by calling `mcpwm_capture_timer_get_resolution()`.

The callback function is called within the ISR context, so it should **not** attempt to block. For example, you may make sure that only FreeRTOS APIs with the `_ISR` suffix are called within the function.

The parameter `user_data` of `mcpwm_capture_channel_register_event_callbacks()` function is used to save your context. It is passed to the callback function directly.

This function will lazy install interrupt service for the MCPWM capture channel, whereas the service can only be removed in `mcpwm_del_capture_channel`.

Enable and Disable Capture Channel The capture channel is not enabled after allocation by `mcpwm_new_capture_channel()`. You should call `mcpwm_capture_channel_enable()` and `mcpwm_capture_channel_disable()` accordingly to enable or disable the channel. If the interrupt service is lazy installed during registering event callbacks for the channel in `mcpwm_capture_channel_register_event_callbacks()`, `mcpwm_capture_channel_enable()` will enable the interrupt service as well.

Enable and Disable Capture Timer Before doing IO control to the capture timer, you need to enable the timer first, by calling `mcpwm_capture_timer_enable()`. Internally, this function:

- switches the capture timer state from **init** to **enable**.
- acquires a proper power management lock if a specific clock source (e.g., APB clock) is selected. See also [Power management](#) for more information.

On the contrary, calling `mcpwm_capture_timer_disable()` will put the timer driver back to **init** state, and release the power management lock.

Start and Stop Capture Timer The basic IO operation of a capture timer is to start and stop. Calling `mcpwm_capture_timer_start()` can start the timer and calling `mcpwm_capture_timer_stop()` can stop the timer immediately.

Trigger a Software Capture Event Sometimes, the software also wants to trigger a "fake" capture event. The `mcpwm_capture_channel_trigger_soft_catch()` is provided for that purpose. Please note that, even though it is a "fake" capture event, it can still cause an interrupt, thus your capture event callback function gets invoked as well.

ETM Event and Task MCPWM comparator is able to generate events that can interact with the [ETM](#) module. The supported events are listed in the `mcpwm_comparator_etm_event_type_t`. You can call `mcpwm_comparator_new_etm_event()` to get the corresponding ETM event handle.

For how to connect the event and task to an ETM channel, please refer to the [ETM](#) documentation.

Power Management When power management is enabled (i.e., `CONFIG_PM_ENABLE` is on), the system will adjust the PLL and APB frequency before going into Light-sleep, thus potentially changing the period of an MCPWM timers' counting step and leading to inaccurate time-keeping.

However, the driver can prevent the system from going into Light-sleep by acquiring a power management lock of type `ESP_PM_NO_LIGHT_SLEEP`. Whenever the driver creates an MCPWM timer instance that has selected PLL as its clock source, the driver guarantees that the power management lock is acquired when enabling the timer by `mcpwm_timer_enable()`. On the contrary, the driver releases the lock when `mcpwm_timer_disable()` is called for that timer.

Likewise, whenever the driver creates an MCPWM capture timer instance, the driver guarantees that the power management lock is acquired when enabling the timer by `mcpwm_capture_timer_enable()`. And releases the lock in `mcpwm_capture_timer_disable()`.

IRAM Safe By default, the MCPWM interrupt will be deferred when the Cache is disabled for reasons like writing/erasing Flash. Thus the event callback functions will not get executed in time, which is not expected in a real-time application.

There is a Kconfig option `CONFIG_MCPWM_ISR_IRAM_SAFE` that:

- enables the interrupt to be serviced even when the cache is disabled
- places all functions used by the ISR into IRAM²
- places the driver object into DRAM (in case it is mapped to PSRAM by accident)

This allows the interrupt to run while the cache is disabled but comes at the cost of increased IRAM consumption.

There is another Kconfig option `CONFIG_MCPWM_CTRL_FUNC_IN_IRAM` that can put commonly used IO control functions into IRAM as well. So, these functions can also be executable when the cache is disabled. The IO control function is as follows:

- `mcpwm_comparator_set_compare_value()`
- `mcpwm_timer_set_period()`

Thread Safety The factory functions like `mcpwm_new_timer()` are guaranteed to be thread-safe by the driver, which means, you can call it from different RTOS tasks without protection by extra locks.

The following function is allowed to run under the ISR context, as the driver uses a critical section to prevent them from being called concurrently in the task and ISR.

- `mcpwm_comparator_set_compare_value()`
- `mcpwm_timer_set_period()`

Other functions that are not related to *Resource Allocation and Initialization*, are not thread-safe. Thus, you should avoid calling them in different tasks without mutex protection.

Kconfig Options

- `CONFIG_MCPWM_ISR_IRAM_SAFE` controls whether the default ISR handler can work when the cache is disabled, see *IRAM Safe* for more information.
- `CONFIG_MCPWM_CTRL_FUNC_IN_IRAM` controls where to place the MCPWM control functions (IRAM or flash), see *IRAM Safe* for more information.
- `CONFIG_MCPWM_ENABLE_DEBUG_LOG` is used to enable the debug log output. Enabling this option will increase the firmware binary size.

Application Examples

- Brushed DC motor speed control by PID algorithm: [peripherals/mcpwm/mcpwm_bdc_speed_control](#)
- BLDC motor control with hall sensor feedback: [peripherals/mcpwm/mcpwm_bldc_hall_control](#)
- Ultrasonic sensor (HC-SR04) distance measurement: [peripherals/mcpwm/mcpwm_capture_hc_sr04](#)

² The callback function and the sub-functions invoked by itself should also be placed in IRAM. You need to take care of this by yourself.

- Servo motor angle control: [peripherals/mcpwm/mcpwm_servo_control](#)
- MCPWM synchronization between timers: [peripherals/mcpwm/mcpwm_sync](#)

API Reference

Header File

- [components/driver/mcpwm/include/driver/mcpwm_timer.h](#)
- This header file can be included with:

```
#include "driver/mcpwm_timer.h"
```

- This header file is a part of the API provided by the driver component. To declare that your component depends on driver, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

[esp_err_t mcpwm_new_timer](#) (const [mcpwm_timer_config_t](#) *config, [mcpwm_timer_handle_t](#) *ret_timer)

Create MCPWM timer.

Parameters

- **config** -- [in] MCPWM timer configuration
- **ret_timer** -- [out] Returned MCPWM timer handle

Returns

- ESP_OK: Create MCPWM timer successfully
- ESP_ERR_INVALID_ARG: Create MCPWM timer failed because of invalid argument
- ESP_ERR_NO_MEM: Create MCPWM timer failed because out of memory
- ESP_ERR_NOT_FOUND: Create MCPWM timer failed because all hardware timers are used up and no more free one
- ESP_FAIL: Create MCPWM timer failed because of other error

[esp_err_t mcpwm_del_timer](#) ([mcpwm_timer_handle_t](#) timer)

Delete MCPWM timer.

Parameters **timer** -- [in] MCPWM timer handle, allocated by [mcpwm_new_timer](#) ()

Returns

- ESP_OK: Delete MCPWM timer successfully
- ESP_ERR_INVALID_ARG: Delete MCPWM timer failed because of invalid argument
- ESP_ERR_INVALID_STATE: Delete MCPWM timer failed because timer is not in init state
- ESP_FAIL: Delete MCPWM timer failed because of other error

[esp_err_t mcpwm_timer_set_period](#) ([mcpwm_timer_handle_t](#) timer, uint32_t period_ticks)

Set a new period for MCPWM timer.

Note: If [mcpwm_timer_config_t::update_period_on_empty](#) and [mcpwm_timer_config_t::update_period_on_sync](#) are not set, the new period will take effect immediately. Otherwise, the new period will take effect when timer counts to zero or on sync event.

Note: You may need to use [mcpwm_comparator_set_compare_value](#) to set a new compare value for MCPWM comparator in order to keep the same PWM duty cycle.

Parameters

- **timer** -- [in] MCPWM timer handle, allocated by `mcpwm_new_timer`
- **period_ticks** -- [in] New period in count ticks

Returns

- ESP_OK: Set new period for MCPWM timer successfully
- ESP_ERR_INVALID_ARG: Set new period for MCPWM timer failed because of invalid argument
- ESP_FAIL: Set new period for MCPWM timer failed because of other error

esp_err_t **mcpwm_timer_enable** (*mcpwm_timer_handle_t* timer)

Enable MCPWM timer.

Parameters **timer** -- [in] MCPWM timer handle, allocated by `mcpwm_new_timer()`

Returns

- ESP_OK: Enable MCPWM timer successfully
- ESP_ERR_INVALID_ARG: Enable MCPWM timer failed because of invalid argument
- ESP_ERR_INVALID_STATE: Enable MCPWM timer failed because timer is enabled already
- ESP_FAIL: Enable MCPWM timer failed because of other error

esp_err_t **mcpwm_timer_disable** (*mcpwm_timer_handle_t* timer)

Disable MCPWM timer.

Parameters **timer** -- [in] MCPWM timer handle, allocated by `mcpwm_new_timer()`

Returns

- ESP_OK: Disable MCPWM timer successfully
- ESP_ERR_INVALID_ARG: Disable MCPWM timer failed because of invalid argument
- ESP_ERR_INVALID_STATE: Disable MCPWM timer failed because timer is disabled already
- ESP_FAIL: Disable MCPWM timer failed because of other error

esp_err_t **mcpwm_timer_start_stop** (*mcpwm_timer_handle_t* timer, *mcpwm_timer_start_stop_cmd_t* command)

Send specific start/stop commands to MCPWM timer.

Parameters

- **timer** -- [in] MCPWM timer handle, allocated by `mcpwm_new_timer()`
- **command** -- [in] Supported command list for MCPWM timer

Returns

- ESP_OK: Start or stop MCPWM timer successfully
- ESP_ERR_INVALID_ARG: Start or stop MCPWM timer failed because of invalid argument
- ESP_ERR_INVALID_STATE: Start or stop MCPWM timer failed because timer is not enabled
- ESP_FAIL: Start or stop MCPWM timer failed because of other error

esp_err_t **mcpwm_timer_register_event_callbacks** (*mcpwm_timer_handle_t* timer, const *mcpwm_timer_event_callbacks_t* *cbs, void *user_data)

Set event callbacks for MCPWM timer.

Note: The first call to this function needs to be before the call to `mcpwm_timer_enable`

Note: User can deregister a previously registered callback by calling this function and setting the callback member in the `cbs` structure to NULL.

Parameters

- **timer** -- [in] MCPWM timer handle, allocated by `mcpwm_new_timer()`
- **cbs** -- [in] Group of callback functions
- **user_data** -- [in] User data, which will be passed to callback functions directly

Returns

- ESP_OK: Set event callbacks successfully
- ESP_ERR_INVALID_ARG: Set event callbacks failed because of invalid argument
- ESP_ERR_INVALID_STATE: Set event callbacks failed because timer is not in init state
- ESP_FAIL: Set event callbacks failed because of other error

esp_err_t **mcpwm_timer_set_phase_on_sync** (*mcpwm_timer_handle_t* timer, const *mcpwm_timer_sync_phase_config_t* *config)

Set sync phase for MCPWM timer.

Parameters

- **timer** -- [in] MCPWM timer handle, allocated by `mcpwm_new_timer()`
- **config** -- [in] MCPWM timer sync phase configuration

Returns

- ESP_OK: Set sync phase for MCPWM timer successfully
- ESP_ERR_INVALID_ARG: Set sync phase for MCPWM timer failed because of invalid argument
- ESP_FAIL: Set sync phase for MCPWM timer failed because of other error

Structures

struct **mcpwm_timer_event_callbacks_t**

Group of supported MCPWM timer event callbacks.

Note: The callbacks are all running under ISR environment

Public Members

mcpwm_timer_event_cb_t **on_full**

callback function when MCPWM timer counts to peak value

mcpwm_timer_event_cb_t **on_empty**

callback function when MCPWM timer counts to zero

mcpwm_timer_event_cb_t **on_stop**

callback function when MCPWM timer stops

struct **mcpwm_timer_config_t**

MCPWM timer configuration.

Public Members

int **group_id**

Specify from which group to allocate the MCPWM timer

mcpwm_timer_clock_source_t **clk_src**

MCPWM timer clock source

uint32_t **resolution_hz**

Counter resolution in Hz The step size of each count tick equals to (1 / resolution_hz) seconds

mcpwm_timer_count_mode_t **count_mode**

Count mode

uint32_t **period_ticks**

Number of count ticks within a period

int **intr_priority**

MCPWM timer interrupt priority, if set to 0, the driver will try to allocate an interrupt with a relative low priority (1,2,3)

uint32_t **update_period_on_empty**

Whether to update period when timer counts to zero

uint32_t **update_period_on_sync**

Whether to update period on sync event

struct *mcpwm_timer_config_t*::[anonymous] **flags**

Extra configuration flags for timer

struct **mcpwm_timer_sync_phase_config_t**

MCPWM Timer sync phase configuration.

Public Members

mcpwm_sync_handle_t **sync_src**

The sync event source. Set to NULL will disable the timer being synced by others

uint32_t **count_value**

The count value that should lock to upon sync event

mcpwm_timer_direction_t **direction**

The count direction that should lock to upon sync event

Header File

- [components/driver/mcpwm/include/driver/mcpwm_oper.h](#)
- This header file can be included with:

```
#include "driver/mcpwm_oper.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t **mcpwm_new_operator** (const *mcpwm_operator_config_t* *config, *mcpwm_oper_handle_t* *ret_oper)

Create MCPWM operator.

Parameters

- **config** -- [in] MCPWM operator configuration
- **ret_oper** -- [out] Returned MCPWM operator handle

Returns

- ESP_OK: Create MCPWM operator successfully
- ESP_ERR_INVALID_ARG: Create MCPWM operator failed because of invalid argument
- ESP_ERR_NO_MEM: Create MCPWM operator failed because out of memory
- ESP_ERR_NOT_FOUND: Create MCPWM operator failed because can't find free resource
- ESP_FAIL: Create MCPWM operator failed because of other error

esp_err_t **mcpwm_del_operator** (*mcpwm_oper_handle_t* oper)

Delete MCPWM operator.

Parameters **oper** -- [in] MCPWM operator, allocated by `mcpwm_new_operator()`

Returns

- ESP_OK: Delete MCPWM operator successfully
- ESP_ERR_INVALID_ARG: Delete MCPWM operator failed because of invalid argument
- ESP_FAIL: Delete MCPWM operator failed because of other error

esp_err_t **mcpwm_operator_connect_timer** (*mcpwm_oper_handle_t* oper, *mcpwm_timer_handle_t* timer)

Connect MCPWM operator and timer, so that the operator can be driven by the timer.

Parameters

- **oper** -- [in] MCPWM operator handle, allocated by `mcpwm_new_operator()`
- **timer** -- [in] MCPWM timer handle, allocated by `mcpwm_new_timer()`

Returns

- ESP_OK: Connect MCPWM operator and timer successfully
- ESP_ERR_INVALID_ARG: Connect MCPWM operator and timer failed because of invalid argument
- ESP_FAIL: Connect MCPWM operator and timer failed because of other error

esp_err_t **mcpwm_operator_set_brake_on_fault** (*mcpwm_oper_handle_t* oper, const *mcpwm_brake_config_t* *config)

Set brake method for MCPWM operator.

Parameters

- **oper** -- [in] MCPWM operator, allocated by `mcpwm_new_operator()`
- **config** -- [in] MCPWM brake configuration

Returns

- ESP_OK: Set trip for operator successfully
- ESP_ERR_INVALID_ARG: Set trip for operator failed because of invalid argument
- ESP_FAIL: Set trip for operator failed because of other error

esp_err_t **mcpwm_operator_recover_from_fault** (*mcpwm_oper_handle_t* oper, *mcpwm_fault_handle_t* fault)

Try to make the operator recover from fault.

Note: To recover from fault or escape from trip, you make sure the fault signal has disappeared already. Otherwise the recovery can't succeed.

Parameters

- **oper** -- [in] MCPWM operator, allocated by `mcpwm_new_operator()`
- **fault** -- [in] MCPWM fault handle

Returns

- ESP_OK: Recover from fault successfully
- ESP_ERR_INVALID_ARG: Recover from fault failed because of invalid argument
- ESP_ERR_INVALID_STATE: Recover from fault failed because the fault source is still active
- ESP_FAIL: Recover from fault failed because of other error

esp_err_t **mcpwm_operator_register_event_callbacks** (*mcpwm_oper_handle_t* oper, const *mcpwm_operator_event_callbacks_t* *cbs, void *user_data)

Set event callbacks for MCPWM operator.

Note: User can deregister a previously registered callback by calling this function and setting the callback member in the `cbs` structure to NULL.

Parameters

- **oper** -- [in] MCPWM operator handle, allocated by `mcpwm_new_operator()`
- **cbs** -- [in] Group of callback functions
- **user_data** -- [in] User data, which will be passed to callback functions directly

Returns

- ESP_OK: Set event callbacks successfully
- ESP_ERR_INVALID_ARG: Set event callbacks failed because of invalid argument
- ESP_FAIL: Set event callbacks failed because of other error

esp_err_t **mcpwm_operator_apply_carrier** (*mcpwm_oper_handle_t* oper, const *mcpwm_carrier_config_t* *config)

Apply carrier feature for MCPWM operator.

Parameters

- **oper** -- [in] MCPWM operator, allocated by `mcpwm_new_operator()`
- **config** -- [in] MCPWM carrier specific configuration

Returns

- ESP_OK: Set carrier for operator successfully
- ESP_ERR_INVALID_ARG: Set carrier for operator failed because of invalid argument
- ESP_FAIL: Set carrier for operator failed because of other error

Structures

struct **mcpwm_operator_config_t**

MCPWM operator configuration.

Public Members

int **group_id**

Specify from which group to allocate the MCPWM operator

int **intr_priority**

MCPWM operator interrupt priority, if set to 0, the driver will try to allocate an interrupt with a relative low priority (1,2,3)

uint32_t **update_gen_action_on_tez**

Whether to update generator action when timer counts to zero

uint32_t **update_gen_action_on_tep**

Whether to update generator action when timer counts to peak

uint32_t **update_gen_action_on_sync**

Whether to update generator action on sync event

uint32_t **update_dead_time_on_tez**

Whether to update dead time when timer counts to zero

uint32_t **update_dead_time_on_tep**

Whether to update dead time when timer counts to peak

uint32_t **update_dead_time_on_sync**

Whether to update dead time on sync event

struct *mcpwm_operator_config_t*::[anonymous] **flags**

Extra configuration flags for operator

struct **mcpwm_brake_config_t**

MCPWM brake configuration structure.

Public Members

mcpwm_fault_handle_t **fault**

Which fault causes the operator to brake

mcpwm_operator_brake_mode_t **brake_mode**

Brake mode

uint32_t **cbc_recover_on_tez**

Recovery CBC brake state on tez event

uint32_t **cbc_recover_on_tep**

Recovery CBC brake state on tep event

struct *mcpwm_brake_config_t*::[anonymous] **flags**

Extra flags for brake configuration

struct **mcpwm_operator_event_callbacks_t**

Group of supported MCPWM operator event callbacks.

Note: The callbacks are all running under ISR environment

Public Members

mcpwm_brake_event_cb_t **on_brake_cbc**

callback function when mcpwm operator brakes in CBC

mcpwm_brake_event_cb_t **on_brake_ost**

callback function when mcpwm operator brakes in OST

struct **mcpwm_carrier_config_t**

MCPWM carrier configuration structure.

Public Members

mcpwm_carrier_clock_source_t **clk_src**

MCPWM carrier clock source

uint32_t **frequency_hz**

Carrier frequency in Hz

uint32_t **first_pulse_duration_us**

The duration of the first PWM pulse, in us

float **duty_cycle**

Carrier duty cycle

uint32_t **invert_before_modulate**

Invert the raw signal

uint32_t **invert_after_modulate**

Invert the modulated signal

struct *mcpwm_carrier_config_t*::[anonymous] **flags**

Extra flags for carrier configuration

Header File

- [components/driver/mcpwm/include/driver/mcpwm_cmpr.h](#)
- This header file can be included with:

```
#include "driver/mcpwm_cmpr.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t **mcpwm_new_comparator** (*mcpwm_oper_handle_t* oper, const *mcpwm_comparator_config_t* *config, *mcpwm_cmpr_handle_t* *ret_cmpr)

Create MCPWM comparator.

Parameters

- **oper** -- [in] MCPWM operator, allocated by `mcpwm_new_operator()`, the new comparator will be allocated from this operator
- **config** -- [in] MCPWM comparator configuration
- **ret_cmpr** -- [out] Returned MCPWM comparator

Returns

- ESP_OK: Create MCPWM comparator successfully
- ESP_ERR_INVALID_ARG: Create MCPWM comparator failed because of invalid argument
- ESP_ERR_NO_MEM: Create MCPWM comparator failed because out of memory
- ESP_ERR_NOT_FOUND: Create MCPWM comparator failed because can't find free resource
- ESP_FAIL: Create MCPWM comparator failed because of other error

esp_err_t **mcpwm_del_comparator** (*mcpwm_cmpr_handle_t* cmpr)

Delete MCPWM comparator.

Parameters **cmpr** -- [in] MCPWM comparator handle, allocated by `mcpwm_new_comparator()`

Returns

- ESP_OK: Delete MCPWM comparator successfully
- ESP_ERR_INVALID_ARG: Delete MCPWM comparator failed because of invalid argument
- ESP_FAIL: Delete MCPWM comparator failed because of other error

esp_err_t **mcpwm_new_event_comparator** (*mcpwm_oper_handle_t* oper, const *mcpwm_event_comparator_config_t* *config, *mcpwm_cmpr_handle_t* *ret_cmpr)

Create MCPWM event comparator.

Parameters

- **oper** -- [in] MCPWM operator, allocated by `mcpwm_new_operator()`, the new event comparator will be allocated from this operator
- **config** -- [in] MCPWM comparator configuration
- **ret_cmpr** -- [out] Returned MCPWM event comparator

Returns

- ESP_OK: Create MCPWM event comparator successfully
- ESP_ERR_INVALID_ARG: Create MCPWM event comparator failed because of invalid argument
- ESP_ERR_NO_MEM: Create MCPWM event comparator failed because out of memory
- ESP_ERR_NOT_FOUND: Create MCPWM event comparator failed because can't find free resource
- ESP_FAIL: Create MCPWM event comparator failed because of other error

esp_err_t **mcpwm_comparator_register_event_callbacks** (*mcpwm_cmpr_handle_t* cmpr, const *mcpwm_comparator_event_callbacks_t* *cbs, void *user_data)

Set event callbacks for MCPWM comparator.

Note: User can deregister a previously registered callback by calling this function and setting the callback member in the `cbs` structure to NULL.

Parameters

- **cmpr** -- [in] MCPWM comparator handle, allocated by `mcpwm_new_comparator()`
- **cbs** -- [in] Group of callback functions
- **user_data** -- [in] User data, which will be passed to callback functions directly

Returns

- ESP_OK: Set event callbacks successfully
- ESP_ERR_INVALID_ARG: Set event callbacks failed because of invalid argument

- ESP_FAIL: Set event callbacks failed because of other error

esp_err_t **mcpwm_comparator_set_compare_value** (*mcpwm_cmpr_handle_t* cmpr, uint32_t cmp_ticks)

Set MCPWM comparator's compare value.

Parameters

- **cmpr** -- [in] MCPWM comparator handle, allocated by `mcpwm_new_comparator()`
- **cmp_ticks** -- [in] The new compare value

Returns

- ESP_OK: Set MCPWM compare value successfully
- ESP_ERR_INVALID_ARG: Set MCPWM compare value failed because of invalid argument (e.g. the `cmp_ticks` is out of range)
- ESP_ERR_INVALID_STATE: Set MCPWM compare value failed because the operator doesn't have a timer connected
- ESP_FAIL: Set MCPWM compare value failed because of other error

Structures

struct **mcpwm_comparator_config_t**

MCPWM comparator configuration.

Public Members

int **intr_priority**

MCPWM comparator interrupt priority, if set to 0, the driver will try to allocate an interrupt with a relative low priority (1,2,3)

uint32_t **update_cmp_on_tez**

Whether to update compare value when timer count equals to zero (tez)

uint32_t **update_cmp_on_tep**

Whether to update compare value when timer count equals to peak (tep)

uint32_t **update_cmp_on_sync**

Whether to update compare value on sync event

struct *mcpwm_comparator_config_t*::[anonymous] **flags**

Extra configuration flags for comparator

struct **mcpwm_event_comparator_config_t**

MCPWM event comparator configuration.

struct **mcpwm_comparator_event_callbacks_t**

Group of supported MCPWM compare event callbacks.

Note: The callbacks are all running under ISR environment

Public Members

mcpwm_compare_event_cb_t on_reach

ISR callback function which would be invoked when counter reaches compare value

Header File

- `components/driver/mcpwm/include/driver/mcpwm_gen.h`
- This header file can be included with:

```
#include "driver/mcpwm_gen.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t **mcpwm_new_generator** (*mcpwm_oper_handle_t* oper, const *mcpwm_generator_config_t* *config, *mcpwm_gen_handle_t* *ret_gen)

Allocate MCPWM generator from given operator.

Parameters

- **oper** -- [in] MCPWM operator, allocated by `mcpwm_new_operator()`
- **config** -- [in] MCPWM generator configuration
- **ret_gen** -- [out] Returned MCPWM generator

Returns

- `ESP_OK`: Create MCPWM generator successfully
- `ESP_ERR_INVALID_ARG`: Create MCPWM generator failed because of invalid argument
- `ESP_ERR_NO_MEM`: Create MCPWM generator failed because out of memory
- `ESP_ERR_NOT_FOUND`: Create MCPWM generator failed because can't find free resource
- `ESP_FAIL`: Create MCPWM generator failed because of other error

esp_err_t **mcpwm_del_generator** (*mcpwm_gen_handle_t* gen)

Delete MCPWM generator.

Parameters **gen** -- [in] MCPWM generator handle, allocated by `mcpwm_new_generator()`

Returns

- `ESP_OK`: Delete MCPWM generator successfully
- `ESP_ERR_INVALID_ARG`: Delete MCPWM generator failed because of invalid argument
- `ESP_FAIL`: Delete MCPWM generator failed because of other error

esp_err_t **mcpwm_generator_set_force_level** (*mcpwm_gen_handle_t* gen, int level, bool hold_on)

Set force level for MCPWM generator.

Note: The force level will be applied to the generator immediately, regardless any other events that would change the generator's behaviour.

Note: If the `hold_on` is true, the force level will retain forever, until user removes the force level by setting the force level to `-1`.

Note: If the `hold_on` is false, the force level can be overridden by the next event action.

Note: The force level set by this function can be inverted by GPIO matrix or dead-time module. So the level set here doesn't equal to the final output level.

Parameters

- **gen** -- [in] MCPWM generator handle, allocated by `mcpwm_new_generator()`
- **level** -- [in] GPIO level to be applied to MCPWM generator, specially, -1 means to remove the force level
- **hold_on** -- [in] Whether the forced PWM level should retain (i.e. will remain unchanged until manually remove the force level)

Returns

- `ESP_OK`: Set force level for MCPWM generator successfully
- `ESP_ERR_INVALID_ARG`: Set force level for MCPWM generator failed because of invalid argument
- `ESP_FAIL`: Set force level for MCPWM generator failed because of other error

esp_err_t `mcpwm_generator_set_action_on_timer_event` (*mcpwm_gen_handle_t* gen,
mcpwm_gen_timer_event_action_t
ev_act)

Set generator action on MCPWM timer event.

Parameters

- **gen** -- [in] MCPWM generator handle, allocated by `mcpwm_new_generator()`
- **ev_act** -- [in] MCPWM timer event action, can be constructed by `MCPWM_GEN_TIMER_EVENT_ACTION` helper macro

Returns

- `ESP_OK`: Set generator action successfully
- `ESP_ERR_INVALID_ARG`: Set generator action failed because of invalid argument
- `ESP_ERR_INVALID_STATE`: Set generator action failed because of timer is not connected to operator
- `ESP_FAIL`: Set generator action failed because of other error

esp_err_t `mcpwm_generator_set_actions_on_timer_event` (*mcpwm_gen_handle_t* gen,
mcpwm_gen_timer_event_action_t
ev_act, ...)

Set generator actions on multiple MCPWM timer events.

Note: This is an aggregation version of `mcpwm_generator_set_action_on_timer_event`, which allows user to set multiple actions in one call.

Parameters

- **gen** -- [in] MCPWM generator handle, allocated by `mcpwm_new_generator()`
- **ev_act** -- [in] MCPWM timer event action list, must be terminated by `MCPWM_GEN_TIMER_EVENT_ACTION_END()`

Returns

- `ESP_OK`: Set generator actions successfully
- `ESP_ERR_INVALID_ARG`: Set generator actions failed because of invalid argument
- `ESP_ERR_INVALID_STATE`: Set generator actions failed because of timer is not connected to operator
- `ESP_FAIL`: Set generator actions failed because of other error

esp_err_t **mcpwm_generator_set_action_on_compare_event** (*mcpwm_gen_handle_t* generator,
mcpwm_gen_compare_event_action_t
ev_act)

Set generator action on MCPWM compare event.

Parameters

- **generator** -- [in] MCPWM generator handle, allocated by `mcpwm_new_generator()`
- **ev_act** -- [in] MCPWM compare event action, can be constructed by `MCPWM_GEN_COMPARE_EVENT_ACTION` helper macro

Returns

- `ESP_OK`: Set generator action successfully
- `ESP_ERR_INVALID_ARG`: Set generator action failed because of invalid argument
- `ESP_FAIL`: Set generator action failed because of other error

esp_err_t **mcpwm_generator_set_actions_on_compare_event** (*mcpwm_gen_handle_t* generator,
mcpwm_gen_compare_event_action_t
ev_act, ...)

Set generator actions on multiple MCPWM compare events.

Note: This is an aggregation version of `mcpwm_generator_set_action_on_compare_event`, which allows user to set multiple actions in one call.

Parameters

- **generator** -- [in] MCPWM generator handle, allocated by `mcpwm_new_generator()`
- **ev_act** -- [in] MCPWM compare event action list, must be terminated by `MCPWM_GEN_COMPARE_EVENT_ACTION_END()`

Returns

- `ESP_OK`: Set generator actions successfully
- `ESP_ERR_INVALID_ARG`: Set generator actions failed because of invalid argument
- `ESP_FAIL`: Set generator actions failed because of other error

esp_err_t **mcpwm_generator_set_action_on_brake_event** (*mcpwm_gen_handle_t* generator,
mcpwm_gen_brake_event_action_t
ev_act)

Set generator action on MCPWM brake event.

Parameters

- **generator** -- [in] MCPWM generator handle, allocated by `mcpwm_new_generator()`
- **ev_act** -- [in] MCPWM brake event action, can be constructed by `MCPWM_GEN_BRAKE_EVENT_ACTION` helper macro

Returns

- `ESP_OK`: Set generator action successfully
- `ESP_ERR_INVALID_ARG`: Set generator action failed because of invalid argument
- `ESP_FAIL`: Set generator action failed because of other error

esp_err_t **mcpwm_generator_set_actions_on_brake_event** (*mcpwm_gen_handle_t* generator,
mcpwm_gen_brake_event_action_t
ev_act, ...)

Set generator actions on multiple MCPWM brake events.

Note: This is an aggregation version of `mcpwm_generator_set_action_on_brake_event`, which allows user to set multiple actions in one call.

Parameters

- **generator** -- [in] MCPWM generator handle, allocated by `mcpwm_new_generator()`
- **ev_act** -- [in] MCPWM brake event action list, must be terminated by `MCPWM_GEN_BRAKE_EVENT_ACTION_END()`

Returns

- `ESP_OK`: Set generator actions successfully
- `ESP_ERR_INVALID_ARG`: Set generator actions failed because of invalid argument
- `ESP_FAIL`: Set generator actions failed because of other error

esp_err_t **mcpwm_generator_set_action_on_fault_event** (*mcpwm_gen_handle_t* generator, *mcpwm_gen_fault_event_action_t* ev_act)

Set generator action on MCPWM Fault event.

Parameters

- **generator** -- [in] MCPWM generator handle, allocated by `mcpwm_new_generator()`
- **ev_act** -- [in] MCPWM trigger event action, can be constructed by `MCPWM_GEN_FAULT_EVENT_ACTION` helper macro

Returns

- `ESP_OK`: Set generator action successfully
- `ESP_ERR_INVALID_ARG`: Set generator action failed because of invalid argument
- `ESP_FAIL`: Set generator action failed because of other error

esp_err_t **mcpwm_generator_set_action_on_sync_event** (*mcpwm_gen_handle_t* generator, *mcpwm_gen_sync_event_action_t* ev_act)

Set generator action on MCPWM Sync event.

Note: The trigger only support one sync action, regardless of the kinds. Should not call this function more than once.

Parameters

- **generator** -- [in] MCPWM generator handle, allocated by `mcpwm_new_generator()`
- **ev_act** -- [in] MCPWM trigger event action, can be constructed by `MCPWM_GEN_SYNC_EVENT_ACTION` helper macro

Returns

- `ESP_OK`: Set generator action successfully
- `ESP_ERR_INVALID_ARG`: Set generator action failed because of invalid argument
- `ESP_FAIL`: Set generator action failed because of other error

esp_err_t **mcpwm_generator_set_dead_time** (*mcpwm_gen_handle_t* in_generator, *mcpwm_gen_handle_t* out_generator, const *mcpwm_dead_time_config_t* *config)

Set dead time for MCPWM generator.

Note: Due to a hardware limitation, you can't set rising edge delay for both MCPWM generator 0 and 1 at the same time, otherwise, there will be a conflict inside the dead time module. The same goes for the falling edge setting. But you can set both the rising edge and falling edge delay for the same MCPWM generator.

Parameters

- **in_generator** -- [in] MCPWM generator, before adding the dead time
- **out_generator** -- [in] MCPWM generator, after adding the dead time
- **config** -- [in] MCPWM dead time configuration

Returns

- `ESP_OK`: Set dead time for MCPWM generator successfully
- `ESP_ERR_INVALID_ARG`: Set dead time for MCPWM generator failed because of invalid argument
- `ESP_ERR_INVALID_STATE`: Set dead time for MCPWM generator failed because of invalid state (e.g. delay module is already in use by other generator)
- `ESP_FAIL`: Set dead time for MCPWM generator failed because of other error

Structures

struct **mcpwm_generator_config_t**

MCPWM generator configuration.

Public Members

int **gen_gpio_num**

The GPIO number used to output the PWM signal

uint32_t **invert_pwm**

Whether to invert the PWM signal (done by GPIO matrix)

uint32_t **io_loop_back**

For debug/test, the signal output from the GPIO will be fed to the input path as well

uint32_t **io_od_mode**

Configure the GPIO as open-drain mode

uint32_t **pull_up**

Whether to pull up internally

uint32_t **pull_down**

Whether to pull down internally

struct *mcpwm_generator_config_t*::[anonymous] **flags**

Extra configuration flags for generator

struct **mcpwm_gen_timer_event_action_t**

Generator action on specific timer event.

Public Members

mcpwm_timer_direction_t **direction**

Timer direction

mcpwm_timer_event_t **event**

Timer event

mcpwm_generator_action_t **action**

Generator action should perform

struct **mcpwm_gen_compare_event_action_t**

Generator action on specific comparator event.

Public Members

mcpwm_timer_direction_t **direction**

Timer direction

mcpwm_cmpr_handle_t **comparator**

Comparator handle

mcpwm_generator_action_t **action**

Generator action should perform

struct **mcpwm_gen_brake_event_action_t**

Generator action on specific brake event.

Public Members

mcpwm_timer_direction_t **direction**

Timer direction

mcpwm_operator_brake_mode_t **brake_mode**

Brake mode

mcpwm_generator_action_t **action**

Generator action should perform

struct **mcpwm_gen_fault_event_action_t**

Generator action on specific fault event.

Public Members

mcpwm_timer_direction_t **direction**

Timer direction

mcpwm_fault_handle_t **fault**

Which fault as the trigger. Only support GPIO fault

mcpwm_generator_action_t **action**

Generator action should perform

struct **mcpwm_gen_sync_event_action_t**

Generator action on specific sync event.

Public Members

mcpwm_timer_direction_t **direction**

Timer direction

mcpwm_sync_handle_t **sync**

Which sync as the trigger

mcpwm_generator_action_t **action**

Generator action should perform

struct **mcpwm_dead_time_config_t**

MCPWM dead time configuration structure.

Public Members

uint32_t **posedge_delay_ticks**

delay time applied to rising edge, 0 means no rising delay time

uint32_t **negedge_delay_ticks**

delay time applied to falling edge, 0 means no falling delay time

uint32_t **invert_output**

Invert the signal after applied the dead time

struct *mcpwm_dead_time_config_t*::[anonymous] **flags**

Extra flags for dead time configuration

Macros

MCPWM_GEN_TIMER_EVENT_ACTION (dir, ev, act)

Help macros to construct a *mcpwm_gen_timer_event_action_t* entry.

MCPWM_GEN_TIMER_EVENT_ACTION_END ()

MCPWM_GEN_COMPARE_EVENT_ACTION (dir, cmp, act)

Help macros to construct a *mcpwm_gen_compare_event_action_t* entry.

MCPWM_GEN_COMPARE_EVENT_ACTION_END ()

MCPWM_GEN_BRAKE_EVENT_ACTION (dir, mode, act)

Help macros to construct a *mcpwm_gen_brake_event_action_t* entry.

MCPWM_GEN_BRAKE_EVENT_ACTION_END ()

MCPWM_GEN_FAULT_EVENT_ACTION (dir, flt, act)

Help macros to construct a *mcpwm_gen_fault_event_action_t* entry.

MCPWM_GEN_SYNC_EVENT_ACTION (dir, syn, act)

Help macros to construct a *mcpwm_gen_sync_event_action_t* entry.

Header File

- [components/driver/mcpwm/include/driver/mcpwm_fault.h](#)
- This header file can be included with:

```
#include "driver/mcpwm_fault.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t **mcpwm_new_gpio_fault** (const *mcpwm_gpio_fault_config_t* *config, *mcpwm_fault_handle_t* *ret_fault)

Create MCPWM GPIO fault.

Parameters

- **config** -- [in] MCPWM GPIO fault configuration
- **ret_fault** -- [out] Returned GPIO fault handle

Returns

- **ESP_OK**: Create MCPWM GPIO fault successfully
- **ESP_ERR_INVALID_ARG**: Create MCPWM GPIO fault failed because of invalid argument
- **ESP_ERR_NO_MEM**: Create MCPWM GPIO fault failed because out of memory
- **ESP_ERR_NOT_FOUND**: Create MCPWM GPIO fault failed because can't find free resource
- **ESP_FAIL**: Create MCPWM GPIO fault failed because of other error

esp_err_t **mcpwm_new_soft_fault** (const *mcpwm_soft_fault_config_t* *config, *mcpwm_fault_handle_t* *ret_fault)

Create MCPWM software fault.

Parameters

- **config** -- [in] MCPWM software fault configuration
- **ret_fault** -- [out] Returned software fault handle

Returns

- **ESP_OK**: Create MCPWM software fault successfully
- **ESP_ERR_INVALID_ARG**: Create MCPWM software fault failed because of invalid argument
- **ESP_ERR_NO_MEM**: Create MCPWM software fault failed because out of memory
- **ESP_FAIL**: Create MCPWM software fault failed because of other error

esp_err_t **mcpwm_del_fault** (*mcpwm_fault_handle_t* fault)

Delete MCPWM fault.

Parameters **fault** -- [in] MCPWM fault handle allocated by `mcpwm_new_gpio_fault()` or `mcpwm_new_soft_fault()`

Returns

- **ESP_OK**: Delete MCPWM fault successfully
- **ESP_ERR_INVALID_ARG**: Delete MCPWM fault failed because of invalid argument
- **ESP_FAIL**: Delete MCPWM fault failed because of other error

esp_err_t **mcpwm_soft_fault_activate** (*mcpwm_fault_handle_t* fault)

Activate the software fault, trigger the fault event for once.

Parameters **fault** -- [in] MCPWM soft fault, allocated by `mcpwm_new_soft_fault()`

Returns

- ESP_OK: Trigger MCPWM software fault event successfully
- ESP_ERR_INVALID_ARG: Trigger MCPWM software fault event failed because of invalid argument
- ESP_FAIL: Trigger MCPWM software fault event failed because of other error

esp_err_t **mcpwm_fault_register_event_callbacks** (*mcpwm_fault_handle_t* fault, const *mcpwm_fault_event_callbacks_t* *cbs, void *user_data)

Set event callbacks for MCPWM fault.

Note: User can deregister a previously registered callback by calling this function and setting the callback member in the *cbs* structure to NULL.

Parameters

- **fault** -- **[in]** MCPWM GPIO fault handle, allocated by *mcpwm_new_gpio_fault()*
- **cbs** -- **[in]** Group of callback functions
- **user_data** -- **[in]** User data, which will be passed to callback functions directly

Returns

- ESP_OK: Set event callbacks successfully
- ESP_ERR_INVALID_ARG: Set event callbacks failed because of invalid argument
- ESP_FAIL: Set event callbacks failed because of other error

Structures

struct **mcpwm_gpio_fault_config_t**

MCPWM GPIO fault configuration structure.

Public Members

int **group_id**

In which MCPWM group that the GPIO fault belongs to

int **intr_priority**

MCPWM GPIO fault interrupt priority, if set to 0, the driver will try to allocate an interrupt with a relative low priority (1,2,3)

int **gpio_num**

GPIO used by the fault signal

uint32_t **active_level**

On which level the fault signal is treated as active

uint32_t **io_loop_back**

For debug/test, the signal output from the GPIO will be fed to the input path as well

uint32_t **pull_up**

Whether to pull up internally

uint32_t **pull_down**

Whether to pull down internally

struct *mcpwm_gpio_fault_config_t*::[anonymous] **flags**

Extra configuration flags for GPIO fault

struct **mcpwm_soft_fault_config_t**

MCPWM software fault configuration structure.

struct **mcpwm_fault_event_callbacks_t**

Group of supported MCPWM fault event callbacks.

Note: The callbacks are all running under ISR environment

Public Members

mcpwm_fault_event_cb_t **on_fault_enter**

ISR callback function that would be invoked when fault signal becomes active

mcpwm_fault_event_cb_t **on_fault_exit**

ISR callback function that would be invoked when fault signal becomes inactive

Header File

- [components/driver/mcpwm/include/driver/mcpwm_sync.h](#)
- This header file can be included with:

```
#include "driver/mcpwm_sync.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t **mcpwm_new_timer_sync_src** (*mcpwm_timer_handle_t* timer, const *mcpwm_timer_sync_src_config_t* *config, *mcpwm_sync_handle_t* *ret_sync)

Create MCPWM timer sync source.

Parameters

- **timer** -- [in] MCPWM timer handle, allocated by `mcpwm_new_timer()`
- **config** -- [in] MCPWM timer sync source configuration
- **ret_sync** -- [out] Returned MCPWM sync handle

Returns

- **ESP_OK**: Create MCPWM timer sync source successfully
- **ESP_ERR_INVALID_ARG**: Create MCPWM timer sync source failed because of invalid argument
- **ESP_ERR_NO_MEM**: Create MCPWM timer sync source failed because out of memory
- **ESP_ERR_INVALID_STATE**: Create MCPWM timer sync source failed because the timer has created a sync source before
- **ESP_FAIL**: Create MCPWM timer sync source failed because of other error

```
esp_err_t mcpwm_new_gpio_sync_src (const mcpwm_gpio_sync_src_config_t *config,  
                                   mcpwm_sync_handle_t *ret_sync)
```

Create MCPWM GPIO sync source.

Parameters

- **config** -- [in] MCPWM GPIO sync source configuration
- **ret_sync** -- [out] Returned MCPWM GPIO sync handle

Returns

- ESP_OK: Create MCPWM GPIO sync source successfully
- ESP_ERR_INVALID_ARG: Create MCPWM GPIO sync source failed because of invalid argument
- ESP_ERR_NO_MEM: Create MCPWM GPIO sync source failed because out of memory
- ESP_ERR_NOT_FOUND: Create MCPWM GPIO sync source failed because can't find free resource
- ESP_FAIL: Create MCPWM GPIO sync source failed because of other error

```
esp_err_t mcpwm_new_soft_sync_src (const mcpwm_soft_sync_config_t *config, mcpwm_sync_handle_t  
                                   *ret_sync)
```

Create MCPWM software sync source.

Parameters

- **config** -- [in] MCPWM software sync source configuration
- **ret_sync** -- [out] Returned software sync handle

Returns

- ESP_OK: Create MCPWM software sync successfully
- ESP_ERR_INVALID_ARG: Create MCPWM software sync failed because of invalid argument
- ESP_ERR_NO_MEM: Create MCPWM software sync failed because out of memory
- ESP_FAIL: Create MCPWM software sync failed because of other error

```
esp_err_t mcpwm_del_sync_src (mcpwm_sync_handle_t sync)
```

Delete MCPWM sync source.

Parameters **sync** -- [in] MCPWM sync handle, allocated by `mcpwm_new_timer_sync_src()` or `mcpwm_new_gpio_sync_src()` or `mcpwm_new_soft_sync_src()`

Returns

- ESP_OK: Delete MCPWM sync source successfully
- ESP_ERR_INVALID_ARG: Delete MCPWM sync source failed because of invalid argument
- ESP_FAIL: Delete MCPWM sync source failed because of other error

```
esp_err_t mcpwm_soft_sync_activate (mcpwm_sync_handle_t sync)
```

Activate the software sync, trigger the sync event for once.

Parameters **sync** -- [in] MCPWM soft sync handle, allocated by `mcpwm_new_soft_sync_src()`

Returns

- ESP_OK: Trigger MCPWM software sync event successfully
- ESP_ERR_INVALID_ARG: Trigger MCPWM software sync event failed because of invalid argument
- ESP_FAIL: Trigger MCPWM software sync event failed because of other error

Structures

```
struct mcpwm_timer_sync_src_config_t
```

MCPWM timer sync source configuration.

Public Members

mcpwm_timer_event_t **timer_event**

Timer event, upon which MCPWM timer will generate the sync signal

uint32_t **propagate_input_sync**

The input sync signal would be routed to its sync output

struct *mcpwm_timer_sync_src_config_t*::[anonymous] **flags**

Extra configuration flags for timer sync source

struct **mcpwm_gpio_sync_src_config_t**

MCPWM GPIO sync source configuration.

Public Members

int **group_id**

MCPWM group ID

int **gpio_num**

GPIO used by sync source

uint32_t **active_neg**

Whether the sync signal is active on negedge, by default, the sync signal's posedge is treated as active

uint32_t **io_loop_back**

For debug/test, the signal output from the GPIO will be fed to the input path as well

uint32_t **pull_up**

Whether to pull up internally

uint32_t **pull_down**

Whether to pull down internally

struct *mcpwm_gpio_sync_src_config_t*::[anonymous] **flags**

Extra configuration flags for GPIO sync source

struct **mcpwm_soft_sync_config_t**

MCPWM software sync configuration structure.

Header File

- [components/driver/mcpwm/include/driver/mcpwm_cap.h](#)
- This header file can be included with:

```
#include "driver/mcpwm_cap.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

REQUIRES driver

or

PRIV_REQUIRES driver

Functions

esp_err_t **mcpwm_new_capture_timer** (const *mcpwm_capture_timer_config_t* *config, *mcpwm_cap_timer_handle_t* *ret_cap_timer)

Create MCPWM capture timer.

Parameters

- **config** -- [in] MCPWM capture timer configuration
- **ret_cap_timer** -- [out] Returned MCPWM capture timer handle

Returns

- ESP_OK: Create MCPWM capture timer successfully
- ESP_ERR_INVALID_ARG: Create MCPWM capture timer failed because of invalid argument
- ESP_ERR_NO_MEM: Create MCPWM capture timer failed because out of memory
- ESP_ERR_NOT_FOUND: Create MCPWM capture timer failed because can't find free resource
- ESP_FAIL: Create MCPWM capture timer failed because of other error

esp_err_t **mcpwm_del_capture_timer** (*mcpwm_cap_timer_handle_t* cap_timer)

Delete MCPWM capture timer.

Parameters **cap_timer** -- [in] MCPWM capture timer, allocated by `mcpwm_new_capture_timer()`

Returns

- ESP_OK: Delete MCPWM capture timer successfully
- ESP_ERR_INVALID_ARG: Delete MCPWM capture timer failed because of invalid argument
- ESP_FAIL: Delete MCPWM capture timer failed because of other error

esp_err_t **mcpwm_capture_timer_enable** (*mcpwm_cap_timer_handle_t* cap_timer)

Enable MCPWM capture timer.

Parameters **cap_timer** -- [in] MCPWM capture timer handle, allocated by `mcpwm_new_capture_timer()`

Returns

- ESP_OK: Enable MCPWM capture timer successfully
- ESP_ERR_INVALID_ARG: Enable MCPWM capture timer failed because of invalid argument
- ESP_ERR_INVALID_STATE: Enable MCPWM capture timer failed because timer is enabled already
- ESP_FAIL: Enable MCPWM capture timer failed because of other error

esp_err_t **mcpwm_capture_timer_disable** (*mcpwm_cap_timer_handle_t* cap_timer)

Disable MCPWM capture timer.

Parameters **cap_timer** -- [in] MCPWM capture timer handle, allocated by `mcpwm_new_capture_timer()`

Returns

- ESP_OK: Disable MCPWM capture timer successfully
- ESP_ERR_INVALID_ARG: Disable MCPWM capture timer failed because of invalid argument
- ESP_ERR_INVALID_STATE: Disable MCPWM capture timer failed because timer is disabled already
- ESP_FAIL: Disable MCPWM capture timer failed because of other error

esp_err_t **mcpwm_capture_timer_start** (*mcpwm_cap_timer_handle_t* cap_timer)

Start MCPWM capture timer.

Parameters **cap_timer** -- **[in]** MCPWM capture timer, allocated by `mcpwm_new_capture_timer()`

Returns

- ESP_OK: Start MCPWM capture timer successfully
- ESP_ERR_INVALID_ARG: Start MCPWM capture timer failed because of invalid argument
- ESP_FAIL: Start MCPWM capture timer failed because of other error

esp_err_t **mcpwm_capture_timer_stop** (*mcpwm_cap_timer_handle_t* cap_timer)

Stop MCPWM capture timer.

Parameters **cap_timer** -- **[in]** MCPWM capture timer, allocated by `mcpwm_new_capture_timer()`

Returns

- ESP_OK: Stop MCPWM capture timer successfully
- ESP_ERR_INVALID_ARG: Stop MCPWM capture timer failed because of invalid argument
- ESP_FAIL: Stop MCPWM capture timer failed because of other error

esp_err_t **mcpwm_capture_timer_get_resolution** (*mcpwm_cap_timer_handle_t* cap_timer, uint32_t *out_resolution)

Get MCPWM capture timer resolution, in Hz.

Parameters

- **cap_timer** -- **[in]** MCPWM capture timer, allocated by `mcpwm_new_capture_timer()`
- **out_resolution** -- **[out]** Returned capture timer resolution, in Hz

Returns

- ESP_OK: Get capture timer resolution successfully
- ESP_ERR_INVALID_ARG: Get capture timer resolution failed because of invalid argument
- ESP_FAIL: Get capture timer resolution failed because of other error

esp_err_t **mcpwm_capture_timer_set_phase_on_sync** (*mcpwm_cap_timer_handle_t* cap_timer, const *mcpwm_capture_timer_sync_phase_config_t* *config)

Set sync phase for MCPWM capture timer.

Parameters

- **cap_timer** -- **[in]** MCPWM capture timer, allocated by `mcpwm_new_capture_timer()`
- **config** -- **[in]** MCPWM capture timer sync phase configuration

Returns

- ESP_OK: Set sync phase for MCPWM capture timer successfully
- ESP_ERR_INVALID_ARG: Set sync phase for MCPWM capture timer failed because of invalid argument
- ESP_FAIL: Set sync phase for MCPWM capture timer failed because of other error

esp_err_t **mcpwm_new_capture_channel** (*mcpwm_cap_timer_handle_t* cap_timer, const *mcpwm_capture_channel_config_t* *config, *mcpwm_cap_channel_handle_t* *ret_cap_channel)

Create MCPWM capture channel.

Note: The created capture channel won't be enabled until calling `mcpwm_capture_channel_enable`

Parameters

- **cap_timer** -- [in] MCPWM capture timer, allocated by `mcpwm_new_capture_timer()`, will be connected to the new capture channel
- **config** -- [in] MCPWM capture channel configuration
- **ret_cap_channel** -- [out] Returned MCPWM capture channel

Returns

- ESP_OK: Create MCPWM capture channel successfully
- ESP_ERR_INVALID_ARG: Create MCPWM capture channel failed because of invalid argument
- ESP_ERR_NO_MEM: Create MCPWM capture channel failed because out of memory
- ESP_ERR_NOT_FOUND: Create MCPWM capture channel failed because can't find free resource
- ESP_FAIL: Create MCPWM capture channel failed because of other error

esp_err_t **mcpwm_del_capture_channel** (*mcpwm_cap_channel_handle_t* cap_channel)

Delete MCPWM capture channel.

Parameters **cap_channel** -- [in] MCPWM capture channel handle, allocated by `mcpwm_new_capture_channel()`

Returns

- ESP_OK: Delete MCPWM capture channel successfully
- ESP_ERR_INVALID_ARG: Delete MCPWM capture channel failed because of invalid argument
- ESP_FAIL: Delete MCPWM capture channel failed because of other error

esp_err_t **mcpwm_capture_channel_enable** (*mcpwm_cap_channel_handle_t* cap_channel)

Enable MCPWM capture channel.

Note: This function will transit the channel state from init to enable.

Note: This function will enable the interrupt service, if it's lazy installed in `mcpwm_capture_channel_register_event_callbacks()`.

Parameters **cap_channel** -- [in] MCPWM capture channel handle, allocated by `mcpwm_new_capture_channel()`

Returns

- ESP_OK: Enable MCPWM capture channel successfully
- ESP_ERR_INVALID_ARG: Enable MCPWM capture channel failed because of invalid argument
- ESP_ERR_INVALID_STATE: Enable MCPWM capture channel failed because the channel is already enabled
- ESP_FAIL: Enable MCPWM capture channel failed because of other error

esp_err_t **mcpwm_capture_channel_disable** (*mcpwm_cap_channel_handle_t* cap_channel)

Disable MCPWM capture channel.

Parameters **cap_channel** -- [in] MCPWM capture channel handle, allocated by `mcpwm_new_capture_channel()`

Returns

- ESP_OK: Disable MCPWM capture channel successfully
- ESP_ERR_INVALID_ARG: Disable MCPWM capture channel failed because of invalid argument
- ESP_ERR_INVALID_STATE: Disable MCPWM capture channel failed because the channel is not enabled yet
- ESP_FAIL: Disable MCPWM capture channel failed because of other error

esp_err_t **mcpwm_capture_channel_register_event_callbacks** (*mcpwm_cap_channel_handle_t* cap_channel, const *mcpwm_capture_event_callbacks_t* *cbs, void *user_data)

Set event callbacks for MCPWM capture channel.

Note: The first call to this function needs to be before the call to `mcpwm_capture_channel_enable`

Note: User can deregister a previously registered callback by calling this function and setting the callback member in the `cbs` structure to NULL.

Parameters

- **cap_channel** -- [in] MCPWM capture channel handle, allocated by `mcpwm_new_capture_channel()`
- **cbs** -- [in] Group of callback functions
- **user_data** -- [in] User data, which will be passed to callback functions directly

Returns

- ESP_OK: Set event callbacks successfully
- ESP_ERR_INVALID_ARG: Set event callbacks failed because of invalid argument
- ESP_ERR_INVALID_STATE: Set event callbacks failed because the channel is not in init state
- ESP_FAIL: Set event callbacks failed because of other error

esp_err_t **mcpwm_capture_channel_trigger_soft_catch** (*mcpwm_cap_channel_handle_t* cap_channel)

Trigger a catch by software.

Parameters **cap_channel** -- [in] MCPWM capture channel handle, allocated by `mcpwm_new_capture_channel()`

Returns

- ESP_OK: Trigger software catch successfully
- ESP_ERR_INVALID_ARG: Trigger software catch failed because of invalid argument
- ESP_ERR_INVALID_STATE: Trigger software catch failed because the channel is not enabled yet
- ESP_FAIL: Trigger software catch failed because of other error

Structures

struct **mcpwm_capture_timer_config_t**

MCPWM capture timer configuration structure.

Public Members

int **group_id**

Specify from which group to allocate the capture timer

mcpwm_capture_clock_source_t **clk_src**

MCPWM capture timer clock source

uint32_t **resolution_hz**

Resolution of capture timer

struct **mcpwm_capture_timer_sync_phase_config_t**

MCPWM Capture timer sync phase configuration.

Public Members

mcpwm_sync_handle_t **sync_src**

The sync event source

uint32_t **count_value**

The count value that should lock to upon sync event

mcpwm_timer_direction_t **direction**

The count direction that should lock to upon sync event

struct **mcpwm_capture_channel_config_t**

MCPWM capture channel configuration structure.

Public Members

int **gpio_num**

GPIO used capturing input signal

int **intr_priority**

MCPWM capture interrupt priority, if set to 0, the driver will try to allocate an interrupt with a relative low priority (1,2,3)

uint32_t **prescale**

Prescale of input signal, effective frequency = cap_input_clk/prescale

uint32_t **pos_edge**

Whether to capture on positive edge

uint32_t **neg_edge**

Whether to capture on negative edge

uint32_t **pull_up**

Whether to pull up internally

uint32_t **pull_down**

Whether to pull down internally

uint32_t **invert_cap_signal**

Invert the input capture signal

uint32_t **io_loop_back**

For debug/test, the signal output from the GPIO will be fed to the input path as well

uint32_t **keep_io_conf_at_exit**

For debug/test, whether to keep the GPIO configuration when capture channel is deleted. By default, driver will reset the GPIO pin at exit.

struct *mcpwm_capture_channel_config_t*::[anonymous] **flags**

Extra configuration flags for capture channel

struct **mcpwm_capture_event_callbacks_t**

Group of supported MCPWM capture event callbacks.

Note: The callbacks are all running under ISR environment

Public Members

mcpwm_capture_event_cb_t **on_cap**

Callback function that would be invoked when capture event occurred

Header File

- [components/driver/mcpwm/include/driver/mcpwm_etm.h](#)
- This header file can be included with:

```
#include "driver/mcpwm_etm.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t **mcpwm_comparator_new_etm_event** (*mcpwm_cmpr_handle_t* cmpr, const *mcpwm_cmpr_etm_event_config_t* *config, *esp_etm_event_handle_t* *out_event)

Get the ETM event for MCPWM comparator.

Note: The created ETM event object can be deleted later by calling `esp_etm_del_event`

Parameters

- **cmpr** -- [in] MCPWM comparator, allocated by `mcpwm_new_comparator()` or `mcpwm_new_event_comparator()`
- **config** -- [in] MCPWM ETM comparator event configuration
- **out_event** -- [out] Returned ETM event handle

Returns

- **ESP_OK**: Get ETM event successfully
- **ESP_ERR_INVALID_ARG**: Get ETM event failed because of invalid argument
- **ESP_FAIL**: Get ETM event failed because of other error

Structures

struct **mcpwm_cmpr_etm_event_config_t**
MCPWM event comparator ETM event configuration.

Public Members

mcpwm_comparator_etm_event_type_t **event_type**
MCPWM comparator ETM event type

Header File

- [components/driver/mcpwm/include/driver/mcpwm_types.h](#)
- This header file can be included with:

```
#include "driver/mcpwm_types.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Structures

struct **mcpwm_timer_event_data_t**
MCPWM timer event data.

Public Members

uint32_t **count_value**
MCPWM timer count value

mcpwm_timer_direction_t **direction**
MCPWM timer count direction

struct **mcpwm_brake_event_data_t**
MCPWM brake event data.

struct **mcpwm_fault_event_data_t**
MCPWM fault event data.

struct **mcpwm_compare_event_data_t**
MCPWM compare event data.

Public Members

uint32_t **compare_ticks**

Compare value

mcpwm_timer_direction_t **direction**

Count direction

struct **mcpwm_capture_event_data_t**

MCPWM capture event data.

Public Members

uint32_t **cap_value**

Captured value

mcpwm_capture_edge_t **cap_edge**

Capture edge

Type Definitions

typedef struct mcpwm_timer_t ***mcpwm_timer_handle_t**

Type of MCPWM timer handle.

typedef struct mcpwm_oper_t ***mcpwm_oper_handle_t**

Type of MCPWM operator handle.

typedef struct mcpwm_cmpr_t ***mcpwm_cmpr_handle_t**

Type of MCPWM comparator handle.

typedef struct mcpwm_gen_t ***mcpwm_gen_handle_t**

Type of MCPWM generator handle.

typedef struct mcpwm_fault_t ***mcpwm_fault_handle_t**

Type of MCPWM fault handle.

typedef struct mcpwm_sync_t ***mcpwm_sync_handle_t**

Type of MCPWM sync handle.

typedef struct mcpwm_cap_timer_t ***mcpwm_cap_timer_handle_t**

Type of MCPWM capture timer handle.

typedef struct mcpwm_cap_channel_t ***mcpwm_cap_channel_handle_t**

Type of MCPWM capture channel handle.

typedef bool (***mcpwm_timer_event_cb_t**)(*mcpwm_timer_handle_t* timer, const *mcpwm_timer_event_data_t* *edata, void *user_ctx)

MCPWM timer event callback function.

Param timer [in] MCPWM timer handle

Param edata [in] MCPWM timer event data, fed by driver

Param user_ctx [in] User data, set in `mcpwm_timer_register_event_callbacks()`

Return Whether a high priority task has been waken up by this function

```
typedef bool (*mcpwm_brake_event_cb_t)(mcpwm_oper_handle_t oper, const mcpwm_brake_event_data_t *edata, void *user_ctx)
```

MCPWM operator brake event callback function.

Param oper [in] MCPWM operator handle

Param edata [in] MCPWM brake event data, fed by driver

Param user_ctx [in] User data, set in `mcpwm_operator_register_event_callbacks()`

Return Whether a high priority task has been waken up by this function

```
typedef bool (*mcpwm_fault_event_cb_t)(mcpwm_fault_handle_t fault, const mcpwm_fault_event_data_t *edata, void *user_ctx)
```

MCPWM fault event callback function.

Param fault MCPWM fault handle

Param edata MCPWM fault event data, fed by driver

Param user_ctx User data, set in `mcpwm_fault_register_event_callbacks()`

Return whether a task switch is needed after the callback returns

```
typedef bool (*mcpwm_compare_event_cb_t)(mcpwm_cmpr_handle_t comparator, const mcpwm_compare_event_data_t *edata, void *user_ctx)
```

MCPWM comparator event callback function.

Param comparator MCPWM comparator handle

Param edata MCPWM comparator event data, fed by driver

Param user_ctx User data, set in `mcpwm_comparator_register_event_callbacks()`

Return Whether a high priority task has been waken up by this function

```
typedef bool (*mcpwm_capture_event_cb_t)(mcpwm_cap_channel_handle_t cap_channel, const mcpwm_capture_event_data_t *edata, void *user_ctx)
```

MCPWM capture event callback function.

Param cap_channel MCPWM capture channel handle

Param edata MCPWM capture event data, fed by driver

Param user_ctx User data, set in `mcpwm_capture_channel_register_event_callbacks()`

Return Whether a high priority task has been waken up by this function

Header File

- `components/hal/include/hal/mcpwm_types.h`
- This header file can be included with:

```
#include "hal/mcpwm_types.h"
```

Type Definitions

```
typedef soc_periph_mcpwm_timer_clk_src_t mcpwm_timer_clock_source_t
```

MCPWM timer clock source.

```
typedef soc_periph_mcpwm_capture_clk_src_t mcpwm_capture_clock_source_t
```

MCPWM capture clock source.

```
typedef soc_periph_mcpwm_carrier_clk_src_t mcpwm_carrier_clock_source_t
```

MCPWM carrier clock source.

Enumerations

enum **mcpwm_timer_direction_t**

MCPWM timer count direction.

Values:

enumerator **MCPWM_TIMER_DIRECTION_UP**

Counting direction: Increase

enumerator **MCPWM_TIMER_DIRECTION_DOWN**

Counting direction: Decrease

enum **mcpwm_timer_event_t**

MCPWM timer events.

Values:

enumerator **MCPWM_TIMER_EVENT_EMPTY**

MCPWM timer counts to zero (i.e. counter is empty)

enumerator **MCPWM_TIMER_EVENT_FULL**

MCPWM timer counts to peak (i.e. counter is full)

enumerator **MCPWM_TIMER_EVENT_INVALID**

MCPWM timer invalid event

enum **mcpwm_timer_count_mode_t**

MCPWM timer count modes.

Values:

enumerator **MCPWM_TIMER_COUNT_MODE_PAUSE**

MCPWM timer paused

enumerator **MCPWM_TIMER_COUNT_MODE_UP**

MCPWM timer counting up

enumerator **MCPWM_TIMER_COUNT_MODE_DOWN**

MCPWM timer counting down

enumerator **MCPWM_TIMER_COUNT_MODE_UP_DOWN**

MCPWM timer counting up and down

enum **mcpwm_timer_start_stop_cmd_t**

MCPWM timer commands, specify the way to start or stop the timer.

Values:

enumerator **MCPWM_TIMER_STOP_EMPTY**

MCPWM timer stops when next count reaches zero

enumerator **MCPWM_TIMER_STOP_FULL**

MCPWM timer stops when next count reaches peak

enumerator **MCPWM_TIMER_START_NO_STOP**

MCPWM timer starts counting, and don't stop until received stop command

enumerator **MCPWM_TIMER_START_STOP_EMPTY**

MCPWM timer starts counting and stops when next count reaches zero

enumerator **MCPWM_TIMER_START_STOP_FULL**

MCPWM timer starts counting and stops when next count reaches peak

enum **mcpwm_generator_action_t**

MCPWM generator actions.

Values:

enumerator **MCPWM_GEN_ACTION_KEEP**

Generator action: Keep the same level

enumerator **MCPWM_GEN_ACTION_LOW**

Generator action: Force to low level

enumerator **MCPWM_GEN_ACTION_HIGH**

Generator action: Force to high level

enumerator **MCPWM_GEN_ACTION_TOGGLE**

Generator action: Toggle level

enum **mcpwm_operator_brake_mode_t**

MCPWM operator brake mode.

Values:

enumerator **MCPWM_OPER_BRAKE_MODE_CBC**

Brake mode: CBC (cycle by cycle)

enumerator **MCPWM_OPER_BRAKE_MODE_OST**

Brake mode: OST (one shot)

enumerator **MCPWM_OPER_BRAKE_MODE_INVALID**

MCPWM operator invalid brake mode

enum **mcpwm_capture_edge_t**

MCPWM capture edge.

Values:

enumerator **MCPWM_CAP_EDGE_POS**

Capture on the positive edge

enumerator **MCPWM_CAP_EDGE_NEG**

Capture on the negative edge

enum **mcpwm_comparator_etm_event_type_t**

MCPWM comparator specific events that supported by the ETM module.

Values:

enumerator **MCPWM_CMPR_ETM_EVENT_EQUAL**

The count value equals the value of comparator

enumerator **MCPWM_CMPR_ETM_EVENT_MAX**

Maximum number of comparator events

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.5.14 Parallel IO

Introduction

The Parallel IO peripheral is a general purpose parallel interface that can be used to connect to external devices such as LED matrix, LCD display, Printer and Camera. The peripheral has independent TX and RX units. Each unit can have up to 8 or 16 data signals plus 1 or 2 clock signals.¹

Warning: At the moment, the Parallel IO driver only supports TX mode. The RX feature is still working in progress.

Application Examples

- Simple REG LED Matrix with HUB75 interface: [peripherals/parlio/simple_rgb_led_matrix](#).

API Reference

Header File

- [components/driver/parlio/include/driver/parlio_tx.h](#)
- This header file can be included with:

```
#include "driver/parlio_tx.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

¹ Different ESP chip series might have different numbers of PARLIO TX/RX instances, and the maximum data bus can also be different. For more details, please refer to [ESP32-P4 Technical Reference Manual](#) > Chapter **Parallel IO (PARLIO)** [PDF]. The driver does not forbid you from applying for more driver objects, but it returns error when all available hardware resources are used up. Please always check the return value when doing resource allocation (e.g., `parlio_new_tx_unit()`).

PRIV_REQUIRES driver

Functions

esp_err_t **parlio_new_tx_unit** (const *parlio_tx_unit_config_t* *config, *parlio_tx_unit_handle_t* *ret_unit)

Create a Parallel IO TX unit.

Parameters

- **config** -- [in] Parallel IO TX unit configuration
- **ret_unit** -- [out] Returned Parallel IO TX unit handle

Returns

- ESP_OK: Create Parallel IO TX unit successfully
- ESP_ERR_INVALID_ARG: Create Parallel IO TX unit failed because of invalid argument
- ESP_ERR_NO_MEM: Create Parallel IO TX unit failed because of out of memory
- ESP_ERR_NOT_FOUND: Create Parallel IO TX unit failed because all TX units are used up and no more free one
- ESP_ERR_NOT_SUPPORTED: Create Parallel IO TX unit failed because some feature is not supported by hardware, e.g. clock gating
- ESP_FAIL: Create Parallel IO TX unit failed because of other error

esp_err_t **parlio_del_tx_unit** (*parlio_tx_unit_handle_t* unit)

Delete a Parallel IO TX unit.

Parameters **unit** -- [in] Parallel IO TX unit that created by `parlio_new_tx_unit`

Returns

- ESP_OK: Delete Parallel IO TX unit successfully
- ESP_ERR_INVALID_ARG: Delete Parallel IO TX unit failed because of invalid argument
- ESP_ERR_INVALID_STATE: Delete Parallel IO TX unit failed because it is still in working
- ESP_FAIL: Delete Parallel IO TX unit failed because of other error

esp_err_t **parlio_tx_unit_enable** (*parlio_tx_unit_handle_t* unit)

Enable the Parallel IO TX unit.

Note: This function will transit the driver state from init to enable

Note: This function will acquire a PM lock that might be installed during channel allocation

Note: If there're transaction pending in the queue, this function will pick up the first one and start the transfer

Parameters **unit** -- [in] Parallel IO TX unit that created by `parlio_new_tx_unit`

Returns

- ESP_OK: Enable Parallel IO TX unit successfully
- ESP_ERR_INVALID_ARG: Enable Parallel IO TX unit failed because of invalid argument
- ESP_ERR_INVALID_STATE: Enable Parallel IO TX unit failed because it is already enabled
- ESP_FAIL: Enable Parallel IO TX unit failed because of other error

esp_err_t **parlio_tx_unit_disable** (*parlio_tx_unit_handle_t* unit)

Disable the Parallel IO TX unit.

Note: This function will transit the driver state from enable to init

Note: This function will release the PM lock that might be installed during channel allocation

Note: If one transaction is undergoing, this function will terminate it immediately

Parameters **unit** -- [in] Parallel IO TX unit that created by `parlio_new_tx_unit`

Returns

- `ESP_OK`: Disable Parallel IO TX unit successfully
- `ESP_ERR_INVALID_ARG`: Disable Parallel IO TX unit failed because of invalid argument
- `ESP_ERR_INVALID_STATE`: Disable Parallel IO TX unit failed because it's not enabled yet
- `ESP_FAIL`: Disable Parallel IO TX unit failed because of other error

esp_err_t **parlio_tx_unit_register_event_callbacks** (*parlio_tx_unit_handle_t* tx_unit, const *parlio_tx_event_callbacks_t* *cbs, void *user_data)

Set event callbacks for Parallel IO TX unit.

Note: User can deregister a previously registered callback by calling this function and setting the callback member in the `cbs` structure to NULL.

Note: When `CONFIG_PARLIO_ISR_IRAM_SAFE` is enabled, the callback itself and functions called by it should be placed in IRAM. The variables used in the function should be in the SRAM as well. The `user_data` should also reside in SRAM.

Parameters

- **tx_unit** -- [in] Parallel IO TX unit that created by `parlio_new_tx_unit`
- **cbs** -- [in] Group of callback functions
- **user_data** -- [in] User data, which will be passed to callback functions directly

Returns

- `ESP_OK`: Set event callbacks successfully
- `ESP_ERR_INVALID_ARG`: Set event callbacks failed because of invalid argument
- `ESP_FAIL`: Set event callbacks failed because of other error

esp_err_t **parlio_tx_unit_transmit** (*parlio_tx_unit_handle_t* tx_unit, const void *payload, size_t payload_bits, const *parlio_transmit_config_t* *config)

Transmit data on by Parallel IO TX unit.

Note: After the function returns, it doesn't mean the transaction is finished. This function only constructs a transaction structure and push into a queue.

Parameters

- **tx_unit** -- [in] Parallel IO TX unit that created by `parlio_new_tx_unit`
- **payload** -- [in] Pointer to the data to be transmitted
- **payload_bits** -- [in] Length of the data to be transmitted, in bits
- **config** -- [in] Transmit configuration

Returns

- ESP_OK: Transmit data successfully
- ESP_ERR_INVALID_ARG: Transmit data failed because of invalid argument
- ESP_ERR_INVALID_STATE: Transmit data failed because the Parallel IO TX unit is not enabled
- ESP_FAIL: Transmit data failed because of other error

esp_err_t **parlio_tx_unit_wait_all_done** (*parlio_tx_unit_handle_t* tx_unit, int timeout_ms)

Wait for all pending TX transactions done.

Parameters

- **tx_unit** -- [in] Parallel IO TX unit that created by `parlio_new_tx_unit`
- **timeout_ms** -- [in] Timeout in milliseconds, -1 means to wait forever

Returns

- ESP_OK: All pending TX transactions is finished and recycled
- ESP_ERR_INVALID_ARG: Wait for all pending TX transactions done failed because of invalid argument
- ESP_ERR_TIMEOUT: Wait for all pending TX transactions done timeout
- ESP_FAIL: Wait for all pending TX transactions done failed because of other error

Structures

struct **parlio_tx_unit_config_t**

Parallel IO TX unit configuration.

Public Members

parlio_clock_source_t **clk_src**

Parallel IO internal clock source

gpio_num_t **clk_in_gpio_num**

If the clock source is input from external, set the corresponding GPIO number. Otherwise, set to -1 and the driver will use the internal `clk_src` as clock source. This option has higher priority than `clk_src`

uint32_t **input_clk_src_freq_hz**

Frequency of the input clock source, valid only if `clk_in_gpio_num` is not -1

uint32_t **output_clk_freq_hz**

Frequency of the output clock. It's divided from either internal `clk_src` or external clock source

size_t **data_width**

Parallel IO data width, can set to 1/2/4/8/..., but can't bigger than `PARLIO_TX_UNIT_MAX_DATA_WIDTH`

gpio_num_t **data_gpio_nums**[`PARLIO_TX_UNIT_MAX_DATA_WIDTH`]

Parallel IO data GPIO numbers, if any GPIO is not used, you can set it to -1

gpio_num_t **clk_out_gpio_num**

GPIO number of the output clock signal, the clock is synced with TX data

gpio_num_t **valid_gpio_num**

GPIO number of the valid signal, which stays high when transferring data. Note that, the valid signal will always occupy the MSB data bit

size_t **trans_queue_depth**

Depth of internal transaction queue

size_t **max_transfer_size**

Maximum transfer size in one transaction, in bytes. This decides the number of DMA nodes will be used for each transaction

parlio_sample_edge_t **sample_edge**

Parallel IO sample edge

parlio_shift_edge_t **shift_edge**

Parallel IO Tx shift edge

parlio_bit_pack_order_t **bit_pack_order**

Set the order of packing the bits into bytes (only works when `data_width < 8`)

uint32_t **clk_gate_en**

Enable TX clock gating, the output clock will be controlled by the MSB bit of the data bus, i.e. by `data_gpio_nums[PARLIO_TX_UNIT_MAX_DATA_WIDTH-1]`. High level to enable the clock output, low to disable

uint32_t **io_loop_back**

For debug/test, the signal output from the GPIO will be fed to the input path as well

struct *parlio_tx_unit_config_t*::[anonymous] **flags**

Extra configuration flags

struct **parlio_tx_done_event_data_t**

Type of Parallel IO TX done event data.

struct **parlio_tx_event_callbacks_t**

Group of Parallel IO TX callbacks.

Note: The callbacks are all running under ISR environment

Note: When `CONFIG_PARLIO_ISR_IRAM_SAFE` is enabled, the callback itself and functions called by it should be placed in IRAM. The variables used in the function should be in the SRAM as well.

Public Members

parlio_tx_done_callback_t **on_trans_done**

Event callback, invoked when one transmission is finished

struct **parlio_transmit_config_t**

Parallel IO transmit configuration.

Public Members

uint32_t **idle_value**

The value on the data line when the parallel IO is in idle state

uint32_t **queue_nonblocking**

If set, when the transaction queue is full, driver will not block the thread but return directly

struct *parlio_transmit_config_t*::[anonymous] **flags**

Transmit specific config flags

Type Definitions

typedef bool (***parlio_tx_done_callback_t**)(*parlio_tx_unit_handle_t* tx_unit, const *parlio_tx_done_event_data_t* *edata, void *user_ctx)

Prototype of parlio tx event callback.

Param tx_unit [in] Parallel IO TX unit that created by `parlio_new_tx_unit`

Param edata [in] Point to Parallel IO TX event data. The lifecycle of this pointer memory is inside this function, user should copy it into static memory if used outside this function.

Param user_ctx [in] User registered context, passed from `parlio_tx_unit_register_event_callbacks`

Return Whether a high priority task has been waken up by this callback function

Header File

- [components/driver/parlio/include/driver/parlio_types.h](#)
- This header file can be included with:

```
#include "driver/parlio_types.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Type Definitions

typedef struct *parlio_tx_unit_t* ***parlio_tx_unit_handle_t**

Type of Parallel IO TX unit handle.

Header File

- [components/hal/include/hal/parlio_types.h](#)
- This header file can be included with:

```
#include "hal/parlio_types.h"
```

Macros

PARLIO_TX_UNIT_MAX_DATA_WIDTH

Maximum data width of TX unit.

Type Definitions

typedef *soc_periph_parlio_clk_src_t* **parlio_clock_source_t**

Parallel IO clock source.

Note: User should select the clock source based on the power and resolution requirement

Enumerations

enum **parlio_sample_edge_t**

Parallel IO Rx sample edge.

Values:

enumerator **PARLIO_SAMPLE_EDGE_NEG**

Sample data on falling edge of clock

enumerator **PARLIO_SAMPLE_EDGE_POS**

Sample data on rising edge of clock

enum **parlio_shift_edge_t**

Parallel IO Tx shift edge.

Values:

enumerator **PARLIO_SHIFT_EDGE_POS**

Shift data on rising edge of clock

enumerator **PARLIO_SHIFT_EDGE_NEG**

Shift data on falling edge of clock

enum **parlio_bit_pack_order_t**

Parallel IO bit packing order.

Data in memory: Byte 0: MSB < B0.7 B0.6 B0.5 B0.4 B0.3 B0.2 B0.1 B0.0 > LSB Byte 1: MSB < B1.7 B1.6 B1.5 B1.4 B1.3 B1.2 B1.1 B1.0 > LSB

Output on line (PARLIO_BIT_PACK_ORDER_LSB): Cycle 0 Cycle 1 Cycle 2 —> time GPIO 0: B0.0 B0.4 B1.0 GPIO 1: B0.1 B0.5 B1.1 GPIO 2: B0.2 B0.6 B1.2 GPIO 3: B0.3 B0.7 B1.3

Output on line (PARLIO_BIT_PACK_ORDER_MSB): Cycle 0 Cycle 1 Cycle 2 —> time GPIO 0: B0.4 B0.0 B1.4 GPIO 1: B0.5 B0.1 B1.5 GPIO 2: B0.6 B0.2 B1.6 GPIO 3: B0.7 B0.3 B1.7

Values:

enumerator **PARLIO_BIT_PACK_ORDER_LSB**

Bit pack order: LSB

enumerator **PARLIO_BIT_PACK_ORDER_MSB**

Bit pack order: MSB

2.5.15 Pulse Counter (PCNT)

Introduction

The PCNT (Pulse Counter) module is designed to count the number of rising and/or falling edges of input signals. The ESP32-P4 contains multiple pulse counter units in the module.¹ Each unit is in effect an independent counter with multiple channels, where each channel can increment/decrement the counter on a rising/falling edge. Furthermore, each channel can be configured separately.

PCNT channels can react to signals of **edge** type and **level** type, however for simple applications, detecting the edge signal is usually sufficient. PCNT channels can be configured react to both pulse edges (i.e., rising and falling edge), and can be configured to increase, decrease or do nothing to the unit's counter on each edge. The level signal is the so-called **control signal**, which is used to control the counting mode of the edge signals that are attached to the same channel. By combining the usage of both edge and level signals, a PCNT unit can act as a **quadrature decoder**.

Besides that, PCNT unit is equipped with a separate glitch filter, which is helpful to remove noise from the signal.

Typically, a PCNT module can be used in scenarios like:

- Calculate periodic signal's frequency by counting the pulse numbers within a time slice
- Decode quadrature signals into speed and direction

Functional Overview

Description of the PCNT functionality is divided into the following sections:

- [Resource Allocation](#) - covers how to allocate PCNT units and channels with properly set of configurations. It also covers how to recycle the resources when they finished working.
- [Set Up Channel Actions](#) - covers how to configure the PCNT channel to behave on different signal edges and levels.
- [Watch Points](#) - describes how to configure PCNT watch points (i.e., tell PCNT unit to trigger an event when the count reaches a certain value).
- [Register Event Callbacks](#) - describes how to hook your specific code to the watch point event callback function.
- [Set Glitch Filter](#) - describes how to enable and set the timing parameters for the internal glitch filter.
- [Use External Clear Signal](#) - describes how to set the parameters for the external clear signal.
- [Enable and Disable Unit](#) - describes how to enable and disable the PCNT unit.
- [Unit IO Control](#) - describes IO control functions of PCNT unit, like enable glitch filter, start and stop unit, get and clear count value.
- [Power Management](#) - describes what functionality will prevent the chip from going into low power mode.
- [IRAM Safe](#) - describes tips on how to make the PCNT interrupt and IO control functions work better along with a disabled cache.
- [Thread Safety](#) - lists which APIs are guaranteed to be thread safe by the driver.
- [Kconfig Options](#) - lists the supported Kconfig options that can be used to make a different effect on driver behavior.

Resource Allocation The PCNT unit and channel are represented by `pcnt_unit_handle_t` and `pcnt_channel_handle_t` respectively. All available units and channels are maintained by the driver in a resource pool, so you do not need to know the exact underlying instance ID.

Install PCNT Unit To install a PCNT unit, there is a configuration structure that needs to be given in advance: `pcnt_unit_config_t`:

¹ Different ESP chip series might have different number of PCNT units and channels. Please refer to the [TRM] for details. The driver does not forbid you from applying for more PCNT units and channels, but it returns error when all available hardware resources are used up. Please always check the return value when doing resource allocation (e.g., `pcnt_new_unit()`).

- `pcnt_unit_config_t::low_limit` and `pcnt_unit_config_t::high_limit` specify the range for the internal hardware counter. The counter will reset to zero automatically when it crosses either the high or low limit.
- `pcnt_unit_config_t::accum_count` sets whether to create an internal accumulator for the counter. This is helpful when you want to extend the counter's width, which by default is 16 bit at most, defined in the hardware. See also [Compensate Overflow Loss](#) for how to use this feature to compensate the overflow loss.
- `pcnt_unit_config_t::intr_priority` sets the priority of the interrupt. If it is set to 0, the driver will allocate an interrupt with a default priority. Otherwise, the driver will use the given priority.

Note: Since all PCNT units share the same interrupt source, when installing multiple PCNT units make sure that the interrupt priority `pcnt_unit_config_t::intr_priority` is the same for each unit.

Unit allocation and initialization is done by calling a function `pcnt_new_unit()` with `pcnt_unit_config_t` as an input parameter. The function will return a PCNT unit handle only when it runs correctly. Specifically, when there are no more free PCNT units in the pool (i.e., unit resources have been used up), then this function will return `ESP_ERR_NOT_FOUND` error. The total number of available PCNT units is recorded by `SOC_PCNT_UNITS_PER_GROUP` for reference.

If a previously created PCNT unit is no longer needed, it is recommended to recycle the resource by calling `pcnt_del_unit()`. Which in return allows the underlying unit hardware to be used for other purposes. Before deleting a PCNT unit, one should ensure the following prerequisites:

- The unit is in the init state, in other words, the unit is either disabled by `pcnt_unit_disable()` or not enabled yet.
- The attached PCNT channels are all removed by `pcnt_del_channel()`.

```
#define EXAMPLE_PCNT_HIGH_LIMIT 100
#define EXAMPLE_PCNT_LOW_LIMIT -100

pcnt_unit_config_t unit_config = {
    .high_limit = EXAMPLE_PCNT_HIGH_LIMIT,
    .low_limit = EXAMPLE_PCNT_LOW_LIMIT,
};
pcnt_unit_handle_t pcnt_unit = NULL;
ESP_ERROR_CHECK(pcnt_new_unit(&unit_config, &pcnt_unit));
```

Install PCNT Channel To install a PCNT channel, you must initialize a `pcnt_chan_config_t` structure in advance, and then call `pcnt_new_channel()`. The configuration fields of the `pcnt_chan_config_t` structure are described below:

- `pcnt_chan_config_t::edge_gpio_num` and `pcnt_chan_config_t::level_gpio_num` specify the GPIO numbers used by **edge** type signal and **level** type signal. Please note, either of them can be assigned to -1 if it is not actually used, and thus it will become a **virtual IO**. For some simple pulse counting applications where one of the level/edge signals is fixed (i.e., never changes), you can reclaim a GPIO by setting the signal as a virtual IO on channel allocation. Setting the level/edge signal as a virtual IO causes that signal to be internally routed to a fixed High/Low logic level, thus allowing you to save a GPIO for other purposes.
- `pcnt_chan_config_t::virt_edge_io_level` and `pcnt_chan_config_t::virt_level_io_level` specify the virtual IO level for **edge** and **level** input signal, to ensure a deterministic state for such control signal. Please note, they are only valid when either `pcnt_chan_config_t::edge_gpio_num` or `pcnt_chan_config_t::level_gpio_num` is assigned to -1.
- `pcnt_chan_config_t::invert_edge_input` and `pcnt_chan_config_t::invert_level_input` are used to decide whether to invert the input signals before they going into PCNT hardware. The invert is done by GPIO matrix instead of PCNT hardware.
- `pcnt_chan_config_t::io_loop_back` is for debug only, which enables both the GPIO's input and output paths. This can help to simulate the pulse signals by function `gpio_set_level()` on the same GPIO.

Channel allocating and initialization is done by calling a function `pcnt_new_channel()` with the above `pcnt_chan_config_t` as an input parameter plus a PCNT unit handle returned from `pcnt_new_unit()`. This function will return a PCNT channel handle if it runs correctly. Specifically, when there are no more free PCNT channel within the unit (i.e., channel resources have been used up), then this function will return `ESP_ERR_NOT_FOUND` error. The total number of available PCNT channels within the unit is recorded by `SOC_PCNT_CHANNELS_PER_UNIT` for reference. Note that, when install a PCNT channel for a specific unit, one should ensure the unit is in the init state, otherwise this function will return `ESP_ERR_INVALID_STATE` error.

If a previously created PCNT channel is no longer needed, it is recommended to recycle the resources by calling `pcnt_del_channel()`. Which in return allows the underlying channel hardware to be used for other purposes.

```
#define EXAMPLE_CHAN_GPIO_A 0
#define EXAMPLE_CHAN_GPIO_B 2

pcnt_chan_config_t chan_config = {
    .edge_gpio_num = EXAMPLE_CHAN_GPIO_A,
    .level_gpio_num = EXAMPLE_CHAN_GPIO_B,
};
pcnt_channel_handle_t pcnt_chan = NULL;
ESP_ERROR_CHECK(pcnt_new_channel(pcnt_unit, &chan_config, &pcnt_chan));
```

Set Up Channel Actions The PCNT will increase/decrease/hold its internal count value when the input pulse signal toggles. You can set different actions for edge signal and/or level signal.

- `pcnt_channel_set_edge_action()` function is to set specific actions for rising and falling edge of the signal attached to the `pcnt_chan_config_t::edge_gpio_num`. Supported actions are listed in `pcnt_channel_edge_action_t`.
- `pcnt_channel_set_level_action()` function is to set specific actions for high and low level of the signal attached to the `pcnt_chan_config_t::level_gpio_num`. Supported actions are listed in `pcnt_channel_level_action_t`. This function is not mandatory if the `pcnt_chan_config_t::level_gpio_num` is set to -1 when allocating PCNT channel by `pcnt_new_channel()`.

```
// decrease the counter on rising edge, increase the counter on falling edge
ESP_ERROR_CHECK(pcnt_channel_set_edge_action(pcnt_chan, PCNT_CHANNEL_EDGE_ACTION_
↪DECREASE, PCNT_CHANNEL_EDGE_ACTION_INCREASE));
// keep the counting mode when the control signal is high level, and reverse the_
↪counting mode when the control signal is low level
ESP_ERROR_CHECK(pcnt_channel_set_level_action(pcnt_chan, PCNT_CHANNEL_LEVEL_ACTION_
↪KEEP, PCNT_CHANNEL_LEVEL_ACTION_INVERSE));
```

Watch Points Each PCNT unit can be configured to watch several different values that you are interested in. The value to be watched is also called **Watch Point**. The watch point itself can not exceed the range set in `pcnt_unit_config_t` by `pcnt_unit_config_t::low_limit` and `pcnt_unit_config_t::high_limit`. When the counter reaches either watch point, a watch event will be triggered and notify you by interrupt if any watch event callback has ever registered in `pcnt_unit_register_event_callbacks()`. See [Register Event Callbacks](#) for how to register event callbacks.

The watch point can be added and removed by `pcnt_unit_add_watch_point()` and `pcnt_unit_remove_watch_point()`. The commonly-used watch points are: **zero cross**, **maximum/minimum count** and other threshold values. The number of available watch point is limited, `pcnt_unit_add_watch_point()` will return error `ESP_ERR_NOT_FOUND` if it can not find any free hardware resource to save the watch point. You can not add the same watch point for multiple times, otherwise it will return error `ESP_ERR_INVALID_STATE`.

It is recommended to remove the unused watch point by `pcnt_unit_remove_watch_point()` to recycle the watch point resources.

```
// add zero across watch point
ESP_ERROR_CHECK(pcmt_unit_add_watch_point(pcmt_unit, 0));
// add high limit watch point
ESP_ERROR_CHECK(pcmt_unit_add_watch_point(pcmt_unit, EXAMPLE_PCNT_HIGH_LIMIT));
```

Register Event Callbacks When PCNT unit reaches any enabled watch point, specific event will be generated and notify the CPU by interrupt. If you have some function that want to get executed when event happens, you should hook your function to the interrupt service routine by calling `pcmt_unit_register_event_callbacks()`. All supported event callbacks are listed in the `pcmt_event_callbacks_t`:

- `pcmt_event_callbacks_t::on_reach` sets a callback function for watch point event. As this function is called within the ISR context, you must ensure that the function does not attempt to block (e.g., by making sure that only FreeRTOS APIs with ISR suffix are called from within the function). The function prototype is declared in `pcmt_watch_cb_t`.

You can save their own context to `pcmt_unit_register_event_callbacks()` as well, via the parameter `user_ctx`. This user data will be directly passed to the callback functions.

In the callback function, the driver will fill in the event data of specific event. For example, the watch point event data is declared as `pcmt_watch_event_data_t`:

- `pcmt_watch_event_data_t::watch_point_value` saves the watch point value that triggers the event.
- `pcmt_watch_event_data_t::zero_cross_mode` saves how the PCNT unit crosses the zero point in the latest time. The possible zero cross modes are listed in the `pcmt_unit_zero_cross_mode_t`. Usually different zero cross mode means different **counting direction** and **counting step size**.

Registering callback function results in lazy installation of interrupt service, thus this function should only be called before the unit is enabled by `pcmt_unit_enable()`. Otherwise, it can return `ESP_ERR_INVALID_STATE` error.

```
static bool example_pcmt_on_reach(pcmt_unit_handle_t unit, const pcmt_watch_event_
↳data_t *edata, void *user_ctx)
{
    BaseType_t high_task_wakeup;
    QueueHandle_t queue = (QueueHandle_t)user_ctx;
    // send watch point to queue, from this interrupt callback
    xQueueSendFromISR(queue, &(edata->watch_point_value), &high_task_wakeup);
    // return whether a high priority task has been waken up by this function
    return (high_task_wakeup == pdTRUE);
}

pcmt_event_callbacks_t cbs = {
    .on_reach = example_pcmt_on_reach,
};
QueueHandle_t queue = xQueueCreate(10, sizeof(int));
ESP_ERROR_CHECK(pcmt_unit_register_event_callbacks(pcmt_unit, &cbs, queue));
```

Set Glitch Filter The PCNT unit features filters to ignore possible short glitches in the signals. The parameters that can be configured for the glitch filter are listed in `pcmt_glitch_filter_config_t`:

- `pcmt_glitch_filter_config_t::max_glitch_ns` sets the maximum glitch width, in nano seconds. If a signal pulse's width is smaller than this value, then it will be treated as noise and will not increase/decrease the internal counter.

You can enable the glitch filter for PCNT unit by calling `pcmt_unit_set_glitch_filter()` with the filter configuration provided above. Particularly, you can disable the glitch filter later by calling `pcmt_unit_set_glitch_filter()` with a NULL filter configuration.

This function should be called when the unit is in the init state. Otherwise, it will return `ESP_ERR_INVALID_STATE` error.

Note: The glitch filter operates using the APB clock. To ensure the counter does not miss any pulses, the maximum glitch width should be longer than one APB_CLK cycle (typically 12.5 ns if APB is 80 MHz). Since the APB frequency can change with Dynamic Frequency Scaling (DFS), the filter may not function as expected in such cases. Therefore, the driver installs a power management lock for each PCNT unit. For more details on the power management strategy used in the PCNT driver, please refer to [Power Management](#).

```
pcnt_glitch_filter_config_t filter_config = {
    .max_glitch_ns = 1000,
};
ESP_ERROR_CHECK(pcnt_unit_set_glitch_filter(pcnt_unit, &filter_config));
```

Use External Clear Signal The PCNT unit can receive a clear signal from the GPIO. The parameters that can be configured for the clear signal are listed in `pcnt_clear_signal_config_t`:

- `pcnt_clear_signal_config_t::clear_signal_gpio_num` specify the GPIO numbers used by **clear** signal. The default active level is high, and the input mode is pull-down enabled.
- `pcnt_clear_signal_config_t::invert_clear_signal` is used to decide whether to invert the input signal before it going into PCNT hardware. The invert is done by GPIO matrix instead of PCNT hardware. The input mode is pull-up enabled when the input signal is inverted.
- `pcnt_clear_signal_config_t::io_loop_back` is for debug only, which enables both the GPIO's input and output paths. This can help to simulate the clear signal by function `gpio_set_level()` for the same GPIO.

This signal acts in the same way as calling `pcnt_unit_clear_count()`, but is not subject to software latency, and is suitable for use in situations with low latency requirements. Also please note, the flip frequency of this signal can not be too high.

```
pcnt_clear_signal_config_t clear_signal_config = {
    .clear_signal_gpio_num = PCNT_CLEAR_SIGNAL_GPIO,
};
ESP_ERROR_CHECK(pcnt_unit_set_clear_signal(pcnt_unit, &clear_signal_config));
```

Enable and Disable Unit Before doing IO control to the PCNT unit, you need to enable it first, by calling `pcnt_unit_enable()`. Internally, this function:

- switches the PCNT driver state from **init** to **enable**.
- enables the interrupt service if it has been lazy installed in `pcnt_unit_register_event_callbacks()`.
- acquires a proper power management lock if it has been installed. See also [Power Management](#) for more information.

On the contrary, calling `pcnt_unit_disable()` will do the opposite, that is, put the PCNT driver back to the **init** state, disable the interrupts service and release the power management lock.

Unit IO Control

Start/Stop and Clear Calling `pcnt_unit_start()` makes the PCNT unit start to work, increase or decrease counter according to pulse signals. On the contrary, calling `pcnt_unit_stop()` will stop the PCNT unit but retain current count value. Instead, clearing counter can only be done by calling `pcnt_unit_clear_count()`.

Note, `pcnt_unit_start()` and `pcnt_unit_stop()` should be called when the unit has been enabled by `pcnt_unit_enable()`. Otherwise, it will return `ESP_ERR_INVALID_STATE` error.

Get Count Value You can read current count value at any time by calling `pcnt_unit_get_count()`. The returned count value is a **signed** integer, where the sign can be used to reflect the direction.

```
int pulse_count = 0;
ESP_ERROR_CHECK(pcnt_unit_get_count(pcnt_unit, &pulse_count));
```

Compensate Overflow Loss The internal hardware counter will be cleared to zero automatically when it reaches high or low limit. If you want to compensate for that count loss and extend the counter's bit-width, you can:

1. Enable `pcnt_unit_config_t::accum_count` when installing the PCNT unit.
2. Add the high/low limit as the [Watch Points](#).
3. Now, the returned count value from the `pcnt_unit_get_count()` function not only reflects the hardware's count value, but also accumulates the high/low overflow loss to it.

Note: `pcnt_unit_clear_count()` resets the accumulated count value as well.

Note: When enabling the count overflow compensation, it is recommended to use as large a high/low count limit as possible, as it can avoid frequent interrupt triggering, improve system performance, and avoid compensation failure due to multiple overflows.

Power Management When power management is enabled (i.e., `CONFIG_PM_ENABLE` is on), the system adjusts the APB frequency before entering light sleep, which can cause the PCNT glitch filter to misinterpret valid signals as noise.

To prevent this, the driver can acquire a power management lock of type `ESP_PM_APB_FREQ_MAX`, ensuring the APB frequency remains constant. This lock is acquired when the PCNT unit is enabled via `pcnt_unit_enable()` and released when the unit is disabled via `pcnt_unit_disable()`.

IRAM Safe By default, the PCNT interrupt will be deferred when the Cache is disabled for reasons like writing/erasing Flash. Thus the alarm interrupt will not get executed in time, which is not expected in a real-time application.

There is a Kconfig option `CONFIG_PCNT_ISR_IRAM_SAFE` that:

1. Enables the interrupt being serviced even when cache is disabled
2. Places all functions that used by the ISR into IRAM²
3. Places driver object into DRAM (in case it is mapped to PSRAM by accident)

This allows the interrupt to run while the cache is disabled but comes at the cost of increased IRAM consumption.

There is another Kconfig option `CONFIG_PCNT_CTRL_FUNC_IN_IRAM` that can put commonly used IO control functions into IRAM as well. So that these functions can also be executable when the cache is disabled. These IO control functions are as follows:

- `pcnt_unit_start()`
- `pcnt_unit_stop()`
- `pcnt_unit_clear_count()`
- `pcnt_unit_get_count()`

Thread Safety The factory functions `pcnt_new_unit()` and `pcnt_new_channel()` are guaranteed to be thread safe by the driver, which means, you can call them from different RTOS tasks without protection by extra locks.

The following functions are allowed to run under ISR context, the driver uses a critical section to prevent them being called concurrently in both task and ISR.

- `pcnt_unit_start()`

² `pcnt_event_callbacks_t::on_reach` callback and the functions invoked by itself should also be placed in IRAM, you need to take care of them by themselves.

- [`pcnt\_unit\_stop\(\)`](#)
- [`pcnt\_unit\_clear\_count\(\)`](#)
- [`pcnt\_unit\_get\_count\(\)`](#)

Other functions that take the [`pcnt\_unit\_handle\_t`](#) and [`pcnt\_channel\_handle\_t`](#) as the first positional parameter, are not treated as thread safe. This means you should avoid calling them from multiple tasks.

Kconfig Options

- [`CONFIG\_PCNT\_CTRL\_FUNC\_IN\_IRAM`](#) controls where to place the PCNT control functions (IRAM or Flash), see [*IRAM Safe*](#) for more information.
- [`CONFIG\_PCNT\_ISR\_IRAM\_SAFE`](#) controls whether the default ISR handler can work when cache is disabled, see [*IRAM Safe*](#) for more information.
- [`CONFIG\_PCNT\_ENABLE\_DEBUG\_LOG`](#) is used to enable the debug log output. Enabling this option increases the firmware binary size.

Application Examples

- Decode the quadrature signals from rotary encoder: [peripherals/pcnt/rotary_encoder](#).

API Reference

Header File

- [components/driver/pcnt/include/driver/pulse_cnt.h](#)
- This header file can be included with:

```
#include "driver/pulse_cnt.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

[`esp\_err\_t pcnt\_new\_unit`](#) (const [`pcnt\_unit\_config\_t`](#) *config, [`pcnt\_unit\_handle\_t`](#) *ret_unit)

Create a new PCNT unit, and return the handle.

Note: The newly created PCNT unit is put in the init state.

Parameters

- **config** -- [in] PCNT unit configuration
- **ret_unit** -- [out] Returned PCNT unit handle

Returns

- `ESP_OK`: Create PCNT unit successfully
- `ESP_ERR_INVALID_ARG`: Create PCNT unit failed because of invalid argument (e.g. high/low limit value out of the range)
- `ESP_ERR_NO_MEM`: Create PCNT unit failed because out of memory
- `ESP_ERR_NOT_FOUND`: Create PCNT unit failed because all PCNT units are used up and no more free one
- `ESP_FAIL`: Create PCNT unit failed because of other error

esp_err_t **pcnt_del_unit** (*pcnt_unit_handle_t* unit)

Delete the PCNT unit handle.

Note: A PCNT unit can't be in the enable state when this function is invoked. See also `pcnt_unit_disable()` for how to disable a unit.

Parameters **unit** -- [in] PCNT unit handle created by `pcnt_new_unit()`

Returns

- **ESP_OK**: Delete the PCNT unit successfully
- **ESP_ERR_INVALID_ARG**: Delete the PCNT unit failed because of invalid argument
- **ESP_ERR_INVALID_STATE**: Delete the PCNT unit failed because the unit is not in init state or some PCNT channel is still in working
- **ESP_FAIL**: Delete the PCNT unit failed because of other error

esp_err_t **pcnt_unit_set_glitch_filter** (*pcnt_unit_handle_t* unit, const *pcnt_glitch_filter_config_t* *config)

Set glitch filter for PCNT unit.

Note: The glitch filter module is clocked from APB, and APB frequency can be changed during DFS, which in return make the filter out of action. So this function will lazy-install a PM lock internally when the power management is enabled. With this lock, the APB frequency won't be changed. The PM lock can be uninstalled in `pcnt_del_unit()`.

Note: This function should be called when the PCNT unit is in the init state (i.e. before calling `pcnt_unit_enable()`)

Parameters

- **unit** -- [in] PCNT unit handle created by `pcnt_new_unit()`
- **config** -- [in] PCNT filter configuration, set config to NULL means disabling the filter function

Returns

- **ESP_OK**: Set glitch filter successfully
- **ESP_ERR_INVALID_ARG**: Set glitch filter failed because of invalid argument (e.g. glitch width is too big)
- **ESP_ERR_INVALID_STATE**: Set glitch filter failed because the unit is not in the init state
- **ESP_FAIL**: Set glitch filter failed because of other error

esp_err_t **pcnt_unit_set_clear_signal** (*pcnt_unit_handle_t* unit, const *pcnt_clear_signal_config_t* *config)

Set clear signal for PCNT unit.

Note: The function of clear signal is the same as `pcnt_unit_clear_count()`. High-level Active

Parameters

- **unit** -- [in] PCNT unit handle created by `pcnt_new_unit()`
- **config** -- [in] PCNT clear signal configuration, set config to NULL means disabling the clear signal

Returns

- **ESP_OK**: Set clear signal successfully
- **ESP_ERR_INVALID_ARG**: Set clear signal failed because of invalid argument

- ESP_ERR_INVALID_STATE: Set clear signal failed because set clear signal repeatedly or disable clear signal before set it
- ESP_FAIL: Set clear signal failed because of other error

esp_err_t **pcnt_unit_enable** (*pcnt_unit_handle_t* unit)

Enable the PCNT unit.

Note: This function will transit the unit state from init to enable.

Note: This function will enable the interrupt service, if it's lazy installed in `pcnt_unit_register_event_callbacks()`.

Note: This function will acquire the PM lock if it's lazy installed in `pcnt_unit_set_glitch_filter()`.

Note: Enable a PCNT unit doesn't mean to start it. See also `pcnt_unit_start()` for how to start the PCNT counter.

Parameters **unit** -- [in] PCNT unit handle created by `pcnt_new_unit()`

Returns

- ESP_OK: Enable PCNT unit successfully
- ESP_ERR_INVALID_ARG: Enable PCNT unit failed because of invalid argument
- ESP_ERR_INVALID_STATE: Enable PCNT unit failed because the unit is already enabled
- ESP_FAIL: Enable PCNT unit failed because of other error

esp_err_t **pcnt_unit_disable** (*pcnt_unit_handle_t* unit)

Disable the PCNT unit.

Note: This function will do the opposite work to the `pcnt_unit_enable()`

Note: Disable a PCNT unit doesn't mean to stop it. See also `pcnt_unit_stop()` for how to stop the PCNT counter.

Parameters **unit** -- [in] PCNT unit handle created by `pcnt_new_unit()`

Returns

- ESP_OK: Disable PCNT unit successfully
- ESP_ERR_INVALID_ARG: Disable PCNT unit failed because of invalid argument
- ESP_ERR_INVALID_STATE: Disable PCNT unit failed because the unit is not enabled yet
- ESP_FAIL: Disable PCNT unit failed because of other error

esp_err_t **pcnt_unit_start** (*pcnt_unit_handle_t* unit)

Start the PCNT unit, the counter will start to count according to the edge and/or level input signals.

Note: This function should be called when the unit is in the enable state (i.e. after calling `pcnt_unit_enable()`)

Note: This function is allowed to run within ISR context

Note: This function will be placed into IRAM if `CONFIG_PCNT_CTRL_FUNC_IN_IRAM` is on, so that it's allowed to be executed when Cache is disabled

Parameters `unit` -- [in] PCNT unit handle created by `pcnt_new_unit()`

Returns

- `ESP_OK`: Start PCNT unit successfully
- `ESP_ERR_INVALID_ARG`: Start PCNT unit failed because of invalid argument
- `ESP_ERR_INVALID_STATE`: Start PCNT unit failed because the unit is not enabled yet
- `ESP_FAIL`: Start PCNT unit failed because of other error

esp_err_t **pcnt_unit_stop** (*pcnt_unit_handle_t* unit)

Stop PCNT from counting.

Note: This function should be called when the unit is in the enable state (i.e. after calling `pcnt_unit_enable()`)

Note: The stop operation won't clear the counter. Also see `pcnt_unit_clear_count()` for how to clear pulse count value.

Note: This function is allowed to run within ISR context

Note: This function will be placed into IRAM if `CONFIG_PCNT_CTRL_FUNC_IN_IRAM`, so that it is allowed to be executed when Cache is disabled

Parameters `unit` -- [in] PCNT unit handle created by `pcnt_new_unit()`

Returns

- `ESP_OK`: Stop PCNT unit successfully
- `ESP_ERR_INVALID_ARG`: Stop PCNT unit failed because of invalid argument
- `ESP_ERR_INVALID_STATE`: Stop PCNT unit failed because the unit is not enabled yet
- `ESP_FAIL`: Stop PCNT unit failed because of other error

esp_err_t **pcnt_unit_clear_count** (*pcnt_unit_handle_t* unit)

Clear PCNT pulse count value to zero.

Note: It's recommended to call this function after adding a watch point by `pcnt_unit_add_watch_point()`, so that the newly added watch point is effective immediately.

Note: This function is allowed to run within ISR context

Note: This function will be placed into IRAM if `CONFIG_PCNT_CTRL_FUNC_IN_IRAM`, so that it's allowed to be executed when Cache is disabled

Parameters **unit** -- [in] PCNT unit handle created by `pcnt_new_unit()`

Returns

- `ESP_OK`: Clear PCNT pulse count successfully
- `ESP_ERR_INVALID_ARG`: Clear PCNT pulse count failed because of invalid argument
- `ESP_FAIL`: Clear PCNT pulse count failed because of other error

esp_err_t **pcnt_unit_get_count** (*pcnt_unit_handle_t* unit, int *value)

Get PCNT count value.

Note: This function is allowed to run within ISR context

Note: This function will be placed into IRAM if `CONFIG_PCNT_CTRL_FUNC_IN_IRAM`, so that it's allowed to be executed when Cache is disabled

Parameters

- **unit** -- [in] PCNT unit handle created by `pcnt_new_unit()`
- **value** -- [out] Returned count value

Returns

- `ESP_OK`: Get PCNT pulse count successfully
- `ESP_ERR_INVALID_ARG`: Get PCNT pulse count failed because of invalid argument
- `ESP_FAIL`: Get PCNT pulse count failed because of other error

esp_err_t **pcnt_unit_register_event_callbacks** (*pcnt_unit_handle_t* unit, const *pcnt_event_callbacks_t* *cbs, void *user_data)

Set event callbacks for PCNT unit.

Note: User registered callbacks are expected to be runnable within ISR context

Note: The first call to this function needs to be before the call to `pcnt_unit_enable`

Note: User can deregister a previously registered callback by calling this function and setting the callback member in the `cbs` structure to NULL.

Parameters

- **unit** -- [in] PCNT unit handle created by `pcnt_new_unit()`
- **cbs** -- [in] Group of callback functions
- **user_data** -- [in] User data, which will be passed to callback functions directly

Returns

- `ESP_OK`: Set event callbacks successfully
- `ESP_ERR_INVALID_ARG`: Set event callbacks failed because of invalid argument
- `ESP_ERR_INVALID_STATE`: Set event callbacks failed because the unit is not in init state
- `ESP_FAIL`: Set event callbacks failed because of other error

esp_err_t **pcnt_unit_add_watch_point** (*pcnt_unit_handle_t* unit, int watch_point)

Add a watch point for PCNT unit, PCNT will generate an event when the counter value reaches the watch point value.

Parameters

- **unit** -- [in] PCNT unit handle created by `pcnt_new_unit()`

- **watch_point** -- [in] Value to be watched

Returns

- ESP_OK: Add watch point successfully
- ESP_ERR_INVALID_ARG: Add watch point failed because of invalid argument (e.g. the value to be watched is out of the limitation set in [pcnt_unit_config_t](#))
- ESP_ERR_INVALID_STATE: Add watch point failed because the same watch point has already been added
- ESP_ERR_NOT_FOUND: Add watch point failed because no more hardware watch point can be configured
- ESP_FAIL: Add watch point failed because of other error

[esp_err_t](#) **pcnt_unit_remove_watch_point** ([pcnt_unit_handle_t](#) unit, int watch_point)

Remove a watch point for PCNT unit.

Parameters

- **unit** -- [in] PCNT unit handle created by [pcnt_new_unit\(\)](#)
- **watch_point** -- [in] Watch point value

Returns

- ESP_OK: Remove watch point successfully
- ESP_ERR_INVALID_ARG: Remove watch point failed because of invalid argument
- ESP_ERR_INVALID_STATE: Remove watch point failed because the watch point was not added by [pcnt_unit_add_watch_point\(\)](#) yet
- ESP_FAIL: Remove watch point failed because of other error

[esp_err_t](#) **pcnt_new_channel** ([pcnt_unit_handle_t](#) unit, const [pcnt_chan_config_t](#) *config, [pcnt_channel_handle_t](#) *ret_chan)

Create PCNT channel for specific unit, each PCNT has several channels associated with it.

Note: This function should be called when the unit is in init state (i.e. before calling [pcnt_unit_enable\(\)](#))

Parameters

- **unit** -- [in] PCNT unit handle created by [pcnt_new_unit\(\)](#)
- **config** -- [in] PCNT channel configuration
- **ret_chan** -- [out] Returned channel handle

Returns

- ESP_OK: Create PCNT channel successfully
- ESP_ERR_INVALID_ARG: Create PCNT channel failed because of invalid argument
- ESP_ERR_NO_MEM: Create PCNT channel failed because of insufficient memory
- ESP_ERR_NOT_FOUND: Create PCNT channel failed because all PCNT channels are used up and no more free one
- ESP_ERR_INVALID_STATE: Create PCNT channel failed because the unit is not in the init state
- ESP_FAIL: Create PCNT channel failed because of other error

[esp_err_t](#) **pcnt_del_channel** ([pcnt_channel_handle_t](#) chan)

Delete the PCNT channel.

Parameters **chan** -- [in] PCNT channel handle created by [pcnt_new_channel\(\)](#)

Returns

- ESP_OK: Delete the PCNT channel successfully
- ESP_ERR_INVALID_ARG: Delete the PCNT channel failed because of invalid argument
- ESP_FAIL: Delete the PCNT channel failed because of other error

[esp_err_t](#) **pcnt_channel_set_edge_action** ([pcnt_channel_handle_t](#) chan, [pcnt_channel_edge_action_t](#) pos_act, [pcnt_channel_edge_action_t](#) neg_act)

Set channel actions when edge signal changes (e.g. falling or rising edge occurred). The edge signal is input

from the `edge_gpio_num` configured in `pcnt_chan_config_t`. We use these actions to control when and how to change the counter value.

Parameters

- **chan** -- [in] PCNT channel handle created by `pcnt_new_channel()`
- **pos_act** -- [in] Action on posedge signal
- **neg_act** -- [in] Action on negedge signal

Returns

- ESP_OK: Set edge action for PCNT channel successfully
- ESP_ERR_INVALID_ARG: Set edge action for PCNT channel failed because of invalid argument
- ESP_FAIL: Set edge action for PCNT channel failed because of other error

`esp_err_t pcnt_channel_set_level_action` (`pcnt_channel_handle_t` chan, `pcnt_channel_level_action_t` high_act, `pcnt_channel_level_action_t` low_act)

Set channel actions when level signal changes (e.g. signal level goes from high to low). The level signal is input from the `level_gpio_num` configured in `pcnt_chan_config_t`. We use these actions to control when and how to change the counting mode.

Parameters

- **chan** -- [in] PCNT channel handle created by `pcnt_new_channel()`
- **high_act** -- [in] Action on high level signal
- **low_act** -- [in] Action on low level signal

Returns

- ESP_OK: Set level action for PCNT channel successfully
- ESP_ERR_INVALID_ARG: Set level action for PCNT channel failed because of invalid argument
- ESP_FAIL: Set level action for PCNT channel failed because of other error

Structures

struct `pcnt_watch_event_data_t`

PCNT watch event data.

Public Members

int `watch_point_value`

Watch point value that triggered the event

`pcnt_unit_zero_cross_mode_t` `zero_cross_mode`

Zero cross mode

struct `pcnt_event_callbacks_t`

Group of supported PCNT callbacks.

Note: The callbacks are all running under ISR environment

Note: When `CONFIG_PCNT_ISR_IRAM_SAFE` is enabled, the callback itself and functions called by it should be placed in IRAM.

Public Members

pcnt_watch_cb_t on_reach

Called when PCNT unit counter reaches any watch point

struct **pcnt_unit_config_t**

PCNT unit configuration.

Public Members

int **low_limit**

Low limitation of the count unit, should be lower than 0

int **high_limit**

High limitation of the count unit, should be higher than 0

int **intr_priority**

PCNT interrupt priority, if set to 0, the driver will try to allocate an interrupt with a relative low priority (1,2,3)

uint32_t **accum_count**

Whether to accumulate the count value when overflows at the high/low limit

struct *pcnt_unit_config_t*::[anonymous] **flags**

Extra flags

struct **pcnt_chan_config_t**

PCNT channel configuration.

Public Members

int **edge_gpio_num**

GPIO number used by the edge signal, input mode with pull up enabled. Set to -1 if unused

int **level_gpio_num**

GPIO number used by the level signal, input mode with pull up enabled. Set to -1 if unused

uint32_t **invert_edge_input**

Invert the input edge signal

uint32_t **invert_level_input**

Invert the input level signal

uint32_t **virt_edge_io_level**

Virtual edge IO level, 0: low, 1: high. Only valid when edge_gpio_num is set to -1

uint32_t **virt_level_io_level**

Virtual level IO level, 0: low, 1: high. Only valid when level_gpio_num is set to -1

uint32_t **io_loop_back**

For debug/test, the signal output from the GPIO will be fed to the input path as well

struct *pcnt_chan_config_t*::[anonymous] **flags**

Channel config flags

struct **pcnt_glitch_filter_config_t**

PCNT glitch filter configuration.

Public Members

uint32_t **max_glitch_ns**

Pulse width smaller than this threshold will be treated as glitch and ignored, in the unit of ns

struct **pcnt_clear_signal_config_t**

PCNT clear signal configuration.

Public Members

int **clear_signal_gpio_num**

GPIO number used by the clear signal, the default active level is high, input mode with pull down enabled

uint32_t **invert_clear_signal**

Invert the clear input signal and set input mode with pull up

uint32_t **io_loop_back**

For debug/test, the signal output from the GPIO will be fed to the input path as well

struct *pcnt_clear_signal_config_t*::[anonymous] **flags**

clear signal config flags

Type Definitions

typedef struct pcnt_unit_t ***pcnt_unit_handle_t**

Type of PCNT unit handle.

typedef struct pcnt_chan_t ***pcnt_channel_handle_t**

Type of PCNT channel handle.

typedef bool (***pcnt_watch_cb_t**)(*pcnt_unit_handle_t* unit, const *pcnt_watch_event_data_t* *edata, void *user_ctx)

PCNT watch event callback prototype.

Note: The callback function is invoked from an ISR context, so it should meet the restrictions of not calling any blocking APIs when implementing the callback. e.g. must use ISR version of FreeRTOS APIs.

Param unit [in] PCNT unit handle

Param edata [in] PCNT event data, fed by the driver

Param `user_ctx` [in] User data, passed from `pcnt_unit_register_event_callbacks()`
Return Whether a high priority task has been woken up by this function

Header File

- [components/hal/include/hal/pcnt_types.h](#)
- This header file can be included with:

```
#include "hal/pcnt_types.h"
```

Enumerations

enum **pcnt_channel_level_action_t**

PCNT channel action on control level.

Values:

enumerator **PCNT_CHANNEL_LEVEL_ACTION_KEEP**

Keep current count mode

enumerator **PCNT_CHANNEL_LEVEL_ACTION_INVERSE**

Invert current count mode (increase -> decrease, decrease -> increase)

enumerator **PCNT_CHANNEL_LEVEL_ACTION_HOLD**

Hold current count value

enum **pcnt_channel_edge_action_t**

PCNT channel action on signal edge.

Values:

enumerator **PCNT_CHANNEL_EDGE_ACTION_HOLD**

Hold current count value

enumerator **PCNT_CHANNEL_EDGE_ACTION_INCREASE**

Increase count value

enumerator **PCNT_CHANNEL_EDGE_ACTION_DECREASE**

Decrease count value

enum **pcnt_unit_zero_cross_mode_t**

PCNT unit zero cross mode.

Values:

enumerator **PCNT_UNIT_ZERO_CROSS_POS_ZERO**

start from positive value, end to zero, i.e. +N->0

enumerator **PCNT_UNIT_ZERO_CROSS_NEG_ZERO**

start from negative value, end to zero, i.e. -N->0

enumerator **PCNT_UNIT_ZERO_CROSS_NEG_POS**

start from negative value, end to positive value, i.e. -N->+M

enumerator **PCNT_UNIT_ZERO_CROSS_POS_NEG**

start from positive value, end to negative value, i.e. +N->-M

2.5.16 Remote Control Transceiver (RMT)

Introduction

The RMT (Remote Control Transceiver) peripheral was designed to act as an infrared transceiver. However, due to the flexibility of its data format, RMT can be extended to a versatile and general-purpose transceiver, transmitting or receiving many other types of signals. From the perspective of network layering, the RMT hardware contains both physical and data link layers. The physical layer defines the communication media and bit signal representation. The data link layer defines the format of an RMT frame. The minimal data unit in the frame is called the **RMT symbol**, which is represented by `rmt_symbol_word_t` in the driver.

ESP32-P4 contains multiple channels in the RMT peripheral¹. Each channel can be independently configured as either transmitter or receiver.

Typically, the RMT peripheral can be used in the following scenarios:

- Transmit or receive infrared signals, with any IR protocols, e.g., NEC
- General-purpose sequence generator
- Transmit signals in a hardware-controlled loop, with a finite or infinite number of times
- Multi-channel simultaneous transmission
- Modulate the carrier to the output signal or demodulate the carrier from the input signal

Layout of RMT Symbols The RMT hardware defines data in its own pattern -- the **RMT symbol**. The diagram below illustrates the bit fields of an RMT symbol. Each symbol consists of two pairs of two values. The first value in the pair is a 15-bit value representing the signal's duration in units of RMT ticks. The second in the pair is a 1-bit value representing the signal's logic level, i.e., high or low.




Fig. 17: Structure of RMT symbols (L - signal level)

RMT Transmitter Overview The data path and control path of an RMT TX channel is illustrated in the figure below:

¹ Different ESP chip series might have different numbers of RMT channels. Please refer to [TRM] for details. The driver does not forbid you from applying for more RMT channels, but it returns an error when there are no hardware resources available. Please always check the return value when doing [Resource Allocation](#).




Fig. 18: RMT Transmitter Overview

The driver encodes the user's data into RMT data format, then the RMT transmitter can generate the waveforms according to the encoding artifacts. It is also possible to modulate a high-frequency carrier signal before being routed to a GPIO pad.

RMT Receiver Overview The data path and control path of an RMT RX channel is illustrated in the figure below:




Fig. 19: RMT Receiver Overview

The RMT receiver can sample incoming signals into RMT data format, and store the data in memory. It is also possible to tell the receiver the basic characteristics of the incoming signal, so that the signal's stop condition can be recognized, and signal glitches and noise can be filtered out. The RMT peripheral also supports demodulating the high-frequency carrier from the base signal.

Functional Overview

The description of the RMT functionality is divided into the following sections:

- [Resource Allocation](#) - covers how to allocate and properly configure RMT channels. It also covers how to recycle channels and other resources when they are no longer used.
- [Carrier Modulation and Demodulation](#) - describes how to modulate and demodulate the carrier signals for TX and RX channels respectively.
- [Register Event Callbacks](#) - covers how to register user-provided event callbacks to receive RMT channel events.
- [Enable and Disable Channel](#) - shows how to enable and disable the RMT channel.
- [Initiate TX Transaction](#) - describes the steps to initiate a transaction for a TX channel.
- [Initiate RX Transaction](#) - describes the steps to initiate a transaction for an RX channel.
- [Multiple Channels Simultaneous Transmission](#) - describes how to collect multiple channels into a sync group so that their transmissions can be started simultaneously.

- *RMT Encoder* - focuses on how to write a customized encoder by combining multiple primitive encoders that are provided by the driver.
- *Power Management* - describes how different clock sources affects power consumption.
- *IRAM Safe* - describes how disabling the cache affects the RMT driver, and tips to mitigate it.
- *Thread Safety* - lists which APIs are guaranteed to be thread-safe by the driver.
- *Kconfig Options* - describes the various Kconfig options supported by the RMT driver.

Resource Allocation Both RMT TX and RX channels are represented by `rmt_channel_handle_t` in the driver. The driver internally manages which channels are available and hands out a free channel on request.

Install RMT TX Channel To install an RMT TX channel, there is a configuration structure that needs to be given in advance `rmt_tx_channel_config_t`. The following list describes each member of the configuration structure.

- `rmt_tx_channel_config_t::gpio_num` sets the GPIO number used by the transmitter.
- `rmt_tx_channel_config_t::clk_src` selects the source clock for the RMT channel. The available clocks are listed in `rmt_clock_source_t`. Note that, the selected clock is also used by other channels, which means the user should ensure this configuration is the same when allocating other channels, regardless of TX or RX. For the effect on the power consumption of different clock sources, please refer to the *Power Management* section.
- `rmt_tx_channel_config_t::resolution_hz` sets the resolution of the internal tick counter. The timing parameter of the RMT signal is calculated based on this **tick**.
- `rmt_tx_channel_config_t::mem_block_symbols` has a slightly different meaning based on if the DMA backend is enabled or not.
 - If the DMA is enabled via `rmt_tx_channel_config_t::with_dma`, then this field controls the size of the internal DMA buffer. To achieve a better throughput and smaller CPU overhead, you can set a larger value, e.g., 1024.
 - If DMA is not used, this field controls the size of the dedicated memory block owned by the channel, which should be at least 48.
- `rmt_tx_channel_config_t::trans_queue_depth` sets the depth of the internal transaction queue, the deeper the queue, the more transactions can be prepared in the backlog.
- `rmt_tx_channel_config_t::invert_out` is used to decide whether to invert the RMT signal before sending it to the GPIO pad.
- `rmt_tx_channel_config_t::with_dma` enables the DMA backend for the channel. Using the DMA allows a significant amount of the channel's workload to be offloaded from the CPU. However, the DMA backend is not available on all ESP chips, please refer to [TRM] before you enable this option. Or you might encounter a `ESP_ERR_NOT_SUPPORTED` error.
- `rmt_tx_channel_config_t::io_loop_back` enables both input and output capabilities on the channel's assigned GPIO. Thus, by binding a TX and RX channel to the same GPIO, loopback can be achieved.
- `rmt_tx_channel_config_t::io_od_mode` configures the channel's assigned GPIO as open-drain. When combined with `rmt_tx_channel_config_t::io_loop_back`, a bi-directional bus (e.g., 1-wire) can be achieved.
- `rmt_tx_channel_config_t::intr_priority` Set the priority of the interrupt. If set to 0, then the driver will use a interrupt with low or medium priority (priority level may be one of 1,2 or 3), otherwise use the priority indicated by `rmt_tx_channel_config_t::intr_priority`. Please use the number form (1,2,3), not the bitmask form ((1<<1),(1<<2),(1<<3)). Please pay attention that once the interrupt priority is set, it cannot be changed until `rmt_del_channel()` is called.

Once the `rmt_tx_channel_config_t` structure is populated with mandatory parameters, users can call `rmt_new_tx_channel()` to allocate and initialize a TX channel. This function returns an RMT channel handle if it runs correctly. Specifically, when there are no more free channels in the RMT resource pool, this function returns `ESP_ERR_NOT_FOUND` error. If some feature (e.g., DMA backend) is not supported by the hardware, it returns `ESP_ERR_NOT_SUPPORTED` error.

```
rmt_channel_handle_t tx_chan = NULL;
rmt_tx_channel_config_t tx_chan_config = {
    .clk_src = RMT_CLK_SRC_DEFAULT,    // select source clock
```

(continues on next page)

(continued from previous page)

```

        .gpio_num = 0,                // GPIO number
        .mem_block_symbols = 64,      // memory block size, 64 * 4 = 256 Bytes
        .resolution_hz = 1 * 1000 * 1000, // 1 MHz tick resolution, i.e., 1 tick = 1 μs
        .trans_queue_depth = 4,       // set the number of transactions that can
        ↪pend in the background
        .flags.invert_out = false,     // do not invert output signal
        .flags.with_dma = false,      // do not need DMA backend
    };
    ESP_ERROR_CHECK(rmt_new_tx_channel(&tx_chan_config, &tx_chan));

```

Install RMT RX Channel To install an RMT RX channel, there is a configuration structure that needs to be given in advance `rmt_rx_channel_config_t`. The following list describes each member of the configuration structure.

- `rmt_rx_channel_config_t::gpio_num` sets the GPIO number used by the receiver.
- `rmt_rx_channel_config_t::clk_src` selects the source clock for the RMT channel. The available clocks are listed in `rmt_clock_source_t`. Note that, the selected clock is also used by other channels, which means the user should ensure this configuration is the same when allocating other channels, regardless of TX or RX. For the effect on the power consumption of different clock sources, please refer to the [Power Management](#) section.
- `rmt_rx_channel_config_t::resolution_hz` sets the resolution of the internal tick counter. The timing parameter of the RMT signal is calculated based on this **tick**.
- `rmt_rx_channel_config_t::mem_block_symbols` has a slightly different meaning based on whether the DMA backend is enabled.
 - If the DMA is enabled via `rmt_rx_channel_config_t::with_dma`, this field controls the maximum size of the DMA buffer.
 - If DMA is not used, this field controls the size of the dedicated memory block owned by the channel, which should be at least 48.
- `rmt_rx_channel_config_t::invert_in` is used to invert the input signals before it is passed to the RMT receiver. The inversion is done by the GPIO matrix instead of by the RMT peripheral.
- `rmt_rx_channel_config_t::with_dma` enables the DMA backend for the channel. Using the DMA allows a significant amount of the channel's workload to be offloaded from the CPU. However, the DMA backend is not available on all ESP chips, please refer to [TRM] before you enable this option. Or you might encounter a `ESP_ERR_NOT_SUPPORTED` error.
- `rmt_rx_channel_config_t::io_loop_back` enables both input and output capabilities on the channel's assigned GPIO. Thus, by binding a TX and RX channel to the same GPIO, loopback can be achieved.
- `rmt_rx_channel_config_t::intr_priority` Set the priority of the interrupt. If set to 0, then the driver will use an interrupt with low or medium priority (priority level may be one of 1,2 or 3), otherwise use the priority indicated by `rmt_rx_channel_config_t::intr_priority`. Please use the number form (1,2,3), not the bitmask form ((1<<1),(1<<2),(1<<3)). Please pay attention that once the interrupt priority is set, it cannot be changed until `rmt_del_channel()` is called.

Once the `rmt_rx_channel_config_t` structure is populated with mandatory parameters, users can call `rmt_new_rx_channel()` to allocate and initialize an RX channel. This function returns an RMT channel handle if it runs correctly. Specifically, when there are no more free channels in the RMT resource pool, this function returns `ESP_ERR_NOT_FOUND` error. If some feature (e.g., DMA backend) is not supported by the hardware, it returns `ESP_ERR_NOT_SUPPORTED` error.

```

rmt_channel_handle_t rx_chan = NULL;
rmt_rx_channel_config_t rx_chan_config = {
    .clk_src = RMT_CLK_SRC_DEFAULT,    // select source clock
    .resolution_hz = 1 * 1000 * 1000, // 1 MHz tick resolution, i.e., 1 tick = 1 μs
    .mem_block_symbols = 64,           // memory block size, 64 * 4 = 256 Bytes
    .gpio_num = 2,                    // GPIO number
    .flags.invert_in = false,         // do not invert input signal
    .flags.with_dma = false,          // do not need DMA backend
};
ESP_ERROR_CHECK(rmt_new_rx_channel(&rx_chan_config, &rx_chan));

```

Note: Due to a software limitation in the GPIO driver, when both TX and RX channels are bound to the same GPIO, ensure the RX Channel is initialized before the TX Channel. If the TX Channel was set up first, then during the RX Channel setup, the previous RMT TX Channel signal will be overridden by the GPIO control signal.

Uninstall RMT Channel If a previously installed RMT channel is no longer needed, it is recommended to recycle the resources by calling `rmt_del_channel()`, which in return allows the underlying software and hardware resources to be reused for other purposes.

Carrier Modulation and Demodulation The RMT transmitter can generate a carrier wave and modulate it onto the message signal. Compared to the message signal, the carrier signal's frequency is significantly higher. In addition, the user can only set the frequency and duty cycle for the carrier signal. The RMT receiver can demodulate the carrier signal from the incoming signal. Note that, carrier modulation and demodulation are not supported on all ESP chips, please refer to [TRM] before configuring the carrier, or you might encounter a `ESP_ERR_NOT_SUPPORTED` error.

Carrier-related configurations lie in `rmt_carrier_config_t`:

- `rmt_carrier_config_t::frequency_hz` sets the carrier frequency, in Hz.
- `rmt_carrier_config_t::duty_cycle` sets the carrier duty cycle.
- `rmt_carrier_config_t::polarity_active_low` sets the carrier polarity, i.e., on which level the carrier is applied.
- `rmt_carrier_config_t::always_on` sets whether to output the carrier even when the data transmission has finished. This configuration is only valid for the TX channel.

Note: For the RX channel, we should not set the carrier frequency exactly to the theoretical value. It is recommended to leave a tolerance for the carrier frequency. For example, in the snippet below, we set the frequency to 25 KHz, instead of the 38 KHz configured on the TX side. The reason is that reflection and refraction occur when a signal travels through the air, leading to distortion on the receiver side.

```
rmt_carrier_config_t tx_carrier_cfg = {
    .duty_cycle = 0.33,           // duty cycle 33%
    .frequency_hz = 38000,       // 38 KHz
    .flags.polarity_active_low = false, // carrier should be modulated to high
    ↪level
};
// modulate carrier to TX channel
ESP_ERROR_CHECK(rmt_apply_carrier(tx_chan, &tx_carrier_cfg));

rmt_carrier_config_t rx_carrier_cfg = {
    .duty_cycle = 0.33,           // duty cycle 33%
    .frequency_hz = 25000,       // 25 KHz carrier, should be smaller than
    ↪the transmitter's carrier frequency
    .flags.polarity_active_low = false, // the carrier is modulated to high level
};
// demodulate carrier from RX channel
ESP_ERROR_CHECK(rmt_apply_carrier(rx_chan, &rx_carrier_cfg));
```

Register Event Callbacks When an event occurs on an RMT channel (e.g., transmission or receiving is completed), the CPU is notified of this event via an interrupt. If you have some function that needs to be called when a particular events occur, you can register a callback for that event to the RMT driver's ISR (Interrupt Service Routine) by calling `rmt_tx_register_event_callbacks()` and `rmt_rx_register_event_callbacks()` for TX and RX channel respectively. Since the registered callback functions are called in the interrupt context, the user should ensure the callback function does not block, e.g., by making sure that only FreeRTOS APIs with the `FromISR` suffix are called from within the function. The callback function has a boolean return value used to indicate whether a higher priority task has been unblocked by the callback.

The TX channel-supported event callbacks are listed in the `rmt_tx_event_callbacks_t`:

- `rmt_tx_event_callbacks_t::on_trans_done` sets a callback function for the "trans-done" event. The function prototype is declared in `rmt_tx_done_callback_t`.

The RX channel-supported event callbacks are listed in the `rmt_rx_event_callbacks_t`:

- `rmt_rx_event_callbacks_t::on_recv_done` sets a callback function for "receive-done" event. The function prototype is declared in `rmt_rx_done_callback_t`.

Users can save their own context in `rmt_tx_register_event_callbacks()` and `rmt_rx_register_event_callbacks()` as well, via the parameter `user_data`. The user data is directly passed to each callback function.

In the callback function, users can fetch the event-specific data that is filled by the driver in the `edata`. Note that the `edata` pointer is only valid during the callback.

The TX-done event data is defined in `rmt_tx_done_event_data_t`:

- `rmt_tx_done_event_data_t::num_symbols` indicates the number of transmitted RMT symbols. This also reflects the size of the encoding artifacts. Please note, this value accounts for the EOF symbol as well, which is appended by the driver to mark the end of one transaction.

The RX-complete event data is defined in `rmt_rx_done_event_data_t`:

- `rmt_rx_done_event_data_t::received_symbols` points to the received RMT symbols. These symbols are saved in the `buffer` parameter of the `rmt_receive()` function. Users should not free this receive buffer before the callback returns.
- `rmt_rx_done_event_data_t::num_symbols` indicates the number of received RMT symbols. This value is not larger than the `buffer_size` parameter of `rmt_receive()` function. If the `buffer_size` is not sufficient to accommodate all the received RMT symbols, the driver only keeps the maximum number of symbols that the buffer can hold, and excess symbols are discarded or ignored.

Enable and Disable Channel `rmt_enable()` must be called in advance before transmitting or receiving RMT symbols. For TX channels, enabling a channel enables a specific interrupt and prepares the hardware to dispatch transactions. For RX channels, enabling a channel enables an interrupt, but the receiver is not started during this time, as the characteristics of the incoming signal have yet to be specified. The receiver is started in `rmt_receive()`.

`rmt_disable()` does the opposite by disabling the interrupt and clearing any pending interrupts. The transmitter and receiver are disabled as well.

```
ESP_ERROR_CHECK(rmt_enable(tx_chan));
ESP_ERROR_CHECK(rmt_enable(rx_chan));
```

Initiate TX Transaction RMT is a special communication peripheral, as it is unable to transmit raw byte streams like SPI and I2C. RMT can only send data in its own format `rmt_symbol_word_t`. However, the hardware does not help to convert the user data into RMT symbols, this can only be done in software by the so-called **RMT Encoder**. The encoder is responsible for encoding user data into RMT symbols and then writing to the RMT memory block or the DMA buffer. For how to create an RMT encoder, please refer to [RMT Encoder](#).

Once you created an encoder, you can initiate a TX transaction by calling `rmt_transmit()`. This function takes several positional parameters like channel handle, encoder handle, and payload buffer. Besides, you also need to provide a transmission-specific configuration in `rmt_transmit_config_t`:

- `rmt_transmit_config_t::loop_count` sets the number of transmission loops. After the transmitter has finished one round of transmission, it can restart the same transmission again if this value is not set to zero. As the loop is controlled by hardware, the RMT channel can be used to generate many periodic sequences with minimal CPU intervention.
 - Setting `rmt_transmit_config_t::loop_count` to `-1` means an infinite loop transmission. In this case, the channel does not stop until `rmt_disable()` is called. The "trans-done" event is not generated as well.

- Setting `rmt_transmit_config_t::loop_count` to a positive number means finite number of iterations. In this case, the "trans-done" event is when the specified number of iterations have completed.

Note: The **loop transmit** feature is not supported on all ESP chips, please refer to [TRM] before you configure this option, or you might encounter `ESP_ERR_NOT_SUPPORTED` error.

- `rmt_transmit_config_t::eot_level` sets the output level when the transmitter finishes working or stops working by calling `rmt_disable()`.
- `rmt_transmit_config_t::queue_nonblocking` sets whether to wait for a free slot in the transaction queue when it is full. If this value is set to `true`, then the function will return with an error code `ESP_ERR_INVALID_STATE` when the queue is full. Otherwise, the function will block until a free slot is available in the queue.

Note: There is a limitation in the transmission size if the `rmt_transmit_config_t::loop_count` is set to non-zero, i.e., to enable the loop feature. The total amount of symbols returned by the encoder should not exceed the capacity of `SOC_RMT_MEM_WORDS_PER_CHANNEL`, or you might see an error message like encoding artifacts can't exceed hw memory block for loop transmission. If you have to start a large transaction by loop, you can try either of the following methods.

- Increase the `rmt_tx_channel_config_t::mem_block_symbols`. This approach does not work if the DMA backend is also enabled.
- Customize an encoder and construct an infinite loop in the encoding function. See also *RMT Encoder*.

Internally, `rmt_transmit()` constructs a transaction descriptor and sends it to a job queue, which is dispatched in the ISR. So it is possible that the transaction is not started yet when `rmt_transmit()` returns. To ensure all pending transactions to complete, the user can use `rmt_tx_wait_all_done()`.

Multiple Channels Simultaneous Transmission In some real-time control applications (e.g., to make two robotic arms move simultaneously), you do not want any time drift between different channels. The RMT driver can help to manage this by creating a so-called **Sync Manager**. The sync manager is represented by `rmt_sync_manager_handle_t` in the driver. The procedure of RMT sync transmission is shown as follows:




Fig. 20: RMT TX Sync

Install RMT Sync Manager To create a sync manager, the user needs to tell which channels are going to be managed in the `rmt_sync_manager_config_t`:

- `rmt_sync_manager_config_t::tx_channel_array` points to the array of TX channels to be managed.
- `rmt_sync_manager_config_t::array_size` sets the number of channels to be managed.

`rmt_new_sync_manager()` can return a manager handle on success. This function could also fail due to various errors such as invalid arguments, etc. Especially, when the sync manager has been installed before, and there are no hardware resources to create another manager, this function reports `ESP_ERR_NOT_FOUND` error. In addition, if the sync manager is not supported by the hardware, it reports a `ESP_ERR_NOT_SUPPORTED` error. Please refer to [TRM] before using the sync manager feature.

Start Transmission Simultaneously For any managed TX channel, it does not start the machine until `rmt_transmit()` has been called on all channels in `rmt_sync_manager_config_t::tx_channel_array`. Before that, the channel is just put in a waiting state. TX channels will usually complete their transactions at different times due to differing transactions, thus resulting in a loss of sync. So before restarting a simultaneous transmission, the user needs to call `rmt_sync_reset()` to synchronize all channels again.

Calling `rmt_del_sync_manager()` can recycle the sync manager and enable the channels to initiate transactions independently afterward.

```
rmt_channel_handle_t tx_channels[2] = {NULL}; // declare two channels
int tx_gpio_number[2] = {0, 2};
// install channels one by one
for (int i = 0; i < 2; i++) {
    rmt_tx_channel_config_t tx_chan_config = {
        .clk_src = RMT_CLK_SRC_DEFAULT, // select source clock
        .gpio_num = tx_gpio_number[i], // GPIO number
        .mem_block_symbols = 64, // memory block size, 64 * 4 = 256 Bytes
        .resolution_hz = 1 * 1000 * 1000, // 1 MHz resolution
        .trans_queue_depth = 1, // set the number of transactions that
        ↪ can pend in the background
    };
    ESP_ERROR_CHECK(rmt_new_tx_channel(&tx_chan_config, &tx_channels[i]));
}
// install sync manager
rmt_sync_manager_handle_t synchro = NULL;
rmt_sync_manager_config_t synchro_config = {
    .tx_channel_array = tx_channels,
    .array_size = sizeof(tx_channels) / sizeof(tx_channels[0]),
};
ESP_ERROR_CHECK(rmt_new_sync_manager(&synchro_config, &synchro));

ESP_ERROR_CHECK(rmt_transmit(tx_channels[0], led_strip_encoders[0], led_data, led_
↪ num * 3, &transmit_config));
// tx_channels[0] does not start transmission until call of `rmt_transmit()` for
↪ tx_channels[1] returns
ESP_ERROR_CHECK(rmt_transmit(tx_channels[1], led_strip_encoders[1], led_data, led_
↪ num * 3, &transmit_config));
```

Initiate RX Transaction As also discussed in the *Enable and Disable Channel*, calling `rmt_enable()` does not prepare an RX to receive RMT symbols. The user needs to specify the basic characteristics of the incoming signals in `rmt_receive_config_t`:

- `rmt_receive_config_t::signal_range_min_ns` specifies the minimal valid pulse duration in either high or low logic levels. A pulse width that is smaller than this value is treated as a glitch, and ignored by the hardware.
- `rmt_receive_config_t::signal_range_max_ns` specifies the maximum valid pulse duration in either high or low logic levels. A pulse width that is bigger than this value is treated as **Stop Signal**, and the receiver generates receive-complete event immediately.

The RMT receiver starts the RX machine after the user calls `rmt_receive()` with the provided configuration above. Note that, this configuration is transaction specific, which means, to start a new round of reception, the user needs to set the `rmt_receive_config_t` again. The receiver saves the incoming signals into its internal memory block or DMA buffer, in the format of `rmt_symbol_word_t`.

Due to the limited size of the memory block, the RMT receiver notifies the driver to copy away the accumulated symbols in a ping-pong way.

The copy destination should be provided in the `buffer` parameter of `rmt_receive()` function. If this buffer overflows due to an insufficient buffer size, the receiver can continue to work, but overflowed symbols are dropped and the following error message is reported: `user buffer too small, received symbols truncated`. Please take care of the lifecycle of the `buffer` parameter, ensuring that the buffer is not recycled before the receiver is finished or stopped.

The receiver is stopped by the driver when it finishes working, i.e., receive a signal whose duration is bigger than `rmt_receive_config_t::signal_range_max_ns`. The user needs to call `rmt_receive()` again to restart the receiver, if necessary. The user can get the received data in the `rmt_rx_event_callbacks_t::on_rcv_done` callback. See also [Register Event Callbacks](#) for more information.

```
static bool example_rmt_rx_done_callback(rmt_channel_handle_t channel, const rmt_
    rx_done_event_data_t *edata, void *user_data)
{
    BaseType_t high_task_wakeup = pdFALSE;
    QueueHandle_t receive_queue = (QueueHandle_t)user_data;
    // send the received RMT symbols to the parser task
    xQueueSendFromISR(receive_queue, edata, &high_task_wakeup);
    // return whether any task is woken up
    return high_task_wakeup == pdTRUE;
}

QueueHandle_t receive_queue = xQueueCreate(1, sizeof(rmt_rx_done_event_data_t));
rmt_rx_event_callbacks_t cbs = {
    .on_rcv_done = example_rmt_rx_done_callback,
};
ESP_ERROR_CHECK(rmt_rx_register_event_callbacks(rx_channel, &cbs, receive_queue));

// the following timing requirement is based on NEC protocol
rmt_receive_config_t receive_config = {
    .signal_range_min_ns = 1250, // the shortest duration for NEC signal is
    // 560 µs, 1250 ns < 560 µs, valid signal is not treated as noise
    .signal_range_max_ns = 12000000, // the longest duration for NEC signal is
    // 9000 µs, 12000000 ns > 9000 µs, the receive does not stop early
};

rmt_symbol_word_t raw_symbols[64]; // 64 symbols should be sufficient for a
    standard NEC frame
// ready to receive
ESP_ERROR_CHECK(rmt_receive(rx_channel, raw_symbols, sizeof(raw_symbols), &receive_
    config));
// wait for the RX-done signal
rmt_rx_done_event_data_t rx_data;
xQueueReceive(receive_queue, &rx_data, portMAX_DELAY);
// parse the received symbols
example_parse_nec_frame(rx_data.received_symbols, rx_data.num_symbols);
```

RMT Encoder An RMT encoder is part of the RMT TX transaction, whose responsibility is to generate and write the correct RMT symbols into hardware memory or DMA buffer at a specific time. There are some special restrictions for an encoding function:

- During a single transaction, the encoding function may be called multiple times. This is necessary because the target RMT memory block cannot hold all the artifacts at once. To overcome this limitation, the driver utilizes

a **ping-pong** approach, where the encoding session is divided into multiple parts. This means that the encoder needs to **keep track of its state** to continue encoding from where it left off in the previous part.

- The encoding function is running in the ISR context. To speed up the encoding session, it is highly recommended to put the encoding function into IRAM. This can also avoid the cache miss during encoding.

To help get started with the RMT driver faster, some commonly used encoders are provided out-of-the-box. They can either work alone or be chained together into a new encoder. See also [Composite Pattern](#) for the principle behind it. The driver has defined the encoder interface in `rmt_encoder_t`, it contains the following functions:

- `rmt_encoder_t::encode` is the fundamental function of an encoder. This is where the encoding session happens.
 - The function might be called multiple times within a single transaction. The encode function should return the state of the current encoding session.
 - The supported states are listed in the `rmt_encode_state_t`. If the result contains `RMT_ENCODING_COMPLETE`, it means the current encoder has finished work.
 - If the result contains `RMT_ENCODING_MEM_FULL`, the program needs to yield from the current session, as there is no space to save more encoding artifacts.
- `rmt_encoder_t::reset` should reset the encoder state back to the initial state (the RMT encoder is stateful).
 - If the RMT transmitter is manually stopped without resetting its corresponding encoder, subsequent encoding session can be erroneous.
 - This function is also called implicitly in `rmt_disable()`.
- `rmt_encoder_t::del` should free the resources allocated by the encoder.

Copy Encoder A copy encoder is created by calling `rmt_new_copy_encoder()`. A copy encoder's main functionality is to copy the RMT symbols from user space into the driver layer. It is usually used to encode `const` data, i.e., data does not change at runtime after initialization such as the leading code in the IR protocol.

A configuration structure `rmt_copy_encoder_config_t` should be provided in advance before calling `rmt_new_copy_encoder()`. Currently, this configuration is reserved for future expansion, and has no specific use or setting items for now.

Bytes Encoder A bytes encoder is created by calling `rmt_new_bytes_encoder()`. The bytes encoder's main functionality is to convert the user space byte stream into RMT symbols dynamically. It is usually used to encode dynamic data, e.g., the address and command fields in the IR protocol.

A configuration structure `rmt_bytes_encoder_config_t` should be provided in advance before calling `rmt_new_bytes_encoder()`:

- `rmt_bytes_encoder_config_t::bit0` and `rmt_bytes_encoder_config_t::bit1` are necessary to specify the encoder how to represent bit zero and bit one in the format of `rmt_symbol_word_t`.
- `rmt_bytes_encoder_config_t::msb_first` sets the bit endianness of each byte. If it is set to true, the encoder encodes the **Most Significant Bit** first. Otherwise, it encodes the **Least Significant Bit** first.

Besides the primitive encoders provided by the driver, the user can implement his own encoder by chaining the existing encoders together. A common encoder chain is shown as follows:

Customize RMT Encoder for NEC Protocol This section demonstrates how to write an NEC encoder. The NEC IR protocol uses pulse distance encoding of the message bits. Each pulse burst is $562.5\ \mu\text{s}$ in length, logical bits are transmitted as follows. It is worth mentioning that the least significant bit of each byte is sent first.

- Logical 0: a $562.5\ \mu\text{s}$ pulse burst followed by a $562.5\ \mu\text{s}$ space, with a total transmit time of $1.125\ \text{ms}$
- Logical 1: a $562.5\ \mu\text{s}$ pulse burst followed by a $1.6875\ \text{ms}$ space, with a total transmit time of $2.25\ \text{ms}$

When a key is pressed on the remote controller, the transmitted message includes the following elements in the specified order:




Fig. 21: RMT Encoder Chain




Fig. 22: IR NEC Frame

- 9 ms leading pulse burst, also called the "AGC pulse"
- 4.5 ms space
- 8-bit address for the receiving device
- 8-bit logical inverse of the address
- 8-bit command
- 8-bit logical inverse of the command
- a final 562.5 μ s pulse burst to signify the end of message transmission

Then you can construct the NEC `rmt_encoder_t::encode` function in the same order, for example:

```
// IR NEC scan code representation
typedef struct {
    uint16_t address;
    uint16_t command;
} ir_nec_scan_code_t;

// construct an encoder by combining primitive encoders
typedef struct {
    rmt_encoder_t base; // the base "class" declares the standard
    encoder interface
    rmt_encoder_t *copy_encoder; // use the copy_encoder to encode the leading
    and ending pulse
    rmt_encoder_t *bytes_encoder; // use the bytes_encoder to encode the address
    and command data
    rmt_symbol_word_t nec_leading_symbol; // NEC leading code with RMT
    representation
    rmt_symbol_word_t nec_ending_symbol; // NEC ending code with RMT
    representation
    int state; // record the current encoding state, i.e., we are in which
    encoding phase
} rmt_ir_nec_encoder_t;
```

(continues on next page)

(continued from previous page)

```

static size_t rmt_encode_ir_nec(rmt_encoder_t *encoder, rmt_channel_handle_t
↪channel, const void *primary_data, size_t data_size, rmt_encode_state_t *ret_
↪state)
{
    rmt_ir_nec_encoder_t *nec_encoder = __containerof(encoder, rmt_ir_nec_encoder_
↪t, base);
    rmt_encode_state_t session_state = RMT_ENCODING_RESET;
    rmt_encode_state_t state = RMT_ENCODING_RESET;
    size_t encoded_symbols = 0;
    ir_nec_scan_code_t *scan_code = (ir_nec_scan_code_t *)primary_data;
    rmt_encoder_handle_t copy_encoder = nec_encoder->copy_encoder;
    rmt_encoder_handle_t bytes_encoder = nec_encoder->bytes_encoder;
    switch (nec_encoder->state) {
        case 0: // send leading code
            encoded_symbols += copy_encoder->encode(copy_encoder, channel, &nec_
↪encoder->nec_leading_symbol,
                                                                    sizeof(rmt_symbol_word_t), &
↪session_state);
            if (session_state & RMT_ENCODING_COMPLETE) {
                nec_encoder->state = 1; // we can only switch to the next state when_
↪the current encoder finished
            }
            if (session_state & RMT_ENCODING_MEM_FULL) {
                state |= RMT_ENCODING_MEM_FULL;
                goto out; // yield if there is no free space to put other encoding_
↪artifacts
            }
            // fall-through
            case 1: // send address
                encoded_symbols += bytes_encoder->encode(bytes_encoder, channel, &scan_
↪code->address, sizeof(uint16_t), &session_state);
                if (session_state & RMT_ENCODING_COMPLETE) {
                    nec_encoder->state = 2; // we can only switch to the next state when_
↪the current encoder finished
                }
                if (session_state & RMT_ENCODING_MEM_FULL) {
                    state |= RMT_ENCODING_MEM_FULL;
                    goto out; // yield if there is no free space to put other encoding_
↪artifacts
                }
                // fall-through
                case 2: // send command
                    encoded_symbols += bytes_encoder->encode(bytes_encoder, channel, &scan_
↪code->command, sizeof(uint16_t), &session_state);
                    if (session_state & RMT_ENCODING_COMPLETE) {
                        nec_encoder->state = 3; // we can only switch to the next state when_
↪the current encoder finished
                    }
                    if (session_state & RMT_ENCODING_MEM_FULL) {
                        state |= RMT_ENCODING_MEM_FULL;
                        goto out; // yield if there is no free space to put other encoding_
↪artifacts
                    }
                    // fall-through
                    case 3: // send ending code
                        encoded_symbols += copy_encoder->encode(copy_encoder, channel, &nec_
↪encoder->nec_ending_symbol,
                                                                    sizeof(rmt_symbol_word_t), &
↪session_state);
                        if (session_state & RMT_ENCODING_COMPLETE) {

```

(continues on next page)

(continued from previous page)

```

        nec_encoder->state = RMT_ENCODING_RESET; // back to the initial_
↳encoding session
        state |= RMT_ENCODING_COMPLETE; // telling the caller the NEC encoding_
↳has finished
    }
    if (session_state & RMT_ENCODING_MEM_FULL) {
        state |= RMT_ENCODING_MEM_FULL;
        goto out; // yield if there is no free space to put other encoding_
↳artifacts
    }
}
out:
    *ret_state = state;
    return encoded_symbols;
}

```

A full sample code can be found in [peripherals/rmt/ir_nec_transceiver](#). In the above snippet, we use a switch-case and several goto statements to implement a [Finite-state machine](#). With this pattern, users can construct much more complex IR protocols.

Power Management When power management is enabled, i.e., [CONFIG_PM_ENABLE](#) is on, the system adjusts the APB frequency before going into Light-sleep, thus potentially changing the resolution of the RMT internal counter.

However, the driver can prevent the system from changing APB frequency by acquiring a power management lock of type [ESP_PM_APB_FREQ_MAX](#). Whenever the user creates an RMT channel that has selected [RMT_CLK_SRC_APB](#) as the clock source, the driver guarantees that the power management lock is acquired after the channel enabled by [rmt_enable\(\)](#). Likewise, the driver releases the lock after [rmt_disable\(\)](#) is called for the same channel. This also reveals that the [rmt_enable\(\)](#) and [rmt_disable\(\)](#) should appear in pairs.

If the channel clock source is selected to others like [RMT_CLK_SRC_XTAL](#), then the driver does not install a power management lock for it, which is more suitable for a low-power application as long as the source clock can still provide sufficient resolution.

IRAM Safe By default, the RMT interrupt is deferred when the Cache is disabled for reasons like writing or erasing the main Flash. Thus the transaction-done interrupt does not get handled in time, which is not acceptable in a real-time application. What is worse, when the RMT transaction relies on **ping-pong** interrupt to successively encode or copy RMT symbols, a delayed interrupt can lead to an unpredictable result.

There is a Kconfig option [CONFIG_RMT_ISR_IRAM_SAFE](#) that has the following features:

1. Enable the interrupt being serviced even when the cache is disabled
2. Place all functions used by the ISR into IRAM²
3. Place the driver object into DRAM in case it is mapped to PSRAM by accident

This Kconfig option allows the interrupt handler to run while the cache is disabled but comes at the cost of increased IRAM consumption.

Another Kconfig option [CONFIG_RMT_RECV_FUNC_IN_IRAM](#) can place [rmt_receive\(\)](#) into the IRAM as well. So that the receive function can be used even when the flash cache is disabled.

Thread Safety The factory function [rmt_new_tx_channel\(\)](#), [rmt_new_rx_channel\(\)](#) and [rmt_new_sync_manager\(\)](#) are guaranteed to be thread-safe by the driver, which means, user can call them from different RTOS tasks without protection by extra locks. Other functions that take the [rmt_channel_handle_t](#) and [rmt_sync_manager_handle_t](#) as the first positional parameter, are not thread-safe. which means the user should avoid calling them from multiple tasks.

² The callback function, e.g., [rmt_tx_event_callbacks_t::on_trans_done](#), and the functions invoked by itself should also reside in IRAM, users need to take care of this by themselves.

The following functions are allowed to use under ISR context as well.

- [rmt_receive\(\)](#)

Kconfig Options

- [CONFIG_RMT_ISR_IRAM_SAFE](#) controls whether the default ISR handler can work when cache is disabled, see also [IRAM Safe](#) for more information.
- [CONFIG_RMT_ENABLE_DEBUG_LOG](#) is used to enable the debug log at the cost of increased firmware binary size.
- [CONFIG_RMT_RECV_FUNC_IN_IRAM](#) controls where to place the RMT receive function (IRAM or Flash), see [IRAM Safe](#) for more information.

Application Examples

- RMT-based RGB LED strip customized encoder: [peripherals/rmt/led_strip](#)
- RMT IR NEC protocol encoding and decoding: [peripherals/rmt/ir_nec_transceiver](#)
- RMT transactions in queue: [peripherals/rmt/musical_buzzer](#)
- RMT-based stepper motor with S-curve algorithm: [peripherals/rmt/stepper_motor](#)
- RMT infinite loop for driving DShot ESC: [peripherals/rmt/dshot_esc](#)
- RMT simulate 1-wire protocol (take DS18B20 as example): [peripherals/rmt/onewire](#)

FAQ

- Why the RMT encoder results in more data than expected?

The RMT encoding takes place in the ISR context. If your RMT encoding session takes a long time (e.g., by logging debug information) or the encoding session is deferred somehow because of interrupt latency, then it is possible the transmitting becomes **faster** than the encoding. As a result, the encoder can not prepare the next data in time, leading to the transmitter sending the previous data again. There is no way to ask the transmitter to stop and wait. You can mitigate the issue by combining the following ways:

- Increase the [rmt_tx_channel_config_t::mem_block_symbols](#), in steps of 48.
- Place the encoding function in the IRAM.
- Enables the [rmt_tx_channel_config_t::with_dma](#) if it is available for your chip.

API Reference

Header File

- [components/driver/rmt/include/driver/rmt_tx.h](#)
- This header file can be included with:

```
#include "driver/rmt_tx.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t **rmt_new_tx_channel** (const *rmt_tx_channel_config_t* *config, *rmt_channel_handle_t* *ret_chan)

Create a RMT TX channel.

Parameters

- **config** -- [in] TX channel configurations
- **ret_chan** -- [out] Returned generic RMT channel handle

Returns

- ESP_OK: Create RMT TX channel successfully
- ESP_ERR_INVALID_ARG: Create RMT TX channel failed because of invalid argument
- ESP_ERR_NO_MEM: Create RMT TX channel failed because out of memory
- ESP_ERR_NOT_FOUND: Create RMT TX channel failed because all RMT channels are used up and no more free one
- ESP_ERR_NOT_SUPPORTED: Create RMT TX channel failed because some feature is not supported by hardware, e.g. DMA feature is not supported by hardware
- ESP_FAIL: Create RMT TX channel failed because of other error

esp_err_t **rmt_transmit** (*rmt_channel_handle_t* tx_channel, *rmt_encoder_handle_t* encoder, const void *payload, size_t payload_bytes, const *rmt_transmit_config_t* *config)

Transmit data by RMT TX channel.

Note: This function constructs a transaction descriptor then pushes to a queue. The transaction will not start immediately if there's another one under processing. Based on the setting of *rmt_transmit_config_t::queue_nonblocking*, if there're too many transactions pending in the queue, this function can block until it has free slot, otherwise just return quickly.

Note: The data to be transmitted will be encoded into RMT symbols by the specific encoder.

Parameters

- **tx_channel** -- [in] RMT TX channel that created by *rmt_new_tx_channel()*
- **encoder** -- [in] RMT encoder that created by various factory APIs like *rmt_new_bytes_encoder()*
- **payload** -- [in] The raw data to be encoded into RMT symbols
- **payload_bytes** -- [in] Size of the payload in bytes
- **config** -- [in] Transmission specific configuration

Returns

- ESP_OK: Transmit data successfully
- ESP_ERR_INVALID_ARG: Transmit data failed because of invalid argument
- ESP_ERR_INVALID_STATE: Transmit data failed because channel is not enabled
- ESP_ERR_NOT_SUPPORTED: Transmit data failed because some feature is not supported by hardware, e.g. unsupported loop count
- ESP_FAIL: Transmit data failed because of other error

esp_err_t **rmt_tx_wait_all_done** (*rmt_channel_handle_t* tx_channel, int timeout_ms)

Wait for all pending TX transactions done.

Note: This function will block forever if the pending transaction can't be finished within a limited time (e.g. an infinite loop transaction). See also *rmt_disable()* for how to terminate a working channel.

Parameters

- **tx_channel** -- [in] RMT TX channel that created by *rmt_new_tx_channel()*
- **timeout_ms** -- [in] Wait timeout, in ms. Specially, -1 means to wait forever.

Returns

- ESP_OK: Flush transactions successfully
- ESP_ERR_INVALID_ARG: Flush transactions failed because of invalid argument

- ESP_ERR_TIMEOUT: Flush transactions failed because of timeout
- ESP_FAIL: Flush transactions failed because of other error

esp_err_t **rmt_tx_register_event_callbacks** (*rmt_channel_handle_t* tx_channel, const *rmt_tx_event_callbacks_t* *cbs, void *user_data)

Set event callbacks for RMT TX channel.

Note: User can deregister a previously registered callback by calling this function and setting the callback member in the `cbs` structure to NULL.

Note: When CONFIG_RMT_ISR_IRAM_SAFE is enabled, the callback itself and functions called by it should be placed in IRAM. The variables used in the function should be in the SRAM as well. The `user_data` should also reside in SRAM.

Parameters

- **tx_channel** -- [in] RMT generic channel that created by `rmt_new_tx_channel()`
- **cbs** -- [in] Group of callback functions
- **user_data** -- [in] User data, which will be passed to callback functions directly

Returns

- ESP_OK: Set event callbacks successfully
- ESP_ERR_INVALID_ARG: Set event callbacks failed because of invalid argument
- ESP_FAIL: Set event callbacks failed because of other error

esp_err_t **rmt_new_sync_manager** (const *rmt_sync_manager_config_t* *config, *rmt_sync_manager_handle_t* *ret_syncro)

Create a synchronization manager for multiple TX channels, so that the managed channel can start transmitting at the same time.

Note: All the channels to be managed should be enabled by `rmt_enable()` before put them into sync manager.

Parameters

- **config** -- [in] Synchronization manager configuration
- **ret_syncro** -- [out] Returned synchronization manager handle

Returns

- ESP_OK: Create sync manager successfully
- ESP_ERR_INVALID_ARG: Create sync manager failed because of invalid argument
- ESP_ERR_NOT_SUPPORTED: Create sync manager failed because it is not supported by hardware
- ESP_ERR_INVALID_STATE: Create sync manager failed because not all channels are enabled
- ESP_ERR_NO_MEM: Create sync manager failed because out of memory
- ESP_ERR_NOT_FOUND: Create sync manager failed because all sync controllers are used up and no more free one
- ESP_FAIL: Create sync manager failed because of other error

esp_err_t **rmt_del_sync_manager** (*rmt_sync_manager_handle_t* syncro)

Delete synchronization manager.

Parameters **syncro** -- [in] Synchronization manager handle returned from `rmt_new_sync_manager()`

Returns

- ESP_OK: Delete the synchronization manager successfully
- ESP_ERR_INVALID_ARG: Delete the synchronization manager failed because of invalid argument
- ESP_FAIL: Delete the synchronization manager failed because of other error

esp_err_t **rmt_sync_reset** (*rmt_sync_manager_handle_t* synchro)

Reset synchronization manager.

Parameters **synchro** -- **[in]** Synchronization manager handle returned from `rmt_new_sync_manager()`

Returns

- ESP_OK: Reset the synchronization manager successfully
- ESP_ERR_INVALID_ARG: Reset the synchronization manager failed because of invalid argument
- ESP_FAIL: Reset the synchronization manager failed because of other error

esp_err_t **rmt_tx_switch_gpio** (*rmt_channel_handle_t* channel, *gpio_num_t* gpio_num, bool invert_out)

Switch GPIO for RMT TX channel.

Parameters

- **channel** -- **[in]** RMT TX channel handle
- **gpio_num** -- **[in]** New GPIO number to be used
- **invert_out** -- **[in]** Whether to invert the output signal

Returns

- ESP_OK: Switch GPIO successfully
- ESP_ERR_INVALID_ARG: Switch GPIO failed because of invalid argument
- ESP_ERR_INVALID_STATE: Switch GPIO failed because channel is not disabled
- ESP_FAIL: Switch GPIO failed because of other error

Structures

struct **rmt_tx_event_callbacks_t**

Group of RMT TX callbacks.

Note: The callbacks are all running under ISR environment

Note: When CONFIG_RMT_ISR_IRAM_SAFE is enabled, the callback itself and functions called by it should be placed in IRAM. The variables used in the function should be in the SRAM as well.

Public Members

rmt_tx_done_callback_t **on_trans_done**

Event callback, invoked when transmission is finished

struct **rmt_tx_channel_config_t**

RMT TX channel specific configuration.

Public Members

gpio_num_t **gpio_num**

GPIO number used by RMT TX channel. Set to -1 if unused

rmt_clock_source_t **clk_src**

Clock source of RMT TX channel, channels in the same group must use the same clock source

uint32_t **resolution_hz**

Channel clock resolution, in Hz

size_t **mem_block_symbols**

Size of memory block, in number of *rmt_symbol_word_t*, must be an even. In the DMA mode, this field controls the DMA buffer size, it can be set to a large value; In the normal mode, this field controls the number of RMT memory block that will be used by the channel.

size_t **trans_queue_depth**

Depth of internal transfer queue, increase this value can support more transfers pending in the background

int **intr_priority**

RMT interrupt priority, if set to 0, the driver will try to allocate an interrupt with a relative low priority (1,2,3)

uint32_t **invert_out**

Whether to invert the RMT channel signal before output to GPIO pad

uint32_t **with_dma**

If set, the driver will allocate an RMT channel with DMA capability

uint32_t **io_loop_back**

The signal output from the GPIO will be fed to the input path as well

uint32_t **io_od_mode**

Configure the GPIO as open-drain mode

uint32_t **init_level**

Set the initial level of the RMT channel signal

struct *rmt_tx_channel_config_t*::[anonymous] **flags**

TX channel config flags

struct **rmt_transmit_config_t**

RMT transmit specific configuration.

Public Members

int **loop_count**

Specify the times of transmission in a loop, -1 means transmitting in an infinite loop

uint32_t **eot_level**

Set the output level for the "End Of Transmission"

uint32_t **queue_nonblocking**

If set, when the transaction queue is full, driver will not block the thread but return directly

struct *rmt_transmit_config_t*::[anonymous] **flags**

Transmit specific config flags

struct **rmt_sync_manager_config_t**

Synchronous manager configuration.

Public Members

const *rmt_channel_handle_t* ***tx_channel_array**

Array of TX channels that are about to be managed by a synchronous controller

size_t **array_size**

Size of the tx_channel_array

Header File

- [components/driver/rmt/include/driver/rmt_rx.h](#)
- This header file can be included with:

```
#include "driver/rmt_rx.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t **rmt_new_rx_channel** (const *rmt_rx_channel_config_t* *config, *rmt_channel_handle_t* *ret_chan)

Create a RMT RX channel.

Parameters

- **config** -- [in] RX channel configurations
- **ret_chan** -- [out] Returned generic RMT channel handle

Returns

- **ESP_OK**: Create RMT RX channel successfully
- **ESP_ERR_INVALID_ARG**: Create RMT RX channel failed because of invalid argument
- **ESP_ERR_NO_MEM**: Create RMT RX channel failed because out of memory
- **ESP_ERR_NOT_FOUND**: Create RMT RX channel failed because all RMT channels are used up and no more free one
- **ESP_ERR_NOT_SUPPORTED**: Create RMT RX channel failed because some feature is not supported by hardware, e.g. DMA feature is not supported by hardware
- **ESP_FAIL**: Create RMT RX channel failed because of other error

esp_err_t **rmt_receive** (*rmt_channel_handle_t* rx_channel, void *buffer, size_t buffer_size, const *rmt_receive_config_t* *config)

Initiate a receive job for RMT RX channel.

Note: This function is non-blocking, it initiates a new receive job and then returns. User should check the received data from the `on_recv_done` callback that registered by `rmt_rx_register_event_callbacks()`.

Note: This function can also be called in ISR context.

Note: If you want this function to work even when the flash cache is disabled, please enable the `CONFIG_RMT_RECV_FUNC_IN_IRAM` option.

Parameters

- **rx_channel** -- [in] RMT RX channel that created by `rmt_new_rx_channel()`
- **buffer** -- [in] The buffer to store the received RMT symbols
- **buffer_size** -- [in] size of the `buffer`, in bytes
- **config** -- [in] Receive specific configurations

Returns

- `ESP_OK`: Initiate receive job successfully
- `ESP_ERR_INVALID_ARG`: Initiate receive job failed because of invalid argument
- `ESP_ERR_INVALID_STATE`: Initiate receive job failed because channel is not enabled
- `ESP_FAIL`: Initiate receive job failed because of other error

esp_err_t **rmt_rx_register_event_callbacks** (*rmt_channel_handle_t* rx_channel, const *rmt_rx_event_callbacks_t* *cbs, void *user_data)

Set callbacks for RMT RX channel.

Note: User can deregister a previously registered callback by calling this function and setting the callback member in the `cbs` structure to `NULL`.

Note: When `CONFIG_RMT_ISR_IRAM_SAFE` is enabled, the callback itself and functions called by it should be placed in IRAM. The variables used in the function should be in the SRAM as well. The `user_data` should also reside in SRAM.

Parameters

- **rx_channel** -- [in] RMT generic channel that created by `rmt_new_rx_channel()`
- **cbs** -- [in] Group of callback functions
- **user_data** -- [in] User data, which will be passed to callback functions directly

Returns

- `ESP_OK`: Set event callbacks successfully
- `ESP_ERR_INVALID_ARG`: Set event callbacks failed because of invalid argument
- `ESP_FAIL`: Set event callbacks failed because of other error

Structures

struct **rmt_rx_event_callbacks_t**

Group of RMT RX callbacks.

Note: The callbacks are all running under ISR environment

Note: When `CONFIG_RMT_ISR_IRAM_SAFE` is enabled, the callback itself and functions called by it should be placed in IRAM. The variables used in the function should be in the SRAM as well.

Public Members

rmt_rx_done_callback_t **on_recv_done**

Event callback, invoked when one RMT channel receiving transaction completes

struct **rmt_rx_channel_config_t**

RMT RX channel specific configuration.

Public Members

gpio_num_t **gpio_num**

GPIO number used by RMT RX channel. Set to -1 if unused

rmt_clock_source_t **clk_src**

Clock source of RMT RX channel, channels in the same group must use the same clock source

uint32_t **resolution_hz**

Channel clock resolution, in Hz

size_t **mem_block_symbols**

Size of memory block, in number of *rmt_symbol_word_t*, must be an even. In the DMA mode, this field controls the DMA buffer size, it can be set to a large value (e.g. 1024); In the normal mode, this field controls the number of RMT memory block that will be used by the channel.

uint32_t **invert_in**

Whether to invert the incoming RMT channel signal

uint32_t **with_dma**

If set, the driver will allocate an RMT channel with DMA capability

uint32_t **io_loop_back**

For debug/test, the signal output from the GPIO will be fed to the input path as well

struct *rmt_rx_channel_config_t*::[anonymous] **flags**

RX channel config flags

int **intr_priority**

RMT interrupt priority, if set to 0, the driver will try to allocate an interrupt with a relative low priority (1,2,3)

struct **rmt_receive_config_t**

RMT receive specific configuration.

Public Members

`uint32_t signal_range_min_ns`

A pulse whose width is smaller than this threshold will be treated as glitch and ignored

`uint32_t signal_range_max_ns`

RMT will stop receiving if one symbol level has kept more than `signal_range_max_ns`

Header File

- [components/driver/rmt/include/driver/rmt_common.h](#)
- This header file can be included with:

```
#include "driver/rmt_common.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

`esp_err_t rmt_del_channel(rmt_channel_handle_t channel)`

Delete an RMT channel.

Parameters `channel` -- **[in]** RMT generic channel that created by `rmt_new_tx_channel()` or `rmt_new_rx_channel()`

Returns

- `ESP_OK`: Delete RMT channel successfully
- `ESP_ERR_INVALID_ARG`: Delete RMT channel failed because of invalid argument
- `ESP_ERR_INVALID_STATE`: Delete RMT channel failed because it is still in working
- `ESP_FAIL`: Delete RMT channel failed because of other error

`esp_err_t rmt_apply_carrier(rmt_channel_handle_t channel, const rmt_carrier_config_t *config)`

Apply modulation feature for TX channel or demodulation feature for RX channel.

Parameters

- **channel** -- **[in]** RMT generic channel that created by `rmt_new_tx_channel()` or `rmt_new_rx_channel()`
- **config** -- **[in]** Carrier configuration. Specially, a `NULL` config means to disable the carrier modulation or demodulation feature

Returns

- `ESP_OK`: Apply carrier configuration successfully
- `ESP_ERR_INVALID_ARG`: Apply carrier configuration failed because of invalid argument
- `ESP_FAIL`: Apply carrier configuration failed because of other error

`esp_err_t rmt_enable(rmt_channel_handle_t channel)`

Enable the RMT channel.

Note: This function will acquire a PM lock that might be installed during channel allocation

Parameters `channel` -- **[in]** RMT generic channel that created by `rmt_new_tx_channel()` or `rmt_new_rx_channel()`

Returns

- ESP_OK: Enable RMT channel successfully
- ESP_ERR_INVALID_ARG: Enable RMT channel failed because of invalid argument
- ESP_ERR_INVALID_STATE: Enable RMT channel failed because it's enabled already
- ESP_FAIL: Enable RMT channel failed because of other error

esp_err_t **rmt_disable**(*rmt_channel_handle_t* channel)

Disable the RMT channel.

Note: This function will release a PM lock that might be installed during channel allocation

Parameters **channel** -- [in] RMT generic channel that created by `rmt_new_tx_channel()` or `rmt_new_rx_channel()`

Returns

- ESP_OK: Disable RMT channel successfully
- ESP_ERR_INVALID_ARG: Disable RMT channel failed because of invalid argument
- ESP_ERR_INVALID_STATE: Disable RMT channel failed because it's not enabled yet
- ESP_FAIL: Disable RMT channel failed because of other error

Structures

struct **rmt_carrier_config_t**

RMT carrier wave configuration (for either modulation or demodulation)

Public Members

uint32_t **frequency_hz**

Carrier wave frequency, in Hz, 0 means disabling the carrier

float **duty_cycle**

Carrier wave duty cycle (0~100%)

uint32_t **polarity_active_low**

Specify the polarity of carrier, by default it's modulated to base signal's high level

uint32_t **always_on**

If set, the carrier can always exist even there's not transfer undergoing

struct *rmt_carrier_config_t*::[anonymous] **flags**

Carrier config flags

Header File

- [components/driver/rmt/include/driver/rmt_encoder.h](#)
- This header file can be included with:

```
#include "driver/rmt_encoder.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

REQUIRES driver

or

PRIV_REQUIRES driver

Functions

esp_err_t **rmt_new_bytes_encoder** (const *rmt_bytes_encoder_config_t* *config, *rmt_encoder_handle_t* *ret_encoder)

Create RMT bytes encoder, which can encode byte stream into RMT symbols.

Parameters

- **config** -- [in] Bytes encoder configuration
- **ret_encoder** -- [out] Returned encoder handle

Returns

- ESP_OK: Create RMT bytes encoder successfully
- ESP_ERR_INVALID_ARG: Create RMT bytes encoder failed because of invalid argument
- ESP_ERR_NO_MEM: Create RMT bytes encoder failed because out of memory
- ESP_FAIL: Create RMT bytes encoder failed because of other error

esp_err_t **rmt_bytes_encoder_update_config** (*rmt_encoder_handle_t* bytes_encoder, const *rmt_bytes_encoder_config_t* *config)

Update the configuration of the bytes encoder.

Note: The configurations of the bytes encoder is also set up by `rmt_new_bytes_encoder()`. This function is used to update the configuration of the bytes encoder at runtime.

Parameters

- **bytes_encoder** -- [in] Bytes encoder handle, created by e.g `rmt_new_bytes_encoder()`
- **config** -- [in] Bytes encoder configuration

Returns

- ESP_OK: Update RMT bytes encoder successfully
- ESP_ERR_INVALID_ARG: Update RMT bytes encoder failed because of invalid argument
- ESP_FAIL: Update RMT bytes encoder failed because of other error

esp_err_t **rmt_new_copy_encoder** (const *rmt_copy_encoder_config_t* *config, *rmt_encoder_handle_t* *ret_encoder)

Create RMT copy encoder, which copies the given RMT symbols into RMT memory.

Parameters

- **config** -- [in] Copy encoder configuration
- **ret_encoder** -- [out] Returned encoder handle

Returns

- ESP_OK: Create RMT copy encoder successfully
- ESP_ERR_INVALID_ARG: Create RMT copy encoder failed because of invalid argument
- ESP_ERR_NO_MEM: Create RMT copy encoder failed because out of memory
- ESP_FAIL: Create RMT copy encoder failed because of other error

esp_err_t **rmt_del_encoder** (*rmt_encoder_handle_t* encoder)

Delete RMT encoder.

Parameters **encoder** -- [in] RMT encoder handle, created by e.g `rmt_new_bytes_encoder()`

Returns

- ESP_OK: Delete RMT encoder successfully
- ESP_ERR_INVALID_ARG: Delete RMT encoder failed because of invalid argument
- ESP_FAIL: Delete RMT encoder failed because of other error

esp_err_t **rmt_encoder_reset** (*rmt_encoder_handle_t* encoder)

Reset RMT encoder.

Parameters **encoder** -- [in] RMT encoder handle, created by e.g. `rmt_new_bytes_encoder()`

Returns

- ESP_OK: Reset RMT encoder successfully
- ESP_ERR_INVALID_ARG: Reset RMT encoder failed because of invalid argument
- ESP_FAIL: Reset RMT encoder failed because of other error

void ***rmt_alloc_encoder_mem** (size_t size)

A helper function to allocate a proper memory for RMT encoder.

Parameters **size** -- Size of memory to be allocated

Returns Pointer to the allocated memory if the allocation is successful, NULL otherwise

Structures

struct **rmt_encoder_t**

Interface of RMT encoder.

Public Members

size_t (***encode**)(*rmt_encoder_t* *encoder, *rmt_channel_handle_t* tx_channel, const void *primary_data, size_t data_size, *rmt_encode_state_t* *ret_state)

Encode the user data into RMT symbols and write into RMT memory.

Note: The encoding function will also be called from an ISR context, thus the function must not call any blocking API.

Note: It's recommended to put this function implementation in the IRAM, to achieve a high performance and less interrupt latency.

Param encoder [in] Encoder handle

Param tx_channel [in] RMT TX channel handle, returned from `rmt_new_tx_channel()`

Param primary_data [in] App data to be encoded into RMT symbols

Param data_size [in] Size of primary_data, in bytes

Param ret_state [out] Returned current encoder's state

Return Number of RMT symbols that the primary data has been encoded into

esp_err_t (***reset**)(*rmt_encoder_t* *encoder)

Reset encoding state.

Param encoder [in] Encoder handle

Return

- ESP_OK: reset encoder successfully
- ESP_FAIL: reset encoder failed

```
esp_err_t (*del)(rmt_encoder_t *encoder)
```

Delete encoder object.

Param encoder [in] Encoder handle

Return

- ESP_OK: delete encoder successfully
- ESP_FAIL: delete encoder failed

```
struct rmt_bytes_encoder_config_t
```

Bytes encoder configuration.

Public Members

```
rmt_symbol_word_t bit0
```

How to represent BIT0 in RMT symbol

```
rmt_symbol_word_t bit1
```

How to represent BIT1 in RMT symbol

```
uint32_t msb_first
```

Whether to encode MSB bit first

```
struct rmt_bytes_encoder_config_t::[anonymous] flags
```

Encoder config flag

```
struct rmt_copy_encoder_config_t
```

Copy encoder configuration.

Enumerations

```
enum rmt_encode_state_t
```

RMT encoding state.

Values:

enumerator **RMT_ENCODING_RESET**

The encoding session is in reset state

enumerator **RMT_ENCODING_COMPLETE**

The encoding session is finished, the caller can continue with subsequent encoding

enumerator **RMT_ENCODING_MEM_FULL**

The encoding artifact memory is full, the caller should return from current encoding session

Header File

- `components/driver/rmt/include/driver/rmt_types.h`
- This header file can be included with:

```
#include "driver/rmt_types.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

REQUIRES driver

or

PRIV_REQUIRES driver

Structures

struct **rmt_tx_done_event_data_t**

Type of RMT TX done event data.

Public Members

size_t **num_symbols**

The number of transmitted RMT symbols, including one EOF symbol, which is appended by the driver to mark the end of a transmission. For a loop transmission, this value only counts for one round.

struct **rmt_rx_done_event_data_t**

Type of RMT RX done event data.

Public Members

rmt_symbol_word_t ***received_symbols**

Point to the received RMT symbols

size_t **num_symbols**

The number of received RMT symbols

Type Definitions

typedef struct rmt_channel_t ***rmt_channel_handle_t**

Type of RMT channel handle.

typedef struct rmt_sync_manager_t ***rmt_sync_manager_handle_t**

Type of RMT synchronization manager handle.

typedef struct *rmt_encoder_t* ***rmt_encoder_handle_t**

Type of RMT encoder handle.

typedef bool (***rmt_tx_done_callback_t**)(*rmt_channel_handle_t* tx_chan, const *rmt_tx_done_event_data_t* *edata, void *user_ctx)

Prototype of RMT event callback.

Param tx_chan [in] RMT channel handle, created from `rmt_new_tx_channel()`

Param edata [in] Point to RMT event data. The lifecycle of this pointer memory is inside this function, user should copy it into static memory if used outside this function.

Param user_ctx [in] User registered context, passed from `rmt_tx_register_event_callbacks()`

Return Whether a high priority task has been waken up by this callback function

```
typedef bool (*rmt_rx_done_callback_t)(rmt_channel_handle_t rx_chan, const  
rmt_rx_done_event_data_t *edata, void *user_ctx)
```

Prototype of RMT event callback.

Param rx_chan [in] RMT channel handle, created from `rmt_new_rx_channel()`

Param edata [in] Point to RMT event data. The lifecycle of this pointer memory is inside this function, user should copy it into static memory if used outside this function.

Param user_ctx [in] User registered context, passed from
`rmt_rx_register_event_callbacks()`

Return Whether a high priority task has been waken up by this function

Header File

- [components/hal/include/hal/rmt_types.h](#)
- This header file can be included with:

```
#include "hal/rmt_types.h"
```

Unions

union **rmt_symbol_word_t**

#include <rmt_types.h> The layout of RMT symbol stored in memory, which is decided by the hardware design.

Public Members

uint16_t **duration0**

Duration of level0

uint16_t **level0**

Level of the first part

uint16_t **duration1**

Duration of level1

uint16_t **level1**

Level of the second part

struct *rmt_symbol_word_t*::[anonymous] [**anonymous**]

uint32_t **val**

Equivalent unsigned value for the RMT symbol

Type Definitions

```
typedef soc_periph_rmt_clk_src_t rmt_clock_source_t
```

RMT group clock source.

Note: User should select the clock source based on the power and resolution requirement

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.5.17 SD Pull-up Requirements

Espressif hardware products are designed for multiple use cases which may require different pull states on pins. For this reason, the pull state of particular pins on certain products needs to be adjusted to provide the pull-ups required in the SD bus.

SD pull-up requirements apply to cases where ESP32-P4 uses the SPI or SDMMC controller to communicate with SD cards. When an SD card is operating in SPI mode or 1-bit SD mode, the CMD and DATA (DAT0 - DAT3) lines of the SD bus must be pulled up by 10 kOhm resistors. SD cards and SDIO devices should also have pull-ups on all above-mentioned lines (regardless of whether these lines are connected to the host) in order to prevent them from entering a wrong state.

This document has the following structure:

- [Overview of compatibility](#) between the default pull states on pins of Espressif's products and the states required by the SD bus
- [Solutions](#) - ideas on how to resolve compatibility issues
- [Related information](#) - other relevant information

Overview of Compatibility

This section provides an overview of compatibility issues that might occur when using SDIO (secure digital input output). Since the SD bus needs to be connected to pull-ups, these issues should be resolved regardless of whether they are related to master (host) or slave (device). Each issue has links to its respective solution. A solution for a host and device may differ.

Systems on a Chip (SoCs) ESP32-P4 SDMMC host controller allows using any of GPIOs for any of SD interface signals. However, it is recommended to avoid using strapping GPIOs, GPIOs with internal weak pull-downs and GPIOs commonly used for other purposes to prevent conflicts:

Systems in Packages (SIP)

Modules

Development Boards

Solutions

No Pull-ups If you use a development board without pull-ups, you can do the following:

- If your host and slave device are on separate boards, replace one of them with a board that has pull-ups. For the list of Espressif's development boards with pull-ups, go to [Development Boards](#).
- Attach external pull-ups by connecting each pin which requires a pull-up to VDD via a 10 kOhm resistor.

Related Information

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.5.18 SDMMC Host Driver

Overview

ESP32-P4's SDMMC host peripheral has two slots. Each slot can be used independently to connect to an SD card, SDIO device, or eMMC chip.

Both slots SDMMC_HOST_SLOT_0 and SDMMC_HOST_SLOT_1 support 1-, 4- and 8-line SD interfaces. The slots are connected to ESP32-P4 GPIOs using the GPIO matrix. This means that any GPIO may be used for each of the SD card signals.

Supported Speed Modes

SDMMC Host driver supports the following speed modes:

- Default Speed (20 MHz): 1-line or 4-line with SD cards, and 1-line, 4-line, or 8-line with 3.3 V eMMC
- High Speed (40 MHz): 1-line or 4-line with SD cards, and 1-line, 4-line, or 8-line with 3.3 V eMMC
- High Speed DDR (40 MHz): 4-line with 3.3 V eMMC

Speed modes not supported at present:

- High Speed DDR mode: 8-line eMMC
- UHS-I 1.8 V modes: 4-line SD cards

Using the SDMMC Host Driver

Of all the functions listed below, only the following ones will be used directly by most applications:

- `sdmmc_host_init()`
- `sdmmc_host_init_slot()`
- `sdmmc_host_deinit()`

Other functions, such as the ones given below, will be called by the SD/MMC protocol layer via function pointers in the `sdmmc_host_t` structure:

- `sdmmc_host_set_bus_width()`
- `sdmmc_host_set_card_clk()`
- `sdmmc_host_do_transaction()`

Configuring Bus Width and Frequency

With the default initializers for `sdmmc_host_t` and `sdmmc_slot_config_t`, i.e., `SDMMC_HOST_DEFAULT` and `SDMMC_SLOT_CONFIG_DEFAULT`, SDMMC Host driver will attempt to use the widest bus supported by the card (4 lines for SD, 8 lines for eMMC) and the frequency of 20 MHz.

In the designs where communication at 40 MHz frequency can be achieved, it is possible to increase the bus frequency by changing the `max_freq_khz` field of `sdmmc_host_t`:

```
sdmmc_host_t host = SDMMC_HOST_DEFAULT();
host.max_freq_khz = SDMMC_FREQ_HIGHSPEED;
```

If you need a specific frequency other than standard speeds, you are free to use any value from within an appropriate range of the SD interface given (SDMMC or SDSPI). However, the real clock frequency shall be calculated by the underlying driver and the value can be different from the one required.

For the SDMMC, `max_freq_khz` works as the upper limit so the final frequency value shall be always lower or equal. For the SDSPI, the nearest fitting frequency is supplied and thus the value can be greater than/equal to/lower than `max_freq_khz`.

To configure the bus width, set the `width` field of `sdmmc_slot_config_t`. For example, to set 1-line mode:

```
sdmmc_slot_config_t slot = SDMMC_SLOT_CONFIG_DEFAULT();
slot.width = 1;
```

Configuring GPIOs

ESP32-P4 SDMMC Host can be configured to use arbitrary GPIOs for each of the signals. Configuration is performed by setting members of `sdmmc_slot_config_t` structure.

For example, to use GPIOs 1-6 for CLK, CMD, and D0-D3 signals respectively:

```
sdmmc_slot_config_t slot = SDMMC_SLOT_CONFIG_DEFAULT();
slot.clk = GPIO_NUM_1;
slot.cmd = GPIO_NUM_2;
slot.d0 = GPIO_NUM_3;
slot.d1 = GPIO_NUM_4;
slot.d2 = GPIO_NUM_5;
slot.d3 = GPIO_NUM_6;
```

It is also possible to configure Card Detect and Write Protect pins. Similar to other signals, set `cd` and `wp` members of the same structure:

```
slot.cd = GPIO_NUM_7;
slot.wp = GPIO_NUM_8;
```

`SDMMC_SLOT_CONFIG_DEFAULT` sets both to `GPIO_NUM_NC`, meaning that by default the signals are not used.

Once `sdmmc_slot_config_t` structure is initialized this way, you can use it when calling `sdmmc_host_init_slot()` or one of the higher level functions (such as `esp_vfs_fat_sdmmc_mount()`).

DDR Mode for eMMC Chips

By default, DDR mode will be used if:

- SDMMC host frequency is set to `SDMMC_FREQ_HIGHSPEED` in `sdmmc_host_t` structure, and
- eMMC chip reports DDR mode support in its CSD register

DDR mode places higher requirements for signal integrity. To disable DDR mode while keeping the `SDMMC_FREQ_HIGHSPEED` frequency, clear the `SDMMC_HOST_FLAG_DDR` bit in `sdmmc_host_t::flags` field of the `sdmmc_host_t`:

```
sdmmc_host_t host = SDMMC_HOST_DEFAULT();
host.max_freq_khz = SDMMC_FREQ_HIGHSPEED;
host.flags &= ~SDMMC_HOST_FLAG_DDR;
```

See also

- [SD/SDIO/MMC Driver](#): introduces the higher-level driver which implements the protocol layer.
- [SD SPI Host Driver](#): introduces a similar driver that uses the SPI controller and is limited to SD protocol's SPI mode.
- [SD Pull-up Requirements](#): introduces pull-up support and compatibilities of modules and development kits.

API Reference

Header File

- [components/driver/sdmmc/include/driver/sdmmc_host.h](#)
- This header file can be included with:

```
#include "driver/sdmmc_host.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t **sdmmc_host_init** (void)

Initialize SDMMC host peripheral.

Note: This function is not thread safe

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE if `sdmmc_host_init` was already called
- ESP_ERR_NO_MEM if memory can not be allocated

esp_err_t **sdmmc_host_init_slot** (int slot, const *sdmmc_slot_config_t* *slot_config)

Initialize given slot of SDMMC peripheral.

On the ESP32, SDMMC peripheral has two slots:

- Slot 0: 8-bit wide, maps to HS1_* signals in PIN MUX
- Slot 1: 4-bit wide, maps to HS2_* signals in PIN MUX

Card detect and write protect signals can be routed to arbitrary GPIOs using GPIO matrix.

Note: This function is not thread safe

Parameters

- **slot** -- slot number (SDMMC_HOST_SLOT_0 or SDMMC_HOST_SLOT_1)
- **slot_config** -- additional configuration for the slot

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE if host has not been initialized using `sdmmc_host_init`

esp_err_t **sdmmc_host_set_bus_width** (int slot, size_t width)

Select bus width to be used for data transfer.

SD/MMC card must be initialized prior to this command, and a command to set bus width has to be sent to the card (e.g. SD_APP_SET_BUS_WIDTH)

Note: This function is not thread safe

Parameters

- **slot** -- slot number (SDMMC_HOST_SLOT_0 or SDMMC_HOST_SLOT_1)
- **width** -- bus width (1, 4, or 8 for slot 0; 1 or 4 for slot 1)

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if slot number or width is not valid

size_t **sdmmc_host_get_slot_width** (int slot)

Get bus width configured in `sdmmc_host_init_slot` to be used for data transfer.

Parameters **slot** -- slot number (SDMMC_HOST_SLOT_0 or SDMMC_HOST_SLOT_1)

Returns configured bus width of the specified slot.

esp_err_t **sdmmc_host_set_card_clk** (int slot, uint32_t freq_khz)

Set card clock frequency.

Currently only integer fractions of 40MHz clock can be used. For High Speed cards, 40MHz can be used. For Default Speed cards, 20MHz can be used.

Note: This function is not thread safe

Parameters

- **slot** -- slot number (SDMMC_HOST_SLOT_0 or SDMMC_HOST_SLOT_1)
- **freq_khz** -- card clock frequency, in kHz

Returns

- ESP_OK on success
- other error codes may be returned in the future

esp_err_t **sdmmc_host_set_bus_ddr_mode** (int slot, bool ddr_enabled)

Enable or disable DDR mode of SD interface.

Parameters

- **slot** -- slot number (SDMMC_HOST_SLOT_0 or SDMMC_HOST_SLOT_1)
- **ddr_enabled** -- enable or disable DDR mode

Returns

- ESP_OK on success
- ESP_ERR_NOT_SUPPORTED if DDR mode is not supported on this slot

esp_err_t **sdmmc_host_set_cclk_always_on** (int slot, bool cclk_always_on)

Enable or disable always-on card clock When `cclk_always_on` is false, the host controller is allowed to shut down the card clock between the commands. When `cclk_always_on` is true, the clock is generated even if no command is in progress.

Parameters

- **slot** -- slot number
- **cclk_always_on** -- enable or disable always-on clock

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if the slot number is invalid

esp_err_t **sdmmc_host_do_transaction** (int slot, *sdmmc_command_t* *cmdinfo)

Send command to the card and get response.

This function returns when command is sent and response is received, or data is transferred, or timeout occurs.

Attention Data buffer passed in `cmdinfo->data` must be in DMA capable memory

Note: This function is not thread safe w.r.t. init/deinit functions, and bus width/clock speed configuration functions. Multiple tasks can call `sdmmc_host_do_transaction` as long as other `sdmmc_host_*` functions are not called.

Parameters

- **slot** -- slot number (SDMMC_HOST_SLOT_0 or SDMMC_HOST_SLOT_1)
- **cmdinfo** -- pointer to structure describing command and data to transfer

Returns

- ESP_OK on success
- ESP_ERR_TIMEOUT if response or data transfer has timed out
- ESP_ERR_INVALID_CRC if response or data transfer CRC check has failed
- ESP_ERR_INVALID_RESPONSE if the card has sent an invalid response
- ESP_ERR_INVALID_SIZE if the size of data transfer is not valid in SD protocol
- ESP_ERR_INVALID_ARG if the data buffer is not in DMA capable memory

esp_err_t **sdmmc_host_io_int_enable** (int slot)

Enable IO interrupts.

This function configures the host to accept SDIO interrupts.

Parameters **slot** -- slot number (SDMMC_HOST_SLOT_0 or SDMMC_HOST_SLOT_1)

Returns returns ESP_OK, other errors possible in the future

esp_err_t **sdmmc_host_io_int_wait** (int slot, TickType_t timeout_ticks)

Block until an SDIO interrupt is received, or timeout occurs.

Parameters

- **slot** -- slot number (SDMMC_HOST_SLOT_0 or SDMMC_HOST_SLOT_1)
- **timeout_ticks** -- number of RTOS ticks to wait for the interrupt

Returns

- ESP_OK on success (interrupt received)
- ESP_ERR_TIMEOUT if the interrupt did not occur within timeout_ticks

esp_err_t **sdmmc_host_deinit** (void)

Disable SDMMC host and release allocated resources.

Note: This function is not thread safe

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE if `sdmmc_host_init` function has not been called

esp_err_t **sdmmc_host_get_real_freq** (int slot, int *real_freq_khz)

Provides a real frequency used for an SD card installed on specific slot of SD/MMC host controller.

This function calculates real working frequency given by current SD/MMC host controller setup for required slot: it reads associated host and card dividers from corresponding SDMMC registers, calculates respective frequency and stores the value into the 'real_freq_khz' parameter

Parameters

- **slot** -- slot number (SDMMC_HOST_SLOT_0 or SDMMC_HOST_SLOT_1)
- **real_freq_khz** -- [out] output parameter for the result frequency (in kHz)

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG on `real_freq_khz == NULL` or invalid slot number used

esp_err_t **sdmmc_host_set_input_delay** (int slot, *sdmmc_delay_phase_t* delay_phase)

set input delay

- This API sets delay when the SDMMC Host samples the signal from the SD Slave.
- This API will check if the given `delay_phase` is valid or not.
- This API will print out the delay time, in picosecond (ps)

Note: ESP32 doesn't support this feature, you will get an `ESP_ERR_NOT_SUPPORTED`

Parameters

- **slot** -- slot number (SDMMC_HOST_SLOT_0 or SDMMC_HOST_SLOT_1)
- **delay_phase** -- delay phase, this API will convert the phase into picoseconds and print it out

Returns

- `ESP_OK`: ON success.
- `ESP_ERR_INVALID_ARG`: Invalid argument.
- `ESP_ERR_NOT_SUPPORTED`: ESP32 doesn't support this feature.

Structures

struct **sdmmc_slot_config_t**

Extra configuration for SDMMC peripheral slot

Public Members

gpio_num_t **clk**

GPIO number of CLK signal.

gpio_num_t **cmd**

GPIO number of CMD signal.

gpio_num_t **d0**

GPIO number of D0 signal.

gpio_num_t **d1**

GPIO number of D1 signal.

gpio_num_t **d2**

GPIO number of D2 signal.

gpio_num_t **d3**

GPIO number of D3 signal.

gpio_num_t **d4**

GPIO number of D4 signal. Ignored in 1- or 4- line mode.

gpio_num_t **d5**

GPIO number of D5 signal. Ignored in 1- or 4- line mode.

`gpio_num_t d6`

GPIO number of D6 signal. Ignored in 1- or 4- line mode.

`gpio_num_t d7`

GPIO number of D7 signal. Ignored in 1- or 4- line mode.

`gpio_num_t gpio_cd`

GPIO number of card detect signal.

`gpio_num_t cd`

GPIO number of card detect signal; shorter name.

`gpio_num_t gpio_wp`

GPIO number of write protect signal.

`gpio_num_t wp`

GPIO number of write protect signal; shorter name.

`uint8_t width`

Bus width used by the slot (might be less than the max width supported)

`uint32_t flags`

Features used by this slot.

Macros

SDMMC_SLOT_FLAG_INTERNAL_PULLUP

Enable internal pullups on enabled pins. The internal pullups are insufficient however, please make sure external pullups are connected on the bus. This is for debug / example purpose only.

SDMMC_SLOT_FLAG_WP_ACTIVE_HIGH

GPIO write protect polarity. 0 means "active low", i.e. card is protected when the GPIO is low; 1 means "active high", i.e. card is protected when GPIO is high.

2.5.19 SD SPI Host Driver

Overview

The SD SPI host driver allows communication with one or more SD cards using the SPI Master driver, which utilizes the SPI host. Each card is accessed through an SD SPI device, represented by an SD SPI handle `sdspi_dev_handle_t`, which returns when the device is attached to an SPI bus by calling `sdspi_host_init_device()`. It is important to note that the SPI bus should be initialized beforehand by `spi_bus_initialize()`.

With the help of *SPI Master Driver* the SD SPI host driver based on, the SPI bus can be shared among SD cards and other SPI devices. The SPI Master driver will handle exclusive access from different tasks.

The SD SPI driver uses software-controlled CS signal.

How to Use

Firstly, use the macro `SDSPI_DEVICE_CONFIG_DEFAULT` to initialize the structure `sdspi_device_config_t`, which is used to initialize an SD SPI device. This macro will also fill in the default pin mappings, which are the same as the pin mappings of the SDMMC host driver. Modify the host and pins of the structure to desired value. Then call `sdspi_host_init_device` to initialize the SD SPI device and attach to its bus.

Then use the `SDSPI_HOST_DEFAULT` macro to initialize the `sdmmc_host_t` structure, which is used to store the state and configurations of the upper layer (SD/SDIO/MMC driver). Modify the `slot` parameter of the structure to the SD SPI device SD SPI handle just returned from `sdspi_host_init_device`. Call `sdmmc_card_init` with the `sdmmc_host_t` to probe and initialize the SD card.

Now you can use SD/SDIO/MMC driver functions to access your card!

Other Details

Only the following driver's API functions are normally used by most applications:

- `sdspi_host_init()`
- `sdspi_host_init_device()`
- `sdspi_host_remove_device()`
- `sdspi_host_deinit()`

Other functions are mostly used by the protocol level SD/SDIO/MMC driver via function pointers in the `sdmmc_host_t` structure. For more details, see [SD/SDIO/MMC Driver](#).

Note: SD over SPI does not support speeds above `SDMMC_FREQ_DEFAULT` due to the limitations of the SPI driver.

Warning: If you want to share the SPI bus among SD card and other SPI devices, there are some restrictions, see [Sharing the SPI Bus Among SD Cards and Other SPI Devices](#).

Related Docs

Sharing the SPI Bus Among SD Cards and Other SPI Devices

The SD card has an SPI mode, enabling it to function as an SPI device, but there are some restrictions that we need to pay attention to.

Pin Loading of Other Devices When adding more devices onto the same bus, the overall pin loading increases. The loading consists of AC loading (pin capacitor) and DC loading (pull-ups).

AC Loading SD cards, designed for high-speed communications, have small pin capacitors (AC loading) to work until 50 MHz. However, the other attached devices will increase the pin's AC loading.

Heavy AC loading of a pin may prevent the pin from being toggled quickly. By using an oscilloscope, you will see the edges of the pin become smoother, i.e., the gradient of the edge is smaller. The setup timing requirements of an SD card may be violated when the card is connected to a bus with a high AC load. Even worse, high AC loads may cause the SD card and other SPI devices to fail to properly resolve clock signals from the host, affecting communication stability.

This issue may be more obvious if other attached devices are not designed to work at the same frequency as the SD card, because they may have larger pin capacitors. The larger the pin capacity, the greater the pin response time, the smaller the max frequency the SD bus can work.

To see if your pin AC loading is too heavy, you can try the following tests:

Terminology:

- **launch edge**: at which clock edge the data starts to toggle;
- **latch edge**: at which clock edge the data is supposed to be sampled by the receiver. For SD card, it is the rising edge.

1. Use an oscilloscope to see the clock and compare the data line to the clock.
 - If you see the clock is not fast enough, e.g., the rising/falling edge is longer than 1/4 of the clock cycle, it means the clock is skewed too much.
 - If you see the data line unstable before the latch edge of the clock, it means the load of the data line is too large.

You may also observe the corresponding phenomenon that data delayed largely from the launching edge of the clock with logic analyzers. But it is not as obvious as with an oscilloscope.

2. Try to use a slower clock frequency.

If the lower frequency can work while the higher frequency cannot, it is an indication that the AC loading on the pins is too large.

If the AC loading of the pins is too large, you can either use other faster devices with lower pin load or slow down the clock speed.

DC Loading The pull-ups required by SD cards are usually around 10 kOhm to 50 kOhm, which may be too strong for some other SPI devices.

Check the specification of your device about its DC output current, it should be larger than 700 μ A, otherwise, the device output may not be read correctly.

Initialization Sequence

Note: If you see any problem in the following steps, please make sure the timing is correct first. You can try to slow down the clock speed, such as setting `SDMMC_FREQ_PROBING` to 400 kHz for SD card, to avoid the influence of pin AC loading, as discussed in the previous section.

When using an SD card with other SPI devices on the same SPI bus, due to the restrictions of the SD card startup flow, the following initialization sequence should be followed. Refer to [storage/sd_card](#) for further details.

1. Initialize the SPI bus properly by `spi_bus_initialize()`.
2. Tie the CS lines of all other devices than the SD card to idle state (by default it's high). This is to avoid conflicts with the SD card in the following step.

You can do this by either:

1. Attach devices to the SPI bus by calling `spi_bus_add_device()`. This function will by default initialize the GPIO that is used as CS to the idle level: high.
2. Initialize GPIO on the CS pin that needs to be tied up before actually adding a new device.
3. Rely on the internal/external pull-up (**not recommended**) to pull up all the CS pins when the GPIOs of ESP are not initialized yet. You need to check carefully the pull-up is strong enough and there are no other pull-downs that will influence the pull-up. For example, internal pull-down should be enabled.
3. Mount the card to the filesystem by calling `esp_vfs_fat_sdspi_mount()`.

This step will put the SD card into the SPI mode, which **should** be done before all other SPI communications on the same bus. Otherwise, the card will stay in the SD mode, in which mode it may randomly respond to any SPI communications on the bus, even when its CS line is not addressed.

If you want to test this behavior, please also note that, once the card is put into SPI mode, it will not return to SD mode before the next power cycle, i.e., powered down and powered up again.
4. Now you can talk to other SPI devices freely!

API Reference

Header File

- [components/driver/spi/include/driver/sdspi_host.h](#)

- This header file can be included with:

```
#include "driver/sdspi_host.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t **sdspi_host_init** (void)

Initialize SD SPI driver.

Note: This function is not thread safe

Returns

- `ESP_OK` on success
- other error codes may be returned in future versions

esp_err_t **sdspi_host_init_device** (const *sdspi_device_config_t* *dev_config, *sdspi_dev_handle_t* *out_handle)

Attach and initialize an SD SPI device on the specific SPI bus.

Note: This function is not thread safe

Note: Initialize the SPI bus by `spi_bus_initialize()` before calling this function.

Note: The SDIO over `sdspi` needs an extra interrupt line. Call `gpio_install_isr_service()` before this function.

Parameters

- **dev_config** -- pointer to device configuration structure
- **out_handle** -- Output of the handle to the `sdspi` device.

Returns

- `ESP_OK` on success
- `ESP_ERR_INVALID_ARG` if `sdspi_host_init_device` has invalid arguments
- `ESP_ERR_NO_MEM` if memory can not be allocated
- other errors from the underlying `spi_master` and `gpio` drivers

esp_err_t **sdspi_host_remove_device** (*sdspi_dev_handle_t* handle)

Remove an SD SPI device.

Parameters **handle** -- Handle of the SD SPI device

Returns Always `ESP_OK`

esp_err_t **sdspi_host_do_transaction** (*sdspi_dev_handle_t* handle, *sdmmc_command_t* *cmdinfo)

Send command to the card and get response.

This function returns when command is sent and response is received, or data is transferred, or timeout occurs.

Note: This function is not thread safe w.r.t. init/deinit functions, and bus width/clock speed configuration functions. Multiple tasks can call `sdspi_host_do_transaction` as long as other `sdspi_host_*` functions are not called.

Parameters

- **handle** -- Handle of the sdspi device
- **cmdinfo** -- pointer to structure describing command and data to transfer

Returns

- ESP_OK on success
- ESP_ERR_TIMEOUT if response or data transfer has timed out
- ESP_ERR_INVALID_CRC if response or data transfer CRC check has failed
- ESP_ERR_INVALID_RESPONSE if the card has sent an invalid response

esp_err_t **sdspi_host_set_card_clk** (*sdspi_dev_handle_t* host, uint32_t freq_khz)

Set card clock frequency.

Currently only integer fractions of 40MHz clock can be used. For High Speed cards, 40MHz can be used. For Default Speed cards, 20MHz can be used.

Note: This function is not thread safe

Parameters

- **host** -- Handle of the sdspi device
- **freq_khz** -- card clock frequency, in kHz

Returns

- ESP_OK on success
- other error codes may be returned in the future

esp_err_t **sdspi_host_get_real_freq** (*sdspi_dev_handle_t* handle, int *real_freq_khz)

Calculate working frequency for specific device.

Parameters

- **handle** -- SDSPI device handle
- **real_freq_khz** -- [out] output parameter to hold the calculated frequency (in kHz)

Returns

- ESP_ERR_INVALID_ARG : handle is NULL or invalid or real_freq_khz parameter is NULL
- ESP_OK : Success

esp_err_t **sdspi_host_deinit** (void)

Release resources allocated using `sdspi_host_init`.

Note: This function is not thread safe

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE if `sdspi_host_init` function has not been called

esp_err_t **sdspi_host_io_int_enable** (*sdspi_dev_handle_t* handle)

Enable SDIO interrupt.

Parameters **handle** -- Handle of the sdspi device

Returns

- ESP_OK on success

esp_err_t **sdspi_host_io_int_wait** (*sdspi_dev_handle_t* handle, TickType_t timeout_ticks)

Wait for SDIO interrupt until timeout.

Parameters

- **handle** -- Handle of the sdspi device
- **timeout_ticks** -- Ticks to wait before timeout.

Returns

- ESP_OK on success

Structures

struct **sdspi_device_config_t**

Extra configuration for SD SPI device.

Public Members

spi_host_device_t **host_id**

SPI host to use, SPIx_HOST (see spi_types.h).

gpio_num_t **gpio_cs**

GPIO number of CS signal.

gpio_num_t **gpio_cd**

GPIO number of card detect signal.

gpio_num_t **gpio_wp**

GPIO number of write protect signal.

gpio_num_t **gpio_int**

GPIO number of interrupt line (input) for SDIO card.

bool **gpio_wp_polarity**

GPIO write protect polarity 0 means "active low", i.e. card is protected when the GPIO is low; 1 means "active high", i.e. card is protected when GPIO is high.

Macros

SDSPI_DEFAULT_HOST

SDSPI_DEFAULT_DMA

SDSPI_HOST_DEFAULT ()

Default *sdmmc_host_t* structure initializer for SD over SPI driver.

Uses SPI mode and max frequency set to 20MHz

'slot' should be set to an sdspi device initialized by *sdspi_host_init_device* ().

SDSPI_SLOT_NO_CS

indicates that card select line is not used

SDSPI_SLOT_NO_CD

indicates that card detect line is not used

SDSPI_SLOT_NO_WP

indicates that write protect line is not used

SDSPI_SLOT_NO_INT

indicates that interrupt line is not used

SDSPI_IO_ACTIVE_LOW**SDSPI_DEVICE_CONFIG_DEFAULT()**

Macro defining default configuration of SD SPI device.

Type Definitions

typedef int **sdspi_dev_handle_t**

Handle representing an SD SPI device.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.5.20 SPI Flash API

Overview

The `spi_flash` component contains API functions related to reading, writing, erasing, and memory mapping for data in the external flash.

For higher-level API functions which work with partitions defined in the *partition table*, see *Partitions API*

Note: `esp_partition_*` APIs are recommended to be used instead of the lower level `esp_flash_*` API functions when accessing the main SPI flash chip, since they conduct bounds checking and are guaranteed to calculate correct offsets in flash based on the information in the partition table. `esp_flash_*` functions can still be used directly when accessing an external (secondary) SPI flash chip.

Different from the API before ESP-IDF v4.0, the functionality of `esp_flash_*` APIs is not limited to the "main" SPI flash chip (the same SPI flash chip from which program runs). With different chip pointers, you can access external flash chips connected to not only SPI0/1 but also other SPI buses like SPI2.

Note: Instead of going through the cache connected to the SPI0 peripheral, most `esp_flash_*` APIs go through other SPI peripherals like SPI1, SPI2, etc. This makes them able to access not only the main flash, but also external (secondary) flash.

However, due to the limitations of the cache, operations through the cache are limited to the main flash. The address range limitation for these operations is also on the cache side. The cache is not able to access external flash chips or address range above its capabilities. These cache operations include: mmap, encrypted read/write, executing code or access to variables in the flash.

Note: Flash APIs after ESP-IDF v4.0 are no longer **atomic**. If a write operation occurs during another on-going read operation, and the flash addresses of both operations overlap, the data returned from the read operation may contain both old data and new data (that was updated written by the write operation).

Note: Encrypted flash operations are only supported with the main flash chip (and not with other flash chips, that is on SPI1 with different CS, or on other SPI buses). Reading through cache is only supported on the main flash, which is determined by the HW.

Support for Features of Flash Chips

Quad/Dual Mode Chips Features of different flashes are implemented in different ways and thus need special support. The fast/slow read and Dual mode (DOUT/DIO) of almost all flashes with 24-bit address are supported, because they do not need any vendor-specific commands.

Quad mode (QIO/QOUT) is supported on the following chip types:

1. ISSI
2. GD
3. MXIC
4. FM
5. Winbond
6. XMC
7. BOYA

Note: Only when one flash series listed above is supported by ESP32-P4, this flash series is supported by the chip driver by default. You can use `Component config > SPI Flash driver > Auto-detect flash chips` in `menuconfig` to enable/disable a flash series.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Optional Features

Optional Features for Flash Some features are not supported on all ESP chips and Flash chips. You can check the list below for more information.

- *Auto Suspend & Resume*
- *Flash unique ID*
- *High performance mode*
- *OPI flash support*
- *32-bit Address Flash Chips*

Note: When Flash optional features listed in this page are used, aside from the capability of ESP chips, and ESP-IDF version you are using, you will also need to make sure these features are supported by flash chips used.

- If you are using an official Espressif modules/SiP. Some of the modules/SiPs always support the feature, in this case you can see these features listed in the datasheet. Otherwise please contact [Espressif's business team](#) to know if we can supply such products for you.

- If you are making your own modules with your own bought flash chips, and you need features listed above. Please contact your vendor if they support the those features, and make sure that the chips can be supplied continuously.
-

Attention: This document only shows that ESP-IDF code has supported the features of those flash chips. It is not a list of stable flash chips certified by Espressif. If you build your own hardware from flash chips with your own bought flash chips (even with flash listed in this page), you need to validate the reliability of flash chips yourself.

Auto Suspend & Resume This feature is only supported on ESP32-S3, ESP32-C2, ESP32-C3, ESP32-C6, ESP32-H2 for now.

The support for ESP32-P4 may be added in the future.

Flash Unique ID This feature is supported on all Espressif chips.

Unique ID is not flash id, which means flash has 64-Bit unique ID for each device. The instruction to read the unique ID (4Bh) accesses a factory-set read-only 64-bit number that is unique to each flash device. This ID number helps you to recognize each single device. Not all flash vendors support this feature. If you try to read the unique ID on a chip which does not have this feature, the behavior is not determined. The support list is as follows.

List of Flash chips that support this feature:

1. ISSI
2. GD
3. TH
4. FM
5. Winbond
6. XMC
7. BOYA

High Performance Mode This featurer is only supported on ESP32-S3 for now.

The support for ESP32-S2, ESP32-C3, ESP32-C6, ESP32-H2, ESP32-P4 may be added in the future.

OPI flash Support This feature is only suppoetrd on ESP32-S3 for now.

OPI flash means that the flash chip supports octal peripheral interface, which has octal I/O pins. Different octal flash has different configurations and different commands. Hence, it is necessary to carefully check the support list.

32-bit Address Flash Chips This feature is supported on all Espressif chips (with various restrictions to application).

Most NOR flash chips used by Espressif chips use 24-bits address, which can cover 16 MBytes memory. However, for larger memory (usually equal to or larger than 16 MBytes), flash uses a 32-bits address to address larger memory. Regretfully, 32-bits address chips have vendor-specific commands, so we need to support the chips one by one.

List of Flash chips that support this feature:

1. W25Q256
2. GD25Q256

Important: Over 16 MBytes space on flash mentioned above can be only used for data saving, like file system. If your data/instructions over 16 MBytes spaces need to be mapped to MMU (so as to be accessed by the CPU), please enable the config `IDF_EXPERIMENTAL_FEATURES` and `BOOTLOADER_CACHE_32BIT_ADDR_FLASH` and read the limitations following:

1. This feature is valid only for 4-line flash. Octal flash supports 32-bit-addr by default
 2. This feature needs the MMU on ESP chip to be able to map to ≥ 16 MB physical address on the Flash. (Only ESP32S3 supports this up to now)
 3. This option is experimental, which means it can not use on all flash chips stable, for more information, please contact Espressif Business support.
-

There are some features that are not supported by all flash chips, or not supported by all Espressif chips. These features include:

- 32-bit address flash - usually means that the flash has higher capacity (equal to or larger than 16 MB) that needs longer addresses.
- Flash unique ID - means that flash supports its unique 64-bit ID.

If you want to use these features, please ensure both ESP32-P4 and ALL flash chips in your product support these features. For more details, refer to [Optional Features for Flash](#).

You may also customise your own flash chip driver. See [Overriding Default Chip Drivers](#) for more details.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Warning: Customizing SPI Flash Chip Drivers is considered an "expert" feature. Users should only do so at their own risk. (See the notes below)

Overriding Default Chip Drivers During the SPI Flash driver's initialization (i.e., `esp_flash_init()`), there is a chip detection step during which the driver iterates through a Default Chip Driver List and determine which chip driver can properly support the currently connected flash chip. The Default Chip Drivers are provided by the ESP-IDF, thus are updated in together with each ESP-IDF version. However ESP-IDF also allows users to customize their own chip drivers.

Users should note the following when customizing chip drivers:

1. You may need to rely on some non-public ESP-IDF functions, which have slight possibility to change between ESP-IDF versions. On the one hand, these changes may be useful bug fixes for your driver, on the other hand, they may also be breaking changes (i.e., breaks your code).
2. Some ESP-IDF bug fixes to other chip drivers are not automatically applied to your own custom chip drivers.
3. If the protection of flash is not handled properly, there may be some random reliability issues.
4. If you update to a newer ESP-IDF version that has support for more chips, you will have to manually add those new chip drivers into your custom chip driver list. Otherwise the driver will only search for the drivers in custom list you provided.

Steps For Creating Custom Chip Drivers and Overriding the ESP-IDF Default Driver List

1. Enable the `CONFIG_SPI_FLASH_OVERRIDE_CHIP_DRIVER_LIST` config option. This prevents compilation and linking of the Default Chip Driver List (`default_registered_chips`) provided by ESP-IDF. Instead, the linker searches for the structure of the same name (`default_registered_chips`) that must be provided by the user.
2. Add a new component in your project, e.g., `custom_chip_driver`.
3. Copy the necessary chip driver files from the `spi_flash` component in ESP-IDF. This may include:
 - `spi_flash_chip_drivers.c` (to provide the `default_registered_chips` structure)
 - Any of the `spi_flash_chip_*.c` files that matches your own flash model best
 - `CMakeLists.txt` and `linker.lf` files

Modify the files above properly. Including:

- Change the `default_registered_chips` variable to non-static and remove the `#ifdef` logic around it.

- Update `linker.lf` file to rename the fragment header and the library name to match the new component.
- If reusing other drivers, some header names need prefixing with `spi_flash/` when included from outside `spi_flash` component.

Note:

- When writing your own flash chip driver, you can set your flash chip capabilities through `spi_flash_chip_***(vendor)_get_caps` and points the function pointer `get_chip_caps` for protection to the `spi_flash_chip_***_get_caps` function. The steps are as follows.
 1. Please check whether your flash chip have the capabilities listed in `spi_flash_caps_t` by checking the flash datasheet.
 2. Write a function named `spi_flash_chip_***(vendor)_get_caps`. Take the example below as a reference. (if the flash support suspend and read unique id).
 3. Points the pointer `get_chip_caps` (in `spi_flash_chip_t`) to the function mentioned above.

```
spi_flash_caps_t spi_flash_chip_***(vendor)_get_caps(esp_flash_t *chip)
{
    spi_flash_caps_t caps_flags = 0;
    // 32-bit-address flash is not supported
    flash_suspend is supported
    caps_flags |= SPI_FLASHS_CHIP_CAP_SUSPEND;
    // flash read unique id.
    caps_flags |= SPI_FLASH_CHIP_CAP_UNIQUE_ID;
    return caps_flags;
}
```

```
const spi_flash_chip_t esp_flash_chip_eon = {
    // Other function pointers
    .get_chip_caps = spi_flash_chip_eon_get_caps,
};
```

- You also can see how to implement this in the example [storage/custom_flash_driver](#).

4. Write a new `CMakeLists.txt` file for the `custom_chip_driver` component, including an additional line to add a linker dependency from `spi_flash` to `custom_chip_driver`:

```
idf_component_register(SRCS "spi_flash_chip_drivers.c"
                        "spi_flash_chip_mychip.c" # modify as needed
                        REQUIRES hal
                        PRIV_REQUIRES spi_flash
                        LDFRAGMENTS linker.lf)
idf_component_add_link_dependency(FROM spi_flash)
```

- An example of this component `CMakeLists.txt` can be found in [storage/custom_flash_driver/components/custom_chip_driver/CMakeLists.txt](#)
5. The `linker.lf` is used to put every chip driver that you are going to use whilst cache is disabled into internal RAM. See [Linker Script Generation](#) for more details. Make sure this file covers all the source files that you add.
 6. Build your project, and you will see the new flash driver is used.

Example See also [storage/custom_flash_driver](#).

Initializing a Flash Device

To use the `esp_flash_*` APIs, you need to initialise a flash chip on a certain SPI bus, as shown below:

1. Call `spi_bus_initialize()` to properly initialize an SPI bus. This function initializes the resources (I/O, DMA, interrupts) shared among devices attached to this bus.

2. Call `spi_bus_add_flash_device()` to attach the flash device to the bus. This function allocates memory and fills the members for the `esp_flash_t` structure. The CS I/O is also initialized here.
3. Call `esp_flash_init()` to actually communicate with the chip. This also detects the chip type, and influence the following operations.

Note: Multiple flash chips can be attached to the same bus now.

SPI Flash Access API

This is the set of API functions for working with data in flash:

- `esp_flash_read()` reads data from flash to RAM
- `esp_flash_write()` writes data from RAM to flash
- `esp_flash_erase_region()` erases specific region of flash
- `esp_flash_erase_chip()` erases the whole flash
- `esp_flash_get_chip_size()` returns flash chip size, in bytes, as configured in menuconfig

Generally, try to avoid using the raw SPI flash functions to the "main" SPI flash chip in favour of *partition-specific functions*.

SPI Flash Size

The SPI flash size is configured by writing a field in the software bootloader image header, flashed at offset 0x1000.

By default, the SPI flash size is detected by `esptool.py` when this bootloader is written to flash, and the header is updated with the correct size. Alternatively, it is possible to generate a fixed flash size by setting [CONFIG_ESPTOOLPY_FLASHSIZE](#) in the project configuration.

If it is necessary to override the configured flash size at runtime, it is possible to set the `chip_size` member of the `g_rom_flashchip` structure. This size is used by `esp_flash_*` functions (in both software & ROM) to check the bounds.

Concurrency Constraints for Flash on SPI1

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Concurrency Constraints for Flash on SPI1

The SPI0/1 bus is shared between the instruction & data cache (for firmware execution) and the SPI1 peripheral (controlled by the drivers including this SPI Flash driver). Hence, operations to SPI1 will cause significant influence to the whole system. This kind of operations include calling SPI Flash API or other drivers on SPI1 bus, any operations like read/write/erase or other user defined SPI operations, regardless to the main flash or other SPI slave devices.

On ESP32-P4, these caches must be disabled while reading/writing/erasing.

When the Caches Are Disabled Under this condition, all CPUs should always execute code and access data from internal RAM. The APIs documented in this file will disable the caches automatically and transparently.

The way that these APIs disable the caches suspends all the other tasks. Besides, all non-IRAM-safe interrupts will be disabled. The other core will be polling in a busy loop. These will be restored until the Flash operation completes.

See also [OS Functions](#) and [SPI Bus Lock](#).

There are no such constraints and impacts for flash chips on other SPI buses than SPI0/1.

For differences between internal RAM (e.g., IRAM, DRAM) and flash cache, please refer to the [application memory layout](#) documentation.

IRAM-Safe Interrupt Handlers For interrupt handlers which need to execute when the cache is disabled (e.g., for low latency operations), set the `ESP_INTR_FLAG_IRAM` flag when the [interrupt handler is registered](#).

You must ensure that all data and functions accessed by these interrupt handlers, including the ones that handlers call, are located in IRAM or DRAM. See [How to Place Code in IRAM](#).

If a function or symbol is not correctly put into IRAM/DRAM, and the interrupt handler reads from the flash cache during a flash operation, it will cause a crash due to Illegal Instruction exception (for code which should be in IRAM) or garbage data to be read (for constant data which should be in DRAM).

Note: When working with strings in ISRs, it is not advised to use `printf` and other output functions. For debugging purposes, use `ESP_DRAM_LOGE()` and similar macros when logging from ISRs. Make sure that both TAG and format string are placed into DRAM in that case.

Non-IRAM-Safe Interrupt Handlers If the `ESP_INTR_FLAG_IRAM` flag is not set when registering, the interrupt handler will not get executed when the caches are disabled. Once the caches are restored, the non-IRAM-safe interrupts will be re-enabled. After this moment, the interrupt handler will run normally again. This means that as long as caches are disabled, users will not see the corresponding hardware event happening.

Attention: The SPI0/1 bus is shared between the instruction & data cache (for firmware execution) and the SPI1 peripheral (controlled by the drivers including this SPI flash driver). Hence, calling SPI Flash API on SPI1 bus (including the main flash) causes significant influence to the whole system. See [Concurrency Constraints for Flash on SPI1](#) for more details.

SPI Flash Encryption

It is possible to encrypt the contents of SPI flash and have it transparently decrypted by hardware.

Refer to the [Flash Encryption documentation](#) for more details.

Memory Mapping API

ESP32-P4 features memory hardware which allows regions of flash memory to be mapped into instruction and data address spaces. This mapping works only for read operations. It is not possible to modify contents of flash memory by writing to a mapped memory region.

Mapping happens in 64 KB pages. Memory mapping hardware can map flash into the data address space and the instruction address space. See the technical reference manual for more details and limitations about memory mapping hardware.

Note that some pages are used to map the application itself into memory, so the actual number of available pages may be less than the capability of the hardware.

Reading data from flash using a memory mapped region is the only way to decrypt contents of flash when [flash encryption](#) is enabled. Decryption is performed at the hardware level.

Memory mapping API are declared in `spi_flash_mmap.h` and `esp_partition.h`:

- `spi_flash_mmap()` maps a region of physical flash addresses into instruction space or data space of the CPU.
- `spi_flash_munmap()` unmaps previously mapped region.
- `esp_partition_mmap()` maps part of a partition into the instruction space or data space of the CPU.

Differences between `spi_flash_mmap()` and `esp_partition_mmap()` are as follows:

- `spi_flash_mmap()` must be given a 64 KB aligned physical address.
- `esp_partition_mmap()` may be given any arbitrary offset within the partition. It adjusts the returned pointer to mapped memory as necessary.

Note that since memory mapping happens in pages, it may be possible to read data outside of the partition provided to `esp_partition_mmap`, regardless of the partition boundary.

Note: `mmap` is supported by cache, so it can only be used on main flash.

SPI Flash Implementation

The `esp_flash_t` structure holds chip data as well as three important parts of this API:

1. The host driver, which provides the hardware support to access the chip;
2. The chip driver, which provides compatibility service to different chips;
3. The OS functions, provide support of some OS functions (e.g., lock, delay) in different stages (1st/2nd boot, or the app).

Host Driver The host driver relies on an interface (`spi_flash_host_driver_t`) defined in the `spi_flash_types.h` (in the `hal/include/hal` folder). This interface provides some common functions to communicate with the chip.

In other files of the SPI HAL, some of these functions are implemented with existing ESP32-P4 memory-spi functionalities. However, due to the speed limitations of ESP32-P4, the HAL layer cannot provide high-speed implementations to some reading commands (so the support for it was dropped). The files (`memspi_host_driver.h` and `.c`) implement the high-speed version of these commands with the `common_command` function provided in the HAL, and wrap these functions as `spi_flash_host_driver_t` for upper layer to use.

You can also implement your own host driver, even with the GPIO. As long as all the functions in the `spi_flash_host_driver_t` are implemented, the `esp_flash` API can access the flash regardless of the low-level hardware.

Chip Driver The chip driver, defined in `spi_flash_chip_driver.h`, wraps basic functions provided by the host driver for the API layer to use.

Some operations need some commands to be sent first, or read some status afterwards. Some chips need different commands or values, or need special communication ways.

There is a type of chip called `generic_chip` which stands for common chips. Other special chip drivers can be developed on the base of the generic chip.

The chip driver relies on the host driver.

OS Functions Currently the OS function layer provides entries of a lock and delay.

The lock (see [SPI Bus Lock](#)) is used to resolve the conflicts among the access of devices on the same SPI bus, and the SPI Flash chip access. E.g.

1. On SPI1 bus, the cache (used to fetch the data (code) in the Flash and PSRAM) should be disabled when the flash chip on the SPI0/1 is being accessed.
2. On the other buses, the flash driver needs to disable the ISR registered by SPI Master driver, to avoid conflicts.
3. Some devices of SPI Master driver may require to use the bus monopolized during a period (especially when the device does not have a CS wire, or the wire is controlled by software like SDSPI driver).

The delay is used by some long operations which requires the master to wait or polling periodically.

The top API wraps these the chip driver and OS functions into an entire component, and also provides some argument checking.

OS functions can also help to avoid a watchdog timeout when erasing large flash areas. During this time, the CPU is occupied with the flash erasing task. This stops other tasks from being executed. Among these tasks is the idle task to feed the watchdog timer (WDT). If the configuration option [CONFIG_ESP_TASK_WDT_PANIC](#) is selected and the flash operation time is longer than the watchdog timeout period, the system will reboot.

It is pretty hard to totally eliminate this risk, because the erasing time varies with different flash chips, making it hard to be compatible in flash drivers. Therefore, users need to pay attention to it. Please use the following guidelines:

1. It is recommended to enable the [CONFIG_SPI_FLASH_YIELD_DURING_ERASE](#) option to allow the scheduler to re-schedule during erasing flash memory. Besides, following parameters can also be used.
 - Increase [CONFIG_SPI_FLASH_ERASE_YIELD_TICKS](#) or decrease [CONFIG_SPI_FLASH_ERASE_YIELD_DURATION_MS](#) in menuconfig.
 - You can also increase [CONFIG_ESP_TASK_WDT_TIMEOUT_S](#) in menuconfig for a larger watchdog timeout period. However, with larger watchdog timeout period, previously detected timeouts may no longer be detected.
2. Please be aware of the consequences of enabling the [CONFIG_ESP_TASK_WDT_PANIC](#) option when doing long-running SPI flash operations which triggers the panic handler when it times out. However, this option can also help dealing with unexpected exceptions in your application. Please decide whether this is needed to be enabled according to actual condition.
3. During your development, please carefully review the actual flash operation according to the specific requirements and time limits on erasing flash memory of your projects. Always allow reasonable redundancy based on your specific product requirements when configuring the flash erasing timeout threshold, thus improving the reliability of your product.

Implementation Details

In order to perform some flash operations, it is necessary to make sure that both CPUs are not running any code from flash for the duration of the flash operation:

- In a single-core setup, the SDK needs to disable interrupts or scheduler before performing the flash operation.
- In a dual-core setup, the SDK needs to make sure that both CPUs are not running any code from flash.

When SPI flash API is called on CPU A (can be PRO or APP), start the `spi_flash_op_block_func` function on CPU B using the `esp_ipc_call` API. This API wakes up a high priority task on CPU B and tells it to execute a given function, in this case, `spi_flash_op_block_func`. This function disables cache on CPU B and signals that the cache is disabled by setting the `s_flash_op_can_start` flag. Then the task on CPU A disables cache as well and proceeds to execute flash operation.

While a flash operation is running, interrupts can still run on CPUs A and B. It is assumed that all interrupt code is placed into RAM. Once the interrupt allocation API is added, a flag should be added to request the interrupt to be disabled for the duration of a flash operations.

Once the flash operation is complete, the function on CPU A sets another flag, `s_flash_op_complete`, to let the task on CPU B know that it can re-enable cache and release the CPU. Then the function on CPU A re-enables the cache on CPU A as well and returns control to the calling code.

Additionally, all API functions are protected with a mutex (`s_flash_op_mutex`).

In a single core environment ([CONFIG_FREERTOS_UNICORE](#) enabled), you need to disable both caches, so that no inter-CPU communication can take place.

Related Documents

- [Optional Features for Flash](#)
- [Concurrency Constraints for Flash on SPI](#)

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

SPI Flash API ESP-IDF Version vs Chip-ROM Version There is a set of SPI Flash drivers in Chip-ROM which you can use by enabling `CONFIG_SPI_FLASH_ROM_IMPL`. Most of the ESP-IDF SPI Flash driver code are in internal RAM, therefore enabling this option frees some internal RAM usage. Note if you enable this option, this means some SPI Flash driver features and bugfixes that are done in ESP-IDF might not be included in the Chip-ROM version.

Feature Supported by ESP-IDF but Not in Chip-ROM

- Octal Flash chip support. See [OPI flash Support](#) for details.
- 32-bit-address support for GD25Q256. Note this feature is an optional feature, please do read [32-bit Address Flash Chips](#) for details.
- TH Flash chip support.
- Kconfig option `CONFIG_SPI_FLASH_CHECK_ERASE_TIMEOUT_DISABLED`.
- `CONFIG_SPI_FLASH_VERIFY_WRITE`, enabling this option helps you detect bad writing.
- `CONFIG_SPI_FLASH_LOG_FAILED_WRITE`, enabling this option prints the bad writing.
- `CONFIG_SPI_FLASH_WARN_SETTING_ZERO_TO_ONE`, enabling this option checks if you are writing zero to one.
- `CONFIG_SPI_FLASH_DANGEROUS_WRITE`, enabling this option checks for flash programming to certain protected regions like bootloader, partition table or application itself.
- `CONFIG_SPI_FLASH_ENABLE_COUNTERS`, enabling this option to collect performance data for ESP-IDF SPI Flash driver APIs.
- `CONFIG_SPI_FLASH_AUTO_SUSPEND`, enabling this option to automatically suspend / resume a long Flash operation when short Flash operation happens. Note this feature is an optional feature, please do read [Auto Suspend & Resume](#) for more limitations.

Bugfixes Introduced in ESP-IDF but Not in Chip-ROM

- Detected Flash physical size correctly, for larger than 256MBit Flash chips. (Commit ID: b4964279d44f73cce7cfd5cf684567fbd6fd9e)
- Fixed issue that address range may escape from checking for erasing and writing function when their sum overflows 32-bit boundary.

API Reference - SPI Flash

Header File

- `components/spi_flash/include/esp_flash_spi_init.h`
- This header file can be included with:

```
#include "esp_flash_spi_init.h"
```

- This header file is a part of the API provided by the `spi_flash` component. To declare that your component depends on `spi_flash`, add the following to your `CMakeLists.txt`:

```
REQUIRES spi_flash
```

or

```
PRIV_REQUIRES spi_flash
```

Functions

esp_err_t **spi_bus_add_flash_device** (*esp_flash_t* **out_chip, const *esp_flash_spi_device_config_t* *config)

Add a SPI Flash device onto the SPI bus.

The bus should be already initialized by `spi_bus_initialization`.

Parameters

- **out_chip** -- Pointer to hold the initialized chip.
- **config** -- Configuration of the chips to initialize.

Returns

- ESP_ERR_INVALID_ARG: out_chip is NULL, or some field in the config is invalid.
- ESP_ERR_NO_MEM: failed to allocate memory for the chip structures.
- ESP_OK: success.

esp_err_t **spi_bus_remove_flash_device** (*esp_flash_t* *chip)

Remove a SPI Flash device from the SPI bus.

Parameters **chip** -- The flash device to remove.

Returns

- ESP_ERR_INVALID_ARG: The chip is invalid.
- ESP_OK: success.

Structures

struct **esp_flash_spi_device_config_t**

Configurations for the SPI Flash to init.

Public Members

spi_host_device_t **host_id**

Bus to use.

int **cs_io_num**

GPIO pin to output the CS signal.

esp_flash_io_mode_t **io_mode**

IO mode to read from the Flash.

enum *esp_flash_speed_s* **speed**

Speed of the Flash clock. Replaced by `freq_mhz`.

int **input_delay_ns**

Input delay of the data pins, in ns. Set to 0 if unknown.

int **cs_id**

CS line ID, ignored when not `host_id` is not `SPI1_HOST`, or `CONFIG_SPI_FLASH_SHARE_SPI1_BUS` is enabled. In this case, the CS line used is automatically assigned by the SPI bus lock.

int **freq_mhz**

The frequency of flash chip(MHZ)

Header File

- [components/spi_flash/include/esp_flash.h](#)
- This header file can be included with:

```
#include "esp_flash.h"
```

- This header file is a part of the API provided by the `spi_flash` component. To declare that your component depends on `spi_flash`, add the following to your `CMakeLists.txt`:

```
REQUIRES spi_flash
```

or

```
PRIV_REQUIRES spi_flash
```

Functions

esp_err_t **esp_flash_init** (*esp_flash_t* *chip)

Initialise SPI flash chip interface.

This function must be called before any other API functions are called for this chip.

Note: Only the `host` and `read_mode` fields of the chip structure must be initialised before this function is called. Other fields may be auto-detected if left set to zero or `NULL`.

Note: If the `chip->drv` pointer is `NULL`, chip `chip_drv` will be auto-detected based on its manufacturer & product IDs. See `esp_flash_registered_flash_drivers` pointer for details of this process.

Parameters **chip** -- Pointer to SPI flash chip to use. If `NULL`, `esp_flash_default_chip` is substituted.

Returns `ESP_OK` on success, or a flash error code if initialisation fails.

bool **esp_flash_chip_driver_initialized** (const *esp_flash_t* *chip)

Check if appropriate chip driver is set.

Parameters **chip** -- Pointer to SPI flash chip to use. If `NULL`, `esp_flash_default_chip` is substituted.

Returns `true` if set, otherwise `false`.

esp_err_t **esp_flash_read_id** (*esp_flash_t* *chip, uint32_t *out_id)

Read flash ID via the common "RDID" SPI flash command.

ID is a 24-bit value. Lower 16 bits of 'id' are the chip ID, upper 8 bits are the manufacturer ID.

Parameters

- **chip** -- Pointer to identify flash chip. Must have been successfully initialised via `esp_flash_init()`
- **out_id** -- [out] Pointer to receive ID value.

Returns `ESP_OK` on success, or a flash error code if operation failed.

esp_err_t **esp_flash_get_size** (*esp_flash_t* *chip, uint32_t *out_size)

Detect flash size based on flash ID.

Note: 1. Most flash chips use a common format for flash ID, where the lower 4 bits specify the size as a power of 2. If the manufacturer doesn't follow this convention, the size may be incorrectly detected.

- a. The `out_size` returned only stands for The `out_size` stands for the size in the binary image header. If you want to get the real size of the chip, please call `esp_flash_get_physical_size` instead.

Parameters

- **chip** -- Pointer to identify flash chip. Must have been successfully initialised via `esp_flash_init()`
- **out_size** -- [out] Detected size in bytes, standing for the size in the binary image header.

Returns ESP_OK on success, or a flash error code if operation failed.

esp_err_t **esp_flash_get_physical_size** (*esp_flash_t* *chip, uint32_t *flash_size)

Detect flash size based on flash ID.

Note: Most flash chips use a common format for flash ID, where the lower 4 bits specify the size as a power of 2. If the manufacturer doesn't follow this convention, the size may be incorrectly detected.

Parameters

- **chip** -- Pointer to identify flash chip. Must have been successfully initialised via `esp_flash_init()`
- **flash_size** -- [out] Detected size in bytes.

Returns ESP_OK on success, or a flash error code if operation failed.

esp_err_t **esp_flash_read_unique_chip_id** (*esp_flash_t* *chip, uint64_t *out_id)

Read flash unique ID via the common "RDUID" SPI flash command.

ID is a 64-bit value.

Note: This is an optional feature, which is not supported on all flash chips. READ PROGRAMMING GUIDE FIRST!

Parameters

- **chip** -- Pointer to identify flash chip. Must have been successfully initialised via `esp_flash_init()`.
- **out_id** -- [out] Pointer to receive unique ID value.

Returns

- ESP_OK on success, or a flash error code if operation failed.
- ESP_ERR_NOT_SUPPORTED if the chip doesn't support read id.

esp_err_t **esp_flash_erase_chip** (*esp_flash_t* *chip)

Erase flash chip contents.

Parameters **chip** -- Pointer to identify flash chip. Must have been successfully initialised via `esp_flash_init()`

Returns

- ESP_OK on success,
- ESP_ERR_NOT_SUPPORTED if the chip is not able to perform the operation. This is indicated by WREN = 1 after the command is sent.
- ESP_ERR_NOT_ALLOWED if a read-only partition is present.
- Other flash error code if operation failed.

esp_err_t **esp_flash_erase_region** (*esp_flash_t* *chip, uint32_t start, uint32_t len)

Erase a region of the flash chip.

Sector size is specified in `chip->drv->sector_size` field (typically 4096 bytes.) `ESP_ERR_INVALID_ARG` will be returned if the start & length are not a multiple of this size.

Erase is performed using block (multi-sector) erases where possible (block size is specified in `chip->drv->block_erase_size` field, typically 65536 bytes). Remaining sectors are erased using individual sector erase commands.

Parameters

- **chip** -- Pointer to identify flash chip. If NULL, `esp_flash_default_chip` is substituted. Must have been successfully initialised via `esp_flash_init()`
- **start** -- Address to start erasing flash. Must be sector aligned.
- **len** -- Length of region to erase. Must also be sector aligned.

Returns

- `ESP_OK` on success,
- `ESP_ERR_NOT_SUPPORTED` if the chip is not able to perform the operation. This is indicated by `WREN = 1` after the command is sent.
- `ESP_ERR_NOT_ALLOWED` if the address range (`start` — `start + len`) overlaps with a read-only partition address space
- Other flash error code if operation failed.

esp_err_t **esp_flash_get_chip_write_protect** (*esp_flash_t* *chip, bool *write_protected)

Read if the entire chip is write protected.

Note: A correct result for this flag depends on the SPI flash chip model and `chip_drv` in use (via the 'chip->drv' field).

Parameters

- **chip** -- Pointer to identify flash chip. If NULL, `esp_flash_default_chip` is substituted. Must have been successfully initialised via `esp_flash_init()`
- **write_protected** -- [out] Pointer to boolean, set to the value of the write protect flag.

Returns `ESP_OK` on success, or a flash error code if operation failed.

esp_err_t **esp_flash_set_chip_write_protect** (*esp_flash_t* *chip, bool write_protect)

Set write protection for the SPI flash chip.

Some SPI flash chips may require a power cycle before write protect status can be cleared. Otherwise, write protection can be removed via a follow-up call to this function.

Note: Correct behaviour of this function depends on the SPI flash chip model and `chip_drv` in use (via the 'chip->drv' field).

Parameters

- **chip** -- Pointer to identify flash chip. If NULL, `esp_flash_default_chip` is substituted. Must have been successfully initialised via `esp_flash_init()`
- **write_protect** -- Boolean value for the write protect flag

Returns `ESP_OK` on success, or a flash error code if operation failed.

esp_err_t **esp_flash_get_protectable_regions** (const *esp_flash_t* *chip, const *esp_flash_region_t* **out_regions, uint32_t *out_num_regions)

Read the list of individually protectable regions of this SPI flash chip.

Note: Correct behaviour of this function depends on the SPI flash chip model and `chip_drv` in use (via the

'chip->drv' field).

Parameters

- **chip** -- Pointer to identify flash chip. Must have been successfully initialised via `esp_flash_init()`
- **out_regions** -- [out] Pointer to receive a pointer to the array of protectable regions of the chip.
- **out_num_regions** -- [out] Pointer to an integer receiving the count of protectable regions in the array returned in 'regions'.

Returns ESP_OK on success, or a flash error code if operation failed.

esp_err_t **esp_flash_get_protected_region** (*esp_flash_t* *chip, const *esp_flash_region_t* *region, bool *out_protected)

Detect if a region of the SPI flash chip is protected.

Note: It is possible for this result to be false and write operations to still fail, if protection is enabled for the entire chip.

Note: Correct behaviour of this function depends on the SPI flash chip model and chip_drv in use (via the 'chip->drv' field).

Parameters

- **chip** -- Pointer to identify flash chip. Must have been successfully initialised via `esp_flash_init()`
- **region** -- Pointer to a struct describing a protected region. This must match one of the regions returned from `esp_flash_get_protectable_regions(...)`.
- **out_protected** -- [out] Pointer to a flag which is set based on the protected status for this region.

Returns ESP_OK on success, or a flash error code if operation failed.

esp_err_t **esp_flash_set_protected_region** (*esp_flash_t* *chip, const *esp_flash_region_t* *region, bool protect)

Update the protected status for a region of the SPI flash chip.

Note: It is possible for the region protection flag to be cleared and write operations to still fail, if protection is enabled for the entire chip.

Note: Correct behaviour of this function depends on the SPI flash chip model and chip_drv in use (via the 'chip->drv' field).

Parameters

- **chip** -- Pointer to identify flash chip. Must have been successfully initialised via `esp_flash_init()`
- **region** -- Pointer to a struct describing a protected region. This must match one of the regions returned from `esp_flash_get_protectable_regions(...)`.
- **protect** -- Write protection flag to set.

Returns ESP_OK on success, or a flash error code if operation failed.

esp_err_t **esp_flash_read** (*esp_flash_t* *chip, void *buffer, uint32_t address, uint32_t length)

Read data from the SPI flash chip.

There are no alignment constraints on buffer, address or length.

Note: If on-chip flash encryption is used, this function returns raw (ie encrypted) data. Use the flash cache to transparently decrypt data.

Parameters

- **chip** -- Pointer to identify flash chip. If NULL, `esp_flash_default_chip` is substituted. Must have been successfully initialised via `esp_flash_init()`
- **buffer** -- Pointer to a buffer where the data will be read. To get better performance, this should be in the DRAM and word aligned.
- **address** -- Address on flash to read from. Must be less than `chip->size` field.
- **length** -- Length (in bytes) of data to read.

Returns

- `ESP_OK`: success
- `ESP_ERR_NO_MEM`: Buffer is in external PSRAM which cannot be concurrently accessed, and a temporary internal buffer could not be allocated.
- or a flash error code if operation failed.

esp_err_t **esp_flash_write** (*esp_flash_t* *chip, const void *buffer, uint32_t address, uint32_t length)

Write data to the SPI flash chip.

There are no alignment constraints on buffer, address or length.

Parameters

- **chip** -- Pointer to identify flash chip. If NULL, `esp_flash_default_chip` is substituted. Must have been successfully initialised via `esp_flash_init()`
- **address** -- Address on flash to write to. Must be previously erased (SPI NOR flash can only write bits 1->0).
- **buffer** -- Pointer to a buffer with the data to write. To get better performance, this should be in the DRAM and word aligned.
- **length** -- Length (in bytes) of data to write.

Returns

- `ESP_OK` on success
- `ESP_FAIL`, bad write, this will be detected only when `CONFIG_SPI_FLASH_VERIFY_WRITE` is enabled
- `ESP_ERR_NOT_SUPPORTED` if the chip is not able to perform the operation. This is indicated by `WREN = 1` after the command is sent.
- `ESP_ERR_NOT_ALLOWED` if the address range (`address` — `address + length`) overlaps with a read-only partition address space
- Other flash error code if operation failed.

esp_err_t **esp_flash_write_encrypted** (*esp_flash_t* *chip, uint32_t address, const void *buffer, uint32_t length)

Encrypted and write data to the SPI flash chip using on-chip hardware flash encryption.

Note: Both address & length must be 16 byte aligned, as this is the encryption block size

Parameters

- **chip** -- Pointer to identify flash chip. Must be NULL (the main flash chip). For other chips, encrypted write is not supported.

- **address** -- Address on flash to write to. 16 byte aligned. Must be previously erased (SPI NOR flash can only write bits 1->0).
- **buffer** -- Pointer to a buffer with the data to write.
- **length** -- Length (in bytes) of data to write. 16 byte aligned.

Returns

- ESP_OK: on success
- ESP_FAIL: bad write, this will be detected only when CONFIG_SPI_FLASH_VERIFY_WRITE is enabled
- ESP_ERR_NOT_SUPPORTED: encrypted write not supported for this chip.
- ESP_ERR_INVALID_ARG: Either the address, buffer or length is invalid.
- ESP_ERR_NOT_ALLOWED if the address range (address — address + length) overlaps with a read-only partition address space

esp_err_t **esp_flash_read_encrypted** (*esp_flash_t* *chip, uint32_t address, void *out_buffer, uint32_t length)

Read and decrypt data from the SPI flash chip using on-chip hardware flash encryption.

Parameters

- **chip** -- Pointer to identify flash chip. Must be NULL (the main flash chip). For other chips, encrypted read is not supported.
- **address** -- Address on flash to read from.
- **out_buffer** -- Pointer to a buffer for the data to read to.
- **length** -- Length (in bytes) of data to read.

Returns

- ESP_OK: on success
- ESP_ERR_NOT_SUPPORTED: encrypted read not supported for this chip.

static inline bool **esp_flash_is_quad_mode** (const *esp_flash_t* *chip)

Returns true if chip is configured for Quad I/O or Quad Fast Read.

Parameters **chip** -- Pointer to SPI flash chip to use. If NULL, esp_flash_default_chip is substituted.

Returns true if flash works in quad mode, otherwise false

Structures

struct **esp_flash_region_t**

Structure for describing a region of flash.

Public Members

uint32_t **offset**

Start address of this region.

uint32_t **size**

Size of the region.

struct **esp_flash_os_functions_t**

OS-level integration hooks for accessing flash chips inside a running OS.

It's in the public header because some instances should be allocated statically in the startup code. May be updated according to hardware version and new flash chip feature requirements, shouldn't be treated as public API.

For advanced developers, you may replace some of them with your implementations at your own risk.

Public Members

esp_err_t (***start**)(void *arg)

Called before commencing any flash operation. Does not need to be recursive (ie is called at most once for each call to 'end').

esp_err_t (***end**)(void *arg)

Called after completing any flash operation.

esp_err_t (***region_protected**)(void *arg, size_t start_addr, size_t size)

Called before any erase/write operations to check whether the region is limited by the OS

esp_err_t (***delay_us**)(void *arg, uint32_t us)

Delay for at least 'us' microseconds. Called in between 'start' and 'end'.

void (***get_temp_buffer**)(void *arg, size_t request_size, size_t *out_size)

Called for get temp buffer when buffer from application cannot be directly read into/write from.

void (***release_temp_buffer**)(void *arg, void *temp_buf)

Called for release temp buffer.

esp_err_t (***check_yield**)(void *arg, uint32_t chip_status, uint32_t *out_request)

Yield to other tasks. Called during erase operations.

Return ESP_OK means yield needs to be called (got an event to handle), while ESP_ERR_TIMEOUT means skip yield.

esp_err_t (***yield**)(void *arg, uint32_t *out_status)

Yield to other tasks. Called during erase operations.

int64_t (***get_system_time**)(void *arg)

Called for get system time.

void (***set_flash_op_status**)(uint32_t op_status)

Call to set flash operation status

struct **esp_flash_t**

Structure to describe a SPI flash chip connected to the system.

Structure must be initialized before use (passed to esp_flash_init()). It's in the public header because some instances should be allocated statically in the startup code. May be updated according to hardware version and new flash chip feature requirements, shouldn't be treated as public API.

For advanced developers, you may replace some of them with your implementations at your own risk.

Public Members

spi_flash_host_inst_t ***host**

Pointer to hardware-specific "host_driver" structure. Must be initialized before used.

const *spi_flash_chip_t* ***chip_drv**

Pointer to chip-model-specific "adapter" structure. If NULL, will be detected during initialisation.

const *esp_flash_os_functions_t* ***os_func**

Pointer to os-specific hook structure. Call `esp_flash_init_os_functions()` to setup this field, after the host is properly initialized.

void ***os_func_data**

Pointer to argument for os-specific hooks. Left NULL and will be initialized with `os_func`.

esp_flash_io_mode_t **read_mode**

Configured SPI flash read mode. Set before `esp_flash_init` is called.

uint32_t **size**

Size of SPI flash in bytes. If 0, size will be detected during initialisation. Note: this stands for the size in the binary image header. If you want to get the flash physical size, please call `esp_flash_get_physical_size`.

uint32_t **chip_id**

Detected chip id.

uint32_t **busy**

This flag is used to verify chip's status.

uint32_t **hpm_dummy_ena**

This flag is used to verify whether flash works under HPM status.

uint32_t **reserved_flags**

reserved.

Macros

SPI_FLASH_YIELD_REQ_YIELD

SPI_FLASH_YIELD_REQ_SUSPEND

SPI_FLASH_YIELD_STA_RESUME

SPI_FLASH_OS_IS_ERASING_STATUS_FLAG

Type Definitions

typedef struct *spi_flash_chip_t* **spi_flash_chip_t**

Header File

- [components/spi_flash/include/spi_flash_mmap.h](#)
- This header file can be included with:

```
#include "spi_flash_mmap.h"
```

- This header file is a part of the API provided by the `spi_flash` component. To declare that your component depends on `spi_flash`, add the following to your `CMakeLists.txt`:

```
REQUIRES spi_flash
```

or

```
PRIV_REQUIRES spi_flash
```

Functions

`esp_err_t spi_flash_mmap` (`size_t src_addr`, `size_t size`, `spi_flash_mmap_memory_t memory`, `const void **out_ptr`, `spi_flash_mmap_handle_t *out_handle`)

Map region of flash memory into data or instruction address space.

This function allocates sufficient number of 64kB MMU pages and configures them to map the requested region of flash memory into the address space. It may reuse MMU pages which already provide the required mapping.

As with any allocator, if `mmap/munmap` are heavily used then the address space may become fragmented. To troubleshoot issues with page allocation, use `spi_flash_mmap_dump()` function.

Parameters

- **`src_addr`** -- Physical address in flash where requested region starts. This address *must* be aligned to 64kB boundary (`SPI_FLASH_MMU_PAGE_SIZE`)
- **`size`** -- Size of region to be mapped. This size will be rounded up to a 64kB boundary
- **`memory`** -- Address space where the region should be mapped (data or instruction)
- **`out_ptr`** -- [out] Output, pointer to the mapped memory region
- **`out_handle`** -- [out] Output, handle which should be used for `spi_flash_munmap` call

Returns `ESP_OK` on success, `ESP_ERR_NO_MEM` if pages can not be allocated

`esp_err_t spi_flash_mmap_pages` (`const int *pages`, `size_t page_count`, `spi_flash_mmap_memory_t memory`, `const void **out_ptr`, `spi_flash_mmap_handle_t *out_handle`)

Map sequences of pages of flash memory into data or instruction address space.

This function allocates sufficient number of 64kB MMU pages and configures them to map the indicated pages of flash memory contiguously into address space. In this respect, it works in a similar way as `spi_flash_mmap()` but it allows mapping a (maybe non-contiguous) set of pages into a contiguous region of memory.

Parameters

- **`pages`** -- An array of numbers indicating the 64kB pages in flash to be mapped contiguously into memory. These indicate the indexes of the 64kB pages, not the byte-size addresses as used in other functions. Array must be located in internal memory.
- **`page_count`** -- Number of entries in the `pages` array
- **`memory`** -- Address space where the region should be mapped (instruction or data)
- **`out_ptr`** -- [out] Output, pointer to the mapped memory region
- **`out_handle`** -- [out] Output, handle which should be used for `spi_flash_munmap` call

Returns

- `ESP_OK` on success
- `ESP_ERR_NO_MEM` if pages can not be allocated
- `ESP_ERR_INVALID_ARG` if `pagecount` is zero or `pages` array is not in internal memory

void **`spi_flash_munmap`** (`spi_flash_mmap_handle_t handle`)

Release region previously obtained using `spi_flash_mmap`.

Note: Calling this function will not necessarily unmap memory region. Region will only be unmapped when there are no other handles which reference this region. In case of partially overlapping regions it is possible that memory will be unmapped partially.

Parameters **`handle`** -- Handle obtained from `spi_flash_mmap`

void **spi_flash_mmap_dump** (void)

Display information about mapped regions.

This function lists handles obtained using `spi_flash_mmap`, along with range of pages allocated to each handle. It also lists all non-zero entries of MMU table and corresponding reference counts.

uint32_t **spi_flash_mmap_get_free_pages** (*spi_flash_mmap_memory_t* memory)

get free pages number which can be mmap

This function will return number of free pages available in mmu table. This could be useful before calling actual `spi_flash_mmap` (maps flash range to DCache or ICache memory) to check if there is sufficient space available for mapping.

Parameters **memory** -- memory type of MMU table free page

Returns number of free pages which can be mmaped

size_t **spi_flash_cache2phys** (const void *cached)

Given a memory address where flash is mapped, return the corresponding physical flash offset.

Cache address does not have have been assigned via `spi_flash_mmap()`, any address in memory mapped flash space can be looked up.

Parameters **cached** -- Pointer to flashed cached memory.

Returns

- **SPI_FLASH_CACHE2PHYS_FAIL** If cache address is outside flash cache region, or the address is not mapped.
- Otherwise, returns physical offset in flash

const void ***spi_flash_phys2cache** (size_t phys_offs, *spi_flash_mmap_memory_t* memory)

Given a physical offset in flash, return the address where it is mapped in the memory space.

Physical address does not have to have been assigned via `spi_flash_mmap()`, any address in flash can be looked up.

Note: Only the first matching cache address is returned. If MMU flash cache table is configured so multiple entries point to the same physical address, there may be more than one cache address corresponding to that physical address. It is also possible for a single physical address to be mapped to both the IROM and DROM regions.

Note: This function doesn't impose any alignment constraints, but if memory argument is **SPI_FLASH_MMAP_INST** and `phys_offs` is not 4-byte aligned, then reading from the returned pointer will result in a crash.

Parameters

- **phys_offs** -- Physical offset in flash memory to look up.
- **memory** -- Address space type to look up a flash cache address mapping for (instruction or data)

Returns

- NULL if the physical address is invalid or not mapped to flash cache of the specified memory type.
- Cached memory address (in IROM or DROM space) corresponding to `phys_offs`.

Macros

ESP_ERR_FLASH_OP_FAIL

This file contains `spi_flash_mmap_xx` APIs, mainly for doing memory mapping to an SPI0-connected external Flash, as well as some helper functions to convert between virtual and physical address

ESP_ERR_FLASH_OP_TIMEOUT

SPI_FLASH_SEC_SIZE

SPI Flash sector size

SPI_FLASH_MMU_PAGE_SIZE

Flash cache MMU mapping page size

SPI_FLASH_CACHE2PHYS_FAIL

Type Definitions

typedef uint32_t **spi_flash_mmap_handle_t**

Opaque handle for memory region obtained from spi_flash_mmap.

Enumerations

enum **spi_flash_mmap_memory_t**

Enumeration which specifies memory space requested in an mmap call.

Values:

enumerator **SPI_FLASH_MMAP_DATA**

map to data memory, allows byte-aligned access

enumerator **SPI_FLASH_MMAP_INST**

map to instruction memory, allows only 4-byte-aligned access

Header File

- [components/hal/include/hal/spi_flash_types.h](#)
- This header file can be included with:

```
#include "hal/spi_flash_types.h"
```

Structures

struct **spi_flash_trans_t**

Definition of a common transaction. Also holds the return value.

Public Members

uint8_t **reserved**

Reserved, must be 0.

uint8_t **mosi_len**

Output data length, in bytes.

uint8_t **miso_len**

Input data length, in bytes.

uint8_t **address_bitlen**

Length of address in bits, set to 0 if command does not need an address.

uint32_t **address**

Address to perform operation on.

const uint8_t ***mosi_data**

Output data to slave.

uint8_t ***miso_data**

[out] Input data from slave, little endian

uint32_t **flags**

Flags for this transaction. Set to 0 for now.

uint16_t **command**

Command to send.

uint8_t **dummy_bitlen**

Basic dummy bits to use.

uint32_t **io_mode**

Flash working mode when SPI_FLASH_IGNORE_BASEIO is specified.

struct **spi_flash_sus_cmd_conf**

Configuration structure for the flash chip suspend feature.

Public Members

uint32_t **sus_mask**

SUS/SUS1/SUS2 bit in flash register.

uint32_t **cmd_rdsr**

Read flash status register(2) command.

uint32_t **sus_cmd**

Flash suspend command.

uint32_t **res_cmd**

Flash resume command.

uint32_t **reserved**

Reserved, set to 0.

struct **spi_flash_encryption_t**

Structure for flash encryption operations.

Public Members

void (***flash_encryption_enable**)(void)

Enable the flash encryption.

void (***flash_encryption_disable**)(void)

Disable the flash encryption.

void (***flash_encryption_data_prepare**)(uint32_t address, const uint32_t *buffer, uint32_t size)

Prepare flash encryption before operation.

Note: address and buffer must be 8-word aligned.

Param address The destination address in flash for the write operation.

Param buffer Data for programming

Param size Size to program.

void (***flash_encryption_done**)(void)

flash data encryption operation is done.

void (***flash_encryption_destroy**)(void)

Destroy encrypted result

bool (***flash_encryption_check**)(uint32_t address, uint32_t length)

Check if is qualified to encrypt the buffer

Param address the address of written flash partition.

Param length Buffer size.

struct **spi_flash_host_inst_t**

SPI Flash Host driver instance

Public Members

const struct *spi_flash_host_driver_s* ***driver**

Pointer to the implementation function table.

struct **spi_flash_host_driver_s**

Host driver configuration and context structure.

Public Members

esp_err_t (***dev_config**)(*spi_flash_host_inst_t* *host)

Configure the device-related register before transactions. This saves some time to re-configure those registers when we send continuously

esp_err_t (***common_command**)(*spi_flash_host_inst_t* *host, *spi_flash_trans_t* *t)

Send an user-defined spi transaction to the device.

esp_err_t (***read_id**)(*spi_flash_host_inst_t* *host, uint32_t *id)

Read flash ID.

void (***erase_chip**)(*spi_flash_host_inst_t* *host)

Erase whole flash chip.

void (***erase_sector**)(*spi_flash_host_inst_t* *host, uint32_t start_address)

Erase a specific sector by its start address.

void (***erase_block**)(*spi_flash_host_inst_t* *host, uint32_t start_address)

Erase a specific block by its start address.

esp_err_t (***read_status**)(*spi_flash_host_inst_t* *host, uint8_t *out_sr)

Read the status of the flash chip.

esp_err_t (***set_write_protect**)(*spi_flash_host_inst_t* *host, bool wp)

Disable write protection.

void (***program_page**)(*spi_flash_host_inst_t* *host, const void *buffer, uint32_t address, uint32_t length)

Program a page of the flash. Check `max_write_bytes` for the maximum allowed writing length.

bool (***supports_direct_write**)(*spi_flash_host_inst_t* *host, const void *p)

Check whether the SPI host supports direct write.

When cache is disabled, SPI1 doesn't support directly write when buffer isn't internal.

int (***write_data_slicer**)(*spi_flash_host_inst_t* *host, uint32_t address, uint32_t len, uint32_t *align_addr, uint32_t page_size)

Slicer for write data. The `program_page` should be called iteratively with the return value of this function.

Param address Beginning flash address to write

Param len Length request to write

Param align_addr Output of the aligned address to write to

Param page_size Physical page size of the flash chip

Return Length that can be actually written in one `program_page` call

esp_err_t (***read**)(*spi_flash_host_inst_t* *host, void *buffer, uint32_t address, uint32_t read_len)

Read data from the flash. Check `max_read_bytes` for the maximum allowed reading length.

bool (***supports_direct_read**)(*spi_flash_host_inst_t* *host, const void *p)

Check whether the SPI host supports direct read.

When cache is disabled, SPI1 doesn't support directly read when the given buffer isn't internal.

int (***read_data_slicer**)(*spi_flash_host_inst_t* *host, uint32_t address, uint32_t len, uint32_t *align_addr, uint32_t page_size)

Slicer for read data. The `read` should be called iteratively with the return value of this function.

Param address Beginning flash address to read

Param len Length request to read

Param align_addr Output of the aligned address to read

Param page_size Physical page size of the flash chip

Return Length that can be actually read in one `read` call

`uint32_t (*host_status)(spi_flash_host_inst_t *host)`

Check the host status, 0:busy, 1:idle, 2:suspended.

`esp_err_t (*configure_host_io_mode)(spi_flash_host_inst_t *host, uint32_t command, uint32_t addr_bitlen, int dummy_bitlen_base, esp_flash_io_mode_t io_mode)`

Configure the host to work at different read mode. Responsible to compensate the timing and set IO mode.

`void (*poll_cmd_done)(spi_flash_host_inst_t *host)`

Internal use, poll the HW until the last operation is done.

`esp_err_t (*flush_cache)(spi_flash_host_inst_t *host, uint32_t addr, uint32_t size)`

For some host (SPI1), they are shared with a cache. When the data is modified, the cache needs to be flushed. Left NULL if not supported.

`void (*check_suspend)(spi_flash_host_inst_t *host)`

Suspend check erase/program operation, reserved for ESP32-C3 and ESP32-S3 spi flash ROM IMPL.

`void (*resume)(spi_flash_host_inst_t *host)`

Resume flash from suspend manually

`void (*suspend)(spi_flash_host_inst_t *host)`

Set flash in suspend status manually

`esp_err_t (*sus_setup)(spi_flash_host_inst_t *host, const spi_flash_sus_cmd_conf *sus_conf)`

Suspend feature setup for setting cmd and status register mask.

Macros

SPI_FLASH_TRANS_FLAG_CMD16

Send command of 16 bits.

SPI_FLASH_TRANS_FLAG_IGNORE_BASEIO

Not applying the basic io mode configuration for this transaction.

SPI_FLASH_TRANS_FLAG_BYTE_SWAP

Used for DTR mode, to swap the bytes of a pair of rising/falling edge.

SPI_FLASH_TRANS_FLAG_PE_CMD

Indicates that this transaction is to erase/program flash chip.

SPI_FLASH_CONFIG_CONF_BITS

OR the io_mode with this mask, to enable the dummy output feature or replace the first several dummy bits into address to meet the requirements of conf bits. (Used in DIO/QIO/OIO mode)

SPI_FLASH_OPI_FLAG

A flag for flash work in opi mode, the io mode below are opi, above are SPI/QSPI mode. DO NOT use this value in any API.

SPI_FLASH_READ_MODE_MIN

Slowest io mode supported by ESP32, currently SlowRd.

Type Definitions

typedef enum *esp_flash_speed_s* **esp_flash_speed_t**

SPI flash clock speed values, always refer to them by the enum rather than the actual value (more speed may be appended into the list).

A strategy to select the maximum allowed speed is to enumerate from the `ESP_FLASH_SPEED_MAX-1` or highest frequency supported by your flash, and decrease the speed until the probing success.

typedef struct *spi_flash_host_driver_s* **spi_flash_host_driver_t**

Enumerations

enum **esp_flash_speed_s**

SPI flash clock speed values, always refer to them by the enum rather than the actual value (more speed may be appended into the list).

A strategy to select the maximum allowed speed is to enumerate from the `ESP_FLASH_SPEED_MAX-1` or highest frequency supported by your flash, and decrease the speed until the probing success.

Values:

enumerator **ESP_FLASH_5MHZ**

The flash runs under 5MHz.

enumerator **ESP_FLASH_10MHZ**

The flash runs under 10MHz.

enumerator **ESP_FLASH_20MHZ**

The flash runs under 20MHz.

enumerator **ESP_FLASH_26MHZ**

The flash runs under 26MHz.

enumerator **ESP_FLASH_40MHZ**

The flash runs under 40MHz.

enumerator **ESP_FLASH_80MHZ**

The flash runs under 80MHz.

enumerator **ESP_FLASH_120MHZ**

The flash runs under 120MHz, 120MHz can only be used by main flash after timing tuning in system. Do not use this directly in any API.

enumerator **ESP_FLASH_SPEED_MAX**

The maximum frequency supported by the host is `ESP_FLASH_SPEED_MAX-1`.

enum **esp_flash_io_mode_t**

Mode used for reading from SPI flash.

Values:

enumerator **SPI_FLASH_SLOWRD**

Data read using single I/O, some limits on speed.

enumerator **SPI_FLASH_FASTRD**

Data read using single I/O, no limit on speed.

enumerator **SPI_FLASH_DOUT**

Data read using dual I/O.

enumerator **SPI_FLASH_DIO**

Both address & data transferred using dual I/O.

enumerator **SPI_FLASH_QOUT**

Data read using quad I/O.

enumerator **SPI_FLASH_QIO**

Both address & data transferred using quad I/O.

enumerator **SPI_FLASH_OPI_STR**

Only support on OPI flash, flash read and write under STR mode.

enumerator **SPI_FLASH_OPI_DTR**

Only support on OPI flash, flash read and write under DTR mode.

enumerator **SPI_FLASH_READ_MODE_MAX**

The fastest io mode supported by the host is `ESP_FLASH_READ_MODE_MAX-1`.

Header File

- [components/hal/include/hal/esp_flash_err.h](#)
- This header file can be included with:

```
#include "hal/esp_flash_err.h"
```

Macros

ESP_ERR_FLASH_NOT_INITIALISED

`esp_flash_chip_t` structure not correctly initialised by `esp_flash_init()`.

ESP_ERR_FLASH_UNSUPPORTED_HOST

Requested operation isn't supported via this host SPI bus (`chip->spi` field).

ESP_ERR_FLASH_UNSUPPORTED_CHIP

Requested operation isn't supported by this model of SPI flash chip.

ESP_ERR_FLASH_PROTECTED

Write operation failed due to chip's write protection being enabled.

Enumerations

enum **[anonymous]**

Values:

enumerator **ESP_ERR_FLASH_SIZE_NOT_MATCH**

The chip doesn't have enough space for the current partition table.

enumerator **ESP_ERR_FLASH_NO_RESPONSE**

Chip did not respond to the command, or timed out.

Header File

- [components/spi_flash/include/esp_spi_flash_counters.h](#)
- This header file can be included with:

```
#include "esp_spi_flash_counters.h"
```

- This header file is a part of the API provided by the `spi_flash` component. To declare that your component depends on `spi_flash`, add the following to your `CMakeLists.txt`:

```
REQUIRES spi_flash
```

or

```
PRIV_REQUIRES spi_flash
```

Functions

void **esp_flash_reset_counters** (void)

Reset SPI flash operation counters.

void **spi_flash_reset_counters** (void)

void **esp_flash_dump_counters** (FILE *stream)

Print SPI flash operation counters.

void **spi_flash_dump_counters** (void)

const [esp_flash_counters_t](#) ***esp_flash_get_counters** (void)

Return current SPI flash operation counters.

Returns pointer to the [esp_flash_counters_t](#) structure holding values of the operation counters

const [spi_flash_counters_t](#) ***spi_flash_get_counters** (void)

Structures

struct **esp_flash_counter_t**

Structure holding statistics for one type of operation

Public Members

uint32_t **count**

number of times operation was executed

uint32_t **time**

total time taken, in microseconds

uint32_t **bytes**

total number of bytes

struct **esp_flash_counters_t**

Structure for counters of flash actions

Public Members

esp_flash_counter_t **read**

counters for read action, like `esp_flash_read`

esp_flash_counter_t **write**

counters for write action, like `esp_flash_write`

esp_flash_counter_t **erase**

counters for erase action, like `esp_flash_erase`

Type Definitions

typedef *esp_flash_counter_t* **spi_flash_counter_t**

typedef *esp_flash_counters_t* **spi_flash_counters_t**

API Reference - Flash Encrypt

Header File

- [components/bootloader_support/include/esp_flash_encrypt.h](#)
- This header file can be included with:

```
#include "esp_flash_encrypt.h"
```

- This header file is a part of the API provided by the `bootloader_support` component. To declare that your component depends on `bootloader_support`, add the following to your `CMakeLists.txt`:

```
REQUIRES bootloader_support
```

or

```
PRIV_REQUIRES bootloader_support
```

Functions

bool **esp_flash_encryption_enabled** (void)

Is flash encryption currently enabled in hardware?

Flash encryption is enabled if the `FLASH_CRYPT_CNT` efuse has an odd number of bits set.

Returns true if flash encryption is enabled.

esp_err_t **esp_flash_encrypt_check_and_update** (void)

bool **esp_flash_encrypt_state** (void)

Returns the Flash Encryption state and prints it.

Returns True - Flash Encryption is enabled False - Flash Encryption is not enabled

bool **esp_flash_encrypt_initialized_once** (void)

Checks if the first initialization was done.

If the first initialization was done then FLASH_CRYPT_CNT != 0

Returns true - the first initialization was done false - the first initialization was NOT done

esp_err_t **esp_flash_encrypt_init** (void)

The first initialization of Flash Encryption key and related eFuses.

Returns ESP_OK if all operations succeeded

esp_err_t **esp_flash_encrypt_contents** (void)

Encrypts flash content.

Returns ESP_OK if all operations succeeded

esp_err_t **esp_flash_encrypt_enable** (void)

Activates Flash encryption on the chip.

It burns FLASH_CRYPT_CNT eFuse based on the CONFIG_SECURE_FLASH_ENCRYPTION_MODE_RELEASE option.

Returns ESP_OK if all operations succeeded

bool **esp_flash_encrypt_is_write_protected** (bool print_error)

Returns True if the write protection of FLASH_CRYPT_CNT is set.

Parameters **print_error** -- Print error if it is write protected

Returns true - if FLASH_CRYPT_CNT is write protected

esp_err_t **esp_flash_encrypt_region** (uint32_t src_addr, size_t data_length)

Encrypt-in-place a block of flash sectors.

Note: This function resets RTC_WDT between operations with sectors.

Parameters

- **src_addr** -- Source offset in flash. Should be multiple of 4096 bytes.
- **data_length** -- Length of data to encrypt in bytes. Will be rounded up to next multiple of 4096 bytes.

Returns ESP_OK if all operations succeeded, ESP_ERR_FLASH_OP_FAIL if SPI flash fails, ESP_ERR_FLASH_OP_TIMEOUT if flash times out.

void **esp_flash_write_protect_crypt_cnt** (void)

Write protect FLASH_CRYPT_CNT.

Intended to be called as a part of boot process if flash encryption is enabled but secure boot is not used. This should protect against serial re-flashing of an unauthorised code in absence of secure boot.

Note: On ESP32 V3 only, write protecting FLASH_CRYPT_CNT will also prevent disabling UART Down-load Mode. If both are wanted, call esp_efuse_disable_rom_download_mode() before calling this function.

esp_flash_enc_mode_t **esp_get_flash_encryption_mode** (void)

Return the flash encryption mode.

The API is called during boot process but can also be called by application to check the current flash encryption mode of ESP32

Returns

void **esp_flash_encryption_init_checks** (void)

Check the flash encryption mode during startup.

Verifies the flash encryption config during startup:

- Correct any insecure flash encryption settings if hardware Secure Boot is enabled.
- Log warnings if the efuse config doesn't match the project config in any way

Note: This function is called automatically during app startup, it doesn't need to be called from the app.

esp_err_t **esp_flash_encryption_enable_secure_features** (void)

Set all secure eFuse features related to flash encryption.

Returns

- ESP_OK - Successfully

bool **esp_flash_encryption_cfg_verify_release_mode** (void)

Returns the verification status for all physical security features of flash encryption in release mode.

If the device has flash encryption feature configured in the release mode, then it is highly recommended to call this API in the application startup code. This API verifies the sanity of the eFuse configuration against the release (production) mode of the flash encryption feature.

Returns

- True - all eFuses are configured correctly
- False - not all eFuses are configured correctly.

void **esp_flash_encryption_set_release_mode** (void)

Switches Flash Encryption from "Development" to "Release".

If already in "Release" mode, the function will do nothing. If flash encryption efuse is not enabled yet then abort. It burns:

- "disable encrypt in dl mode"
- set FLASH_CRYPT_CNT efuse to max

In case of the targets that support the XTS-AES peripheral's pseudo rounds function, this API would configure the pseudo rounds level efuse bit to level low if the efuse bit is not set already.

Enumerations

enum **esp_flash_enc_mode_t**

Values:

enumerator **ESP_FLASH_ENC_MODE_DISABLED**

enumerator **ESP_FLASH_ENC_MODE_DEVELOPMENT**

enumerator **ESP_FLASH_ENC_MODE_RELEASE**

2.5.21 SPI Master Driver

SPI Master driver is a program that controls ESP32-P4's General Purpose SPI (GP-SPI) peripheral(s) when it functions as a master.

For more hardware information about the GP-SPI peripheral(s), see [ESP32-P4 Technical Reference Manual > SPI Controller \[PDF\]](#).

Terminology

The terms used in relation to the SPI Master driver are given in the table below.

Term	Definition
Host	The SPI controller peripheral inside ESP32-P4 initiates SPI transmissions over the bus and acts as an SPI Master.
Device	SPI slave Device. An SPI bus may be connected to one or more Devices. Each Device shares the MOSI, MISO, and SCLK signals but is only active on the bus when the Host asserts the Device's individual CS line.
Bus	A signal bus, common to all Devices connected to one Host. In general, a bus includes the following lines: MISO, MOSI, SCLK, one or more CS lines, and, optionally, QUADWP and QUADHD. So Devices are connected to the same lines, with the exception that each Device has its own CS line. Several Devices can also share one CS line if connected in a daisy-chain manner.
MOSI	Master Out, Slave In, a.k.a. D. Data transmission from a Host to Device. Also data0 signal in Octal/OPI mode.
MISO	Master In, Slave Out, a.k.a. Q. Data transmission from a Device to Host. Also data1 signal in Octal/OPI mode.
SCLK	Serial Clock. The oscillating signal generated by a Host keeps the transmission of data bits in sync.
CS	Chip Select. Allows a Host to select individual Device(s) connected to the bus in order to send or receive data.
QUADWP	Write Protect signal. Used for 4-bit (qio/qout) transactions. Also for the data2 signal in Octal/OPI mode.
QUADHD	Hold signal. Used for 4-bit (qio/qout) transactions. Also for the data3 signal in Octal/OPI mode.
DATA4	Data4 signal in Octal/OPI mode.
DATA5	Data5 signal in Octal/OPI mode.
DATA6	Data6 signal in Octal/OPI mode.
DATA7	Data7 signal in Octal/OPI mode.
Assertion	The action of activating a line.
De-assertion	The action of returning the line back to inactive (back to idle) status.
Transaction	One instance of a Host asserting a CS line, transferring data to and from a Device, and de-asserting the CS line. Transactions are atomic, which means they can never be interrupted by another transaction.
Launch Edge	Edge of the clock at which the source register launches the signal onto the line.
Latch Edge	Edge of the clock at which the destination register latches in the signal.

Driver Features

The SPI Master driver governs the communications between Hosts and Devices. The driver supports the following features:

- Multi-threaded environments
- Transparent handling of DMA transfers while reading and writing data
- Automatic time-division multiplexing of data coming from different Devices on the same signal bus, see [SPI Bus Lock](#).

Warning: The SPI Master driver allows multiple Devices to be connected on a same SPI bus (sharing a single ESP32-P4 SPI peripheral). As long as each Device is accessed by only one task, the driver is thread-safe.

However, if multiple tasks try to access the same SPI Device, the driver is **not thread-safe**. In this case, it is recommended to either:

- Refactor your application so that each SPI peripheral is only accessed by a single task at a time. You can use `spi_bus_config_t::isr_cpu_id` to register the SPI ISR to the same core as SPI peripheral-related tasks to ensure thread safety.
- Add a mutex lock around the shared Device using `xSemaphoreCreateMutex`.

SPI Features

SPI Master

SPI Bus Lock To realize the multiplexing of different devices from different drivers, including SPI Master, SPI Flash, etc., an SPI bus lock is applied on each SPI bus. Drivers can attach their devices to the bus with the arbitration of the lock.

Each bus lock is initialized with a BG (background) service registered. All devices that request transactions on the bus should wait until the BG is successfully disabled.

- For the SPI1 bus, the BG is the cache. The bus lock disables the cache before device operations start, and enables it again after the device releases the lock. No devices on SPI1 are allowed to use ISR, since it is meaningless for the task to yield to other tasks when the cache is disabled. The SPI Master driver has not supported SPI1 bus. Only the SPI Flash driver can attach to the bus.
- For other buses, the driver can register the ISR as a BG. If a device task requests exclusive bus access, the bus lock will block the task, disable the ISR, and then unblock the task. After the task releases the lock, the lock will try to re-enable the ISR if there are still pending transactions in the ISR.

SPI Transactions

An SPI bus transaction consists of five phases which can be found in the table below. Any of these phases can be skipped.

Phase	Description
Command	In this phase, a command (0-16 bit) is written to the bus by the Host.
Address	In this phase, an address (0-32 bit) is transmitted over the bus by the Host.
Dummy	This phase is configurable and is used to meet the timing requirements.
Write	Host sends data to a Device. This data follows the optional command and address phases and is indistinguishable from them at the electrical level.
Read	Device sends data to its Host.

The attributes of a transaction are determined by the bus configuration structure `spi_bus_config_t`, Device configuration structure `spi_device_interface_config_t`, and transaction configuration structure `spi_transaction_t`.

An SPI Host can send full-duplex transactions, during which the Read and Write phases occur simultaneously. The total transaction length is determined by the sum of the following members:

- `spi_device_interface_config_t::command_bits`
- `spi_device_interface_config_t::address_bits`
- `spi_transaction_t::length`

While the member `spi_transaction_t::rxlength` only determines the length of data received into the buffer.

In half-duplex transactions, the Read and Write phases are not simultaneous (one direction at a time). The lengths of the Write and Read phases are determined by `spi_transaction_t::length` and `spi_transaction_t::rxlength` respectively.

The Command and Address phases are optional, as not every SPI Device requires a command and/or address. This is reflected in the Device's configuration: if `spi_device_interface_config_t::command_bits` and/or `spi_device_interface_config_t::address_bits` are set to zero, no Command or Address phase will occur.

The Read and Write phases can also be optional, as not every transaction requires both writing and reading data. If `spi_transaction_t::rx_buffer` is NULL and `SPI_TRANS_USE_RXDATA` is not set, the Read phase is skipped. If `spi_transaction_t::tx_buffer` is NULL and `SPI_TRANS_USE_TXDATA` is not set, the Write phase is skipped.

The driver supports two types of transactions: interrupt transactions and polling transactions. The programmer can choose to use a different transaction type per Device. If your Device requires both transaction types, see [Notes on Sending Mixed Transactions to the Same Device](#).

Interrupt Transactions Interrupt transactions blocks the transaction routine until the transaction completes, thus allowing the CPU to run other tasks.

An application task can queue multiple transactions, and the driver automatically handles them one by one in the interrupt service routine (ISR). It allows the task to switch to other procedures until all the transactions are complete.

Polling Transactions Polling transactions do not use interrupts. The routine keeps polling the SPI Host's status bit until the transaction is finished.

All the tasks that use interrupt transactions can be blocked by the queue. At this point, they need to wait for the ISR to run twice before the transaction is finished. Polling transactions save time otherwise spent on queue handling and context switching, which results in smaller transaction duration. The disadvantage is that the CPU is busy while these transactions are in progress.

The `spi_device_polling_end()` routine needs an overhead of at least 1 μ s to unblock other tasks when the transaction is finished. It is strongly recommended to wrap a series of polling transactions using the functions `spi_device_acquire_bus()` and `spi_device_release_bus()` to avoid the overhead. For more information, see [Bus Acquiring](#).

Transaction Line Mode Supported line modes for ESP32-P4 are listed as follows, to make use of these modes, set the member `flags` in the struct `spi_transaction_t` as shown in the Transaction Flag column. If you want to check if corresponding IO pins are set or not, set the member `flags` in the `spi_bus_config_t` as shown in the Bus IO setting Flag column.

Mode name	Command Line Width	Address Line Width	Data Line Width	Transaction Flag	Bus IO Setting Flag
Normal SPI	1	1	1	0	0
Dual Output	1	1	2	<code>SPI_TRANS_MODE_DUAL</code>	<code>SPI_COM-MON_BUSFLAG_DUAL</code>
Dual I/O	1	2	2	<code>SPI_TRANS_MODE_DUAL</code> <code>SPI_TRANS_MULTILINE_CMD</code>	<code>SPI_COM-MON_BUSFLAG_DUAL</code> <code>SPI_COM-MON_BUSFLAG_DUAL</code>
Quad Output	1	1	4	<code>SPI_TRANS_MODE_QUAD</code>	<code>SPI_COM-MON_BUSFLAG_QUAD</code>
Quad I/O	1	4	4	<code>SPI_TRANS_MODE_QUAD</code> <code>SPI_TRANS_MULTILINE_CMD</code>	<code>SPI_COM-MON_BUSFLAG_QUAD</code> <code>SPI_COM-MON_BUSFLAG_QUAD</code>
Octal Output	1	1	8	<code>SPI_TRANS_MODE_OCTAL</code>	<code>SPI_COM-MON_BUSFLAG_OCTAL</code>
OPI	8	8	8	<code>SPI_TRANS_MODE_OCTAL</code> <code>SPI_TRANS_MULTILINE_CMD</code>	<code>SPI_COM-MON_BUSFLAG_OCTAL</code> <code>SPI_COM-MON_BUSFLAG_OCTAL</code>

Command and Address Phases During the Command and Address phases, the members `spi_transaction_t::cmd` and `spi_transaction_t::addr` are sent to the bus, nothing is read at this time. The default lengths of the Command and Address phases are set in `spi_device_interface_config_t` by calling `spi_bus_add_device()`. If the flags `SPI_TRANS_VARIABLE_CMD` and `SPI_TRANS_VARIABLE_ADDR` in the member `spi_transaction_t::flags` are not set, the driver automatically sets the length of these phases to default values during Device initialization.

If the lengths of the Command and Address phases need to be variable, declare the struct `spi_transaction_ext_t`, set the flags `SPI_TRANS_VARIABLE_CMD` and/or `SPI_TRANS_VARIABLE_ADDR` in the member `spi_transaction_ext_t::base` and configure the rest of base as usual. Then the length of each phase will be equal to `spi_transaction_ext_t::command_bits` and `spi_transaction_ext_t::address_bits` set in the struct `spi_transaction_ext_t`.

If the Command and Address phase need to have the same number of lines as the data phase, you need to set `SPI_TRANS_MULTILINE_CMD` and/or `SPI_TRANS_MULTILINE_ADDR` to the flags member in the struct `spi_transaction_t`. Also see *Transaction Line Mode*.

Write and Read Phases Normally, the data that needs to be transferred to or from a Device is read from or written to a chunk of memory indicated by the members `spi_transaction_t::rx_buffer` and `spi_transaction_t::tx_buffer`. If DMA is enabled for transfers, the buffers are required to be:

1. Allocated in DMA-capable internal memory (MALLOC_CAP_DMA), see *DMA-Capable Memory*.
2. 32-bit aligned (starting from a 32-bit boundary and having a length of multiples of 4 bytes).

If these requirements are not satisfied, the transaction efficiency will be affected due to the allocation and copying of temporary buffers.

If using more than one data line to transmit, please set `SPI_DEVICE_HALFDUPLEX` flag for the member flags in the struct `spi_device_interface_config_t`. And the member flags in the struct `spi_transaction_t` should be set as described in *Transaction Line Mode*.

Note: Half-duplex transactions with both Read and Write phases are not supported. Please use full duplex mode.

Bus Acquiring Sometimes you might want to send SPI transactions exclusively and continuously so that it takes as little time as possible. For this, you can use bus acquiring, which helps to suspend transactions (both polling or interrupt) to other Devices until the bus is released. To acquire and release a bus, use the functions `spi_device_acquire_bus()` and `spi_device_release_bus()`.

Driver Usage

- Initialize an SPI bus by calling the function `spi_bus_initialize()`. Make sure to set the correct I/O pins in the struct `spi_bus_config_t`. Set the signals that are not needed to -1.
- Register a Device connected to the bus with the driver by calling the function `spi_bus_add_device()`. Make sure to configure any timing requirements the Device might need with the parameter `dev_config`. You should now have obtained the Device's handle which will be used when sending a transaction to it.
- To interact with the Device, fill one or more `spi_transaction_t` structs with any transaction parameters required. Then send the structs either using a polling transaction or an interrupt transaction:
 - **Interrupt** Either queue all transactions by calling the function `spi_device_queue_trans()` and, at a later time, query the result using the function `spi_device_get_trans_result()`, or handle all requests synchronously by feeding them into `spi_device_transmit()`.
 - **Polling** Call the function `spi_device_polling_transmit()` to send polling transactions. Alternatively, if you want to insert something in between, send the transactions by using `spi_device_polling_start()` and `spi_device_polling_end()`.
- (Optional) To perform back-to-back transactions with a Device, call the function `spi_device_acquire_bus()` before sending transactions and `spi_device_release_bus()` after the transactions have been sent.

- (Optional) To remove a certain Device from the bus, call `sbi_bus_remove_device()` with the Device handle as an argument.
- (Optional) To remove the driver from the bus, make sure no more devices are attached and call `sbi_bus_free()`.

The example code for the SPI Master driver can be found in the [peripherals/sbi_master](#) directory of ESP-IDF examples.

Transactions with Data Not Exceeding 32 Bits When the transaction data size is equal to or less than 32 bits, it will be sub-optimal to allocate a buffer for the data. The data can be directly stored in the transaction struct instead. For transmitted data, it can be achieved by using the `sbi_transaction_t::tx_data` member and setting the `SBI_TRANS_USE_TXDATA` flag on the transmission. For received data, use `sbi_transaction_t::rx_data` and set `SBI_TRANS_USE_RXDATA`. In both cases, do not touch the `sbi_transaction_t::tx_buffer` or `sbi_transaction_t::rx_buffer` members, because they use the same memory locations as `sbi_transaction_t::tx_data` and `sbi_transaction_t::rx_data`.

Transactions with Integers Other than `uint8_t` An SPI Host reads and writes data into memory byte by byte. By default, data is sent with the most significant bit (MSB) first, as LSB is first used in rare cases. If a value of fewer than 8 bits needs to be sent, the bits should be written into memory in the MSB first manner.

For example, if `0b00010` needs to be sent, it should be written into a `uint8_t` variable, and the length for reading should be set to 5 bits. The Device will still receive 8 bits with 3 additional "random" bits, so the reading must be performed correctly.

On top of that, ESP32-P4 is a little-endian chip, which means that the least significant byte of `uint16_t` and `uint32_t` variables is stored at the smallest address. Hence, if `uint16_t` is stored in memory, bits [7:0] are sent first, followed by bits [15:8].

For cases when the data to be transmitted has a size differing from `uint8_t` arrays, the following macros can be used to transform data to the format that can be sent by the SPI driver directly:

- `SBI_SWAP_DATA_TX` for data to be transmitted
- `SBI_SWAP_DATA_RX` for data received

Notes on Sending Mixed Transactions to the Same Device To reduce coding complexity, send only one type of transaction (interrupt or polling) to one Device. However, you still can send both interrupt and polling transactions alternately. The notes below explain how to do this.

The polling transactions should be initiated only after all the polling and interrupt transactions are finished.

Since an unfinished polling transaction blocks other transactions, please do not forget to call the function `sbi_device_polling_end()` after `sbi_device_polling_start()` to allow other transactions or to allow other Devices to use the bus. Remember that if there is no need to switch to other tasks during your polling transaction, you can initiate a transaction with `sbi_device_polling_transmit()` so that it will be ended automatically.

In-flight polling transactions are disturbed by the ISR operation to accommodate interrupt transactions. Always make sure that all the interrupt transactions sent to the ISR are finished before you call `sbi_device_polling_start()`. To do that, you can keep calling `sbi_device_get_trans_result()` until all the transactions are returned.

To have better control of the calling sequence of functions, send mixed transactions to the same Device only within a single task.

GPIO Matrix and IO_MUX Most of the chip's peripheral signals have a direct connection to their dedicated IO_MUX pins. However, the signals can also be routed to any other available pins using the less direct GPIO matrix. If at least one signal is routed through the GPIO matrix, then all signals will be routed through it.

When an SPI Host is set to 40 MHz or lower frequencies, routing SPI pins via the GPIO matrix will behave the same compared to routing them via IOMUX.

The IO_MUX pins for SPI buses are given below.

Pin Name	GPIO Number (SPI2)
CS0 ¹	7
SCLK	9
MISO	10
MOSI	8
QUADWP	11
QUADHD	6

Transfer Speed Considerations

There are three factors limiting the transfer speed:

- Transaction interval
- SPI clock frequency
- Cache miss of SPI functions, including callbacks

The main parameter that determines the transfer speed for large transactions is clock frequency. For multiple small transactions, the transfer speed is mostly determined by the length of transaction intervals.

Transaction Duration Transaction duration includes setting up SPI peripheral registers, copying data to FIFOs or setting up DMA links, and the time for SPI transactions.

Interrupt transactions allow appending extra overhead to accommodate the cost of FreeRTOS queues and the time needed for switching between tasks and the ISR.

For **interrupt transactions**, the CPU can switch to other tasks when a transaction is in progress. This saves CPU time but increases the transaction duration. See [Interrupt Transactions](#). For **polling transactions**, it does not block the task but allows to do polling when the transaction is in progress. For more information, see [Polling Transactions](#).

If DMA is enabled, setting up the linked list requires about 2 µs per transaction. When a master is transferring data, it automatically reads the data from the linked list. If DMA is not enabled, the CPU has to write and read each byte from the FIFO by itself. Usually, this is faster than 2 µs, but the transaction length is limited to 64 bytes for both write and read.

The typical transaction duration for one byte of data is given below.

- Interrupt Transaction via DMA: N/A µs.
- Interrupt Transaction via CPU: N/A µs.
- Polling Transaction via DMA: N/A µs.
- Polling Transaction via CPU: N/A µs.

Note that these data are tested with [CONFIG_SPI_MASTER_ISR_IN_IRAM](#) enabled. SPI transaction related code are placed in the internal memory. If this option is turned off (for example, for internal memory optimization), the transaction duration may be affected.

SPI Clock Frequency The clock source of the GPSPI peripherals can be selected by setting `spi_device_handle_t::cfg::clock_source`. You can refer to [spi_clock_source_t](#) to know the supported clock sources.

By default driver sets `spi_device_handle_t::cfg::clock_source` to `SPI_CLK_SRC_DEFAULT`. This usually stands for the highest frequency among GPSPI clock sources. Its value is different among chips.

¹ Only the first Device attached to the bus can use the CS0 pin.

The actual clock frequency of a Device may not be exactly equal to the number you set, it is re-calculated by the driver to the nearest hardware-compatible number, and not larger than the clock frequency of the clock source. You can call `spi_device_get_actual_freq()` to know the actual frequency computed by the driver.

The theoretical maximum transfer speed of the Write or Read phase can be calculated according to the table below:

Line Width of Write/Read phase	Speed (Bps)
1-Line	<i>SPI Frequency / 8</i>
2-Line	<i>SPI Frequency / 4</i>
4-Line	<i>SPI Frequency / 2</i>
8-Line	<i>SPI Frequency</i>

The transfer speed calculation of other phases (Command, Address, Dummy) is similar.

Cache Missing The default config puts only the ISR into the IRAM. Other SPI-related functions, including the driver itself and the callback, might suffer from cache misses and need to wait until the code is read from flash. Select `CONFIG_SPI_MASTER_IN_IRAM` to put the whole SPI driver into IRAM and put the entire callback(s) and its callee functions into IRAM to prevent cache missing.

Note: SPI driver implementation is based on FreeRTOS APIs, to use `CONFIG_SPI_MASTER_IN_IRAM`, you should not enable `CONFIG_FREERTOS_PLACE_FUNCTIONS_INTO_FLASH`.

For an interrupt transaction, the overall cost is $20+8n/F_{spi}[MHz]$ [μs] for n bytes transferred in one transaction. Hence, the transferring speed is: $n/(20+8n/F_{spi})$. An example of transferring speed at 8 MHz clock speed is given in the following table.

Frequency (MHz)	Transaction Interval (μs)	Transaction Length (bytes)	Total Time (μs)	Total Speed (KBps)
8	25	1	26	38.5
8	25	8	33	242.4
8	25	16	41	490.2
8	25	64	89	719.1
8	25	128	153	836.6

When a transaction length is short, the cost of the transaction interval is high. If possible, try to squash several short transactions into one transaction to achieve a higher transfer speed.

Please note that the ISR is disabled during flash operation by default. To keep sending transactions during flash operations, enable `CONFIG_SPI_MASTER_ISR_IN_IRAM` and set `ESP_INTR_FLAG_IRAM` in the member `spi_bus_config_t::intr_flags`. In this case, all the transactions queued before starting flash operations are handled by the ISR in parallel. Also note that the callback of each Device and their callee functions should be in IRAM, or your callback will crash due to cache missing. For more details, see [IRAM-Safe Interrupt Handlers](#).

Application Example

The code example for using the SPI master half duplex mode to read/write an AT93C46D EEPROM (8-bit mode) can be found in the `peripherals/spi_master/hd_eeprom` directory of ESP-IDF examples.

The code example for using the SPI master full duplex mode to drive a SPI_LCD (e.g. ST7789V or ILI9341) can be found in the `peripherals/spi_master/lcd` directory of ESP-IDF examples.

API Reference - SPI Common

Header File

- `components/hal/include/hal/spi_types.h`

- This header file can be included with:

```
#include "hal/spi_types.h"
```

Structures

struct **spi_line_mode_t**

Line mode of SPI transaction phases: CMD, ADDR, DOUT/DIN.

Public Members

uint8_t **cmd_lines**

The line width of command phase, e.g. 2-line-cmd-phase.

uint8_t **addr_lines**

The line width of address phase, e.g. 1-line-addr-phase.

uint8_t **data_lines**

The line width of data phase, e.g. 4-line-data-phase.

Type Definitions

typedef *soc_periph_spi_clk_src_t* **spi_clock_source_t**

Type of SPI clock source.

Enumerations

enum **spi_host_device_t**

Enum with the three SPI peripherals that are software-accessible in it.

Values:

enumerator **SPI1_HOST**

SPI1.

enumerator **SPI2_HOST**

SPI2.

enumerator **SPI3_HOST**

SPI3.

enumerator **SPI_HOST_MAX**

invalid host value

enum **spi_event_t**

SPI Events.

Values:

enumerator **SPI_EV_BUF_TX**

The buffer has sent data to master.

enumerator **SPI_EV_BUF_RX**

The buffer has received data from master.

enumerator **SPI_EV_SEND_DMA_READY**

Slave has loaded its TX data buffer to the hardware (DMA).

enumerator **SPI_EV_SEND**

Master has received certain number of the data, the number is determined by Master.

enumerator **SPI_EV_RECV_DMA_READY**

Slave has loaded its RX data buffer to the hardware (DMA).

enumerator **SPI_EV_RECV**

Slave has received certain number of data from master, the number is determined by Master.

enumerator **SPI_EV_CMD9**

Received CMD9 from master.

enumerator **SPI_EV_CMDA**

Received CMDA from master.

enumerator **SPI_EV_TRANS**

A transaction has done.

enum **spi_command_t**

SPI command.

Values:

enumerator **SPI_CMD_HD_WRBUF**

enumerator **SPI_CMD_HD_RDBUF**

enumerator **SPI_CMD_HD_WRDMA**

enumerator **SPI_CMD_HD_RDDMA**

enumerator **SPI_CMD_HD_SEG_END**

enumerator **SPI_CMD_HD_EN_QPI**

enumerator **SPI_CMD_HD_WR_END**

enumerator **SPI_CMD_HD_INT0**

enumerator **SPI_CMD_HD_INT1**

enumerator **SPI_CMD_HD_INT2**

enum **spi_sampling_point_t**

SPI master RX sample point mode configuration.

Values:

enumerator **SPI_SAMPLING_POINT_PHASE_0**

Data sampling point at 50% cycle delayed then standard timing, (default).

enumerator **SPI_SAMPLING_POINT_PHASE_1**

Data sampling point follows standard SPI timing in master mode.

Header File

- [components/driver/spi/include/driver/spi_common.h](#)
- This header file can be included with:

```
#include "driver/spi_common.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t **spi_bus_initialize** (*spi_host_device_t* host_id, const *spi_bus_config_t* *bus_config, *spi_dma_chan_t* dma_chan)

Initialize a SPI bus.

Warning: SPI0/1 is not supported

Warning: If a DMA channel is selected, any transmit and receive buffer used should be allocated in DMA-capable memory.

Warning: The ISR of SPI is always executed on the core which calls this function. Never starve the ISR on this core or the SPI transactions will not be handled.

Parameters

- **host_id** -- SPI peripheral that controls this bus
- **bus_config** -- Pointer to a *spi_bus_config_t* struct specifying how the host should be initialized
- **dma_chan** -- - Selecting a DMA channel for an SPI bus allows transactions on the bus with size only limited by the amount of internal memory.
 - Selecting `SPI_DMA_DISABLED` limits the size of transactions.
 - Set to `SPI_DMA_DISABLED` if only the SPI flash uses this bus.
 - Set to `SPI_DMA_CH_AUTO` to let the driver to allocate the DMA channel.

Returns

- `ESP_ERR_INVALID_ARG` if configuration is invalid
- `ESP_ERR_INVALID_STATE` if host already is in use

- ESP_ERR_NOT_FOUND if there is no available DMA channel
- ESP_ERR_NO_MEM if out of memory
- ESP_OK on success

esp_err_t **spi_bus_free** (*spi_host_device_t* host_id)

Free a SPI bus.

Warning: In order for this to succeed, all devices have to be removed first.

Parameters *host_id* -- SPI peripheral to free

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_ERR_INVALID_STATE if bus hasn't been initialized before, or not all devices on the bus are freed
- ESP_OK on success

Structures

struct **spi_bus_config_t**

This is a configuration structure for a SPI bus.

You can use this structure to specify the GPIO pins of the bus. Normally, the driver will use the GPIO matrix to route the signals. An exception is made when all signals either can be routed through the IO_MUX or are -1. In that case, the IO_MUX is used, allowing for >40MHz speeds.

Note: Be advised that the slave driver does not use the quadwp/quadhd lines and fields in *spi_bus_config_t* referring to these lines will be ignored and can thus safely be left uninitialized.

Public Members

int **mosi_io_num**

GPIO pin for Master Out Slave In (=spi_d) signal, or -1 if not used.

int **data0_io_num**

GPIO pin for spi data0 signal in quad/octal mode, or -1 if not used.

int **miso_io_num**

GPIO pin for Master In Slave Out (=spi_q) signal, or -1 if not used.

int **data1_io_num**

GPIO pin for spi data1 signal in quad/octal mode, or -1 if not used.

int **sclk_io_num**

GPIO pin for SPI Clock signal, or -1 if not used.

int **quadwp_io_num**

GPIO pin for WP (Write Protect) signal, or -1 if not used.

int **data2_io_num**

GPIO pin for spi data2 signal in quad/octal mode, or -1 if not used.

int **quadhd_io_num**

GPIO pin for HD (Hold) signal, or -1 if not used.

int **data3_io_num**

GPIO pin for spi data3 signal in quad/octal mode, or -1 if not used.

int **data4_io_num**

GPIO pin for spi data4 signal in octal mode, or -1 if not used.

int **data5_io_num**

GPIO pin for spi data5 signal in octal mode, or -1 if not used.

int **data6_io_num**

GPIO pin for spi data6 signal in octal mode, or -1 if not used.

int **data7_io_num**

GPIO pin for spi data7 signal in octal mode, or -1 if not used.

int **max_transfer_sz**

Maximum transfer size, in bytes. Defaults to 4092 if 0 when DMA enabled, or to SOC_SPI_MAXIMUM_BUFFER_SIZE if DMA is disabled.

uint32_t **flags**

Abilities of bus to be checked by the driver. Or-ed value of SPICOMMON_BUSFLAG_* flags.

esp_intr_cpu_affinity_t **isr_cpu_id**

Select cpu core to register SPI ISR.

int **intr_flags**

Interrupt flag for the bus to set the priority, and IRAM attribute, see `esp_intr_alloc.h`. Note that the EDGE, INTRDISABLED attribute are ignored by the driver. Note that if ESP_INTR_FLAG_IRAM is set, ALL the callbacks of the driver, and their callee functions, should be put in the IRAM.

Macros

SPI_MAX_DMA_LEN

SPI_SWAP_DATA_TX (DATA, LEN)

Transform unsigned integer of length <= 32 bits to the format which can be sent by the SPI driver directly.

E.g. to send 9 bits of data, you can:

```
uint16_t data = SPI_SWAP_DATA_TX(0x145, 9);
```

Then points tx_buffer to &data.

Parameters

- **DATA** -- Data to be sent, can be uint8_t, uint16_t or uint32_t.
- **LEN** -- Length of data to be sent, since the SPI peripheral sends from the MSB, this helps to shift the data to the MSB.

SPI_SWAP_DATA_RX (DATA, LEN)

Transform received data of length <= 32 bits to the format of an unsigned integer.

E.g. to transform the data of 15 bits placed in a 4-byte array to integer:

```
uint16_t data = SPI_SWAP_DATA_RX((uint32_t*)t->rx_data, 15);
```

Parameters

- **DATA** -- Data to be rearranged, can be uint8_t, uint16_t or uint32_t.
- **LEN** -- Length of data received, since the SPI peripheral writes from the MSB, this helps to shift the data to the LSB.

SPICOMMON_BUSFLAG_SLAVE

Initialize I/O in slave mode.

SPICOMMON_BUSFLAG_MASTER

Initialize I/O in master mode.

SPICOMMON_BUSFLAG_IOMUX_PINS

Check using iomux pins. Or indicates the pins are configured through the IO mux rather than GPIO matrix.

SPICOMMON_BUSFLAG_GPIO_PINS

Force the signals to be routed through GPIO matrix. Or indicates the pins are routed through the GPIO matrix.

SPICOMMON_BUSFLAG_SCLK

Check existing of SCLK pin. Or indicates CLK line initialized.

SPICOMMON_BUSFLAG_MISO

Check existing of MISO pin. Or indicates MISO line initialized.

SPICOMMON_BUSFLAG_MOSI

Check existing of MOSI pin. Or indicates MOSI line initialized.

SPICOMMON_BUSFLAG_DUAL

Check MOSI and MISO pins can output. Or indicates bus able to work under DIO mode.

SPICOMMON_BUSFLAG_WPHD

Check existing of WP and HD pins. Or indicates WP & HD pins initialized.

SPICOMMON_BUSFLAG_QUAD

Check existing of MOSI/MISO/WP/HD pins as output. Or indicates bus able to work under QIO mode.

SPICOMMON_BUSFLAG_IO4_IO7

Check existing of IO4~IO7 pins. Or indicates IO4~IO7 pins initialized.

SPICOMMON_BUSFLAG_OCTAL

Check existing of MOSI/MISO/WP/HD/SPIIO4/SPIIO5/SPIIO6/SPIIO7 pins as output. Or indicates bus able to work under octal mode.

SPICOMMON_BUSFLAG_NATIVE_PINS**SPICOMMON_BUSFLAG_SLP_ALLOW_PD**

Allow to power down the peripheral during light sleep, and auto recover then.

Type Definitions

typedef [spi_common_dma_t](#) [spi_dma_chan_t](#)

Enumerations

enum [spi_common_dma_t](#)

SPI DMA channels.

Values:

enumerator [SPI_DMA_DISABLED](#)

Do not enable DMA for SPI.

enumerator [SPI_DMA_CH_AUTO](#)

Enable DMA, channel is automatically selected by driver.

API Reference - SPI Master

Header File

- [components/driver/spi/include/driver/spi_master.h](#)
- This header file can be included with:

```
#include "driver/spi_master.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

[esp_err_t](#) [spi_bus_add_device](#) ([spi_host_device_t](#) host_id, const [spi_device_interface_config_t](#) *dev_config, [spi_device_handle_t](#) *handle)

Allocate a device on a SPI bus.

This initializes the internal structures for a device, plus allocates a CS pin on the indicated SPI master peripheral and routes it to the indicated GPIO. All SPI master devices have three CS pins and can thus control up to three devices.

Note: While in general, speeds up to 80MHz on the dedicated SPI pins and 40MHz on GPIO-matrix-routed pins are supported, full-duplex transfers routed over the GPIO matrix only support speeds up to 26MHz.

Parameters

- **host_id** -- SPI peripheral to allocate device on
- **dev_config** -- SPI interface protocol config for the device
- **handle** -- Pointer to variable to hold the device handle

Returns

- [ESP_ERR_INVALID_ARG](#) if parameter is invalid or configuration combination is not supported (e.g. `dev_config->post_cb` isn't set while flag [SPI_DEVICE_NO_RETURN_RESULT](#) is enabled)

- ESP_ERR_INVALID_STATE if selected clock source is unavailable or spi bus not initialized
- ESP_ERR_NOT_FOUND if host doesn't have any free CS slots
- ESP_ERR_NO_MEM if out of memory
- ESP_OK on success

esp_err_t **spi_bus_remove_device** (*spi_device_handle_t* handle)

Remove a device from the SPI bus.

Parameters **handle** -- Device handle to free

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_ERR_INVALID_STATE if device already is freed
- ESP_OK on success

esp_err_t **spi_device_queue_trans** (*spi_device_handle_t* handle, *spi_transaction_t* *trans_desc, TickType_t ticks_to_wait)

Queue a SPI transaction for interrupt transaction execution. Get the result by `spi_device_get_trans_result`.

Note: Normally a device cannot start (queue) polling and interrupt transactions simultaneously.

Parameters

- **handle** -- Device handle obtained using `spi_host_add_dev`
- **trans_desc** -- Description of transaction to execute
- **ticks_to_wait** -- Ticks to wait until there's room in the queue; use `portMAX_DELAY` to never time out.

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid. This can happen if `SPI_TRANS_CS_KEEP_ACTIVE` flag is specified while the bus was not acquired (`spi_device_acquire_bus()` should be called first) or set flag `SPI_TRANS_DMA_BUFFER_ALIGN_MANUAL` but tx or rx buffer not DMA-capable, or `addr&len` not align to cache line size
- ESP_ERR_TIMEOUT if there was no room in the queue before `ticks_to_wait` expired
- ESP_ERR_NO_MEM if allocating DMA-capable temporary buffer failed
- ESP_ERR_INVALID_STATE if previous transactions are not finished
- ESP_OK on success

esp_err_t **spi_device_get_trans_result** (*spi_device_handle_t* handle, *spi_transaction_t* **trans_desc, TickType_t ticks_to_wait)

Get the result of a SPI transaction queued earlier by `spi_device_queue_trans`.

This routine will wait until a transaction to the given device successfully completed. It will then return the description of the completed transaction so software can inspect the result and e.g. free the memory or re-use the buffers.

Parameters

- **handle** -- Device handle obtained using `spi_host_add_dev`
- **trans_desc** -- Pointer to variable able to contain a pointer to the description of the transaction that is executed. The descriptor should not be modified until the descriptor is returned by `spi_device_get_trans_result`.
- **ticks_to_wait** -- Ticks to wait until there's a returned item; use `portMAX_DELAY` to never time out.

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_ERR_NOT_SUPPORTED if flag `SPI_DEVICE_NO_RETURN_RESULT` is set
- ESP_ERR_TIMEOUT if there was no completed transaction before `ticks_to_wait` expired
- ESP_OK on success

esp_err_t **spi_device_transmit** (*spi_device_handle_t* handle, *spi_transaction_t* *trans_desc)

Send a SPI transaction, wait for it to complete, and return the result.

This function is the equivalent of calling `spi_device_queue_trans()` followed by `spi_device_get_trans_result()`. Do not use this when there is still a transaction separately queued (started) from `spi_device_queue_trans()` or `polling_start/transmit` that hasn't been finalized.

Note: This function is not thread safe when multiple tasks access the same SPI device. Normally a device cannot start (queue) polling and interrupt transactions simultaneously.

Parameters

- **handle** -- Device handle obtained using `spi_host_add_dev`
- **trans_desc** -- Description of transaction to execute

Returns

- `ESP_ERR_INVALID_ARG` if parameter is invalid
- `ESP_OK` on success

esp_err_t **spi_device_polling_start** (*spi_device_handle_t* handle, *spi_transaction_t* *trans_desc, TickType_t ticks_to_wait)

Immediately start a polling transaction.

Note: Normally a device cannot start (queue) polling and interrupt transactions simultaneously. Moreover, a device cannot start a new polling transaction if another polling transaction is not finished.

Parameters

- **handle** -- Device handle obtained using `spi_host_add_dev`
- **trans_desc** -- Description of transaction to execute
- **ticks_to_wait** -- Ticks to wait until there's room in the queue; currently only `portMAX_DELAY` is supported.

Returns

- `ESP_ERR_INVALID_ARG` if parameter is invalid. This can happen if `SPI_TRANS_CS_KEEP_ACTIVE` flag is specified while the bus was not acquired (`spi_device_acquire_bus()` should be called first) or set flag `SPI_TRANS_DMA_BUFFER_ALIGN_MANUAL` but tx or rx buffer not DMA-capable, or `addr&len` not align to cache line size
- `ESP_ERR_TIMEOUT` if the device cannot get control of the bus before `ticks_to_wait` expired
- `ESP_ERR_NO_MEM` if allocating DMA-capable temporary buffer failed
- `ESP_ERR_INVALID_STATE` if previous transactions are not finished
- `ESP_OK` on success

esp_err_t **spi_device_polling_end** (*spi_device_handle_t* handle, TickType_t ticks_to_wait)

Poll until the polling transaction ends.

This routine will not return until the transaction to the given device has successfully completed. The task is not blocked, but actively busy-spins for the transaction to be completed.

Parameters

- **handle** -- Device handle obtained using `spi_host_add_dev`
- **ticks_to_wait** -- Ticks to wait until there's a returned item; use `portMAX_DELAY` to never time out.

Returns

- `ESP_ERR_INVALID_ARG` if parameter is invalid
- `ESP_ERR_TIMEOUT` if the transaction cannot finish before `ticks_to_wait` expired
- `ESP_OK` on success

esp_err_t **spi_device_polling_transmit** (*spi_device_handle_t* handle, *spi_transaction_t* *trans_desc)

Send a polling transaction, wait for it to complete, and return the result.

This function is the equivalent of calling `spi_device_polling_start()` followed by `spi_device_polling_end()`. Do not use this when there is still a transaction that hasn't been finalized.

Note: This function is not thread safe when multiple tasks access the same SPI device. Normally a device cannot start (queue) polling and interrupt transactions simultaneously.

Parameters

- **handle** -- Device handle obtained using `spi_host_add_dev`
- **trans_desc** -- Description of transaction to execute

Returns

- `ESP_ERR_INVALID_ARG` if parameter is invalid
- `ESP_ERR_TIMEOUT` if the device cannot get control of the bus
- `ESP_ERR_NO_MEM` if allocating DMA-capable temporary buffer failed
- `ESP_ERR_INVALID_STATE` if previous transactions of same device are not finished
- `ESP_OK` on success

esp_err_t **spi_device_acquire_bus** (*spi_device_handle_t* device, `TickType_t` wait)

Occupy the SPI bus for a device to do continuous transactions.

Transactions to all other devices will be put off until `spi_device_release_bus` is called.

Note: The function will wait until all the existing transactions have been sent.

Parameters

- **device** -- The device to occupy the bus.
- **wait** -- Time to wait before the the bus is occupied by the device. Currently MUST set to `portMAX_DELAY`.

Returns

- `ESP_ERR_INVALID_ARG` : `wait` is not set to `portMAX_DELAY`.
- `ESP_OK` : Success.

void **spi_device_release_bus** (*spi_device_handle_t* dev)

Release the SPI bus occupied by the device. All other devices can start sending transactions.

Parameters **dev** -- The device to release the bus.

esp_err_t **spi_device_get_actual_freq** (*spi_device_handle_t* handle, int *freq_khz)

Calculate working frequency for specific device.

Parameters

- **handle** -- SPI device handle
- **freq_khz** -- [out] output parameter to hold calculated frequency in kHz

Returns

- `ESP_ERR_INVALID_ARG` : `handle` or `freq_khz` parameter is NULL
- `ESP_OK` : Success

int **spi_get_actual_clock** (int fapb, int hz, int duty_cycle)

Calculate the working frequency that is most close to desired frequency.

Parameters

- **fapb** -- The frequency of apb clock, should be `APB_CLK_FREQ`.
- **hz** -- Desired working frequency
- **duty_cycle** -- Duty cycle of the spi clock

Returns Actual working frequency that most fit.

```
void spi_get_timing (bool gpio_is_used, int input_delay_ns, int eff_clk, int *dummy_o, int *cycles_remain_o)
```

Calculate the timing settings of specified frequency and settings.

Note: If ****dummy_o*** is not zero, it means dummy bits should be applied in half duplex mode, and full duplex mode may not work.

Parameters

- **gpio_is_used** -- True if using GPIO matrix, or False if iomux pins are used.
- **input_delay_ns** -- Input delay from SCLK launch edge to MISO data valid.
- **eff_clk** -- Effective clock frequency (in Hz) from `spi_get_actual_clock()`.
- **dummy_o** -- Address of dummy bits used output. Set to NULL if not needed.
- **cycles_remain_o** -- Address of cycles remaining (after dummy bits are used) output.
 - -1 If too many cycles remaining, suggest to compensate half a clock.
 - 0 If no remaining cycles or dummy bits are not used.
 - positive value: cycles suggest to compensate.

```
int spi_get_freq_limit (bool gpio_is_used, int input_delay_ns)
```

Get the frequency limit of current configurations. SPI master working at this limit is OK, while above the limit, full duplex mode and DMA will not work, and dummy bits will be applied in the half duplex mode.

Parameters

- **gpio_is_used** -- True if using GPIO matrix, or False if native pins are used.
- **input_delay_ns** -- Input delay from SCLK launch edge to MISO data valid.

Returns Frequency limit of current configurations.

```
esp_err_t spi_bus_get_max_transaction_len (spi_host_device_t host_id, size_t *max_bytes)
```

Get max length (in bytes) of one transaction.

Parameters

- **host_id** -- SPI peripheral
- **max_bytes** -- **[out]** Max length of one transaction, in bytes

Returns

- ESP_OK: On success
- ESP_ERR_INVALID_ARG: Invalid argument

Structures

```
struct spi_device_interface_config_t
```

This is a configuration for a SPI slave device that is connected to one of the SPI buses.

Public Members

`uint8_t` **command_bits**

Default amount of bits in command phase (0-16), used when `SPI_TRANS_VARIABLE_CMD` is not used, otherwise ignored.

`uint8_t` **address_bits**

Default amount of bits in address phase (0-64), used when `SPI_TRANS_VARIABLE_ADDR` is not used, otherwise ignored.

`uint8_t` **dummy_bits**

Amount of dummy bits to insert between address and data phase.

uint8_t mode

SPI mode, representing a pair of (CPOL, CPHA) configuration:

- 0: (0, 0)
- 1: (0, 1)
- 2: (1, 0)
- 3: (1, 1)

***spi_clock_source_t* clock_source**

Select SPI clock source, `SPI_CLK_SRC_DEFAULT` by default.

uint16_t duty_cycle_pos

Duty cycle of positive clock, in 1/256th increments (128 = 50%/50% duty). Setting this to 0 (=not setting it) is equivalent to setting this to 128.

uint16_t cs_ena_pretrans

Amount of SPI bit-cycles the cs should be activated before the transmission (0-16). This only works on half-duplex transactions.

uint8_t cs_ena_posttrans

Amount of SPI bit-cycles the cs should stay active after the transmission (0-16)

int clock_speed_hz

SPI clock speed in Hz. Derived from `clock_source`.

int input_delay_ns

Maximum data valid time of slave. The time required between SCLK and MISO valid, including the possible clock delay from slave to master. The driver uses this value to give an extra delay before the MISO is ready on the line. Leave at 0 unless you know you need a delay. For better timing performance at high frequency (over 8MHz), it's suggest to have the right value.

***spi_sampling_point_t* sample_point**

Sample point tuning of spi master receiving bit.

int spics_io_num

CS GPIO pin for this device, or -1 if not used.

uint32_t flags

Bitwise OR of `SPI_DEVICE_*` flags.

int queue_size

Transaction queue size. This sets how many transactions can be 'in the air' (queued using `spi_device_queue_trans` but not yet finished using `spi_device_get_trans_result`) at the same time.

***transaction_cb_t* pre_cb**

Callback to be called before a transmission is started.

This callback is called within interrupt context should be in IRAM for best performance, see "Transferring Speed" section in the SPI Master documentation for full details. If not, the callback may crash during flash operation when the driver is initialized with `ESP_INTR_FLAG_IRAM`.

***transaction_cb_t* post_cb**

Callback to be called after a transmission has completed.

This callback is called within interrupt context should be in IRAM for best performance, see "Transferring Speed" section in the SPI Master documentation for full details. If not, the callback may crash during flash operation when the driver is initialized with ESP_INTR_FLAG_IRAM.

struct spi_transaction_t

This structure describes one SPI transaction. The descriptor should not be modified until the transaction finishes.

Public Members**uint32_t flags**

Bitwise OR of SPI_TRANS_* flags.

uint16_t cmd

Command data, of which the length is set in the `command_bits` of *spi_device_interface_config_t*.

NOTE: this field, used to be "command" in ESP-IDF 2.1 and before, is re-written to be used in a new way in ESP-IDF 3.0.

Example: write 0x0123 and `command_bits=12` to send command 0x12, 0x3_ (in previous version, you may have to write 0x3_12).

uint64_t addr

Address data, of which the length is set in the `address_bits` of *spi_device_interface_config_t*.

NOTE: this field, used to be "address" in ESP-IDF 2.1 and before, is re-written to be used in a new way in ESP-IDF3.0.

Example: write 0x123400 and `address_bits=24` to send address of 0x12, 0x34, 0x00 (in previous version, you may have to write 0x12340000).

size_t length

Total data length, in bits.

size_t rxlength

Total data length received, should be not greater than `length` in full-duplex mode (0 defaults this to the value of `length`).

void *user

User-defined variable. Can be used to store eg transaction ID.

const void *tx_buffer

Pointer to transmit buffer, or NULL for no MOSI phase.

uint8_t tx_data[4]

If SPI_TRANS_USE_TXDATA is set, data set here is sent directly from this variable.

void *rx_buffer

Pointer to receive buffer, or NULL for no MISO phase. Written by 4 bytes-unit if DMA is used.

uint8_t **rx_data**[4]

If SPI_TRANS_USE_RXDATA is set, data is received directly to this variable.

struct **spi_transaction_ext_t**

This struct is for SPI transactions which may change their address and command length. Please do set the flags in base to SPI_TRANS_VARIABLE_CMD_ADR to use the bit length here.

Public Members

struct *spi_transaction_t* **base**

Transaction data, so that pointer to *spi_transaction_t* can be converted into *spi_transaction_ext_t*.

uint8_t **command_bits**

The command length in this transaction, in bits.

uint8_t **address_bits**

The address length in this transaction, in bits.

uint8_t **dummy_bits**

The dummy length in this transaction, in bits.

Macros

SPI_MASTER_FREQ_8M

SPI common used frequency (in Hz)

Note: SPI peripheral only has an integer divider, and the default clock source can be different on other targets, so the actual frequency may be slightly different from the desired frequency. 8MHz

SPI_MASTER_FREQ_9M

8.89MHz

SPI_MASTER_FREQ_10M

10MHz

SPI_MASTER_FREQ_11M

11.43MHz

SPI_MASTER_FREQ_13M

13.33MHz

SPI_MASTER_FREQ_16M

16MHz

SPI_MASTER_FREQ_20M

20MHz

SPI_MASTER_FREQ_26M

26.67MHz

SPI_MASTER_FREQ_40M

40MHz

SPI_MASTER_FREQ_80M

80MHz

SPI_DEVICE_TXBIT_LSBFIRST

Transmit command/address/data LSB first instead of the default MSB first.

SPI_DEVICE_RXBIT_LSBFIRST

Receive data LSB first instead of the default MSB first.

SPI_DEVICE_BIT_LSBFIRST

Transmit and receive LSB first.

SPI_DEVICE_3WIRE

Use MOSI (=spid) for both sending and receiving data.

SPI_DEVICE_POSITIVE_CS

Make CS positive during a transaction instead of negative.

SPI_DEVICE_HALFDUPLEX

Transmit data before receiving it, instead of simultaneously.

SPI_DEVICE_CLK_AS_CS

Output clock on CS line if CS is active.

SPI_DEVICE_NO_DUMMY

There are timing issue when reading at high frequency (the frequency is related to whether iomux pins are used, valid time after slave sees the clock).

- In half-duplex mode, the driver automatically inserts dummy bits before reading phase to fix the timing issue. Set this flag to disable this feature.
- In full-duplex mode, however, the hardware cannot use dummy bits, so there is no way to prevent data being read from getting corrupted. Set this flag to confirm that you're going to work with output only, or read without dummy bits at your own risk.

SPI_DEVICE_DDRCLK**SPI_DEVICE_NO_RETURN_RESULT**

Don't return the descriptor to the host on completion (use post_cb to notify instead)

SPI_TRANS_MODE_DIO

Transmit/receive data in 2-bit mode.

SPI_TRANS_MODE_QIO

Transmit/receive data in 4-bit mode.

SPI_TRANS_USE_RXDATA

Receive into rx_data member of *spi_transaction_t* instead into memory at rx_buffer.

SPI_TRANS_USE_TXDATA

Transmit tx_data member of *spi_transaction_t* instead of data at tx_buffer. Do not set tx_buffer when using this.

SPI_TRANS_MODE_DIOQIO_ADDR

Also transmit address in mode selected by SPI_MODE_DIO/SPI_MODE_QIO.

SPI_TRANS_VARIABLE_CMD

Use the command_bits in *spi_transaction_ext_t* rather than default value in *spi_device_interface_config_t*.

SPI_TRANS_VARIABLE_ADDR

Use the address_bits in *spi_transaction_ext_t* rather than default value in *spi_device_interface_config_t*.

SPI_TRANS_VARIABLE_DUMMY

Use the dummy_bits in *spi_transaction_ext_t* rather than default value in *spi_device_interface_config_t*.

SPI_TRANS_CS_KEEP_ACTIVE

Keep CS active after data transfer.

SPI_TRANS_MULTILINE_CMD

The data lines used at command phase is the same as data phase (otherwise, only one data line is used at command phase)

SPI_TRANS_MODE_OCT

Transmit/receive data in 8-bit mode.

SPI_TRANS_MULTILINE_ADDR

The data lines used at address phase is the same as data phase (otherwise, only one data line is used at address phase)

SPI_TRANS_DMA_BUFFER_ALIGN_MANUAL

By default driver will automatically re-alloc dma buffer if it doesn't meet hardware alignment or dma_capable requirements, this flag is for you to disable this feature, you will need to take care of the alignment otherwise driver will return you error ESP_ERR_INVALID_ARG.

Type Definitions

```
typedef void (*transaction_cb_t)(spi_transaction_t *trans)
```

```
typedef struct spi_device_t *spi_device_handle_t
```

Handle for a device on a SPI bus.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.5.22 SPI Slave Driver

SPI Slave driver is a program that controls ESP32-P4's General Purpose SPI (GP-SPI) peripheral(s) when it functions as a slave.

For more hardware information about the GP-SPI peripheral(s), see **ESP32-P4 Technical Reference Manual > SPI Controller** [PDF].

Terminology

The terms used in relation to the SPI slave driver are given in the table below.

Term	Definition
Host	The SPI controller peripheral external to ESP32-P4 that initiates SPI transmissions over the bus, and acts as an SPI Master.
Device	SPI slave device (general purpose SPI controller). Each Device shares the MOSI, MISO and SCLK signals but is only active on the bus when the Host asserts the Device's individual CS line.
Bus	A signal bus, common to all Devices connected to one Host. In general, a bus includes the following lines: MISO, MOSI, SCLK, one or more CS lines, and, optionally, QUADWP and QUADHD. So Devices are connected to the same lines, with the exception that each Device has its own CS line. Several Devices can also share one CS line if connected in the daisy-chain manner.
MISO	Master In, Slave Out, a.k.a. Q. Data transmission from a Device to Host.
MOSI	Master Out, Slave In, a.k.a. D. Data transmission from a Host to Device.
SCLK	Serial Clock. Oscillating signal generated by a Host that keeps the transmission of data bits in sync.
CS	Chip Select. Allows a Host to select individual Device(s) connected to the bus in order to send or receive data.
QUADWP	Write Protect signal. Only used for 4-bit (qio/qout) transactions.
QUADHD	Hold signal. Only used for 4-bit (qio/qout) transactions.
Assertion	The action of activating a line. The opposite action of returning the line back to inactive (back to idle) is called de-assertion .
Transaction	One instance of a Host asserting a CS line, transferring data to and from a Device, and de-asserting the CS line. Transactions are atomic, which means they can never be interrupted by another transaction.
Launch Edge	Edge of the clock at which the source register launches the signal onto the line.
Latch Edge	Edge of the clock at which the destination register latches in the signal.

Driver Features

The SPI slave driver allows using the SPI peripherals as full-duplex Devices. The driver can send/receive transactions up to 64 bytes in length, or utilize DMA to send/receive longer transactions. However, there are some *known issues* related to DMA.

The SPI slave driver supports registering the SPI ISR to a certain CPU core. If multiple tasks try to access the same SPI Device simultaneously, it is recommended that your application be refactored so that each SPI peripheral is only accessed by a single task at a time. Please also use `spi_bus_config_t::isr_cpu_id` to register the SPI ISR to the same core as SPI peripheral related tasks to ensure thread safety.

SPI Transactions

A full-duplex SPI transaction begins when the Host asserts the CS line and starts sending out clock pulses on the SCLK line. Every clock pulse, a data bit is shifted from the Host to the Device on the MOSI line and back on the

MISO line at the same time. At the end of the transaction, the Host de-asserts the CS line.

The attributes of a transaction are determined by the configuration structure for an SPI peripheral acting as a slave device `spi_slave_interface_config_t`, and transaction configuration structure `spi_slave_transaction_t`.

As not every transaction requires both writing and reading data, you can choose to configure the `spi_transaction_t` structure for TX only, RX only, or TX and RX transactions. If `spi_slave_transaction_t::rx_buffer` is set to NULL, the read phase will be skipped. Similarly, if `spi_slave_transaction_t::tx_buffer` is set to NULL, the write phase will be skipped.

Note: A Host should not start a transaction before its Device is ready for receiving data. It is recommended to use another GPIO pin for a handshake signal to sync the Devices. For more details, see [Transaction Interval](#).

Driver Usage

- Initialize an SPI peripheral as a Device by calling the function `spi_slave_initialize()`. Make sure to set the correct I/O pins in the struct `bus_config`. Set the unused signals to -1.
- Before initiating transactions, fill one or more `spi_slave_transaction_t` structs with the transaction parameters required. Either queue all transactions by calling the function `spi_slave_queue_trans()` and, at a later time, query the result by using the function `spi_slave_get_trans_result()`, or handle all requests individually by feeding them into `spi_slave_transmit()`. The latter two functions will be blocked until the Host has initiated and finished a transaction, causing the queued data to be sent and received.
- (Optional) To unload the SPI slave driver, call `spi_slave_free()`.

Transaction Data and Master/Slave Length Mismatches

Normally, the data that needs to be transferred to or from a Device is read or written to a chunk of memory indicated by the `spi_slave_transaction_t::rx_buffer` and `spi_slave_transaction_t::tx_buffer`. The SPI driver can be configured to use DMA for transfers, in which case these buffers must be allocated in DMA-capable memory using `pvPortMallocCaps(size, MALLOC_CAP_DMA)`.

The amount of data that the driver can read or write to the buffers is limited by `spi_slave_transaction_t::length`. However, this member does not define the actual length of an SPI transaction. A transaction's length is determined by the clock and CS lines driven by the Host. The actual length of the transmission can be read only after a transaction is finished from the member `spi_slave_transaction_t::trans_len`.

If the length of the transmission is greater than the buffer length, only the initial number of bits specified in the `spi_slave_transaction_t::length` member will be sent and received. In this case, `spi_slave_transaction_t::trans_len` is set to `spi_slave_transaction_t::length` instead of the actual transaction length. To meet the actual transaction length requirements, set `spi_slave_transaction_t::length` to a value greater than the maximum `spi_slave_transaction_t::trans_len` expected. If the transmission length is shorter than the buffer length, only the data equal to the length of the buffer will be transmitted.

GPIO Matrix and IO_MUX Most of chip's peripheral signals have direct connection to their dedicated IO_MUX pins. However, the signals can also be routed to any other available pins using the less direct GPIO matrix. If at least one signal is routed through the GPIO matrix, then all signals will be routed through it.

When an SPI Host is set to 80 MHz or lower frequencies, routing SPI pins via GPIO matrix will behave the same compared to routing them via IO_MUX.

The IO_MUX pins for SPI buses are given below.

Pin Name	GPIO Number (SPI2)
CS0	N/A
SCLK	N/A
MISO	N/A
MOSI	N/A
QUADWP	N/A
QUADHD	N/A

Speed and Timing Considerations

Transaction Interval The ESP32-P4 SPI slave peripherals are designed as general purpose Devices controlled by a CPU. As opposed to dedicated slaves, CPU-based SPI Devices have a limited number of pre-defined registers. All transactions must be handled by the CPU, which means that the transfers and responses are not real-time, and there might be noticeable latency.

As a solution, a Device's response rate can be doubled by using the functions `spi_slave_queue_trans()` and then `spi_slave_get_trans_result()` instead of using `spi_slave_transmit()`.

You can also configure a GPIO pin through which the Device will signal to the Host when it is ready for a new transaction. A code example of this can be found in [peripherals/spi_slave](#).

SCLK Frequency Requirements The SPI slaves are designed to operate at up to 60 MHz. The data cannot be recognized or received correctly if the clock is too fast or does not have a 50% duty cycle.

Restrictions and Known Issues

1. If DMA is enabled, the rx buffer should be word-aligned (starting from a 32-bit boundary and having a length of multiples of 4 bytes). Otherwise, DMA may write incorrectly or not in a boundary aligned manner. The driver reports an error if this condition is not satisfied.
Also, a Host should write lengths that are multiples of 4 bytes. The data with inappropriate lengths will be discarded.

Application Example

The code example for Device/Host communication can be found in the [peripherals/spi_slave](#) directory of ESP-IDF examples.

API Reference

Header File

- [components/driver/spi/include/driver/spi_slave.h](#)
- This header file can be included with:

```
#include "driver/spi_slave.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

esp_err_t spi_slave_initialize (*spi_host_device_t* host, const *spi_bus_config_t* *bus_config, const *spi_slave_interface_config_t* *slave_config, *spi_dma_chan_t* dma_chan)

Initialize a SPI bus as a slave interface.

Warning: SPI0/1 is not supported

Warning: If a DMA channel is selected, any transmit and receive buffer used should be allocated in DMA-capable memory.

Warning: The ISR of SPI is always executed on the core which calls this function. Never starve the ISR on this core or the SPI transactions will not be handled.

Parameters

- **host** -- SPI peripheral to use as a SPI slave interface
- **bus_config** -- Pointer to a *spi_bus_config_t* struct specifying how the host should be initialized
- **slave_config** -- Pointer to a *spi_slave_interface_config_t* struct specifying the details for the slave interface
- **dma_chan** -- - Selecting a DMA channel for an SPI bus allows transactions on the bus with size only limited by the amount of internal memory.
 - Selecting SPI_DMA_DISABLED limits the size of transactions.
 - Set to SPI_DMA_DISABLED if only the SPI flash uses this bus.
 - Set to SPI_DMA_CH_AUTO to let the driver to allocate the DMA channel.

Returns

- ESP_ERR_INVALID_ARG if configuration is invalid
- ESP_ERR_INVALID_STATE if host already is in use
- ESP_ERR_NOT_FOUND if there is no available DMA channel
- ESP_ERR_NO_MEM if out of memory
- ESP_OK on success

esp_err_t spi_slave_free (*spi_host_device_t* host)

Free a SPI bus claimed as a SPI slave interface.

Parameters **host** -- SPI peripheral to free

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_ERR_INVALID_STATE if not all devices on the bus are freed
- ESP_OK on success

esp_err_t spi_slave_queue_trans (*spi_host_device_t* host, const *spi_slave_transaction_t* *trans_desc, TickType_t ticks_to_wait)

Queue a SPI transaction for execution.

Queues a SPI transaction to be executed by this slave device. (The transaction queue size was specified when the slave device was initialised via spi_slave_initialize.) This function may block if the queue is full (depending on the ticks_to_wait parameter). No SPI operation is directly initiated by this function, the next queued transaction will happen when the master initiates a SPI transaction by pulling down CS and sending out clock signals.

This function hands over ownership of the buffers in trans_desc to the SPI slave driver; the application is not to access this memory until spi_slave_queue_trans is called to hand ownership back to the application.

Parameters

- **host** -- SPI peripheral that is acting as a slave

- **trans_desc** -- Description of transaction to execute. Not const because we may want to write status back into the transaction description.
- **ticks_to_wait** -- Ticks to wait until there's room in the queue; use portMAX_DELAY to never time out.

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_OK on success

esp_err_t **spi_slave_get_trans_result** (*spi_host_device_t* host, *spi_slave_transaction_t* **trans_desc, TickType_t ticks_to_wait)

Get the result of a SPI transaction queued earlier.

This routine will wait until a transaction to the given device (queued earlier with `spi_slave_queue_trans`) has successfully completed. It will then return the description of the completed transaction so software can inspect the result and e.g. free the memory or re-use the buffers.

It is mandatory to eventually use this function for any transaction queued by `spi_slave_queue_trans`.

Parameters

- **host** -- SPI peripheral to that is acting as a slave
- **trans_desc** -- [out] Pointer to variable able to contain a pointer to the description of the transaction that is executed
- **ticks_to_wait** -- Ticks to wait until there's a returned item; use portMAX_DELAY to never time out.

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_ERR_NOT_SUPPORTED if flag `SPI_SLAVE_NO_RETURN_RESULT` is set
- ESP_OK on success

esp_err_t **spi_slave_transmit** (*spi_host_device_t* host, *spi_slave_transaction_t* *trans_desc, TickType_t ticks_to_wait)

Do a SPI transaction.

Essentially does the same as `spi_slave_queue_trans` followed by `spi_slave_get_trans_result`. Do not use this when there is still a transaction queued that hasn't been finalized using `spi_slave_get_trans_result`.

Parameters

- **host** -- SPI peripheral to that is acting as a slave
- **trans_desc** -- Pointer to variable able to contain a pointer to the description of the transaction that is executed. Not const because we may want to write status back into the transaction description.
- **ticks_to_wait** -- Ticks to wait until there's a returned item; use portMAX_DELAY to never time out.

Returns

- ESP_ERR_INVALID_ARG if parameter is invalid
- ESP_OK on success

Structures

struct **spi_slave_interface_config_t**

This is a configuration for a SPI host acting as a slave device.

Public Members

int **spics_io_num**

CS GPIO pin for this device.

uint32_t **flags**

Bitwise OR of SPI_SLAVE_* flags.

int **queue_size**

Transaction queue size. This sets how many transactions can be 'in the air' (queued using spi_slave_queue_trans but not yet finished using spi_slave_get_trans_result) at the same time.

uint8_t **mode**

SPI mode, representing a pair of (CPOL, CPHA) configuration:

- 0: (0, 0)
- 1: (0, 1)
- 2: (1, 0)
- 3: (1, 1)

slave_transaction_cb_t **post_setup_cb**

Callback called after the SPI registers are loaded with new data.

This callback is called within interrupt context should be in IRAM for best performance, see "Transferring Speed" section in the SPI Master documentation for full details. If not, the callback may crash during flash operation when the driver is initialized with ESP_INTR_FLAG_IRAM.

slave_transaction_cb_t **post_trans_cb**

Callback called after a transaction is done.

This callback is called within interrupt context should be in IRAM for best performance, see "Transferring Speed" section in the SPI Master documentation for full details. If not, the callback may crash during flash operation when the driver is initialized with ESP_INTR_FLAG_IRAM.

struct **spi_slave_transaction_t**

This structure describes one SPI transaction

Public Members

size_t **length**

Total data length, in bits.

size_t **trans_len**

Transaction data length, in bits.

const void ***tx_buffer**

Pointer to transmit buffer, or NULL for no MOSI phase.

void ***rx_buffer**

Pointer to receive buffer, or NULL for no MISO phase. When the DMA is enabled, must start at WORD boundary (`rx_buffer%4==0`), and has length of a multiple of 4 bytes.

void ***user**

User-defined variable. Can be used to store eg transaction ID.

Macros

SPI_SLAVE_TXBIT_LSBFIRST

Transmit command/address/data LSB first instead of the default MSB first.

SPI_SLAVE_RXBIT_LSBFIRST

Receive data LSB first instead of the default MSB first.

SPI_SLAVE_BIT_LSBFIRST

Transmit and receive LSB first.

SPI_SLAVE_NO_RETURN_RESULT

Don't return the descriptor to the host on completion (use `post_trans_cb` to notify instead)

Type Definitions

```
typedef void (*slave_transaction_cb_t)(spi_slave_transaction_t *trans)
```

2.5.23 Universal Asynchronous Receiver/Transmitter (UART)

Introduction

A Universal Asynchronous Receiver/Transmitter (UART) is a hardware feature that handles communication (i.e., timing requirements and data framing) using widely-adopted asynchronous serial communication interfaces, such as RS232, RS422, and RS485. A UART provides a widely adopted and cheap method to realize full-duplex or half-duplex data exchange among different devices.

The ESP32-P4 chip has 5 UART controllers (also referred to as port), each featuring an identical set of registers to simplify programming and for more flexibility.

Each UART controller is independently configurable with parameters such as baud rate, data bit length, bit ordering, number of stop bits, parity bit, etc. All the regular UART controllers are compatible with UART-enabled devices from various manufacturers and can also support Infrared Data Association (IrDA) protocols.

Additionally, the ESP32-P4 chip has one low-power (LP) UART controller. It is the cut-down version of regular UART. Usually, the LP UART controller only support basic UART functionality with a much smaller RAM size, and does not support IrDA or RS485 protocols. For a full list of difference between UART and LP UART, please refer to the **ESP32-P4 Technical Reference Manual > UART Controller (UART) > Features** [PDF].

Functional Overview

The overview describes how to establish communication between an ESP32-P4 and other UART devices using the functions and data types of the UART driver. A typical programming workflow is broken down into the sections provided below:

1. [Set Communication Parameters](#) - Setting baud rate, data bits, stop bits, etc.
2. [Set Communication Pins](#) - Assigning pins for connection to a device
3. [Install Drivers](#) - Allocating ESP32-P4's resources for the UART driver
4. [Run UART Communication](#) - Sending/receiving data
5. [Use Interrupts](#) - Triggering interrupts on specific communication events
6. [Deleting a Driver](#) - Freeing allocated resources if a UART communication is no longer required

Steps 1 to 3 comprise the configuration stage. Step 4 is where the UART starts operating. Steps 5 and 6 are optional.

Additionally, when using the LP UART Controller you need to pay attention to [Use LP UART Controller with HP Core](#).

The UART driver's functions identify each of the UART controllers using `uart_port_t`. This identification is needed for all the following function calls.

Set Communication Parameters UART communication parameters can be configured all in a single step or individually in multiple steps.

Single Step Call the function `uart_param_config()` and pass to it a `uart_config_t` structure. The `uart_config_t` structure should contain all the required parameters. See the example below.

```
const uart_port_t uart_num = UART_NUM_1;
uart_config_t uart_config = {
    .baud_rate = 115200,
    .data_bits = UART_DATA_8_BITS,
    .parity = UART_PARITY_DISABLE,
    .stop_bits = UART_STOP_BITS_1,
    .flow_ctrl = UART_HW_FLOWCTRL_CTS_RTS,
    .rx_flow_ctrl_thresh = 122,
};
// Configure UART parameters
ESP_ERROR_CHECK(uart_param_config(uart_num, &uart_config));
```

For more information on how to configure the hardware flow control options, please refer to [peripherals/uart/uart_echo](#).

Multiple Steps Configure specific parameters individually by calling a dedicated function from the table given below. These functions are also useful if re-configuring a single parameter.

Table 4: Functions for Configuring specific parameters individually

Parameter to Configure	Function
Baud rate	<code>uart_set_baudrate()</code>
Number of transmitted bits	<code>uart_set_word_length()</code> selected out of <code>uart_word_length_t</code>
Parity control	<code>uart_set_parity()</code> selected out of <code>uart_parity_t</code>
Number of stop bits	<code>uart_set_stop_bits()</code> selected out of <code>uart_stop_bits_t</code>
Hardware flow control mode	<code>uart_set_hw_flow_ctrl()</code> selected out of <code>uart_hw_flowcontrol_t</code>
Communication mode	<code>uart_set_mode()</code> selected out of <code>uart_mode_t</code>

Each of the above functions has a `_get_` counterpart to check the currently set value. For example, to check the current baud rate value, call `uart_get_baudrate()`.

Set Communication Pins After setting communication parameters, configure the physical GPIO pins to which the other UART device will be connected. For this, call the function `uart_set_pin()` and specify the GPIO pin numbers to which the driver should route the TX, RX, RTS, and CTS signals. If you want to keep a currently allocated pin number for a specific signal, pass the macro `UART_PIN_NO_CHANGE`.

The same macro `UART_PIN_NO_CHANGE` should be specified for pins that will not be used.

```
// Set UART pins(TX: IO4, RX: IO5, RTS: IO18, CTS: IO19)
ESP_ERROR_CHECK(uart_set_pin(UART_NUM_1, 4, 5, 18, 19));
```

Install Drivers Once the communication pins are set, install the driver by calling `uart_driver_install()` and specify the following parameters:

- Size of TX ring buffer
- Size of RX ring buffer
- Event queue handle and size
- Flags to allocate an interrupt

The function allocates the required internal resources for the UART driver.

```
// Setup UART buffered IO with event queue
const int uart_buffer_size = (1024 * 2);
QueueHandle_t uart_queue;
// Install UART driver using an event queue here
ESP_ERROR_CHECK(uart_driver_install(UART_NUM_1, uart_buffer_size, \
                                   uart_buffer_size, 10, &uart_queue, 0));
```

Once this step is complete, you can connect the external UART device and check the communication.

Run UART Communication Serial communication is controlled by each UART controller's finite state machine (FSM).

The process of sending data involves the following steps:

1. Write data into TX FIFO buffer
2. FSM serializes the data
3. FSM sends the data out

The process of receiving data is similar, but the steps are reversed:

1. FSM processes an incoming serial stream and parallelizes it
2. FSM writes the data into RX FIFO buffer
3. Read the data from RX FIFO buffer

Therefore, an application only writes and reads data from a specific buffer using `uart_write_bytes()` and `uart_read_bytes()` respectively, and the FSM does the rest.

Transmit Data After preparing the data for transmission, call the function `uart_write_bytes()` and pass the data buffer's address and data length to it. The function copies the data to the TX ring buffer (either immediately or after enough space is available), and then exit. When there is free space in the TX FIFO buffer, an interrupt service routine (ISR) moves the data from the TX ring buffer to the TX FIFO buffer in the background. The code below demonstrates the use of this function.

```
// Write data to UART.
char* test_str = "This is a test string.\n";
uart_write_bytes(uart_num, (const char*)test_str, strlen(test_str));
```

The function `uart_write_bytes_with_break()` is similar to `uart_write_bytes()` but adds a serial break signal at the end of the transmission. A 'serial break signal' means holding the TX line low for a period longer than one data frame.

```
// Write data to UART, end with a break signal.
uart_write_bytes_with_break(uart_num, "test break\n", strlen("test break\n"), 100);
```

Another function for writing data to the TX FIFO buffer is `uart_tx_chars()`. Unlike `uart_write_bytes()`, this function does not block until space is available. Instead, it writes all data which can immediately fit into the hardware TX FIFO, and then return the number of bytes that were written.

There is a 'companion' function `uart_wait_tx_done()` that monitors the status of the TX FIFO buffer and returns once it is empty.

```
// Wait for packet to be sent
const uart_port_t uart_num = UART_NUM_1;
ESP_ERROR_CHECK(uart_wait_tx_done(uart_num, 100)); // wait timeout is 100 RTOS_
↳ticks (TickType_t)
```

Receive Data Once the data is received by the UART and saved in the RX FIFO buffer, it needs to be retrieved using the function `uart_read_bytes()`. Before reading data, you can check the number of bytes available in the RX FIFO buffer by calling `uart_get_buffered_data_len()`. An example of using these functions is given below.

```
// Read data from UART.
const uart_port_t uart_num = UART_NUM_1;
uint8_t data[128];
int length = 0;
ESP_ERROR_CHECK(uart_get_buffered_data_len(uart_num, (size_t*)&length));
length = uart_read_bytes(uart_num, data, length, 100);
```

If the data in the RX FIFO buffer is no longer needed, you can clear the buffer by calling `uart_flush()`.

Software Flow Control If the hardware flow control is disabled, you can manually set the RTS and DTR signal levels by using the functions `uart_set_rts()` and `uart_set_dtr()` respectively.

Communication Mode Selection The UART controller supports a number of communication modes. A mode can be selected using the function `uart_set_mode()`. Once a specific mode is selected, the UART driver handles the behavior of a connected UART device accordingly. As an example, it can control the RS485 driver chip using the RTS line to allow half-duplex RS485 communication.

```
// Setup UART in rs485 half duplex mode
ESP_ERROR_CHECK(uart_set_mode(uart_num, UART_MODE_RS485_HALF_DUPLEX));
```

Use Interrupts There are many interrupts that can be generated depending on specific UART states or detected errors. The full list of available interrupts is provided in *ESP32-P4 Technical Reference Manual > UART Controller (UART) > UART Interrupts* and *UHCI Interrupts* [PDF]. You can enable or disable specific interrupts by calling `uart_enable_intr_mask()` or `uart_disable_intr_mask()` respectively.

The `uart_driver_install()` function installs the driver's internal interrupt handler to manage the TX and RX ring buffers and provides high-level API functions like events (see below).

The API provides a convenient way to handle specific interrupts discussed in this document by wrapping them into dedicated functions:

- **Event detection:** There are several events defined in `uart_event_type_t` that may be reported to a user application using the FreeRTOS queue functionality. You can enable this functionality when calling `uart_driver_install()` described in *Install Drivers*. An example of using Event detection can be found in *peripherals/uart/uart_events*.
- **FIFO space threshold or transmission timeout reached:** The TX and RX FIFO buffers can trigger an interrupt when they are filled with a specific number of characters, or on a timeout of sending or receiving data. To use these interrupts, do the following:
 - Configure respective threshold values of the buffer length and timeout by entering them in the structure `uart_intr_config_t` and calling `uart_intr_config()`
 - Enable the interrupts using the functions `uart_enable_tx_intr()` and `uart_enable_rx_intr()`
 - Disable these interrupts using the corresponding functions `uart_disable_tx_intr()` or `uart_disable_rx_intr()`
- **Pattern detection:** An interrupt triggered on detecting a 'pattern' of the same character being received/sent repeatedly. This functionality is demonstrated in the example *peripherals/uart/uart_events*. It can be used, e.g., to detect a command string with a specific number of identical characters (the 'pattern') at the end. The following functions are available:
 - Configure and enable this interrupt using `uart_enable_pattern_det_baud_intr()`
 - Disable the interrupt using `uart_disable_pattern_det_intr()`

Deleting a Driver If the communication established with `uart_driver_install()` is no longer required, the driver can be removed to free allocated resources by calling `uart_driver_delete()`.

Macros The API also defines several macros. For example, `UART_HW_FIFO_LEN` defines the length of hardware FIFO buffers; `UART_BITRATE_MAX` gives the maximum baud rate supported by the UART controllers, etc.

Use LP UART Controller with HP Core The UART driver also supports to control the LP UART controller when the chip is in active mode. The configuration steps for the LP UART are the same as the steps for a normal UART controller, except:

- The port number for the LP UART controller is defined by `LP_UART_NUM_0`.
- The available clock sources for the LP UART controller can be found in [lp_uart_sclk_t](#).
- The size of the hardware FIFO for the LP UART controller is much smaller, which is defined in [SOC_LP_UART_FIFO_LEN](#).
- The GPIO pins for the LP UART controller can only be selected from the LP GPIO pins.

Overview of RS485 Specific Communication Options

Note: The following section uses `[UART_REGISTER_NAME] . [UART_FIELD_BIT]` to refer to UART register fields/bits. For more information on a specific option bit, see **ESP32-P4 Technical Reference Manual > UART Controller (UART) > Register Summary** [\[PDF\]](#). Use the register name to navigate to the register description and then find the field/bit.

- `UART_RS485_CONF_REG.UART_RS485_EN`: setting this bit enables RS485 communication mode support.
- `UART_RS485_CONF_REG.UART_RS485TX_RX_EN`: if this bit is set, the transmitter's output signal loops back to the receiver's input signal.
- `UART_RS485_CONF_REG.UART_RS485RXBY_TX_EN`: if this bit is set, the transmitter will still be sending data if the receiver is busy (remove collisions automatically by hardware).

The ESP32-P4's RS485 UART hardware can detect signal collisions during transmission of a datagram and generate the interrupt `UART_RS485_CLASH_INT` if this interrupt is enabled. The term collision means that a transmitted datagram is not equal to the one received on the other end. Data collisions are usually associated with the presence of other active devices on the bus or might occur due to bus errors.

The collision detection feature allows handling collisions when their interrupts are activated and triggered. The interrupts `UART_RS485_FRM_ERR_INT` and `UART_RS485_PARITY_ERR_INT` can be used with the collision detection feature to control frame errors and parity bit errors accordingly in RS485 mode. This functionality is supported in the UART driver and can be used by selecting the [UART_MODE_RS485_APP_CTRL](#) mode (see the function [uart_set_mode\(\)](#)).

The collision detection feature can work with circuit A and circuit C (see Section [Interface Connection Options](#)). In the case of using circuit A or B, the RTS pin connected to the DE pin of the bus driver should be controlled by the user application. Use the function [uart_get_collision_flag\(\)](#) to check if the collision detection flag has been raised.

The ESP32-P4 UART controllers themselves do not support half-duplex communication as they cannot provide automatic control of the RTS pin connected to the RE/DE input of RS485 bus driver. However, half-duplex communication can be achieved via software control of the RTS pin by the UART driver. This can be enabled by selecting the [UART_MODE_RS485_HALF_DUPLEX](#) mode when calling [uart_set_mode\(\)](#).

Once the host starts writing data to the TX FIFO buffer, the UART driver automatically asserts the RTS pin (logic 1); once the last bit of the data has been transmitted, the driver de-asserts the RTS pin (logic 0). To use this mode, the software would have to disable the hardware flow control function. This mode works with all the used circuits shown below.


Interface Connection Options This section provides example schematics to demonstrate the basic aspects of ESP32-P4's RS485 interface connection.

Note:

- The schematics below do **not** necessarily contain **all required elements**.


- The **analog devices** ADM483 & ADM2483 are examples of common RS485 transceivers and **can be replaced** with other similar transceivers.

Circuit A: Collision Detection Circuit




This circuit is preferable because it allows for collision detection and is quite simple at the same time. The receiver in the line driver is constantly enabled, which allows the UART to monitor the RS485 bus. Echo suppression is performed by the UART peripheral when the bit `UART_RS485_CONF_REG.UART_RS485TX_RX_EN` is enabled.

Circuit B: Manual Switching Transmitter/Receiver Without Collision Detection




This circuit does not allow for collision detection. It suppresses the null bytes that the hardware receives when the bit `UART_RS485_CONF_REG.UART_RS485TX_RX_EN` is set. The bit `UART_RS485_CONF_REG.UART_RS485RXBY_TX_EN` is not applicable in this case.

Circuit C: Auto Switching Transmitter/Receiver



(continues on next page)

(continued from previous page)



This galvanically isolated circuit does not require RTS pin control by a software application or driver because it controls the transceiver direction automatically. However, it requires suppressing null bytes during transmission by setting `UART_RS485_CONF_REG.UART_RS485RXBY_TX_EN` to 1 and `UART_RS485_CONF_REG.UART_RS485TX_RX_EN` to 0. This setup can work in any RS485 UART mode or even in [UART_MODE_UART](#).

Application Examples

The table below describes the code examples available in the directory [peripherals/uart/](#).

Code Example	Description
peripherals/uart/uart_echo	Configuring UART settings, installing the UART driver, and reading/writing over the UART1 interface.
peripherals/uart/uart_events	Reporting various communication events, using pattern detection interrupts.
peripherals/uart/uart_async_rxtxtasks	Transmitting and receiving data in two separate FreeRTOS tasks over the same UART.
peripherals/uart/uart_select	Using synchronous I/O multiplexing for UART file descriptors.
peripherals/uart/uart_echo_rs485	Setting up UART driver to communicate over RS485 interface in half-duplex mode. This example is similar to peripherals/uart/uart_echo but allows communication through an RS485 interface chip connected to ESP32-P4 pins.
peripherals/uart/nmea0183_parser	Obtaining GPS information by parsing NMEA0183 statements received from GPS via the UART peripheral.

API Reference

Header File

- [components/driver/uart/include/driver/uart.h](#)
- This header file can be included with:

```
#include "driver/uart.h"
```

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your CMakeLists.txt:

```
REQUIRES driver
```

or

```
PRIV_REQUIRES driver
```

Functions

`esp_err_t uart_driver_install(uart_port_t uart_num, int rx_buffer_size, int tx_buffer_size, int queue_size, QueueHandle_t *uart_queue, int intr_alloc_flags)`

Install UART driver and set the UART to the default configuration.

UART ISR handler will be attached to the same CPU core that this function is running on.

Note: Rx_buffer_size should be greater than UART_HW_FIFO_LEN(uart_num). Tx_buffer_size should be either zero or greater than UART_HW_FIFO_LEN(uart_num).

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **rx_buffer_size** -- UART RX ring buffer size.
- **tx_buffer_size** -- UART TX ring buffer size. If set to zero, driver will not use TX buffer, TX function will block task until all data have been sent out.
- **queue_size** -- UART event queue size/depth.
- **uart_queue** -- UART event queue handle (out param). On success, a new queue handle is written here to provide access to UART events. If set to NULL, driver will not use an event queue.
- **intr_alloc_flags** -- Flags used to allocate the interrupt. One or multiple (ORred) ESP_INTR_FLAG_* values. See esp_intr_alloc.h for more info. Do not set ESP_INTR_FLAG_IRAM here (the driver's ISR handler is not located in IRAM)

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_driver_delete** (*uart_port_t* uart_num)

Uninstall UART driver.

Parameters **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

bool **uart_is_driver_installed** (*uart_port_t* uart_num)

Checks whether the driver is installed or not.

Parameters **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).

Returns

- true driver is installed
- false driver is not installed

esp_err_t **uart_set_word_length** (*uart_port_t* uart_num, *uart_word_length_t* data_bit)

Set UART data bits.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **data_bit** -- UART data bits

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_get_word_length** (*uart_port_t* uart_num, *uart_word_length_t* *data_bit)

Get the UART data bit configuration.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **data_bit** -- Pointer to accept value of UART data bits.

Returns

- ESP_FAIL Parameter error
- ESP_OK Success, result will be put in (*data_bit)

esp_err_t **uart_set_stop_bits** (*uart_port_t* uart_num, *uart_stop_bits_t* stop_bits)

Set UART stop bits.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **stop_bits** -- UART stop bits

Returns

- ESP_OK Success
- ESP_FAIL Fail

esp_err_t **uart_get_stop_bits** (*uart_port_t* uart_num, *uart_stop_bits_t* *stop_bits)

Get the UART stop bit configuration.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **stop_bits** -- Pointer to accept value of UART stop bits.

Returns

- ESP_FAIL Parameter error
- ESP_OK Success, result will be put in (*stop_bit)

esp_err_t **uart_set_parity** (*uart_port_t* uart_num, *uart_parity_t* parity_mode)

Set UART parity mode.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **parity_mode** -- the enum of uart parity configuration

Returns

- ESP_FAIL Parameter error
- ESP_OK Success

esp_err_t **uart_get_parity** (*uart_port_t* uart_num, *uart_parity_t* *parity_mode)

Get the UART parity mode configuration.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **parity_mode** -- Pointer to accept value of UART parity mode.

Returns

- ESP_FAIL Parameter error
- ESP_OK Success, result will be put in (*parity_mode)

esp_err_t **uart_get_sclk_freq** (*uart_sclk_t* sclk, uint32_t *out_freq_hz)

Get the frequency of a clock source for the HP UART port.

Parameters

- **sclk** -- Clock source
- **out_freq_hz** -- [out] Output of frequency, in Hz

Returns

- ESP_ERR_INVALID_ARG: if the clock source is not supported
- otherwise ESP_OK

esp_err_t **uart_set_baudrate** (*uart_port_t* uart_num, uint32_t baudrate)

Set desired UART baud rate.

Note that the actual baud rate set could have a slight deviation from the user-configured value due to rounding error.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **baudrate** -- UART baud rate.

Returns

- ESP_FAIL Parameter error, such as baud rate unachievable
- ESP_OK Success

esp_err_t **uart_get_baudrate** (*uart_port_t* uart_num, uint32_t *baudrate)

Get the actual UART baud rate.

It returns the real UART rate set in the hardware. It could have a slight deviation from the user-configured baud rate.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **baudrate** -- Pointer to accept value of UART baud rate

Returns

- ESP_FAIL Parameter error or the UART port is not enabled
- ESP_OK Success, result will be put in (*baudrate)

esp_err_t **uart_set_line_inverse** (*uart_port_t* uart_num, uint32_t inverse_mask)

Set UART line inverse mode.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **inverse_mask** -- Choose the wires that need to be inverted. Using the ORred mask of *uart_signal_inv_t*

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_set_hw_flow_ctrl** (*uart_port_t* uart_num, *uart_hw_flowcontrol_t* flow_ctrl, uint8_t rx_thresh)

Set hardware flow control.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **flow_ctrl** -- Hardware flow control mode
- **rx_thresh** -- Threshold of Hardware RX flow control (0 ~ UART_HW_FIFO_LEN(uart_num)). Only when UART_HW_FLOWCTRL_RTS is set, will the rx_thresh value be set.

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_set_sw_flow_ctrl** (*uart_port_t* uart_num, bool enable, uint8_t rx_thresh_xon, uint8_t rx_thresh_xoff)

Set software flow control.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1)
- **enable** -- switch on or off
- **rx_thresh_xon** -- low water mark
- **rx_thresh_xoff** -- high water mark

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_get_hw_flow_ctrl** (*uart_port_t* uart_num, *uart_hw_flowcontrol_t* *flow_ctrl)

Get the UART hardware flow control configuration.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **flow_ctrl** -- Option for different flow control mode.

Returns

- ESP_FAIL Parameter error
- ESP_OK Success, result will be put in (*flow_ctrl)

esp_err_t **uart_clear_intr_status** (*uart_port_t* uart_num, uint32_t clr_mask)

Clear UART interrupt status.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **clr_mask** -- Bit mask of the interrupt status to be cleared.

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_enable_intr_mask** (*uart_port_t* uart_num, uint32_t enable_mask)

Set UART interrupt enable.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **enable_mask** -- Bit mask of the enable bits.

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_disable_intr_mask** (*uart_port_t* uart_num, uint32_t disable_mask)

Clear UART interrupt enable bits.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **disable_mask** -- Bit mask of the disable bits.

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_enable_rx_intr** (*uart_port_t* uart_num)

Enable UART RX interrupt (RX_FULL & RX_TIMEOUT INTERRUPT)

Parameters **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_disable_rx_intr** (*uart_port_t* uart_num)

Disable UART RX interrupt (RX_FULL & RX_TIMEOUT INTERRUPT)

Parameters **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_disable_tx_intr** (*uart_port_t* uart_num)

Disable UART TX interrupt (TX_FIFO_EMPTY INTERRUPT)

Parameters **uart_num** -- UART port number

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_enable_tx_intr** (*uart_port_t* uart_num, int enable, int thresh)

Enable UART TX interrupt (TX_FIFO_EMPTY INTERRUPT)

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **enable** -- Set to 1 to enable the interrupt
- **thresh** -- Threshold of TX FIFO empty interrupt, 0 ~ UART_HW_FIFO_LEN(uart_num). If -1 is given, threshold configuration will be skipped.

Returns

- ESP_OK Success

- ESP_FAIL Parameter error

esp_err_t **uart_set_pin** (*uart_port_t* uart_num, int tx_io_num, int rx_io_num, int rts_io_num, int cts_io_num)

Assign signals of a UART peripheral to GPIO pins.

Note: If the GPIO number configured for a UART signal matches one of the IOMUX signals for that GPIO, the signal will be connected directly via the IOMUX. Otherwise the GPIO and signal will be connected via the GPIO Matrix. For example, if on an ESP32 the call `uart_set_pin(0, 1, 3, -1, -1)` is performed, as GPIO1 is UART0's default TX pin and GPIO3 is UART0's default RX pin, both will be connected to respectively U0TXD and U0RXD through the IOMUX, totally bypassing the GPIO matrix. The check is performed on a per-pin basis. Thus, it is possible to have RX pin binded to a GPIO through the GPIO matrix, whereas TX is binded to its GPIO through the IOMUX.

Note: It is possible to configure TX and RX to share the same IO (single wire mode), but please be aware of output conflict, which could damage the pad. Apply open-drain and pull-up to the pad ahead of time as a protection, or the upper layer protocol must guarantee no output from two ends at the same time.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **tx_io_num** -- UART TX pin GPIO number.
- **rx_io_num** -- UART RX pin GPIO number.
- **rts_io_num** -- UART RTS pin GPIO number.
- **cts_io_num** -- UART CTS pin GPIO number.

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_set_rts** (*uart_port_t* uart_num, int level)

Manually set the UART RTS pin level.

Note: UART must be configured with hardware flow control disabled.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **level** -- 1: RTS output low (active); 0: RTS output high (block)

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_set_dtr** (*uart_port_t* uart_num, int level)

Manually set the UART DTR pin level.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **level** -- 1: DTR output low; 0: DTR output high

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_set_tx_idle_num** (*uart_port_t* uart_num, uint16_t idle_num)

Set UART idle interval after tx FIFO is empty.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).

- **idle_num** -- idle interval after tx FIFO is empty(unit: the time it takes to send one bit under current baudrate)

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_param_config** (*uart_port_t* uart_num, const *uart_config_t* *uart_config)

Set UART configuration parameters.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **uart_config** -- UART parameter settings

Returns

- ESP_OK Success
- ESP_FAIL Parameter error, such as baud rate unachievable

esp_err_t **uart_intr_config** (*uart_port_t* uart_num, const *uart_intr_config_t* *intr_conf)

Configure UART interrupts.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **intr_conf** -- UART interrupt settings

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_wait_tx_done** (*uart_port_t* uart_num, TickType_t ticks_to_wait)

Wait until UART TX FIFO is empty.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **ticks_to_wait** -- Timeout, count in RTOS ticks

Returns

- ESP_OK Success
- ESP_FAIL Parameter error
- ESP_ERR_TIMEOUT Timeout

int **uart_tx_chars** (*uart_port_t* uart_num, const char *buffer, uint32_t len)

Send data to the UART port from a given buffer and length.

This function will not wait for enough space in TX FIFO. It will just fill the available TX FIFO and return when the FIFO is full.

Note: This function should only be used when UART TX buffer is not enabled.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **buffer** -- data buffer address
- **len** -- data length to send

Returns

- (-1) Parameter error
- OTHERS (>=0) The number of bytes pushed to the TX FIFO

int **uart_write_bytes** (*uart_port_t* uart_num, const void *src, size_t size)

Send data to the UART port from a given buffer and length,.

If the UART driver's parameter 'tx_buffer_size' is set to zero: This function will not return until all the data have been sent out, or at least pushed into TX FIFO.

Otherwise, if the 'tx_buffer_size' > 0, this function will return after copying all the data to tx ring buffer, UART ISR will then move data from the ring buffer to TX FIFO gradually.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **src** -- data buffer address
- **size** -- data length to send

Returns

- (-1) Parameter error
- OTHERS (>=0) The number of bytes pushed to the TX FIFO

int **uart_write_bytes_with_break** (*uart_port_t* uart_num, const void *src, size_t size, int brk_len)

Send data to the UART port from a given buffer and length,.

If the UART driver's parameter 'tx_buffer_size' is set to zero: This function will not return until all the data and the break signal have been sent out. After all data is sent out, send a break signal.

Otherwise, if the 'tx_buffer_size' > 0, this function will return after copying all the data to tx ring buffer, UART ISR will then move data from the ring buffer to TX FIFO gradually. After all data sent out, send a break signal.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **src** -- data buffer address
- **size** -- data length to send
- **brk_len** -- break signal duration(unit: the time it takes to send one bit at current baudrate)

Returns

- (-1) Parameter error
- OTHERS (>=0) The number of bytes pushed to the TX FIFO

int **uart_read_bytes** (*uart_port_t* uart_num, void *buf, uint32_t length, TickType_t ticks_to_wait)

UART read bytes from UART buffer.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **buf** -- pointer to the buffer.
- **length** -- data length
- **ticks_to_wait** -- sTimeout, count in RTOS ticks

Returns

- (-1) Error
- OTHERS (>=0) The number of bytes read from UART buffer

esp_err_t **uart_flush** (*uart_port_t* uart_num)

Alias of `uart_flush_input`. UART ring buffer flush. This will discard all data in the UART RX buffer.

Note: Instead of waiting the data sent out, this function will clear UART rx buffer. In order to send all the data in tx FIFO, we can use `uart_wait_tx_done` function.

Parameters **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_flush_input** (*uart_port_t* uart_num)

Clear input buffer, discard all the data in the ring-buffer.

Note: In order to send all the data in tx FIFO, we can use `uart_wait_tx_done` function.

Parameters **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_get_buffered_data_len** (*uart_port_t* uart_num, size_t *size)

UART get RX ring buffer cached data length.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **size** -- Pointer of size_t to accept cached data length

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_get_tx_buffer_free_size** (*uart_port_t* uart_num, size_t *size)

UART get TX ring buffer free space size for the next data to be enqueued.

It returns the tight conservative bound for NOSPLIT ring buffer overall enqueueable payload across up to two chunks.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).
- **size** -- Pointer of size_t to accept the free space size

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **uart_disable_pattern_det_intr** (*uart_port_t* uart_num)

UART disable pattern detect function. Designed for applications like 'AT commands'. When the hardware detects a series of one same character, the interrupt will be triggered.

Parameters **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1).

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

esp_err_t **uart_enable_pattern_det_baud_intr** (*uart_port_t* uart_num, char pattern_chr, uint8_t chr_num, int chr_tout, int post_idle, int pre_idle)

UART enable pattern detect function. Designed for applications like 'AT commands'. When the hardware detect a series of one same character, the interrupt will be triggered.

Parameters

- **uart_num** -- UART port number.
- **pattern_chr** -- character of the pattern.
- **chr_num** -- number of the character, 8bit value.
- **chr_tout** -- timeout of the interval between each pattern characters, 16bit value, unit is the baud-rate cycle you configured. When the duration is more than this value, it will not take this data as at_cmd char.
- **post_idle** -- idle time after the last pattern character, 16bit value, unit is the baud-rate cycle you configured. When the duration is less than this value, it will not take the previous data as the last at_cmd char
- **pre_idle** -- idle time before the first pattern character, 16bit value, unit is the baud-rate cycle you configured. When the duration is less than this value, it will not take this data as the first at_cmd char.

Returns

- ESP_OK Success
- ESP_FAIL Parameter error

int **uart_pattern_pop_pos** (*uart_port_t* uart_num)

Return the nearest detected pattern position in buffer. The positions of the detected pattern are saved in a queue, this function will dequeue the first pattern position and move the pointer to next pattern position.

The following APIs will modify the pattern position info: `uart_flush_input`, `uart_read_bytes`, `uart_driver_delete`, `uart_pop_pattern_pos`. It is the application's responsibility to ensure atomic access to the pattern queue and the rx data buffer when using pattern detect feature.

Note: If the RX buffer is full and flow control is not enabled, the detected pattern may not be found in the rx buffer due to overflow.

Parameters `uart_num` -- UART port number, the max port number is (`UART_NUM_MAX` -1).

Returns

- (-1) No pattern found for current index or parameter error
- others the pattern position in rx buffer.

`int uart_pattern_get_pos (uart_port_t uart_num)`

Return the nearest detected pattern position in buffer. The positions of the detected pattern are saved in a queue, This function do nothing to the queue.

The following APIs will modify the pattern position info: `uart_flush_input`, `uart_read_bytes`, `uart_driver_delete`, `uart_pop_pattern_pos`. It is the application's responsibility to ensure atomic access to the pattern queue and the rx data buffer when using pattern detect feature.

Note: If the RX buffer is full and flow control is not enabled, the detected pattern may not be found in the rx buffer due to overflow.

Parameters `uart_num` -- UART port number, the max port number is (`UART_NUM_MAX` -1).

Returns

- (-1) No pattern found for current index or parameter error
- others the pattern position in rx buffer.

`esp_err_t uart_pattern_queue_reset (uart_port_t uart_num, int queue_length)`

Allocate a new memory with the given length to save record the detected pattern position in rx buffer.

Parameters

- `uart_num` -- UART port number, the max port number is (`UART_NUM_MAX` -1).
- `queue_length` -- Max queue length for the detected pattern. If the queue length is not large enough, some pattern positions might be lost. Set this value to the maximum number of patterns that could be saved in data buffer at the same time.

Returns

- `ESP_ERR_NO_MEM` No enough memory
- `ESP_ERR_INVALID_STATE` Driver not installed
- `ESP_FAIL` Parameter error
- `ESP_OK` Success

`esp_err_t uart_set_mode (uart_port_t uart_num, uart_mode_t mode)`

UART set communication mode.

Note: This function must be executed after `uart_driver_install()`, when the driver object is initialized.

Parameters

- **uart_num** -- Uart number to configure, the max port number is (UART_NUM_MAX -1).
- **mode** -- UART mode to set

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **uart_set_rx_full_threshold** (*uart_port_t* uart_num, int threshold)

Set uart threshold value for RX fifo full.

Note: If application is using higher baudrate and it is observed that bytes in hardware RX fifo are overwritten then this threshold can be reduced

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1)
- **threshold** -- Threshold value above which RX fifo full interrupt is generated

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_ERR_INVALID_STATE Driver is not installed

esp_err_t **uart_set_tx_empty_threshold** (*uart_port_t* uart_num, int threshold)

Set uart threshold values for TX fifo empty.

Parameters

- **uart_num** -- UART port number, the max port number is (UART_NUM_MAX -1)
- **threshold** -- Threshold value below which TX fifo empty interrupt is generated

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_ERR_INVALID_STATE Driver is not installed

esp_err_t **uart_set_rx_timeout** (*uart_port_t* uart_num, const uint8_t tout_thresh)

UART set threshold timeout for TOUT feature.

Parameters

- **uart_num** -- Uart number to configure, the max port number is (UART_NUM_MAX -1).
- **tout_thresh** -- This parameter defines timeout threshold in uart symbol periods. The maximum value of threshold is 126. tout_thresh = 1, defines TOUT interrupt timeout equal to transmission time of one symbol (~11 bit) on current baudrate. If the time is expired the UART_RXFIFO_TOUT_INT interrupt is triggered. If tout_thresh == 0, the TOUT feature is disabled.

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_ERR_INVALID_STATE Driver is not installed

esp_err_t **uart_get_collision_flag** (*uart_port_t* uart_num, bool *collision_flag)

Returns collision detection flag for RS485 mode Function returns the collision detection flag into variable pointed by collision_flag. *collision_flag = true, if collision detected else it is equal to false. This function should be executed when actual transmission is completed (after uart_write_bytes()).

Parameters

- **uart_num** -- Uart number to configure the max port number is (UART_NUM_MAX -1).
- **collision_flag** -- Pointer to variable of type bool to return collision flag.

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

esp_err_t **uart_set_wakeup_threshold** (*uart_port_t* uart_num, int wakeup_threshold)

Set the number of RX pin signal edges for light sleep wakeup.

UART can be used to wake up the system from light sleep. This feature works by counting the number of positive edges on RX pin and comparing the count to the threshold. When the count exceeds the threshold, system is woken up from light sleep. This function allows setting the threshold value.

Stop bit and parity bits (if enabled) also contribute to the number of edges. For example, letter 'a' with ASCII code 97 is encoded as 0100001101 on the wire (with 8n1 configuration), start and stop bits included. This sequence has 3 positive edges (transitions from 0 to 1). Therefore, to wake up the system when 'a' is sent, set wakeup_threshold=3.

The character that triggers wakeup is not received by UART (i.e. it can not be obtained from UART FIFO). Depending on the baud rate, a few characters after that will also not be received. Note that when the chip enters and exits light sleep mode, APB frequency will be changing. To ensure that UART has correct Baud rate all the time, it is necessary to select a source clock which has a fixed frequency and remains active during sleep. For the supported clock sources of the chips, please refer to `uart_sclk_t` or `soc_periph_uart_clk_src_legacy_t`

Note: in ESP32, the wakeup signal can only be input via IO_MUX (i.e. GPIO3 should be configured as function_1 to wake up UART0, GPIO9 should be configured as function_5 to wake up UART1), UART2 does not support light sleep wakeup feature.

Parameters

- **uart_num** -- UART number, the max port number is (UART_NUM_MAX -1).
- **wakeup_threshold** -- number of RX edges for light sleep wakeup, value is 3 .. 0x3ff.

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if uart_num is incorrect or wakeup_threshold is outside of [3, 0x3ff] range.

esp_err_t **uart_get_wakeup_threshold** (*uart_port_t* uart_num, int *out_wakeup_threshold)

Get the number of RX pin signal edges for light sleep wakeup.

See description of `uart_set_wakeup_threshold` for the explanation of UART wakeup feature.

Parameters

- **uart_num** -- UART number, the max port number is (UART_NUM_MAX -1).
- **out_wakeup_threshold** -- [out] output, set to the current value of wakeup threshold for the given UART.

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if out_wakeup_threshold is NULL

esp_err_t **uart_wait_tx_idle_polling** (*uart_port_t* uart_num)

Wait until UART tx memory empty and the last char send ok (polling mode).

•

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_FAIL Driver not installed

Parameters **uart_num** -- UART number

esp_err_t **uart_set_loop_back** (*uart_port_t* uart_num, bool loop_back_en)

Configure TX signal loop back to RX module, just for the test usage.

•

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG Parameter error
- ESP_FAIL Driver not installed

Parameters

- **uart_num** -- UART number
- **loop_back_en** -- Set ture to enable the loop back function, else set it false.

void **uart_set_always_rx_timeout** (*uart_port_t* uart_num, bool always_rx_timeout_en)

Configure behavior of UART RX timeout interrupt.

When always_rx_timeout is true, timeout interrupt is triggered even if FIFO is full. This function can cause extra timeout interrupts triggered only to send the timeout event. Call this function only if you want to ensure timeout interrupt will always happen after a byte stream.

Parameters

- **uart_num** -- UART number
- **always_rx_timeout_en** -- Set to false enable the default behavior of timeout interrupt, set it to true to always trigger timeout interrupt.

Structures

struct **uart_config_t**

UART configuration parameters for uart_param_config function.

Public Members

int **baud_rate**

UART baud rate Note that the actual baud rate set could have a slight deviation from the user-configured value due to rounding error

uart_word_length_t **data_bits**

UART byte size

uart_parity_t **parity**

UART parity mode

uart_stop_bits_t **stop_bits**

UART stop bits

uart_hw_flowcontrol_t **flow_ctrl**

UART HW flow control mode (cts/rts)

uint8_t **rx_flow_ctrl_thresh**

UART HW RTS threshold

uart_sclk_t **source_clk**

UART source clock selection

lp_uart_sclk_t **lp_source_clk**

LP_UART source clock selection

uint32_t **backup_before_sleep**

If set, the driver will backup/restore the HP UART registers before entering/after exiting sleep mode. By this approach, the system can power off HP UART's power domain. This can save power, but at the expense of more RAM being consumed

struct *uart_config_t*::[anonymous] **flags**

Configuration flags

struct **uart_intr_config_t**

UART interrupt configuration parameters for `uart_intr_config` function.

Public Members

uint32_t **intr_enable_mask**

UART interrupt enable mask, choose from `UART_XXXX_INT_ENA_M` under `UART_INT_ENA_REG(i)`, connect with bit-or operator

uint8_t **rx_timeout_thresh**

UART timeout interrupt threshold (unit: time of sending one byte)

uint8_t **txfifo_empty_intr_thresh**

UART TX empty interrupt threshold.

uint8_t **rxfifo_full_thresh**

UART RX full interrupt threshold.

struct **uart_event_t**

Event structure used in UART event queue.

Public Members

uart_event_type_t **type**

UART event type

size_t **size**

UART data size for `UART_DATA` event

bool **timeout_flag**

UART data read timeout flag for `UART_DATA` event (no new data received during configured RX TOUT) If the event is caused by FIFO-full interrupt, then there will be no event with the timeout flag before the next byte coming.

Macros

UART_PIN_NO_CHANGE

UART_FIFO_LEN

Length of the HP UART HW FIFO.

UART_HW_FIFO_LEN (uart_num)

Length of the UART HW FIFO.

UART_BITRATE_MAX

Maximum configurable bitrate.

Type Definitions

typedef [*intr_handle_t*](#) **uart_isr_handle_t**

Enumerations

enum **uart_event_type_t**

UART event types used in the ring buffer.

Values:

enumerator **UART_DATA**

UART data event

enumerator **UART_BREAK**

UART break event

enumerator **UART_BUFFER_FULL**

UART RX buffer full event

enumerator **UART_FIFO_OVF**

UART FIFO overflow event

enumerator **UART_FRAME_ERR**

UART RX frame error event

enumerator **UART_PARITY_ERR**

UART RX parity event

enumerator **UART_DATA_BREAK**

UART TX data and break event

enumerator **UART_PATTERN_DET**

UART pattern detected

enumerator **UART_WAKEUP**

UART wakeup event

enumerator **UART_EVENT_MAX**

UART event max index

Header File

- [components/hal/include/hal/uart_types.h](#)
- This header file can be included with:

```
#include "hal/uart_types.h"
```

Structures

struct **uart_at_cmd_t**

UART AT cmd char configuration parameters Note that this function may different on different chip. Please refer to the TRM at configuration.

Public Members

uint8_t **cmd_char**

UART AT cmd char

uint8_t **char_num**

AT cmd char repeat number

uint32_t **gap_tout**

gap time(in baud-rate) between AT cmd char

uint32_t **pre_idle**

the idle time(in baud-rate) between the non AT char and first AT char

uint32_t **post_idle**

the idle time(in baud-rate) between the last AT char and the none AT char

struct **uart_sw_flowctrl_t**

UART software flow control configuration parameters.

Public Members

uint8_t **xon_char**

Xon flow control char

uint8_t **xoff_char**

Xoff flow control char

uint8_t **xon_thrd**

If the software flow control is enabled and the data amount in rxfifo is less than xon_thrd, an xon_char will be sent

uint8_t **xoff_thrd**

If the software flow control is enabled and the data amount in rxfifo is more than xoff_thrd, an xoff_char will be sent

Type Definitions

typedef *soc_periph_uart_clk_src_legacy_t* **uart_sclk_t**

UART source clock.

typedef *soc_periph_lp_uart_clk_src_t* **lp_uart_sclk_t**

LP_UART source clock.

Enumerations

enum **uart_port_t**

UART port number, can be UART_NUM_0 ~ (UART_NUM_MAX -1).

Values:

enumerator **UART_NUM_0**

UART port 0

enumerator **UART_NUM_1**

UART port 1

enumerator **UART_NUM_2**

UART port 2

enumerator **UART_NUM_3**

UART port 3

enumerator **UART_NUM_4**

UART port 4

enumerator **LP_UART_NUM_0**

LP UART port 0

enumerator **UART_NUM_MAX**

UART port max

enum **uart_mode_t**

UART mode selection.

Values:

enumerator **UART_MODE_UART**

mode: regular UART mode

enumerator **UART_MODE_RS485_HALF_DUPLEX**

mode: half duplex RS485 UART mode control by RTS pin

enumerator **UART_MODE_IRDA**

mode: IRDA UART mode

enumerator **UART_MODE_RS485_COLLISION_DETECT**

mode: RS485 collision detection UART mode (used for test purposes)

enumerator **UART_MODE_RS485_APP_CTRL**

mode: application control RS485 UART mode (used for test purposes)

enum **uart_word_length_t**

UART word length constants.

Values:

enumerator **UART_DATA_5_BITS**

word length: 5bits

enumerator **UART_DATA_6_BITS**

word length: 6bits

enumerator **UART_DATA_7_BITS**

word length: 7bits

enumerator **UART_DATA_8_BITS**

word length: 8bits

enumerator **UART_DATA_BITS_MAX**

enum **uart_stop_bits_t**

UART stop bits number.

Values:

enumerator **UART_STOP_BITS_1**

stop bit: 1bit

enumerator **UART_STOP_BITS_1_5**

stop bit: 1.5bits

enumerator **UART_STOP_BITS_2**

stop bit: 2bits

enumerator **UART_STOP_BITS_MAX**

enum **uart_parity_t**

UART parity constants.

Values:

enumerator **UART_PARITY_DISABLE**

Disable UART parity

enumerator **UART_PARITY_EVEN**

Enable UART even parity

enumerator **UART_PARITY_ODD**

Enable UART odd parity

enum **uart_hw_flowcontrol_t**

UART hardware flow control modes.

Values:

enumerator **UART_HW_FLOWCTRL_DISABLE**

disable hardware flow control

enumerator **UART_HW_FLOWCTRL_RTS**

enable RX hardware flow control (rts)

enumerator **UART_HW_FLOWCTRL_CTS**

enable TX hardware flow control (cts)

enumerator **UART_HW_FLOWCTRL_CTS_RTS**

enable hardware flow control

enumerator **UART_HW_FLOWCTRL_MAX**

enum **uart_signal_inv_t**

UART signal bit map.

Values:

enumerator **UART_SIGNAL_INV_DISABLE**

Disable UART signal inverse

enumerator **UART_SIGNAL_IRDA_TX_INV**

inverse the UART irda_tx signal

enumerator **UART_SIGNAL_IRDA_RX_INV**

inverse the UART irda_rx signal

enumerator **UART_SIGNAL_RXD_INV**

inverse the UART rxd signal

enumerator **UART_SIGNAL_CTS_INV**

inverse the UART cts signal

enumerator **UART_SIGNAL_DSR_INV**

inverse the UART dsr signal

enumerator **UART_SIGNAL_TXD_INV**

inverse the UART txd signal

enumerator **UART_SIGNAL_RTS_INV**

inverse the UART rts signal

enumerator **UART_SIGNAL_DTR_INV**

inverse the UART dtr signal

GPIO Lookup Macros The UART peripherals have dedicated IO_MUX pins to which they are connected directly. However, signals can also be routed to other pins using the less direct GPIO matrix. To use direct routes, you need to know which pin is a dedicated IO_MUX pin for a UART channel. GPIO Lookup Macros simplify the process of finding and assigning IO_MUX pins. You choose a macro based on either the IO_MUX pin number, or a required UART channel name, and the macro returns the matching counterpart for you. See some examples below.

Note: These macros are useful if you need very high UART baud rates (over 40 MHz), which means you will have to use IO_MUX pins only. In other cases, these macros can be ignored, and you can use the GPIO Matrix as it allows you to configure any GPIO pin for any UART function.

1. `UART_NUM_2_TXD_DIRECT_GPIO_NUM` returns the IO_MUX pin number of UART channel 2 TXD pin (pin 17)
2. `UART_GPIO19_DIRECT_CHANNEL` returns the UART number of GPIO 19 when connected to the UART peripheral via IO_MUX (this is `UART_NUM_0`)
3. `UART_CTS_GPIO19_DIRECT_CHANNEL` returns the UART number of GPIO 19 when used as the UART CTS pin via IO_MUX (this is `UART_NUM_0`). It is similar to the above macro but specifies the pin function which is also part of the IO_MUX assignment.

Header File

- `components/soc/esp32p4/include/soc/uart_channel.h`
- This header file can be included with:

```
#include "soc/uart_channel.h"
```

Macros

`UART_GPIO37_DIRECT_CHANNEL`

`UART_NUM_0_TXD_DIRECT_GPIO_NUM`

`UART_GPIO38_DIRECT_CHANNEL`

`UART_NUM_0_RXD_DIRECT_GPIO_NUM`

`UART_GPIO8_DIRECT_CHANNEL`

`UART_NUM_0_RTS_DIRECT_GPIO_NUM`

`UART_GPIO9_DIRECT_CHANNEL`

`UART_NUM_0_CTS_DIRECT_GPIO_NUM`

`UART_TXD_GPIO37_DIRECT_CHANNEL`

`UART_RXD_GPIO38_DIRECT_CHANNEL`

`UART_RTS_GPIO8_DIRECT_CHANNEL`

`UART_CTS_GPIO9_DIRECT_CHANNEL`

`UART_GPIO10_DIRECT_CHANNEL`

`UART_NUM_1_TXD_DIRECT_GPIO_NUM`

`UART_GPIO11_DIRECT_CHANNEL`

`UART_NUM_1_RXD_DIRECT_GPIO_NUM`

`UART_GPIO12_DIRECT_CHANNEL`

`UART_NUM_1_RTS_DIRECT_GPIO_NUM`

`UART_GPIO13_DIRECT_CHANNEL`

`UART_NUM_1_CTS_DIRECT_GPIO_NUM`

`UART_TXD_GPIO10_DIRECT_CHANNEL`

`UART_RXD_GPIO11_DIRECT_CHANNEL`

`UART_RTS_GPIO12_DIRECT_CHANNEL`

`UART_CTS_GPIO13_DIRECT_CHANNEL`

Code examples for this API section are provided in the [peripherals](#) directory of ESP-IDF examples.

2.6 Project Configuration

2.6.1 Introduction

The `esp-idf-kconfig` package that ESP-IDF uses is based on [kconfiglib](#), which is a Python extension to the [Kconfig](#) system. Kconfig provides a compile-time project configuration mechanism and offers configuration options of several types (e.g., integers, strings, and boolens). Kconfig files specify dependencies between options, default values of options, the way options are grouped together, etc.

For the full list of available features, please see [Kconfig](#) and [kconfiglib](#) extentions.

2.6.2 Project Configuration Menu

Application developers can open a terminal-based project configuration menu with the `idf.py menuconfig` build target.

After being updated, this configuration is saved in the `sdkconfig` file under the project root directory. Based on `sdkconfig`, application build targets will generate the `sdkconfig.h` file under the build directory, and will make the `sdkconfig` options available to the project build system and source files.

2.6.3 Using `sdkconfig.defaults`

In some cases, for example, when the `sdkconfig` file is under revision control, it may be inconvenient for the build system to change the `sdkconfig` file. The build system offers a solution to prevent it from happening, which is to create the `sdkconfig.defaults` file. This file is never touched by the build system, and can be created manually or automatically. It contains all the options which matter to the given application and are different from the default ones. The format is the same as that of the `sdkconfig` file. `sdkconfig.defaults` can be created manually when one remembers all the changed configuration, or it can be generated automatically by running the `idf.py save-defconfig` command.

Once `sdkconfig.defaults` is created, `sdkconfig` can be deleted or added to the ignore list of the revision control system (e.g., the `.gitignore` file for `git`). Project build targets will automatically create the `sdkconfig` file, populate it with the settings from the `sdkconfig.defaults` file, and configure the rest of the settings to their default values. Note that during the build process, settings from `sdkconfig.defaults` will not override those already in `sdkconfig`. For more information, see [Custom Sdkconfig Defaults](#).

2.6.4 Kconfig Format Rules

Format rules for Kconfig files are as follows:

- Option names in any menus should have consistent prefixes. The prefix currently should have at least 3 characters.
- The unit of indentation should be 4 spaces. All sub-items belonging to a parent item are indented by one level deeper. For example, `menu` is indented by 0 spaces, `config menu` by 4 spaces, `help in config` by 8 spaces, and the text under `help` by 12 spaces.
- No trailing spaces are allowed at the end of the lines.
- The maximum length of options is 50 characters.
- The maximum length of lines is 120 characters.

Note: The `help` section of each config in the menu is treated as `reStructuredText` to generate the reference documentation for each option.

Format Checker

`kconfcheck` tool in [esp-idf-kconfig](#) package is provided for checking Kconfig files against the above format rules. The checker checks all Kconfig and `Kconfig.projbuild` files given as arguments, and generates a new file with suffix `.new` with some suggestions about how to fix issues (if there are any). Please note that the checker cannot correct all format issues and the responsibility of the developer is to final check and make corrections in order to pass the tests. For example, indentations will be corrected if there is not any misleading formatting, but it cannot come up with a common prefix for options inside a menu.

The `esp-idf-kconfig` package is available in ESP-IDF environments, where the checker tool can be invoked by running command `python -m kconfcheck <path_to_kconfig_file>`.

For more information, please refer to [esp-idf-kconfig package documentation](#).

2.6.5 Backward Compatibility of Kconfig Options

The standard [Kconfig](#) tools ignore unknown options in `sdkconfig`. So if a developer has custom settings for options which are renamed in newer ESP-IDF releases, then the given setting for the option would be silently ignored. Therefore, several features have been adopted to avoid this:

1. `kconfgen` is used by the tool chain to pre-process `sdkconfig` files before anything else. For example, `menuconfig` would read them, so the settings for old options is kept and not ignored.

2. `kconfigen` recursively finds all `sdkconfig.rename` files in ESP-IDF directory which contain old and new `Kconfig` option names. Old options are replaced by new ones in the `sdkconfig` file. Renames that should only appear for a single target can be placed in a target-specific rename file `sdkconfig.rename.TARGET`, where `TARGET` is the target name, e.g., `sdkconfig.rename.esp32s2`.
3. `kconfigen` post-processes `sdkconfig` files and generates all build outputs (`sdkconfig.h`, `sdkconfig.cmake`, and `auto.conf`) by adding a list of compatibility statements, i.e., the values of old options are set for new options after modification. If users still use old options in their code, this will prevent it from breaking.
4. *Deprecated options and their replacements* are automatically generated by `kconfigen`.

2.6.6 Configuration Options Reference

Subsequent sections contain the list of available ESP-IDF options automatically generated from `Kconfig` files. Note that due to dependencies between options, some options listed here may not be visible by default in `menuconfig`.

By convention, all option names are upper-case letters with underscores. When `Kconfig` generates `sdkconfig` and `sdkconfig.h` files, option names are prefixed with `CONFIG_`. So if an option `ENABLE_FOO` is defined in a `Kconfig` file and selected in `menuconfig`, then the `sdkconfig` and `sdkconfig.h` files will have `CONFIG_ENABLE_FOO` defined. In the following sections, option names are also prefixed with `CONFIG_`, same as in the source code.

Build type

Contains:

- `CONFIG_APP_BUILD_TYPE`
- `CONFIG_APP_BUILD_TYPE_PURE_RAM_APP`
- `CONFIG_APP_REPRODUCIBLE_BUILD`
- `CONFIG_APP_NO_BLOBS`

`CONFIG_APP_BUILD_TYPE`

Application build type

Found in: *Build type*

Select the way the application is built.

By default, the application is built as a binary file in a format compatible with the ESP-IDF bootloader. In addition to this application, 2nd stage bootloader is also built. Application and bootloader binaries can be written into flash and loaded/executed from there.

Another option, useful for only very small and limited applications, is to only link the `.elf` file of the application, such that it can be loaded directly into RAM over JTAG or UART. Note that since IRAM and DRAM sizes are very limited, it is not possible to build any complex application this way. However for some kinds of testing and debugging, this option may provide faster iterations, since the application does not need to be written into flash.

Note: when `APP_BUILD_TYPE_RAM` is selected and loaded with JTAG, ESP-IDF does not contain all the startup code required to initialize the CPUs and ROM memory (data/bss). Therefore it is necessary to execute a bit of ROM code prior to executing the application. A `gdbinit` file may look as follows (for ESP32):

```
# Connect to a running instance of OpenOCD target remote :3333 # Reset and halt the target
mon reset halt # Run to a specific point in ROM code, # where most of initialization is
complete. thb *0x40007d54 c # Load the application into RAM load # Run till app_main tb
app_main c
```

Execute this `gdbinit` file as follows:

```
xtensa-esp32-elf-gdb build/app-name.elf -x gdbinit
```

Example gdbinit files for other targets can be found in `tools/test_apps/system/gdb_loadable_elf/`

When loading the BIN with UART, the ROM will jump to ram and run the app after finishing the ROM startup code, so there's no additional startup initialization required. You can use the `load_ram` in `esptool.py` to load the generated .bin file into ram and execute.

Example: `esptool.py --chip {chip} -p {port} -b {baud} --no-stub load_ram {app.bin}`

Recommended `sdkconfig.defaults` for building loadable ELF files is as follows. `CONFIG_APP_BUILD_TYPE_RAM` is required, other options help reduce application memory footprint.

```
CONFIG_APP_BUILD_TYPE_RAM=y CONFIG_VFS_SUPPORT_TERMIOS=
CONFIG_NEWLIB_NANO_FORMAT=y CONFIG_ESP_SYSTEM_PANIC_PRINT_HALT=y
CONFIG_ESP_DEBUG_STUBS_ENABLE= CONFIG_ESP_ERR_TO_NAME_LOOKUP=
```

Available options:

- Default (binary application + 2nd stage bootloader) (`CONFIG_APP_BUILD_TYPE_APP_2NDBOOT`)
- Build app runs entirely in RAM (EXPERIMENTAL) (`CONFIG_APP_BUILD_TYPE_RAM`)

CONFIG_APP_BUILD_TYPE_PURE_RAM_APP

Build app without SPI_FLASH/PSRAM support (saves ram)

Found in: [Build type](#)

If this option is enabled, external memory and related peripherals, such as Cache, MMU, Flash and PSRAM, won't be initialized. Corresponding drivers won't be introduced either. Components that depend on the `spi_flash` component will also be unavailable, such as `app_update`, etc. When this option is enabled, about 26KB of RAM space can be saved.

CONFIG_APP_REPRODUCIBLE_BUILD

Enable reproducible build

Found in: [Build type](#)

If enabled, all date, time, and path information would be eliminated. A .gdbinit file would be created automatically. (or will be appended if you have one already)

Default value:

- No (disabled)

CONFIG_APP_NO_BLOBS

No Binary Blobs

Found in: [Build type](#)

If enabled, this disables the linking of binary libraries in the application build. Note that after enabling this Wi-Fi/Bluetooth will not work.

Default value:

- No (disabled)

Bootloader config

Contains:

- [CONFIG_BOOTLOADER_LOG_LEVEL](#)
- [Bootloader manager](#)

- [CONFIG_BOOTLOADER_COMPILER_OPTIMIZATION](#)
- [CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE](#)
- [CONFIG_BOOTLOADER_REGION_PROTECTION_ENABLE](#)
- [CONFIG_BOOTLOADER_APP_TEST](#)
- [CONFIG_BOOTLOADER_FACTORY_RESET](#)
- [CONFIG_BOOTLOADER_HOLD_TIME_GPIO](#)
- [CONFIG_BOOTLOADER_CUSTOM_RESERVE_RTC](#)
- [Serial Flash Configurations](#)
- [CONFIG_BOOTLOADER_SKIP_VALIDATE_ALWAYS](#)
- [CONFIG_BOOTLOADER_SKIP_VALIDATE_ON_POWER_ON](#)
- [CONFIG_BOOTLOADER_SKIP_VALIDATE_IN_DEEP_SLEEP](#)
- [CONFIG_BOOTLOADER_WDT_ENABLE](#)
- [CONFIG_BOOTLOADER_VDDSDIO_BOOST](#)

Bootloader manager Contains:

- [CONFIG_BOOTLOADER_PROJECT_VER](#)
- [CONFIG_BOOTLOADER_COMPILE_TIME_DATE](#)

CONFIG_BOOTLOADER_COMPILE_TIME_DATE

Use time/date stamp for bootloader

Found in: [Bootloader config](#) > [Bootloader manager](#)

If set, then the bootloader will be built with the current time/date stamp. It is stored in the bootloader description structure. If not set, time/date stamp will be excluded from bootloader image. This can be useful for getting the same binary image files made from the same source, but at different times.

CONFIG_BOOTLOADER_PROJECT_VER

Project version

Found in: [Bootloader config](#) > [Bootloader manager](#)

Project version. It is placed in "version" field of the esp_bootloader_desc structure. The type of this field is "uint32_t".

Range:

- from 0 to 4294967295

Default value:

- 1

CONFIG_BOOTLOADER_COMPILER_OPTIMIZATION

Bootloader optimization Level

Found in: [Bootloader config](#)

This option sets compiler optimization level (gcc -O argument) for the bootloader.

- The default "Size" setting will add the -Os flag to CFLAGS.
- The "Debug" setting will add the -Og flag to CFLAGS.
- The "Performance" setting will add the -O2 flag to CFLAGS.

Note that custom optimization levels may be unsupported.

Available options:

- Size (-Os) ([CONFIG_BOOTLOADER_COMPILER_OPTIMIZATION_SIZE](#))
- Debug (-Og) ([CONFIG_BOOTLOADER_COMPILER_OPTIMIZATION_DEBUG](#))

- Optimize for performance (-O2) (CONFIG_BOOTLOADER_COMPILER_OPTIMIZATION_PERF)
- Debug without optimization (-O0) (Deprecated, will be removed in IDF v6.0) (CONFIG_BOOTLOADER_COMPILER_OPTIMIZATION_NONE)

CONFIG_BOOTLOADER_LOG_LEVEL

Bootloader log verbosity

Found in: [Bootloader config](#)

Specify how much output to see in bootloader logs.

Available options:

- No output (CONFIG_BOOTLOADER_LOG_LEVEL_NONE)
- Error (CONFIG_BOOTLOADER_LOG_LEVEL_ERROR)
- Warning (CONFIG_BOOTLOADER_LOG_LEVEL_WARN)
- Info (CONFIG_BOOTLOADER_LOG_LEVEL_INFO)
- Debug (CONFIG_BOOTLOADER_LOG_LEVEL_DEBUG)
- Verbose (CONFIG_BOOTLOADER_LOG_LEVEL_VERBOSE)

Serial Flash Configurations Contains:

- [CONFIG_BOOTLOADER_FLASH_DC_AWARE](#)
- [CONFIG_BOOTLOADER_FLASH_XMC_SUPPORT](#)

CONFIG_BOOTLOADER_FLASH_DC_AWARE

Allow app adjust Dummy Cycle bits in SPI Flash for higher frequency (READ HELP FIRST)

Found in: [Bootloader config](#) > [Serial Flash Configurations](#)

This will force 2nd bootloader to be loaded by DOUT mode, and will restore Dummy Cycle setting by resetting the Flash

CONFIG_BOOTLOADER_FLASH_XMC_SUPPORT

Enable the support for flash chips of XMC (READ DOCS FIRST)

Found in: [Bootloader config](#) > [Serial Flash Configurations](#)

Perform the startup flow recommended by XMC. Please consult XMC for the details of this flow. XMC chips will be forbidden to be used, when this option is disabled.

DON'T DISABLE THIS UNLESS YOU KNOW WHAT YOU ARE DOING.

comment "Features below require specific hardware (READ DOCS FIRST!)"

CONFIG_BOOTLOADER_VDDSDIO_BOOST

VDDSDIO LDO voltage

Found in: [Bootloader config](#)

If this option is enabled, and VDDSDIO LDO is set to 1.8V (using eFuse or MTDI bootstrapping pin), bootloader will change LDO settings to output 1.9V instead. This helps prevent flash chip from browning out during flash programming operations.

This option has no effect if VDDSDIO is set to 3.3V, or if the internal VDDSDIO regulator is disabled via eFuse.

Available options:

- 1.8V (CONFIG_BOOTLOADER_VDDSDIO_BOOST_1_8V)
- 1.9V (CONFIG_BOOTLOADER_VDDSDIO_BOOST_1_9V)

CONFIG_BOOTLOADER_FACTORY_RESET

GPIO triggers factory reset

Found in: [Bootloader config](#)

Allows to reset the device to factory settings: - clear one or more data partitions; - boot from "factory" partition. The factory reset will occur if there is a GPIO input held at the configured level while device starts up. See settings below.

Default value:

- No (disabled)

CONFIG_BOOTLOADER_NUM_PIN_FACTORY_RESET

Number of the GPIO input for factory reset

Found in: [Bootloader config](#) > [CONFIG_BOOTLOADER_FACTORY_RESET](#)

The selected GPIO will be configured as an input with internal pull-up enabled (note that on some SoCs, not all pins have an internal pull-up, consult the hardware datasheet for details.) To trigger a factory reset, this GPIO must be held high or low (as configured) on startup.

Default value:

- 4 if [CONFIG_BOOTLOADER_FACTORY_RESET](#)

CONFIG_BOOTLOADER_FACTORY_RESET_PIN_LEVEL

Factory reset GPIO level

Found in: [Bootloader config](#) > [CONFIG_BOOTLOADER_FACTORY_RESET](#)

Pin level for factory reset, can be triggered on low or high.

Available options:

- Reset on GPIO low (CONFIG_BOOTLOADER_FACTORY_RESET_PIN_LOW)
- Reset on GPIO high (CONFIG_BOOTLOADER_FACTORY_RESET_PIN_HIGH)

CONFIG_BOOTLOADER_OTA_DATA_ERASE

Clear OTA data on factory reset (select factory partition)

Found in: [Bootloader config](#) > [CONFIG_BOOTLOADER_FACTORY_RESET](#)

The device will boot from "factory" partition (or OTA slot 0 if no factory partition is present) after a factory reset.

CONFIG_BOOTLOADER_DATA_FACTORY_RESET

Comma-separated names of partitions to clear on factory reset

Found in: [Bootloader config](#) > [CONFIG_BOOTLOADER_FACTORY_RESET](#)

Allows customers to select which data partitions will be erased while factory reset.

Specify the names of partitions as a comma-delimited with optional spaces for readability. (Like this: "nvs, phy_init, ...") Make sure that the name specified in the partition table and here are the same. Partitions of type "app" cannot be specified here.

Default value:

- "nvs" if [CONFIG_BOOTLOADER_FACTORY_RESET](#)

CONFIG_BOOTLOADER_APP_TEST

GPIO triggers boot from test app partition

Found in: [Bootloader config](#)

Allows to run the test app from "TEST" partition. A boot from "test" partition will occur if there is a GPIO input pulled low while device starts up. See settings below.

CONFIG_BOOTLOADER_NUM_PIN_APP_TEST

Number of the GPIO input to boot TEST partition

Found in: [Bootloader config](#) > [CONFIG_BOOTLOADER_APP_TEST](#)

The selected GPIO will be configured as an input with internal pull-up enabled. To trigger a test app, this GPIO must be pulled low on reset. After the GPIO input is deactivated and the device reboots, the old application will boot. (factory or OTA[x]). Note that GPIO34-39 do not have an internal pullup and an external one must be provided.

Range:

- from 0 to 39 if [CONFIG_BOOTLOADER_APP_TEST](#)

Default value:

- 18 if [CONFIG_BOOTLOADER_APP_TEST](#)

CONFIG_BOOTLOADER_APP_TEST_PIN_LEVEL

App test GPIO level

Found in: [Bootloader config](#) > [CONFIG_BOOTLOADER_APP_TEST](#)

Pin level for app test, can be triggered on low or high.

Available options:

- Enter test app on GPIO low ([CONFIG_BOOTLOADER_APP_TEST_PIN_LOW](#))
- Enter test app on GPIO high ([CONFIG_BOOTLOADER_APP_TEST_PIN_HIGH](#))

CONFIG_BOOTLOADER_HOLD_TIME_GPIO

Hold time of GPIO for reset/test mode (seconds)

Found in: [Bootloader config](#)

The GPIO must be held low continuously for this period of time after reset before a factory reset or test partition boot (as applicable) is performed.

Default value:

- 5 if [CONFIG_BOOTLOADER_FACTORY_RESET](#) || [CONFIG_BOOTLOADER_APP_TEST](#)

CONFIG_BOOTLOADER_REGION_PROTECTION_ENABLE

Enable protection for unmapped memory regions

Found in: [Bootloader config](#)

Protects the unmapped memory regions of the entire address space from unintended accesses. This will ensure that an exception will be triggered whenever the CPU performs a memory operation on unmapped regions of the address space.

Default value:

- Yes (enabled)

CONFIG_BOOTLOADER_WDT_ENABLE

Use RTC watchdog in start code

Found in: [Bootloader config](#)

Tracks the execution time of startup code. If the execution time is exceeded, the RTC_WDT will restart system. It is also useful to prevent a lock up in start code caused by an unstable power source. NOTE: Tracks the execution time starts from the bootloader code - re-set timeout, while selecting the source for slow_clk - and ends calling app_main. Re-set timeout is needed due to WDT uses a SLOW_CLK clock source. After changing a frequency slow_clk a time of WDT needs to re-set for new frequency. slow_clk depends on RTC_CLK_SRC (INTERNAL_RC or EXTERNAL_CRYSTAL).

Default value:

- Yes (enabled)

CONFIG_BOOTLOADER_WDT_DISABLE_IN_USER_CODE

Allows RTC watchdog disable in user code

Found in: [Bootloader config](#) > [CONFIG_BOOTLOADER_WDT_ENABLE](#)

If this option is set, the ESP-IDF app must explicitly reset, feed, or disable the rtc_wdt in the app's own code. If this option is not set (default), then rtc_wdt will be disabled by ESP-IDF before calling the app_main() function.

Use function rtc_wdt_feed() for resetting counter of rtc_wdt. Use function rtc_wdt_disable() for disabling rtc_wdt.

Default value:

- No (disabled)

CONFIG_BOOTLOADER_WDT_TIME_MS

Timeout for RTC watchdog (ms)

Found in: [Bootloader config](#) > [CONFIG_BOOTLOADER_WDT_ENABLE](#)

Verify that this parameter is correct and more then the execution time. Pay attention to options such as reset to factory, trigger test partition and encryption on boot - these options can increase the execution time. Note: RTC_WDT will reset while encryption operations will be performed.

Range:

- from 0 to 120000

Default value:

- 9000

CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE

Enable app rollback support

Found in: [Bootloader config](#)

After updating the app, the bootloader runs a new app with the "ESP_OTA_IMG_PENDING_VERIFY" state set. This state prevents the re-run of this app. After the first boot of the new app in the user code, the function should be called to confirm the operability of the app or vice versa about its non-operability. If the app is working, then it is marked as valid. Otherwise, it is marked as not valid and rolls back to the previous working app. A reboot is performed, and the app is booted before the software update. Note: If during the first boot a new app

the power goes out or the WDT works, then roll back will happen. Rollback is possible only between the apps with the same security versions.

Default value:

- No (disabled)

CONFIG_BOOTLOADER_APP_ANTI_ROLLBACK

Enable app anti-rollback support

Found in: *Bootloader config* > *CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE*

This option prevents rollback to previous firmware/application image with lower security version.

Default value:

- No (disabled) if *CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE*

CONFIG_BOOTLOADER_APP_SECURE_VERSION

eFuse secure version of app

Found in: *Bootloader config* > *CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE* > *CONFIG_BOOTLOADER_APP_ANTI_ROLLBACK*

The secure version is the sequence number stored in the header of each firmware. The security version is set in the bootloader, version is recorded in the eFuse field as the number of set ones. The allocated number of bits in the efuse field for storing the security version is limited (see *BOOTLOADER_APP_SEC_VER_SIZE_EFUSE_FIELD* option).

Bootloader: When bootloader selects an app to boot, an app is selected that has a security version greater or equal that recorded in eFuse field. The app is booted with a higher (or equal) secure version.

The security version is worth increasing if in previous versions there is a significant vulnerability and their use is not acceptable.

Your partition table should has a scheme with ota_0 + ota_1 (without factory).

Default value:

- 0 if *CONFIG_BOOTLOADER_APP_ANTI_ROLLBACK*

CONFIG_BOOTLOADER_APP_SEC_VER_SIZE_EFUSE_FIELD

Size of the efuse secure version field

Found in: *Bootloader config* > *CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE* > *CONFIG_BOOTLOADER_APP_ANTI_ROLLBACK*

The size of the efuse secure version field. Its length is limited to 32 bits for ESP32 and 16 bits for ESP32-S2. This determines how many times the security version can be increased.

Range:

- from 1 to 16 if *CONFIG_BOOTLOADER_APP_ANTI_ROLLBACK*

Default value:

- 16 if *CONFIG_BOOTLOADER_APP_ANTI_ROLLBACK*

CONFIG_BOOTLOADER_EFUSE_SECURE_VERSION_EMULATE

Emulate operations with efuse secure version(only test)

Found in: *Bootloader config* > *CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE* > *CONFIG_BOOTLOADER_APP_ANTI_ROLLBACK*

This option allows to emulate read/write operations with all eFuses and efuse secure version. It allows to test anti-rollback implementation without permanent write eFuse bits. There should be an entry in partition table with following details: *emul_efuse*, *data*, *efuse*, , *0x2000*.

This option enables: `EFUSE_VIRTUAL` and `EFUSE_VIRTUAL_KEEP_IN_FLASH`.

Default value:

- No (disabled) if [`CONFIG\_BOOTLOADER\_APP\_ANTI\_ROLLBACK`](#)

CONFIG_BOOTLOADER_SKIP_VALIDATE_IN_DEEP_SLEEP

Skip image validation when exiting deep sleep

Found in: [Bootloader config](#)

This option disables the normal validation of an image coming out of deep sleep (checksums, SHA256, and signature). This is a trade-off between wakeup performance from deep sleep, and image integrity checks.

Only enable this if you know what you are doing. It should not be used in conjunction with using `deep_sleep()` entry and changing the active OTA partition as this would skip the validation upon first load of the new OTA partition.

It is possible to enable this option with Secure Boot if "allow insecure options" is enabled, however it's strongly recommended to NOT enable it as it may allow a Secure Boot bypass.

Default value:

- No (disabled) if [`CONFIG\_SECURE\_BOOT`](#) && [`CONFIG\_SECURE\_BOOT\_INSECURE`](#)

CONFIG_BOOTLOADER_SKIP_VALIDATE_ON_POWER_ON

Skip image validation from power on reset (READ HELP FIRST)

Found in: [Bootloader config](#)

Some applications need to boot very quickly from power on. By default, the entire app binary is read from flash and verified which takes up a significant portion of the boot time.

Enabling this option will skip validation of the app when the SoC boots from power on. Note that in this case it's not possible for the bootloader to detect if an app image is corrupted in the flash, therefore it's not possible to safely fall back to a different app partition. Flash corruption of this kind is unlikely but can happen if there is a serious firmware bug or physical damage.

Following other reset types, the bootloader will still validate the app image. This increases the chances that flash corruption resulting in a crash can be detected following soft reset, and the bootloader will fall back to a valid app image. To increase the chances of successfully recovering from a flash corruption event, keep the option `BOOTLOADER_WDT_ENABLE` enabled and consider also enabling `BOOTLOADER_WDT_DISABLE_IN_USER_CODE` - then manually disable the RTC Watchdog once the app is running. In addition, enable both the Task and Interrupt watchdog timers with reset options set.

Default value:

- No (disabled)

CONFIG_BOOTLOADER_SKIP_VALIDATE_ALWAYS

Skip image validation always (READ HELP FIRST)

Found in: [Bootloader config](#)

Selecting this option prevents the bootloader from ever validating the app image before booting it. Any flash corruption of the selected app partition will make the entire SoC unbootable.

Although flash corruption is a very rare case, it is not recommended to select this option. Consider selecting "Skip image validation from power on reset" instead. However, if boot time is the only important factor then it can be enabled.

Default value:

- No (disabled)

CONFIG_BOOTLOADER_CUSTOM_RESERVE_RTC

Reserve RTC FAST memory for custom purposes

Found in: [Bootloader config](#)

This option allows the customer to place data in the RTC FAST memory, this area remains valid when rebooted, except for power loss. This memory is located at a fixed address and is available for both the bootloader and the application. (The application and bootloader must be compiled with the same option). The RTC FAST memory has access only through PRO_CPU.

Default value:

- No (disabled)

CONFIG_BOOTLOADER_CUSTOM_RESERVE_RTC_IN_CRC

Include custom memory in the CRC calculation

Found in: [Bootloader config](#) > [CONFIG_BOOTLOADER_CUSTOM_RESERVE_RTC](#)

This option allows the customer to use the legacy bootloader behavior when the RTC FAST memory CRC calculation takes place. When this option is enabled, the allocated user custom data will be taken into account in the CRC calculation. This means that any change to the custom data would need a CRC update to prevent the bootloader from marking this data as corrupted. If this option is disabled, the custom data will not be taken into account when calculating the RTC FAST memory CRC. The user custom data can be changed freely, without the need to update the CRC. THIS OPTION MUST BE THE SAME FOR BOTH THE BOOTLOADER AND THE APPLICATION BUILDS.

Default value:

- No (disabled) if [CONFIG_BOOTLOADER_CUSTOM_RESERVE_RTC](#)

CONFIG_BOOTLOADER_CUSTOM_RESERVE_RTC_SIZE

Size in bytes for custom purposes

Found in: [Bootloader config](#) > [CONFIG_BOOTLOADER_CUSTOM_RESERVE_RTC](#)

This option reserves in RTC FAST memory the area for custom purposes. If you want to create your own bootloader and save more information in this area of memory, you can increase it. It must be a multiple of 4 bytes. This area (rtc_retain_mem_t) is reserved and has access from the bootloader and an application.

Default value:

- 0 if [CONFIG_BOOTLOADER_CUSTOM_RESERVE_RTC](#)

Security features

Contains:

- [CONFIG_SECURE_BOOT_INSECURE](#)
- [CONFIG_SECURE_SIGNED_APPS_SCHEME](#)
- [CONFIG_SECURE_SIGNED_ON_BOOT_NO_SECURE_BOOT](#)
- [CONFIG_SECURE_FLASH_CHECK_ENC_EN_IN_APP](#)
- [CONFIG_SECURE_BOOT_V2_ALLOW_EFUSE_RD_DIS](#)
- [CONFIG_SECURE_BOOT_ECDSA_KEY_LEN_SIZE](#)
- [CONFIG_SECURE_BOOT_ENABLE_AGGRESSIVE_KEY_REVOKE](#)
- [CONFIG_SECURE_FLASH_ENC_ENABLED](#)
- [CONFIG_SECURE_BOOT](#)
- [CONFIG_SECURE_FLASH_ENCRYPT_ONLY_IMAGE_LEN_IN_APP_PART](#)
- [CONFIG_SECURE_BOOT_FLASH_BOOTLOADER_DEFAULT](#)
- [CONFIG_SECURE_BOOTLOADER_KEY_ENCODING](#)
- [CONFIG_SECURE_FLASH_PSEUDO_ROUND_FUNC](#)

- *Potentially insecure options*
- *CONFIG_SECURE_SIGNED_APPS_NO_SECURE_BOOT*
- *CONFIG_SECURE_BOOT_VERIFICATION_KEY*
- *CONFIG_SECURE_BOOTLOADER_MODE*
- *CONFIG_SECURE_BOOT_BUILD_SIGNED_BINARIES*
- *CONFIG_SECURE_UART_ROM_DL_MODE*
- *CONFIG_SECURE_SIGNED_ON_UPDATE_NO_SECURE_BOOT*

CONFIG_SECURE_SIGNED_APPS_NO_SECURE_BOOT

Require signed app images

Found in: [Security features](#)

Require apps to be signed to verify their integrity.

This option uses the same app signature scheme as hardware secure boot, but unlike hardware secure boot it does not prevent the bootloader from being physically updated. This means that the device can be secured against remote network access, but not physical access. Compared to using hardware Secure Boot this option is much simpler to implement.

CONFIG_SECURE_SIGNED_APPS_SCHEME

App Signing Scheme

Found in: [Security features](#)

Select the Secure App signing scheme. Depends on the Chip Revision. There are two secure boot versions:

1. **Secure boot V1**
 - Legacy custom secure boot scheme. Supported in ESP32 SoC.
2. **Secure boot V2**
 - RSA based secure boot scheme. Supported in ESP32-ECO3 (ESP32 Chip Revision 3 onwards), ESP32-S2, ESP32-C3, ESP32-S3 SoCs.
 - ECDSA based secure boot scheme. Supported in ESP32-C2 SoC.

Available options:

- ECDSA (CONFIG_SECURE_SIGNED_APPS_ECDSA_SCHEME)
Embeds the ECDSA public key in the bootloader and signs the application with an ECDSA key. Refer to the documentation before enabling.
- RSA (CONFIG_SECURE_SIGNED_APPS_RSA_SCHEME)
Appends the RSA-3072 based Signature block to the application. Refer to <Secure Boot Version 2 documentation link> before enabling.
- ECDSA (V2) (CONFIG_SECURE_SIGNED_APPS_ECDSA_V2_SCHEME)
For Secure boot V2 (e.g., ESP32-C2 SoC), appends ECDSA based signature block to the application. Refer to documentation before enabling.

CONFIG_SECURE_BOOT_ECDSA_KEY_LEN_SIZE

ECDSA key size

Found in: [Security features](#)

Select the ECDSA key size. Two key sizes are supported

- 192 bit key using NISTP192 curve
- 256 bit key using NISTP256 curve (Recommended)

The advantage of using 256 bit key is the extra randomness which makes it difficult to be bruteforced compared to 192 bit key. At present, both key sizes are practically implausible to bruteforce.

Available options:

- Using ECC curve NISTP192 (CONFIG_SECURE_BOOT_ECDSA_KEY_LEN_192_BITS)
- Using ECC curve NISTP256 (Recommended) (CONFIG_SECURE_BOOT_ECDSA_KEY_LEN_256_BITS)

CONFIG_SECURE_SIGNED_ON_BOOT_NO_SECURE_BOOT

Bootloader verifies app signatures

Found in: [Security features](#)

If this option is set, the bootloader will be compiled with code to verify that an app is signed before booting it.

If hardware secure boot is enabled, this option is always enabled and cannot be disabled. If hardware secure boot is not enabled, this option doesn't add significant security by itself so most users will want to leave it disabled.

Default value:

- No (disabled) if [CONFIG_SECURE_SIGNED_APPS_NO_SECURE_BOOT](#) && [CONFIG_SECURE_SIGNED_APPS_ECDSA_SCHEME](#)

CONFIG_SECURE_SIGNED_ON_UPDATE_NO_SECURE_BOOT

Verify app signature on update

Found in: [Security features](#)

If this option is set, any OTA updated apps will have the signature verified before being considered valid.

When enabled, the signature is automatically checked whenever the `esp_ota_ops.h` APIs are used for OTA updates, or `esp_image_format.h` APIs are used to verify apps.

If hardware secure boot is enabled, this option is always enabled and cannot be disabled. If hardware secure boot is not enabled, this option still adds significant security against network-based attackers by preventing spoofing of OTA updates.

Default value:

- Yes (enabled) if [CONFIG_SECURE_SIGNED_APPS_NO_SECURE_BOOT](#)

CONFIG_SECURE_BOOT

Enable hardware Secure Boot in bootloader (READ DOCS FIRST)

Found in: [Security features](#)

Build a bootloader which enables Secure Boot on first boot.

Once enabled, Secure Boot will not boot a modified bootloader. The bootloader will only load a partition table or boot an app if the data has a verified digital signature. There are implications for reflashing updated apps once secure boot is enabled.

When enabling secure boot, JTAG and ROM BASIC Interpreter are permanently disabled by default.

Default value:

- No (disabled)

CONFIG_SECURE_BOOT_VERSION

Select secure boot version

Found in: [Security features](#) > CONFIG_SECURE_BOOT

Select the Secure Boot Version. Depends on the Chip Revision. Secure Boot V2 is the new RSA / ECDSA based secure boot scheme.

- RSA based scheme is supported in ESP32 (Revision 3 onwards), ESP32-S2, ESP32-C3 (ECO3), ESP32-S3.
- ECDSA based scheme is supported in ESP32-C2 SoC.

Please note that, RSA or ECDSA secure boot is property of specific SoC based on its HW design, supported crypto accelerators, die-size, cost and similar parameters. Please note that RSA scheme has requirement for bigger key sizes but at the same time it is comparatively faster than ECDSA verification.

Secure Boot V1 is the AES based (custom) secure boot scheme supported in ESP32 SoC.

Available options:

- Enable Secure Boot version 1 (CONFIG_SECURE_BOOT_V1_ENABLED)
Build a bootloader which enables secure boot version 1 on first boot. Refer to the Secure Boot section of the ESP-IDF Programmer's Guide for this version before enabling.
- Enable Secure Boot version 2 (CONFIG_SECURE_BOOT_V2_ENABLED)
Build a bootloader which enables Secure Boot version 2 on first boot. Refer to Secure Boot V2 section of the ESP-IDF Programmer's Guide for this version before enabling.

CONFIG_SECURE_BOOTLOADER_MODE

Secure bootloader mode

Found in: [Security features](#)

Available options:

- One-time flash (CONFIG_SECURE_BOOTLOADER_ONE_TIME_FLASH)
On first boot, the bootloader will generate a key which is not readable externally or by software. A digest is generated from the bootloader image itself. This digest will be verified on each subsequent boot.
Enabling this option means that the bootloader cannot be changed after the first time it is booted.
- Reflashable (CONFIG_SECURE_BOOTLOADER_REFLASHABLE)
Generate a reusable secure bootloader key, derived (via SHA-256) from the secure boot signing key.
This allows the secure bootloader to be re-flashed by anyone with access to the secure boot signing key.
This option is less secure than one-time flash, because a leak of the digest key from one device allows reflashing of any device that uses it.

CONFIG_SECURE_BOOT_BUILD_SIGNED_BINARIES

Sign binaries during build

Found in: [Security features](#)

Once secure boot or signed app requirement is enabled, app images are required to be signed.

If enabled (default), these binary files are signed as part of the build process. The file named in "Secure boot private signing key" will be used to sign the image.

If disabled, unsigned app/partition data will be built. They must be signed manually using `espsecure.py`. Version 1 to enable ECDSA Based Secure Boot and Version 2 to enable RSA based Secure Boot. (for example, on a remote signing server.)

CONFIG_SECURE_BOOT_SIGNING_KEY

Secure boot private signing key

Found in: [Security features](#) > [CONFIG_SECURE_BOOT_BUILD_SIGNED_BINARIES](#)

Path to the key file used to sign app images.

Key file is an ECDSA private key (NIST256p curve) in PEM format for Secure Boot V1. Key file is an RSA private key in PEM format for Secure Boot V2.

Path is evaluated relative to the project directory.

You can generate a new signing key by running the following command: `espsecure.py generate_signing_key secure_boot_signing_key.pem`

See the Secure Boot section of the ESP-IDF Programmer's Guide for this version for details.

Default value:

- "secure_boot_signing_key.pem" if [CONFIG_SECURE_BOOT_BUILD_SIGNED_BINARIES](#)

CONFIG_SECURE_BOOT_VERIFICATION_KEY

Secure boot public signature verification key

Found in: [Security features](#)

Path to a public key file used to verify signed images. Secure Boot V1: This ECDSA public key is compiled into the bootloader and/or app, to verify app images.

Key file is in raw binary format, and can be extracted from a PEM formatted private key using the `espsecure.py extract_public_key` command.

Refer to the Secure Boot section of the ESP-IDF Programmer's Guide for this version before enabling.

CONFIG_SECURE_BOOT_ENABLE_AGGRESSIVE_KEY_REVOKE

Enable Aggressive key revoke strategy

Found in: [Security features](#)

If this option is set, ROM bootloader will revoke the public key digest burned in efuse block if it fails to verify the signature of software bootloader with it. Revocation of keys does not happen when enabling secure boot. Once secure boot is enabled, key revocation checks will be done on subsequent boot-up, while verifying the software bootloader

This feature provides a strong resistance against physical attacks on the device.

NOTE: Once a digest slot is revoked, it can never be used again to verify an image. This can lead to permanent bricking of the device, in case all keys are revoked because of signature verification failure.

Default value:

- No (disabled) if [CONFIG_SECURE_BOOT](#)

CONFIG_SECURE_BOOT_V2_ALLOW_EFUSE_RD_DIS

Do not disable the ability to further read protect eFuses

Found in: [Security features](#)

If not set (default, recommended), on first boot the bootloader will burn the `WR_DIS_RD_DIS` efuse when Secure Boot is enabled. This prevents any more efuses from being read protected.

If this option is set, it will remain possible to write the EFUSE_RD_DIS efuse field after Secure Boot is enabled. This may allow an attacker to read-protect the BLK2 efuse (for ESP32) and BLOCK4-BLOCK10 (i.e. BLOCK_KEY0-BLOCK_KEY5)(for other chips) holding the secure boot public key digest, causing an immediate denial of service and possibly allowing an additional fault injection attack to bypass the signature protection.

The option must be set when you need to program any read-protected key type into the efuses, e.g., HMAC, ECDSA etc. after secure boot has already been enabled on the device. Please refer to secure boot V2 documentation guide for more details.

NOTE: Once a BLOCK is read-protected, the application will read all zeros from that block

NOTE: If "UART ROM download mode (Permanently disabled (recommended))" or "UART ROM download mode (Permanently switch to Secure mode (recommended))" is set, then it is __NOT__ possible to read/write efuses using espfuse.py utility. However, efuse can be read/written from the application

Please refer to the Secure Boot V2 documentation guide for more information.

Default value:

- No (disabled) if `CONFIG_SECURE_BOOT_V2_ENABLED`

CONFIG_SECURE_BOOT_FLASH_BOOTLOADER_DEFAULT

Flash bootloader along with other artifacts when using the default flash command

Found in: [Security features](#)

When Secure Boot V2 is enabled, by default the bootloader is not flashed along with other artifacts like the application and the partition table images, i.e. bootloader has to be separately flashed using the command `idf.py bootloader flash`, whereas, the application and partition table can be flashed using the command `idf.py flash` itself. Enabling this option allows flashing the bootloader along with the other artifacts by invocation of the command `idf.py flash`.

If this option is enabled make sure that even the bootloader is signed using the correct secure boot key, otherwise the bootloader signature verification would fail, as hash of the public key which is present in the bootloader signature would not match with the digest stored into the efuses and thus the device will not be able to boot up.

Default value:

- No (disabled) if `CONFIG_SECURE_BOOT_V2_ENABLED` && `CONFIG_SECURE_BOOT_BUILD_SIGNED_BINARIES`

CONFIG_SECURE_BOOTLOADER_KEY_ENCODING

Hardware Key Encoding

Found in: [Security features](#)

In reflashable secure bootloader mode, a hardware key is derived from the signing key (with SHA-256) and can be written to eFuse with espfuse.py.

Normally this is a 256-bit key, but if 3/4 Coding Scheme is used on the device then the eFuse key is truncated to 192 bits.

This configuration item doesn't change any firmware code, it only changes the size of key binary which is generated at build time.

Available options:

- No encoding (256 bit key) (`CONFIG_SECURE_BOOTLOADER_KEY_ENCODING_256BIT`)
- 3/4 encoding (192 bit key) (`CONFIG_SECURE_BOOTLOADER_KEY_ENCODING_192BIT`)

CONFIG_SECURE_BOOT_INSECURE

Allow potentially insecure options

Found in: [Security features](#)

You can disable some of the default protections offered by secure boot, in order to enable testing or a custom combination of security features.

Only enable these options if you are very sure.

Refer to the Secure Boot section of the ESP-IDF Programmer's Guide for this version before enabling.

Default value:

- No (disabled) if [CONFIG_SECURE_BOOT](#)

CONFIG_SECURE_FLASH_ENC_ENABLED

Enable flash encryption on boot (READ DOCS FIRST)

Found in: [Security features](#)

If this option is set, flash contents will be encrypted by the bootloader on first boot.

Note: After first boot, the system will be permanently encrypted. Re-flashing an encrypted system is complicated and not always possible.

Read [Flash Encryption](#) before enabling.

Default value:

- No (disabled)

CONFIG_SECURE_FLASH_ENCRYPTION_KEYSIZE

Size of generated XTS-AES key

Found in: [Security features](#) > [CONFIG_SECURE_FLASH_ENC_ENABLED](#)

Size of generated XTS-AES key.

- AES-128 uses a 256-bit key (32 bytes) derived from 128 bits (16 bytes) burned in half Efuse key block. Internally, it calculates SHA256(128 bits)
- AES-128 uses a 256-bit key (32 bytes) which occupies one Efuse key block.
- AES-256 uses a 512-bit key (64 bytes) which occupies two Efuse key blocks.

This setting is ignored if either type of key is already burned to Efuse before the first boot. In this case, the pre-burned key is used and no new key is generated.

Available options:

- AES-128 key derived from 128 bits (SHA256(128 bits)) ([CONFIG_SECURE_FLASH_ENCRYPTION_AES128_DERIVED](#))
- AES-128 (256-bit key) ([CONFIG_SECURE_FLASH_ENCRYPTION_AES128](#))
- AES-256 (512-bit key) ([CONFIG_SECURE_FLASH_ENCRYPTION_AES256](#))

CONFIG_SECURE_FLASH_ENCRYPTION_MODE

Enable usage mode

Found in: [Security features](#) > [CONFIG_SECURE_FLASH_ENC_ENABLED](#)

By default Development mode is enabled which allows ROM download mode to perform flash encryption operations (plaintext is sent to the device, and it encrypts it internally and writes ciphertext to flash.) This mode is not secure, it's possible for an attacker to write their own chosen plaintext to flash.

Release mode should always be selected for production or manufacturing. Once enabled it's no longer possible for the device in ROM Download Mode to use the flash encryption hardware.

When EFUSE_VIRTUAL is enabled, SECURE_FLASH_ENCRYPTION_MODE_RELEASE is not available. For CI tests we use IDF_CI_BUILD to bypass it ("export IDF_CI_BUILD=1"). We do not recommend bypassing it for other purposes.

Refer to the Flash Encryption section of the ESP-IDF Programmer's Guide for details.

Available options:

- Development (NOT SECURE) (CONFIG_SECURE_FLASH_ENCRYPTION_MODE_DEVELOPMENT)
- Release (CONFIG_SECURE_FLASH_ENCRYPTION_MODE_RELEASE)

Potentially insecure options Contains:

- [CONFIG_SECURE_BOOT_ALLOW_SHORT_APP_PARTITION](#)
- [CONFIG_SECURE_BOOT_ALLOW_JTAG](#)
- [CONFIG_SECURE_FLASH_UART_BOOTLOADER_ALLOW_ENC](#)
- [CONFIG_SECURE_FLASH_UART_BOOTLOADER_ALLOW_CACHE](#)
- [CONFIG_SECURE_BOOT_ALLOW_UNUSED_DIGEST_SLOTS](#)
- [CONFIG_SECURE_FLASH_REQUIRE_ALREADY_ENABLED](#)
- [CONFIG_SECURE_FLASH_SKIP_WRITE_PROTECTION_CACHE](#)

CONFIG_SECURE_BOOT_ALLOW_JTAG

Allow JTAG Debugging

Found in: [Security features](#) > [Potentially insecure options](#)

If not set (default), the bootloader will permanently disable JTAG (across entire chip) on first boot when either secure boot or flash encryption is enabled.

Setting this option leaves JTAG on for debugging, which negates all protections of flash encryption and some of the protections of secure boot.

Only set this option in testing environments.

Default value:

- No (disabled) if [CONFIG_SECURE_BOOT_INSECURE](#) || [CONFIG_SECURE_FLASH_ENCRYPTION_MODE_DEVELOPMENT](#)

CONFIG_SECURE_BOOT_ALLOW_SHORT_APP_PARTITION

Allow app partition length not 64KB aligned

Found in: [Security features](#) > [Potentially insecure options](#)

If not set (default), app partition size must be a multiple of 64KB. App images are padded to 64KB length, and the bootloader checks any trailing bytes after the signature (before the next 64KB boundary) have not been written. This is because flash cache maps entire 64KB pages into the address space. This prevents an attacker from appending unverified data after the app image in the flash, causing it to be mapped into the address space.

Setting this option allows the app partition length to be unaligned, and disables padding of the app image to this length. It is generally not recommended to set this option, unless you have a legacy partitioning scheme which doesn't support 64KB aligned partition lengths.

CONFIG_SECURE_BOOT_ALLOW_UNUSED_DIGEST_SLOTS

Leave unused digest slots available (not revoke)

Found in: [Security features](#) > [Potentially insecure options](#)

If not set (default), during startup in the app all unused digest slots will be revoked. To revoke unused slot will be called `esp_efuse_set_digest_revoke(num_digest)` for each digest. Revoking unused digest slots makes ensures that no trusted keys can be added later by an attacker. If set, it means that you have a plan to use unused digests slots later.

Default value:

- No (disabled) if [CONFIG_SECURE_BOOT_INSECURE](#)

CONFIG_SECURE_FLASH_UART_BOOTLOADER_ALLOW_ENC

Leave UART bootloader encryption enabled

Found in: [Security features](#) > [Potentially insecure options](#)

If not set (default), the bootloader will permanently disable UART bootloader encryption access on first boot. If set, the UART bootloader will still be able to access hardware encryption.

It is recommended to only set this option in testing environments.

Default value:

- No (disabled) if [CONFIG_SECURE_FLASH_ENCRYPTION_MODE_DEVELOPMENT](#)

CONFIG_SECURE_FLASH_UART_BOOTLOADER_ALLOW_CACHE

Leave UART bootloader flash cache enabled

Found in: [Security features](#) > [Potentially insecure options](#)

If not set (default), the bootloader will permanently disable UART bootloader flash cache access on first boot. If set, the UART bootloader will still be able to access the flash cache.

Only set this option in testing environments.

Default value:

- No (disabled) if [CONFIG_SECURE_FLASH_ENCRYPTION_MODE_DEVELOPMENT](#)

CONFIG_SECURE_FLASH_REQUIRE_ALREADY_ENABLED

Require flash encryption to be already enabled

Found in: [Security features](#) > [Potentially insecure options](#)

If not set (default), and flash encryption is not yet enabled in eFuses, the 2nd stage bootloader will enable flash encryption: generate the flash encryption key and program eFuses. If this option is set, and flash encryption is not yet enabled, the bootloader will error out and reboot. If flash encryption is enabled in eFuses, this option does not change the bootloader behavior.

Only use this option in testing environments, to avoid accidentally enabling flash encryption on the wrong device. The device needs to have flash encryption already enabled using `espefuse.py`.

Default value:

- No (disabled) if [CONFIG_SECURE_FLASH_ENCRYPTION_MODE_DEVELOPMENT](#)

CONFIG_SECURE_FLASH_SKIP_WRITE_PROTECTION_CACHE

Skip write-protection of DIS_CACHE (DIS_ICACHE, DIS_DCACHE)

Found in: [Security features](#) > [Potentially insecure options](#)

If not set (default, recommended), on the first boot the bootloader will burn the write-protection of DIS_CACHE(for ESP32) or DIS_ICACHE/DIS_DCACHE(for other chips) eFuse when Flash Encryption is enabled. Write protection for cache disable efuse prevents the chip from being blocked if it is set by accident. App and bootloader use cache so disabling it makes the chip useless for IDF. Due to other eFuses are linked with the same write protection bit (see the list below) then write-protection will not be done if these SECURE_FLASH_UART_BOOTLOADER_ALLOW_ENC, SECURE_BOOT_ALLOW_JTAG or SECURE_FLASH_UART_BOOTLOADER_ALLOW_CACHE options are selected to give a chance to turn on the chip into the release mode later.

List of eFuses with the same write protection bit: ESP32: MAC, MAC_CRC, DISABLE_APP_CPU, DISABLE_BT, DIS_CACHE, VOL_LEVEL_HP_INV.

ESP32-C3: DIS_ICACHE, DIS_USB_JTAG, DIS_DOWNLOAD_ICACHE, DIS_USB_SERIAL_JTAG, DIS_FORCE_DOWNLOAD, DIS_TWAI, JTAG_SEL_ENABLE, DIS_PAD_JTAG, DIS_DOWNLOAD_MANUAL_ENCRYPT.

ESP32-C6: SWAP_UART_SDIO_EN, DIS_ICACHE, DIS_USB_JTAG, DIS_DOWNLOAD_ICACHE, DIS_USB_SERIAL_JTAG, DIS_FORCE_DOWNLOAD, DIS_TWAI, JTAG_SEL_ENABLE, DIS_PAD_JTAG, DIS_DOWNLOAD_MANUAL_ENCRYPT.

ESP32-H2: DIS_ICACHE, DIS_USB_JTAG, POWERGLITCH_EN, DIS_FORCE_DOWNLOAD, SPI_DOWNLOAD_MSPI_DIS, DIS_TWAI, JTAG_SEL_ENABLE, DIS_PAD_JTAG, DIS_DOWNLOAD_MANUAL_ENCRYPT.

ESP32-S2: DIS_ICACHE, DIS_DCACHE, DIS_DOWNLOAD_ICACHE, DIS_DOWNLOAD_DCACHE, DIS_FORCE_DOWNLOAD, DIS_USB, DIS_TWAI, DIS_BOOT_REMAP, SOFT_DIS_JTAG, HARD_DIS_JTAG, DIS_DOWNLOAD_MANUAL_ENCRYPT.

ESP32-S3: DIS_ICACHE, DIS_DCACHE, DIS_DOWNLOAD_ICACHE, DIS_DOWNLOAD_DCACHE, DIS_FORCE_DOWNLOAD, DIS_USB_OTG, DIS_TWAI, DIS_APP_CPU, DIS_PAD_JTAG, DIS_DOWNLOAD_MANUAL_ENCRYPT, DIS_USB_JTAG, DIS_USB_SERIAL_JTAG, STRAP_JTAG_SEL, USB_PHY_SEL.

CONFIG_SECURE_FLASH_ENCRYPT_ONLY_IMAGE_LEN_IN_APP_PART

Encrypt only the app image that is present in the partition of type app

Found in: [Security features](#)

If set, optimise encryption time for the partition of type APP, by only encrypting the app image that is present in the partition, instead of the whole partition. The image length used for encryption is derived from the image metadata, which includes the size of the app image, checksum, hash and also the signature sector when secure boot is enabled.

If not set (default), the whole partition of type APP would be encrypted, which increases the encryption time but might be useful if there is any custom data appended to the firmware image.

CONFIG_SECURE_FLASH_CHECK_ENC_EN_IN_APP

Check Flash Encryption enabled on app startup

Found in: [Security features](#)

If set (default), in an app during startup code, there is a check of the flash encryption eFuse bit is on (as the bootloader should already have set it). The app requires this bit is on to continue work otherwise abort.

If not set, the app does not care if the flash encryption eFuse bit is set or not.

Default value:

- Yes (enabled) if [CONFIG_SECURE_FLASH_ENC_ENABLED](#)

CONFIG_SECURE_FLASH_PSEUDO_ROUND_FUNC

Permanently enable XTS-AES's pseudo rounds function

Found in: [Security features](#)

If set (default), the bootloader will permanently enable the XTS-AES peripheral's pseudo rounds function. Note: Enabling this config would burn an efuse.

Default value:

- Yes (enabled) if `CONFIG_SECURE_FLASH_ENCRYPTION_MODE_RELEASE` && `CONFIG_SECURE_FLASH_ENC_ENABLED` && `SOC_FLASH_ENCRYPTION_XTS_AES_SUPPORT_PSEUDO_ROUND`
- No (disabled) if `CONFIG_SECURE_FLASH_ENC_ENABLED` && `SOC_FLASH_ENCRYPTION_XTS_AES_SUPPORT_PSEUDO_ROUND`

CONFIG_SECURE_FLASH_PSEUDO_ROUND_FUNC_STRENGTH

Strength of the pseudo rounds function

Found in: [Security features](#) > [CONFIG_SECURE_FLASH_PSEUDO_ROUND_FUNC](#)

The strength of the pseudo rounds functions can be configured to low, medium and high, each denoting the values that would be stored in the efuses field. By default the value to set to low. You can configure the strength of the pseudo rounds functions according to your use cases, for example, increasing the strength would provide higher security but would slow down the flash encryption/decryption operations. For more info regarding the performance impact, please checkout the pseudo round function section of the security guide documentation.

Available options:

- Low (`CONFIG_SECURE_FLASH_PSEUDO_ROUND_FUNC_STRENGTH_LOW`)
- Medium (`CONFIG_SECURE_FLASH_PSEUDO_ROUND_FUNC_STRENGTH_MEDIUM`)
- High (`CONFIG_SECURE_FLASH_PSEUDO_ROUND_FUNC_STRENGTH_HIGH`)

CONFIG_SECURE_UART_ROM_DL_MODE

UART ROM download mode

Found in: [Security features](#)

Available options:

- UART ROM download mode (Permanently disabled (recommended)) (`CONFIG_SECURE_DISABLE_ROM_DL_MODE`)
If set, during startup the app will burn an eFuse bit to permanently disable the UART ROM Download Mode. This prevents any future use of `esptool.py`, `espefuse.py` and similar tools.
Once disabled, if the SoC is booted with strapping pins set for ROM Download Mode then an error is printed instead.
It is recommended to enable this option in any production application where Flash Encryption and/or Secure Boot is enabled and access to Download Mode is not required. It is also possible to permanently disable Download Mode by calling `esp_efuse_disable_rom_download_mode()` at runtime.
- UART ROM download mode (Permanently switch to Secure mode (recommended)) (`CONFIG_SECURE_ENABLE_SECURE_ROM_DL_MODE`)
If set, during startup the app will burn an eFuse bit to permanently switch the UART ROM Download Mode into a separate Secure Download mode. This option can only work if Download Mode is not already disabled by eFuse.

Secure Download mode limits the use of Download Mode functions to update SPI config, changing baud rate, basic flash write and a command to return a summary of currently enabled security features (*get_security_info*).

Secure Download mode is not compatible with the esptool.py flasher stub feature, es-
pfuse.py, read/writing memory or registers, encrypted download, or any other features that interact with unsupported Download Mode commands.

Secure Download mode should be enabled in any application where Flash Encryption and/or Secure Boot is enabled. Disabling this option does not immediately cancel the benefits of the security features, but it increases the potential "attack surface" for an attacker to try and bypass them with a successful physical attack.

It is also possible to enable secure download mode at runtime by calling `esp_efuse_enable_rom_secure_download_mode()`

Note: Secure Download mode is not available for ESP32.

- UART ROM download mode (Enabled (not recommended)) (`CONFIG_SECURE_INSECURE_ALLOW_DL_MODE`)

This is a potentially insecure option. Enabling this option will allow the full UART download mode to stay enabled. This option **SHOULD NOT BE ENABLED** for production use cases.

Application manager

Contains:

- [*CONFIG_APP_EXCLUDE_PROJECT_NAME_VAR*](#)
- [*CONFIG_APP_EXCLUDE_PROJECT_VER_VAR*](#)
- [*CONFIG_APP_PROJECT_VER_FROM_CONFIG*](#)
- [*CONFIG_APP_RETRIEVE_LEN_ELF_SHA*](#)
- [*CONFIG_APP_COMPILE_TIME_DATE*](#)

CONFIG_APP_COMPILE_TIME_DATE

Use time/date stamp for app

Found in: [*Application manager*](#)

If set, then the app will be built with the current time/date stamp. It is stored in the app description structure. If not set, time/date stamp will be excluded from app image. This can be useful for getting the same binary image files made from the same source, but at different times.

CONFIG_APP_EXCLUDE_PROJECT_VER_VAR

Exclude PROJECT_VER from firmware image

Found in: [*Application manager*](#)

The PROJECT_VER variable from the build system will not affect the firmware image. This value will not be contained in the `esp_app_desc` structure.

Default value:

- No (disabled)

CONFIG_APP_EXCLUDE_PROJECT_NAME_VAR

Exclude PROJECT_NAME from firmware image

Found in: [*Application manager*](#)

The PROJECT_NAME variable from the build system will not affect the firmware image. This value will not be contained in the `esp_app_desc` structure.

Default value:

- No (disabled)

CONFIG_APP_PROJECT_VER_FROM_CONFIG

Get the project version from Kconfig

Found in: [Application manager](#)

If this is enabled, then config item APP_PROJECT_VER will be used for the variable PROJECT_VER. Other ways to set PROJECT_VER will be ignored.

Default value:

- No (disabled)

CONFIG_APP_PROJECT_VER

Project version

Found in: [Application manager](#) > [CONFIG_APP_PROJECT_VER_FROM_CONFIG](#)

Project version

Default value:

- 1 if [CONFIG_APP_PROJECT_VER_FROM_CONFIG](#)

CONFIG_APP_RETRIEVE_LEN_ELF_SHA

The length of APP ELF SHA is stored in RAM(chars)

Found in: [Application manager](#)

At startup, the app will read the embedded APP ELF SHA-256 hash value from flash and convert it into a string and store it in a RAM buffer. This ensures the panic handler and core dump will be able to print this string even when cache is disabled. The size of the buffer is APP_RETRIEVE_LEN_ELF_SHA plus the null terminator. Changing this value will change the size of this buffer, in bytes.

Range:

- from 8 to 64

Default value:

- 9

Boot ROM Behavior

Contains:

- [CONFIG_BOOT_ROM_LOG_SCHEME](#)

CONFIG_BOOT_ROM_LOG_SCHEME

Permanently change Boot ROM output

Found in: [Boot ROM Behavior](#)

Controls the Boot ROM log behavior. The rom log behavior can only be changed for once, specific eFuse bit(s) will be burned at app boot stage.

Available options:

- Always Log (CONFIG_BOOT_ROM_LOG_ALWAYS_ON)
Always print ROM logs, this is the default behavior.
- Permanently disable logging (CONFIG_BOOT_ROM_LOG_ALWAYS_OFF)
Don't print ROM logs.

- Log on GPIO High (CONFIG_BOOT_ROM_LOG_ON_GPIO_HIGH)
Print ROM logs when GPIO level is high during start up. The GPIO number is chip dependent, e.g. on ESP32-S2, the control GPIO is GPIO46.
- Log on GPIO Low (CONFIG_BOOT_ROM_LOG_ON_GPIO_LOW)
Print ROM logs when GPIO level is low during start up. The GPIO number is chip dependent, e.g. on ESP32-S2, the control GPIO is GPIO46.

Serial flasher config

Contains:

- [*CONFIG_ESPTOOLPY_AFTER*](#)
- [*CONFIG_ESPTOOLPY_BEFORE*](#)
- [*CONFIG_ESPTOOLPY_HEADER_FLASHSIZE_UPDATE*](#)
- [*CONFIG_ESPTOOLPY_NO_STUB*](#)
- [*CONFIG_ESPTOOLPY_FLASH_SAMPLE_MODE*](#)
- [*CONFIG_ESPTOOLPY_FLASHSIZE*](#)
- [*CONFIG_ESPTOOLPY_FLASHMODE*](#)
- [*CONFIG_ESPTOOLPY_FLASHFREQ*](#)

CONFIG_ESPTOOLPY_NO_STUB

Disable download stub

Found in: [*Serial flasher config*](#)

The flasher tool sends a precompiled download stub first by default. That stub allows things like compressed downloads and more. Usually you should not need to disable that feature

CONFIG_ESPTOOLPY_FLASHMODE

Flash SPI mode

Found in: [*Serial flasher config*](#)

Mode the flash chip is flashed in, as well as the default mode for the binary to run in.

Available options:

- QIO (CONFIG_ESPTOOLPY_FLASHMODE_QIO)
- QOUT (CONFIG_ESPTOOLPY_FLASHMODE_QOUT)
- DIO (CONFIG_ESPTOOLPY_FLASHMODE_DIO)
- DOUT (CONFIG_ESPTOOLPY_FLASHMODE_DOUT)
- OPI (CONFIG_ESPTOOLPY_FLASHMODE_OPI)

CONFIG_ESPTOOLPY_FLASH_SAMPLE_MODE

Flash Sampling Mode

Found in: [*Serial flasher config*](#)

Available options:

- STR Mode (CONFIG_ESPTOOLPY_FLASH_SAMPLE_MODE_STR)
- DTR Mode (CONFIG_ESPTOOLPY_FLASH_SAMPLE_MODE_DTR)

CONFIG_ESPTOOLPY_FLASHFREQ

Flash SPI speed

Found in: [Serial flasher config](#)

Available options:

- 120 MHz (READ DOCS FIRST) (CONFIG_ESPTOOLPY_FLASHFREQ_120M)
 - Optional feature for QSPI Flash. Read docs and enable *CONFIG_SPI_FLASH_HPM_ENA* first!
 - Flash 120 MHz SDR mode is stable.
 - Flash 120 MHz DDR mode is an experimental feature, it works when the temperature is stable.

Risks: If your chip powers on at a certain temperature, then after the temperature increases or decreases by approximately 20 Celsius degrees (depending on the chip), the program will crash randomly.

- 80 MHz (CONFIG_ESPTOOLPY_FLASHFREQ_80M)
- 64 MHz (CONFIG_ESPTOOLPY_FLASHFREQ_64M)
- 60 MHz (CONFIG_ESPTOOLPY_FLASHFREQ_60M)
- 48 MHz (CONFIG_ESPTOOLPY_FLASHFREQ_48M)
- 40 MHz (CONFIG_ESPTOOLPY_FLASHFREQ_40M)
- 32 MHz (CONFIG_ESPTOOLPY_FLASHFREQ_32M)
- 30 MHz (CONFIG_ESPTOOLPY_FLASHFREQ_30M)
- 26 MHz (CONFIG_ESPTOOLPY_FLASHFREQ_26M)
- 24 MHz (CONFIG_ESPTOOLPY_FLASHFREQ_24M)
- 20 MHz (CONFIG_ESPTOOLPY_FLASHFREQ_20M)
- 16 MHz (CONFIG_ESPTOOLPY_FLASHFREQ_16M)
- 15 MHz (CONFIG_ESPTOOLPY_FLASHFREQ_15M)

CONFIG_ESPTOOLPY_FLASHSIZE

Flash size

Found in: [Serial flasher config](#)

SPI flash size, in megabytes

Available options:

- 1 MB (CONFIG_ESPTOOLPY_FLASHSIZE_1MB)
- 2 MB (CONFIG_ESPTOOLPY_FLASHSIZE_2MB)
- 4 MB (CONFIG_ESPTOOLPY_FLASHSIZE_4MB)
- 8 MB (CONFIG_ESPTOOLPY_FLASHSIZE_8MB)
- 16 MB (CONFIG_ESPTOOLPY_FLASHSIZE_16MB)
- 32 MB (CONFIG_ESPTOOLPY_FLASHSIZE_32MB)
- 64 MB (CONFIG_ESPTOOLPY_FLASHSIZE_64MB)
- 128 MB (CONFIG_ESPTOOLPY_FLASHSIZE_128MB)

CONFIG_ESPTOOLPY_HEADER_FLASHSIZE_UPDATE

Detect flash size when flashing bootloader

Found in: [Serial flasher config](#)

If this option is set, flashing the project will automatically detect the flash size of the target chip and update the bootloader image before it is flashed.

Enabling this option turns off the image protection against corruption by a SHA256 digest. Updating the bootloader image before flashing would invalidate the digest.

CONFIG_ESPTOOLPY_BEFORE

Before flashing

Found in: [Serial flasher config](#)

Configure whether esptool.py should reset the ESP32 before flashing.

Automatic resetting depends on the RTS & DTR signals being wired from the serial port to the ESP32. Most USB development boards do this internally.

Available options:

- Reset to bootloader (CONFIG_ESPTOOLPY_BEFORE_RESET)
- No reset (CONFIG_ESPTOOLPY_BEFORE_NORESET)

CONFIG_ESPTOOLPY_AFTER

After flashing

Found in: [Serial flasher config](#)

Configure whether esptool.py should reset the ESP32 after flashing.

Automatic resetting depends on the RTS & DTR signals being wired from the serial port to the ESP32. Most USB development boards do this internally.

Available options:

- Reset after flashing (CONFIG_ESPTOOLPY_AFTER_RESET)
- Stay in bootloader (CONFIG_ESPTOOLPY_AFTER_NORESET)

Partition Table

Contains:

- [CONFIG_PARTITION_TABLE_CUSTOM_FILENAME](#)
- [CONFIG_PARTITION_TABLE_MD5](#)
- [CONFIG_PARTITION_TABLE_OFFSET](#)
- [CONFIG_PARTITION_TABLE_TYPE](#)

CONFIG_PARTITION_TABLE_TYPE

Partition Table

Found in: [Partition Table](#)

The partition table to flash to the ESP32. The partition table determines where apps, data and other resources are expected to be found.

The predefined partition table CSV descriptions can be found in the components/partition_table directory. These are mostly intended for example and development use, it's expected that for production use you will copy one of these CSV files and create a custom partition CSV for your application.

Available options:

- Single factory app, no OTA (CONFIG_PARTITION_TABLE_SINGLE_APP)
This is the default partition table, designed to fit into a 2MB or larger flash with a single 1MB app partition.
The corresponding CSV file in the IDF directory is components/partition_table/partitions_singleapp.csv

This partition table is not suitable for an app that needs OTA (over the air update) capability.

- Single factory app (large), no OTA (CONFIG_PARTITION_TABLE_SINGLE_APP_LARGE)
This is a variation of the default partition table, that expands the 1MB app partition size to 1.5MB to fit more code.

The corresponding CSV file in the IDF directory is `components/partition_table/partitions_singleapp_large.csv`

This partition table is not suitable for an app that needs OTA (over the air update) capability.

- Factory app, two OTA definitions (CONFIG_PARTITION_TABLE_TWO_OTA)
This is a basic OTA-enabled partition table with a factory app partition plus two OTA app partitions. All are 1MB, so this partition table requires 4MB or larger flash size.
The corresponding CSV file in the IDF directory is `components/partition_table/partitions_two_ota.csv`

- Custom partition table CSV (CONFIG_PARTITION_TABLE_CUSTOM)
Specify the path to the partition table CSV to use for your project.
Consult the Partition Table section in the ESP-IDF Programmers Guide for more information.

- Single factory app, no OTA, encrypted NVS (CONFIG_PARTITION_TABLE_SINGLE_APP_ENCRYPTED_NVS)
This is a variation of the default "Single factory app, no OTA" partition table that supports encrypted NVS when using flash encryption. See the Flash Encryption section in the ESP-IDF Programmers Guide for more information.

The corresponding CSV file in the IDF directory is `components/partition_table/partitions_singleapp_encr_nvs.csv`

- Single factory app (large), no OTA, encrypted NVS (CONFIG_PARTITION_TABLE_SINGLE_APP_LARGE_ENC_NVS)
This is a variation of the "Single factory app (large), no OTA" partition table that supports encrypted NVS when using flash encryption. See the Flash Encryption section in the ESP-IDF Programmers Guide for more information.

The corresponding CSV file in the IDF directory is `components/partition_table/partitions_singleapp_large_encr_nvs.csv`

- Factory app, two OTA definitions, encrypted NVS (CONFIG_PARTITION_TABLE_TWO_OTA_ENCRYPTED_NVS)
This is a variation of the "Factory app, two OTA definitions" partition table that supports encrypted NVS when using flash encryption. See the Flash Encryption section in the ESP-IDF Programmers Guide for more information.

The corresponding CSV file in the IDF directory is `components/partition_table/partitions_two_ota_encr_nvs.csv`

CONFIG_PARTITION_TABLE_CUSTOM_FILENAME

Custom partition CSV file

Found in: [Partition Table](#)

Name of the custom partition CSV filename. This path is evaluated relative to the project root directory.

Default value:

- "partitions.csv"

CONFIG_PARTITION_TABLE_OFFSET

Offset of partition table

Found in: [Partition Table](#)

The address of partition table (by default 0x8000). Allows you to move the partition table, it gives more space for the bootloader. Note that the bootloader and app will both need to be compiled with the same PARTITION_TABLE_OFFSET value.

This number should be a multiple of 0x1000.

Note that partition offsets in the partition table CSV file may need to be changed if this value is set to a higher value. To have each partition offset adapt to the configured partition table offset, leave all partition offsets blank in the CSV file.

Default value:

- "0x8000"

CONFIG_PARTITION_TABLE_MD5

Generate an MD5 checksum for the partition table

Found in: [Partition Table](#)

Generate an MD5 checksum for the partition table for protecting the integrity of the table. The generation should be turned off for legacy bootloaders which cannot recognize the MD5 checksum in the partition table.

Default value:

- Yes (enabled)

Compiler options

Contains:

- [CONFIG_COMPILER_OPTIMIZATION_ASSERTION_LEVEL](#)
- [CONFIG_COMPILER_FLOAT_LIB_FROM](#)
- [CONFIG_COMPILER_RT_LIB](#)
- [CONFIG_COMPILER_NO_MERGE_CONSTANTS](#)
- [CONFIG_COMPILER_OPTIMIZATION_CHECKS_SILENT](#)
- [CONFIG_COMPILER_DISABLE_GCC12_WARNINGS](#)
- [CONFIG_COMPILER_DISABLE_GCC13_WARNINGS](#)
- [CONFIG_COMPILER_DUMP_RTL_FILES](#)
- [CONFIG_COMPILER_SAVE_RESTORE_LIBCALLS](#)
- [CONFIG_COMPILER_WARN_WRITE_STRINGS](#)
- [CONFIG_COMPILER_CXX_EXCEPTIONS](#)
- [CONFIG_COMPILER_CXX_RTTI](#)
- [CONFIG_COMPILER_OPTIMIZATION](#)
- [CONFIG_COMPILER_ORPHAN_SECTIONS](#)
- [CONFIG_COMPILER_HIDE_PATHS_MACROS](#)
- [CONFIG_COMPILER_STACK_CHECK_MODE](#)

CONFIG_COMPILER_OPTIMIZATION

Optimization Level

Found in: [Compiler options](#)

This option sets compiler optimization level (gcc -O argument) for the app.

- The "Debug" setting will add the -Og flag to CFLAGS.
- The "Size" setting will add the -Os flag to CFLAGS.
- The "Performance" setting will add the -O2 flag to CFLAGS.
- The "None" setting will add the -O0 flag to CFLAGS.

The "Size" setting cause the compiled code to be smaller and faster, but may lead to difficulties of correlating code addresses to source file lines when debugging.

The "Performance" setting causes the compiled code to be larger and faster, but will be easier to correlated code addresses to source file lines.

"None" with -O0 produces compiled code without optimization.

Note that custom optimization levels may be unsupported.

Compiler optimization for the IDF bootloader is set separately, see the `BOOT-LOADER_COMPILER_OPTIMIZATION` setting.

Available options:

- Debug (-Og) (`CONFIG_COMPILER_OPTIMIZATION_DEBUG`)
- Optimize for size (-Os) (`CONFIG_COMPILER_OPTIMIZATION_SIZE`)
- Optimize for performance (-O2) (`CONFIG_COMPILER_OPTIMIZATION_PERF`)
- Debug without optimization (-O0) (`CONFIG_COMPILER_OPTIMIZATION_NONE`)

CONFIG_COMPILER_OPTIMIZATION_ASSERTION_LEVEL

Assertion level

Found in: *Compiler options*

Assertions can be:

- Enabled. Failure will print verbose assertion details. This is the default.
- Set to "silent" to save code size (failed assertions will abort() but user needs to use the aborting address to find the line number with the failed assertion.)
- Disabled entirely (not recommended for most configurations.) `-DNDEBUG` is added to `CPPFLAGS` in this case.

Available options:

- Enabled (`CONFIG_COMPILER_OPTIMIZATION_ASSERTIONS_ENABLE`)
Enable assertions. Assertion content and line number will be printed on failure.
- Silent (saves code size) (`CONFIG_COMPILER_OPTIMIZATION_ASSERTIONS_SILENT`)
Enable silent assertions. Failed assertions will abort(), user needs to use the aborting address to find the line number with the failed assertion.
- Disabled (sets `-DNDEBUG`) (`CONFIG_COMPILER_OPTIMIZATION_ASSERTIONS_DISABLE`)
If assertions are disabled, `-DNDEBUG` is added to `CPPFLAGS`.

CONFIG_COMPILER_FLOAT_LIB_FROM

Compiler float lib source

Found in: *Compiler options*

In the soft-fp part of libgcc, riscv version is written in C, and handles all edge cases in IEEE754, which makes it larger and performance is slow.

RVfplib is an optimized RISC-V library for FP arithmetic on 32-bit integer processors, for single and double-precision FP. RVfplib is "fast", but it has a few exceptions from IEEE 754 compliance.

Available options:

- libgcc (`CONFIG_COMPILER_FLOAT_LIB_FROM_GCCLIB`)
- librvfp (`CONFIG_COMPILER_FLOAT_LIB_FROM_RVFPLIB`)

CONFIG_COMPILER_OPTIMIZATION_CHECKS_SILENT

Disable messages in `ESP_RETURN_ON_*` and `ESP_EXIT_ON_*` macros

Found in: *Compiler options*

If enabled, the error messages will be discarded in following check macros: -
ESP_RETURN_ON_ERROR - ESP_EXIT_ON_ERROR - ESP_RETURN_ON_FALSE -
ESP_EXIT_ON_FALSE

Default value:

- No (disabled)

CONFIG_COMPILER_HIDE_PATHS_MACROS

Replace ESP-IDF and project paths in binaries

Found in: [Compiler options](#)

When expanding the `__FILE__` and `__BASE_FILE__` macros, replace paths inside ESP-IDF with paths relative to the placeholder string "IDF", and convert paths inside the project directory to relative paths.

This allows building the project with assertions or other code that embeds file paths, without the binary containing the exact path to the IDF or project directories.

This option passes `-fmacro-prefix-map` options to the GCC command line. To replace additional paths in your binaries, modify the project CMakeLists.txt file to pass custom `-fmacro-prefix-map` or `-ffile-prefix-map` arguments.

Default value:

- Yes (enabled)

CONFIG_COMPILER_CXX_EXCEPTIONS

Enable C++ exceptions

Found in: [Compiler options](#)

Enabling this option compiles all IDF C++ files with exception support enabled.

Disabling this option disables C++ exception support in all compiled files, and any libstdc++ code which throws an exception will abort instead.

Enabling this option currently adds an additional ~500 bytes of heap overhead when an exception is thrown in user code for the first time.

Default value:

- No (disabled)

Contains:

- [CONFIG_COMPILER_CXX_EXCEPTIONS_EMG_POOL_SIZE](#)

CONFIG_COMPILER_CXX_EXCEPTIONS_EMG_POOL_SIZE

Emergency Pool Size

Found in: [Compiler options](#) > [CONFIG_COMPILER_CXX_EXCEPTIONS](#)

Size (in bytes) of the emergency memory pool for C++ exceptions. This pool will be used to allocate memory for thrown exceptions when there is not enough memory on the heap.

Default value:

- 0 if [CONFIG_COMPILER_CXX_EXCEPTIONS](#)

CONFIG_COMPILER_CXX_RTTI

Enable C++ run-time type info (RTTI)

Found in: [Compiler options](#)

Enabling this option compiles all C++ files with RTTI support enabled. This increases binary size (typically by tens of kB) but allows using `dynamic_cast` conversion and `typeid` operator.

Default value:

- No (disabled)

CONFIG_COMPILER_STACK_CHECK_MODE

Stack smashing protection mode

Found in: [Compiler options](#)

Stack smashing protection mode. Emit extra code to check for buffer overflows, such as stack smashing attacks. This is done by adding a guard variable to functions with vulnerable objects. The guards are initialized when a function is entered and then checked when the function exits. If a guard check fails, program is halted. Protection has the following modes:

- In NORMAL mode (GCC flag: `-fstack-protector`) only functions that call `alloca`, and functions with buffers larger than 8 bytes are protected.
- STRONG mode (GCC flag: `-fstack-protector-strong`) is like NORMAL, but includes additional functions to be protected -- those that have local array definitions, or have references to local frame addresses.
- In OVERALL mode (GCC flag: `-fstack-protector-all`) all functions are protected.

Modes have the following impact on code performance and coverage:

- performance: NORMAL > STRONG > OVERALL
- coverage: NORMAL < STRONG < OVERALL

The performance impact includes increasing the amount of stack memory required for each task.

Available options:

- None (CONFIG_COMPILER_STACK_CHECK_MODE_NONE)
- Normal (CONFIG_COMPILER_STACK_CHECK_MODE_NORM)
- Strong (CONFIG_COMPILER_STACK_CHECK_MODE_STRONG)
- Overall (CONFIG_COMPILER_STACK_CHECK_MODE_ALL)

CONFIG_COMPILER_NO_MERGE_CONSTANTS

Disable merging const sections

Found in: [Compiler options](#)

Disable merging identical constants (string/floating-point) across compilation units. This helps in better size analysis of the application binary as the rodata section distribution is more uniform across libraries. On downside, it may increase the binary size and hence should be used during development phase only.

CONFIG_COMPILER_WARN_WRITE_STRINGS

Enable `-Wwrite-strings` warning flag

Found in: [Compiler options](#)

Adds `-Wwrite-strings` flag for the C/C++ compilers.

For C, this gives string constants the type `const char[]` so that copying the address of one into a non-const `char *` pointer produces a warning. This warning helps to find at compile time code that tries to write into a string constant.

For C++, this warns about the deprecated conversion from string literals to `char *`.

Default value:

- No (disabled)

CONFIG_COMPILER_SAVE_RESTORE_LIBCALLS

Enable -msave-restore flag to reduce code size

Found in: [Compiler options](#)

Adds -msave-restore to C/C++ compilation flags.

When this flag is enabled, compiler will call library functions to save/restore registers in function prologues/epilogues. This results in lower overall code size, at the expense of slightly reduced performance.

This option can be enabled for RISC-V targets only.

CONFIG_COMPILER_DISABLE_GCC12_WARNINGS

Disable new warnings introduced in GCC 12

Found in: [Compiler options](#)

Enable this option if use GCC 12 or newer, and want to disable warnings which don't appear with GCC 11.

Default value:

- No (disabled)

CONFIG_COMPILER_DISABLE_GCC13_WARNINGS

Disable new warnings introduced in GCC 13

Found in: [Compiler options](#)

Enable this option if use GCC 13 or newer, and want to disable warnings which don't appear with GCC 12.

Default value:

- No (disabled)

CONFIG_COMPILER_DUMP_RTL_FILES

Dump RTL files during compilation

Found in: [Compiler options](#)

If enabled, RTL files will be produced during compilation. These files can be used by other tools, for example to calculate call graphs.

CONFIG_COMPILER_RT_LIB

Compiler runtime library

Found in: [Compiler options](#)

Select runtime library to be used by compiler. - GCC toolchain supports libgcc only. - Clang allows to choose between libgcc or libclang_rt. - For host builds ("linux" target), uses the default library.

Available options:

- libgcc (CONFIG_COMPILER_RT_LIB_GCCLIB)
- libclang_rt (CONFIG_COMPILER_RT_LIB_CLANGRT)
- Host (CONFIG_COMPILER_RT_LIB_HOST)

CONFIG_COMPILER_ORPHAN_SECTIONS

Orphan sections handling

Found in: *Compiler options*

If the linker finds orphan sections, it attempts to place orphan sections after sections of the same attribute such as code vs data, loadable vs non-loadable, etc. That means that orphan sections could be placed between sections defined in IDF linker scripts. This could lead to corruption of the binary image. Configure the linker action here.

Available options:

- Place with warning (CONFIG_COMPILER_ORPHAN_SECTIONS_WARNING)
Places orphan sections with a warning message.
- Place silently (CONFIG_COMPILER_ORPHAN_SECTIONS_PLACE)
Places orphan sections without a warning/error message.

Component config

Contains:

- *ADC and ADC Calibration*
- *Application Level Tracing*
- *Bluetooth*
- *Common ESP-related*
- *Core dump*
- *Driver Configurations*
- *eFuse Bit Manager*
- *CONFIG_BLE_MESH*
- *ESP HID*
- *ESP HTTP client*
- *ESP HTTPS OTA*
- *ESP HTTPS server*
- *ESP NETIF Adapter*
- *ESP PSRAM*
- *ESP Ringbuf*
- *ESP System Settings*
- *ESP-MQTT Configurations*
- *ESP-TLS*
- *Ethernet*
- *Event Loop Library*
- *FAT Filesystem support*
- *FreeRTOS*
- *GDB Stub*
- *Hardware Abstraction Layer (HAL) and Low Level (LL)*
- *Hardware Settings*
- *Heap memory debugging*
- *High resolution timer (esp_timer)*
- *HTTP Server*
- *IEEE 802.15.4*
- *IPC (Inter-Processor Call)*
- *LCD and Touch Panel*
- *Log output*
- *LWIP*
- *Main Flash configuration*
- *mbdTLS*
- *Newlib*
- *NVS*

- *NVS Security Provider*
- *OpenThread*
- *Partition API Configuration*
- *Power Management*
- *Protocomm*
- *PThreads*
- *SoC Settings*
- *SPI Flash driver*
- *SPIFFS Configuration*
- *TCP Transport*
- *Ultra Low Power (ULP) Co-processor*
- *Unity unit testing library*
- *USB-OTG*
- *Virtual file system*
- *Wear Levelling*
- *Wi-Fi*
- *Wi-Fi Provisioning Manager*
- *Wireless Coexistence*

Application Level Tracing Contains:

- *CONFIG_APPTRACE_DESTINATION1*
- *CONFIG_APPTRACE_DESTINATION2*
- *FreeRTOS System View Tracing*
- *CONFIG_APPTRACE_GCOV_ENABLE*
- *CONFIG_APPTRACE_BUF_SIZE*
- *CONFIG_APPTRACE_PENDING_DATA_SIZE_MAX*
- *CONFIG_APPTRACE_POSTMORTEM_FLUSH_THRESH*
- *CONFIG_APPTRACE_ONPANIC_HOST_FLUSH_TMO*
- *CONFIG_APPTRACE_UART_BAUDRATE*
- *CONFIG_APPTRACE_UART_RX_GPIO*
- *CONFIG_APPTRACE_UART_RX_BUFF_SIZE*
- *CONFIG_APPTRACE_UART_TASK_PRIO*
- *CONFIG_APPTRACE_UART_TX_MSG_SIZE*
- *CONFIG_APPTRACE_UART_TX_GPIO*
- *CONFIG_APPTRACE_UART_TX_BUFF_SIZE*

CONFIG_APPTRACE_DESTINATION1

Data Destination 1

Found in: Component config > Application Level Tracing

Select destination for application trace: JTAG or none (to disable).

Available options:

- JTAG (CONFIG_APPTRACE_DEST_JTAG)
- None (CONFIG_APPTRACE_DEST_NONE)

CONFIG_APPTRACE_DESTINATION2

Data Destination 2

Found in: Component config > Application Level Tracing

Select destination for application trace: UART(XX) or none (to disable).

Available options:

- UART0 (CONFIG_APPTRACE_DEST_UART0)
- UART1 (CONFIG_APPTRACE_DEST_UART1)
- UART2 (CONFIG_APPTRACE_DEST_UART2)
- USB_CDC (CONFIG_APPTRACE_DEST_USB_CDC)
- None (CONFIG_APPTRACE_DEST_UART_NONE)

CONFIG_APPTRACE_UART_TX_GPIO

UART TX on GPIO#

Found in: [Component config](#) > [Application Level Tracing](#)

This GPIO is used for UART TX pin.

CONFIG_APPTRACE_UART_RX_GPIO

UART RX on GPIO#

Found in: [Component config](#) > [Application Level Tracing](#)

This GPIO is used for UART RX pin.

CONFIG_APPTRACE_UART_BAUDRATE

UART baud rate

Found in: [Component config](#) > [Application Level Tracing](#)

This baud rate is used for UART.

The app's maximum baud rate depends on the UART clock source. If Power Management is disabled, the UART clock source is the APB clock and all baud rates in the available range will be sufficiently accurate. If Power Management is enabled, REF_TICK clock source is used so the baud rate is divided from 1MHz. Baud rates above 1Mbps are not possible and values between 500Kbps and 1Mbps may not be accurate.

CONFIG_APPTRACE_UART_RX_BUFF_SIZE

UART RX ring buffer size

Found in: [Component config](#) > [Application Level Tracing](#)

Size of the UART input ring buffer. This size related to the baudrate, system tick frequency and amount of data to transfer. The data placed to this buffer before sent out to the interface.

CONFIG_APPTRACE_UART_TX_BUFF_SIZE

UART TX ring buffer size

Found in: [Component config](#) > [Application Level Tracing](#)

Size of the UART output ring buffer. This size related to the baudrate, system tick frequency and amount of data to transfer.

CONFIG_APPTRACE_UART_TX_MSG_SIZE

UART TX message size

Found in: [Component config](#) > [Application Level Tracing](#)

Maximum size of the single message to transfer.

CONFIG_APPTRACE_UART_TASK_PRIO

UART Task Priority

Found in: [Component config](#) > [Application Level Tracing](#)

UART task priority. In case of high events rate, this parameter could be changed up to (config-MAX_PRIORITIES-1).

Range:

- from 1 to 32

Default value:

- 1

CONFIG_APPTRACE_ONPANIC_HOST_FLUSH_TMO

Timeout for flushing last trace data to host on panic

Found in: [Component config](#) > [Application Level Tracing](#)

Timeout for flushing last trace data to host in case of panic. In ms. Use -1 to disable timeout and wait forever.

CONFIG_APPTRACE_POSTMORTEM_FLUSH_THRESH

Threshold for flushing last trace data to host on panic

Found in: [Component config](#) > [Application Level Tracing](#)

Threshold for flushing last trace data to host on panic in post-mortem mode. This is minimal amount of data needed to perform flush. In bytes.

CONFIG_APPTRACE_BUF_SIZE

Size of the apptrace buffer

Found in: [Component config](#) > [Application Level Tracing](#)

Size of the memory buffer for trace data in bytes.

CONFIG_APPTRACE_PENDING_DATA_SIZE_MAX

Size of the pending data buffer

Found in: [Component config](#) > [Application Level Tracing](#)

Size of the buffer for events in bytes. It is useful for buffering events from the time critical code (scheduler, ISRs etc). If this parameter is 0 then events will be discarded when main HW buffer is full.

FreeRTOS SystemView Tracing

 Contains:

- [CONFIG_APPTRACE_SV_CPU](#)
- [CONFIG_APPTRACE_SV_EVT_ISR_ENTER_ENABLE](#)
- [CONFIG_APPTRACE_SV_EVT_ISR_EXIT_ENABLE](#)
- [CONFIG_APPTRACE_SV_EVT_ISR_TO_SCHED_ENABLE](#)
- [CONFIG_APPTRACE_SV_MAX_TASKS](#)
- [CONFIG_APPTRACE_SV_EVT_IDLE_ENABLE](#)
- [CONFIG_APPTRACE_SV_ENABLE](#)
- [CONFIG_APPTRACE_SV_EVT_TASK_CREATE_ENABLE](#)
- [CONFIG_APPTRACE_SV_EVT_TASK_START_EXEC_ENABLE](#)
- [CONFIG_APPTRACE_SV_EVT_TASK_START_READY_ENABLE](#)
- [CONFIG_APPTRACE_SV_EVT_TASK_STOP_EXEC_ENABLE](#)
- [CONFIG_APPTRACE_SV_EVT_TASK_STOP_READY_ENABLE](#)

- [CONFIG_APPTRACE_SV_EVT_TASK_TERMINATE_ENABLE](#)
- [CONFIG_APPTRACE_SV_EVT_TIMER_ENTER_ENABLE](#)
- [CONFIG_APPTRACE_SV_EVT_TIMER_EXIT_ENABLE](#)
- [CONFIG_APPTRACE_SV_TS_SOURCE](#)
- [CONFIG_APPTRACE_SV_EVT_OVERFLOW_ENABLE](#)
- [CONFIG_APPTRACE_SV_BUF_WAIT_TMO](#)

CONFIG_APPTRACE_SV_ENABLE

SystemView Tracing Enable

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Enables support for SEGGER SystemView tracing functionality.

CONFIG_APPTRACE_SV_DEST

SystemView destination

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#) > [CONFIG_APPTRACE_SV_ENABLE](#)

SystemView will transfer data through defined interface.

Available options:

- Data destination JTAG (CONFIG_APPTRACE_SV_DEST_JTAG)
Send SEGGER SystemView events through JTAG interface.
- Data destination UART (CONFIG_APPTRACE_SV_DEST_UART)
Send SEGGER SystemView events through UART interface.

CONFIG_APPTRACE_SV_CPU

CPU to trace

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Define the CPU to trace by SystemView.

Available options:

- CPU0 (CONFIG_APPTRACE_SV_DEST_CPU_0)
Send SEGGER SystemView events for Pro CPU.
- CPU1 (CONFIG_APPTRACE_SV_DEST_CPU_1)
Send SEGGER SystemView events for App CPU.

CONFIG_APPTRACE_SV_TS_SOURCE

Timer to use as timestamp source

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

SystemView needs to use a hardware timer as the source of timestamps when tracing. This option selects the timer for it.

Available options:

- CPU cycle counter (CCOUNT) (CONFIG_APPTRACE_SV_TS_SOURCE_CCOUNT)
- General Purpose Timer (Timer Group) (CONFIG_APPTRACE_SV_TS_SOURCE_GPTIMER)

- esp_timer high resolution timer (CONFIG_APPTRACE_SV_TS_SOURCE_ESP_TIMER)

CONFIG_APPTRACE_SV_MAX_TASKS

Maximum supported tasks

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Configures maximum supported tasks in sysview debug

CONFIG_APPTRACE_SV_BUF_WAIT_TMO

Trace buffer wait timeout

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Configures timeout (in us) to wait for free space in trace buffer. Set to -1 to wait forever and avoid lost events.

CONFIG_APPTRACE_SV_EVT_OVERFLOW_ENABLE

Trace Buffer Overflow Event

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Enables "Trace Buffer Overflow" event.

CONFIG_APPTRACE_SV_EVT_ISR_ENTER_ENABLE

ISR Enter Event

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Enables "ISR Enter" event.

CONFIG_APPTRACE_SV_EVT_ISR_EXIT_ENABLE

ISR Exit Event

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Enables "ISR Exit" event.

CONFIG_APPTRACE_SV_EVT_ISR_TO_SCHED_ENABLE

ISR Exit to Scheduler Event

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Enables "ISR to Scheduler" event.

CONFIG_APPTRACE_SV_EVT_TASK_START_EXEC_ENABLE

Task Start Execution Event

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Enables "Task Start Execution" event.

CONFIG_APPTRACE_SV_EVT_TASK_STOP_EXEC_ENABLE

Task Stop Execution Event

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Enables "Task Stop Execution" event.

CONFIG_APPTRACE_SV_EVT_TASK_START_READY_ENABLE

Task Start Ready State Event

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Enables "Task Start Ready State" event.

CONFIG_APPTRACE_SV_EVT_TASK_STOP_READY_ENABLE

Task Stop Ready State Event

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Enables "Task Stop Ready State" event.

CONFIG_APPTRACE_SV_EVT_TASK_CREATE_ENABLE

Task Create Event

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Enables "Task Create" event.

CONFIG_APPTRACE_SV_EVT_TASK_TERMINATE_ENABLE

Task Terminate Event

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Enables "Task Terminate" event.

CONFIG_APPTRACE_SV_EVT_IDLE_ENABLE

System Idle Event

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Enables "System Idle" event.

CONFIG_APPTRACE_SV_EVT_TIMER_ENTER_ENABLE

Timer Enter Event

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Enables "Timer Enter" event.

CONFIG_APPTRACE_SV_EVT_TIMER_EXIT_ENABLE

Timer Exit Event

Found in: [Component config](#) > [Application Level Tracing](#) > [FreeRTOS SystemView Tracing](#)

Enables "Timer Exit" event.

CONFIG_APPTRACE_GCOV_ENABLE

GCOV to Host Enable

Found in: [Component config](#) > [Application Level Tracing](#)

Enables support for GCOV data transfer to host.

CONFIG_APPTRACE_GCOV_DUMP_TASK_STACK_SIZE

Gcov dump task stack size

Found in: [Component config](#) > [Application Level Tracing](#) > [CONFIG_APPTRACE_GCOV_ENABLE](#)

Configures stack size of Gcov dump task

Default value:

- 2048 if [CONFIG_APPTRACE_GCOV_ENABLE](#)

Bluetooth Contains:

- [Bluedroid Options](#)
- [CONFIG_BT_ENABLED](#)
- [Common Options](#)
- [Controller Options](#)
- [CONFIG_BT_HCI_LOG_DEBUG_EN](#)
- [NimBLE Options](#)
- [CONFIG_BT_RELEASE_IRAM](#)

CONFIG_BT_ENABLED

Bluetooth

Found in: [Component config](#) > [Bluetooth](#)

Select this option to enable Bluetooth and show the submenu with Bluetooth configuration choices.

CONFIG_BT_HOST

Host

Found in: [Component config](#) > [Bluetooth](#) > [CONFIG_BT_ENABLED](#)

This helps to choose Bluetooth host stack

Available options:

- **Bluedroid - Dual-mode** ([CONFIG_BT_BLUEDROID_ENABLED](#))
This option is recommended for classic Bluetooth or for dual-mode usecases
- **NimBLE - BLE only** ([CONFIG_BT_NIMBLE_ENABLED](#))
This option is recommended for BLE only usecases to save on memory
- **Disabled** ([CONFIG_BT_CONTROLLER_ONLY](#))
This option is recommended when you want to communicate directly with the controller (without any host) or when you are using any other host stack not supported by Espressif (not mentioned here).

CONFIG_BT_CONTROLLER

Controller

Found in: *Component config > Bluetooth > CONFIG_BT_ENABLED*

This helps to choose Bluetooth controller stack

Available options:

- Enabled (CONFIG_BT_CONTROLLER_ENABLED)
This option is recommended for Bluetooth controller usecases
- Disabled (CONFIG_BT_CONTROLLER_DISABLED)
This option is recommended for Bluetooth Host only usecases

Bluedroid Options Contains:

- *CONFIG_BT_ABORT_WHEN_ALLOCATION_FAILS*
- *Bluedroid debug option*
- *CONFIG_BT_BTU_TASK_STACK_SIZE*
- *CONFIG_BT_BTC_TASK_STACK_SIZE*
- *CONFIG_BT_BLE_ENABLED*
- *BT DEBUG LOG LEVEL*
- *CONFIG_BT_ACL_CONNECTIONS*
- *CONFIG_BT_SMP_MAX_BONDS*
- *CONFIG_BT_ALLOCATION_FROM_SPIRAM_FIRST*
- *CONFIG_BT_CLASSIC_ENABLED*
- *CONFIG_BT_HID_ENABLED*
- *CONFIG_BT_STACK_NO_LOG*
- *CONFIG_BT_BLE_FEAT_ADV_CODING_SELECTION*
- *CONFIG_BT_BLE_42_FEATURES_SUPPORTED*
- *CONFIG_BT_BLE_50_FEATURES_SUPPORTED*
- *CONFIG_BT_BLE_FEAT_ISO_EN*
- *CONFIG_BT_BLE_FEAT_ADV_MONITOR*
- *CONFIG_BT_BLE_FEAT_CHANNEL_SOUNDING*
- *CONFIG_BT_BLE_FEAT_CONN_SUBRATING*
- *CONFIG_BT_BLE_FEAT_CTE_EN*
- *CONFIG_BT_BLE_HIGH_DUTY_ADV_INTERVAL*
- *CONFIG_BT_MULTI_CONNECTION_ENBALE*
- *CONFIG_BT_BLE_FEAT_PERIODIC_ADV_SYNC_TRANSFER*
- *CONFIG_BT_BLE_FEAT_POWER_CONTROL*
- *CONFIG_BT_BLE_VENDOR_HCI_EN*
- *CONFIG_BT_BLE_FEAT_CREATE_SYNC_ENH*
- *CONFIG_BT_BLUEDROID_ESP_COEX_VSC*
- *CONFIG_BT_BLE_FEAT_PERIODIC_ADV_ENH*
- *CONFIG_BT_BLE_FEAT_PAWR_EN*
- *CONFIG_BT_MAX_DEVICE_NAME_LEN*
- *CONFIG_BT_BLE_FEAT_CREATE_SYNC_RETRY_MAX*
- *CONFIG_BT_BLE_ACT_SCAN_REP_ADV_SCAN*
- *CONFIG_BT_BLUEDROID_PINNED_TO_CORE_CHOICE*
- *CONFIG_BT_BLE_RPA_TIMEOUT*
- *CONFIG_BT_BLE_DYNAMIC_ENV_MEMORY*
- *CONFIG_BT_HFP_WBS_ENABLE*

CONFIG_BT_BTC_TASK_STACK_SIZE

Bluetooth event (callback to application) task stack size

Found in: *Component config > Bluetooth > Bluedroid Options*

This select btc task stack size

Default value:

- 3072 if `CONFIG_BT_BLUEDROID_ENABLED` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLUEDROID_PINNED_TO_CORE_CHOICE

The cpu core which Bluedroid run

Found in: *Component config > Bluetooth > Bluedroid Options*

Which the cpu core to run Bluedroid. Can choose core0 and core1. Can not specify no-affinity.

Available options:

- Core 0 (PRO CPU) (`CONFIG_BT_BLUEDROID_PINNED_TO_CORE_0`)
- Core 1 (APP CPU) (`CONFIG_BT_BLUEDROID_PINNED_TO_CORE_1`)

CONFIG_BT_BTU_TASK_STACK_SIZE

Bluetooth Bluedroid Host Stack task stack size

Found in: *Component config > Bluetooth > Bluedroid Options*

This select btu task stack size

Default value:

- 4352 if `CONFIG_BT_BLUEDROID_ENABLED` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLUEDROID_ESP_COEX_VSC

Enable Espressif Vendor-specific HCI commands for coexist status configuration

Found in: *Component config > Bluetooth > Bluedroid Options*

Enable Espressif Vendor-specific HCI commands for coexist status configuration

Default value:

- Yes (enabled) if (`CONFIG_ESP_COEX_SW_COEXIST_ENABLE` || `CONFIG_BT_CONTROLLER_DISABLED`) && `CONFIG_BT_BLUEDROID_ENABLED` && `CONFIG_BT_BLUEDROID_ENABLED`
- No (disabled) if `CONFIG_BT_BLUEDROID_ENABLED` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_CLASSIC_ENABLED

Classic Bluetooth

Found in: *Component config > Bluetooth > Bluedroid Options*

For now this option needs "SMP_ENABLE" to be set to yes

Default value:

- No (disabled) if `CONFIG_BT_BLUEDROID_ENABLED` && ((`CONFIG_BT_CONTROLLER_ENABLED` && `SOC_BT_CLASSIC_SUPPORTED`) || `CONFIG_BT_CONTROLLER_DISABLED`) && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_CLASSIC_ENABLE_POWER_CTRL_VSC

Enable Espressif Vendor-specific HCI commands for power control of Classic Bluetooth

Found in: [Component config](#) > [Bluetooth](#) > [Bluebird Options](#) > [CONFIG_BT_CLASSIC_ENABLED](#)

Enable Espressif Vendor-specific HCI commands for power control of Classic Bluetooth

Default value:

- No (disabled) if [CONFIG_BT_CLASSIC_ENABLED](#) && [CONFIG_BT_BLUEBIRD_ENABLED](#)

CONFIG_BT_CLASSIC_BQB_ENABLED

Host Qualitification support for Classic Bluetooth

Found in: [Component config](#) > [Bluetooth](#) > [Bluebird Options](#) > [CONFIG_BT_CLASSIC_ENABLED](#)

This enables functionalities of Host qualification for Classic Bluetooth.

Default value:

- No (disabled) if [CONFIG_BT_CLASSIC_ENABLED](#) && [CONFIG_BT_BLUEBIRD_ENABLED](#)

CONFIG_BT_A2DP_ENABLE

A2DP

Found in: [Component config](#) > [Bluetooth](#) > [Bluebird Options](#) > [CONFIG_BT_CLASSIC_ENABLED](#)

Advanced Audio Distribution Profile

Default value:

- No (disabled) if [CONFIG_BT_CLASSIC_ENABLED](#) && [CONFIG_BT_BLUEBIRD_ENABLED](#)

CONFIG_BT_SPP_ENABLED

SPP

Found in: [Component config](#) > [Bluetooth](#) > [Bluebird Options](#) > [CONFIG_BT_CLASSIC_ENABLED](#)

This enables the Serial Port Profile

Default value:

- No (disabled) if [CONFIG_BT_CLASSIC_ENABLED](#) && [CONFIG_BT_BLUEBIRD_ENABLED](#)

CONFIG_BT_L2CAP_ENABLED

BT L2CAP

Found in: [Component config](#) > [Bluetooth](#) > [Bluebird Options](#) > [CONFIG_BT_CLASSIC_ENABLED](#)

This enables the Logical Link Control and Adaptation Layer Protocol. Only supported classic bluetooth.

Default value:

- No (disabled) if [CONFIG_BT_CLASSIC_ENABLED](#) && [CONFIG_BT_BLUEBIRD_ENABLED](#)

CONFIG_BT_HFP_ENABLE

Hands Free/Handset Profile

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_CLASSIC_ENABLED](#)

Hands Free Unit and Audio Gateway can be included simultaneously but they cannot run simultaneously due to internal limitations.

Default value:

- No (disabled) if [CONFIG_BT_CLASSIC_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

Contains:

- [CONFIG_BT_HFP_AG_ENABLE](#)
- [CONFIG_BT_HFP_AUDIO_DATA_PATH](#)
- [CONFIG_BT_HFP_CLIENT_ENABLE](#)

CONFIG_BT_HFP_CLIENT_ENABLE

Hands Free Unit

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_CLASSIC_ENABLED](#) > [CONFIG_BT_HFP_ENABLE](#)

Default value:

- Yes (enabled) if [CONFIG_BT_HFP_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_HFP_AG_ENABLE

Audio Gateway

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_CLASSIC_ENABLED](#) > [CONFIG_BT_HFP_ENABLE](#)

Default value:

- Yes (enabled) if [CONFIG_BT_HFP_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_HFP_AUDIO_DATA_PATH

audio(SCO) data path

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_CLASSIC_ENABLED](#) > [CONFIG_BT_HFP_ENABLE](#)

SCO data path, i.e. HCI or PCM. This option is set using API "esp_bredr_sco_datapath_set" in Bluetooth host. Default SCO data path can also be set in Bluetooth Controller.

Available options:

- PCM ([CONFIG_BT_HFP_AUDIO_DATA_PATH_PCM](#))
- HCI ([CONFIG_BT_HFP_AUDIO_DATA_PATH_HCI](#))

CONFIG_BT_HFP_WBS_ENABLE

Wide Band Speech

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

This enables Wide Band Speech. Should disable it when SCO data path is PCM. Otherwise there will be no data transmitted via GPIOs.

Default value:

- Yes (enabled) if `CONFIG_BT_HFP_AUDIO_DATA_PATH_HCI` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_HID_ENABLED

Classic BT HID

Found in: [Component config > Bluetooth > Bluedroid Options](#)

This enables the BT HID Host

Default value:

- No (disabled) if `CONFIG_BT_CLASSIC_ENABLED` && `CONFIG_BT_BLUEDROID_ENABLED`

Contains:

- `CONFIG_BT_HID_DEVICE_ENABLED`
- `CONFIG_BT_HID_HOST_ENABLED`

CONFIG_BT_HID_HOST_ENABLED

Classic BT HID Host

Found in: [Component config > Bluetooth > Bluedroid Options > CONFIG_BT_HID_ENABLED](#)

This enables the BT HID Host

Default value:

- No (disabled) if `CONFIG_BT_HID_ENABLED` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_HID_DEVICE_ENABLED

Classic BT HID Device

Found in: [Component config > Bluetooth > Bluedroid Options > CONFIG_BT_HID_ENABLED](#)

This enables the BT HID Device

CONFIG_BT_BLE_ENABLED

Bluetooth Low Energy

Found in: [Component config > Bluetooth > Bluedroid Options](#)

This enables Bluetooth Low Energy

Default value:

- Yes (enabled) if `CONFIG_BT_BLUEDROID_ENABLED` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_GATTS_ENABLE

Include GATT server module(GATTS)

Found in: [Component config > Bluetooth > Bluedroid Options > CONFIG_BT_BLE_ENABLED](#)

This option can be disabled when the app work only on gatt client mode

Default value:

- Yes (enabled) if `CONFIG_BT_BLE_ENABLED` && `CONFIG_BT_BLUEDROID_ENABLED`

Contains:

- `CONFIG_BT_GATTS_APPEARANCE_WRITABLE`
- `CONFIG_BT_GATTS_DEVICE_NAME_WRITABLE`

- [CONFIG_BT_GATTS_KEY_MATERIAL_CHAR](#)
- [CONFIG_BT_GATTS_SECURITY_LEVELS_CHAR](#)
- [CONFIG_BT_GATTS_PPCP_CHAR_GAP](#)
- [CONFIG_BT_GATTS_ROBUST_CACHING_ENABLED](#)
- [CONFIG_BT_GATTS_SEND_SERVICE_CHANGE_MODE](#)
- [CONFIG_BT_BLE_BLUFI_ENABLE](#)
- [CONFIG_BT_GATT_MAX_SR_PROFILES](#)
- [CONFIG_BT_GATT_MAX_SR_ATTRIBUTES](#)

CONFIG_BT_GATTS_PPCP_CHAR_GAP

Enable Peripheral Preferred Connection Parameters characteristic in GAP service

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTS_ENABLE](#)

This enables "Peripheral Preferred Connection Parameters" characteristic (UUID: 0x2A04) in GAP service that has connection parameters like min/max connection interval, slave latency and supervision timeout multiplier

Default value:

- No (disabled) if [CONFIG_BT_GATTS_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_BLE_BLUFI_ENABLE

Include blufi function

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTS_ENABLE](#)

This option can be close when the app does not require blufi function.

Default value:

- No (disabled) if [CONFIG_BT_GATTS_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_GATT_MAX_SR_PROFILES

Max GATT Server Profiles

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTS_ENABLE](#)

Maximum GATT Server Profiles Count

Range:

- from 1 to 32 if [CONFIG_BT_GATTS_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

Default value:

- 8 if [CONFIG_BT_GATTS_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_GATT_MAX_SR_ATTRIBUTES

Max GATT Service Attributes

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTS_ENABLE](#)

Maximum GATT Service Attributes Count

Range:

- from 1 to 500 if [CONFIG_BT_GATTS_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

Default value:

- 100 if `CONFIG_BT_GATTS_ENABLE` && `CONFIG_BT_BLUEDROID_ENABLED` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_GATTS_SEND_SERVICE_CHANGE_MODE

GATTS Service Change Mode

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTS_ENABLE](#)

Service change indication mode for GATT Server.

Available options:

- GATTS manually send service change indication (CONFIG_BT_GATTS_SEND_SERVICE_CHANGE_MANUAL)
Manually send service change indication through API `esp_ble_gatts_send_service_change_indication()`
- GATTS automatically send service change indication (CONFIG_BT_GATTS_SEND_SERVICE_CHANGE_AUTO)
Let Bluedroid handle the service change indication internally

CONFIG_BT_GATTS_ROBUST_CACHING_ENABLED

Enable Robust Caching on Server Side

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTS_ENABLE](#)

This option enables the GATT robust caching feature on the server. if turned on, the Client Supported Features characteristic, Database Hash characteristic, and Server Supported Features characteristic will be included in the GAP SERVICE.

Default value:

- No (disabled) if `CONFIG_BT_GATTS_ENABLE` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_GATTS_DEVICE_NAME_WRITABLE

Allow to write device name by GATT clients

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTS_ENABLE](#)

Enabling this option allows remote GATT clients to write device name

Default value:

- No (disabled) if `CONFIG_BT_GATTS_ENABLE` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_GATTS_APPEARANCE_WRITABLE

Allow to write appearance by GATT clients

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTS_ENABLE](#)

Enabling this option allows remote GATT clients to write appearance

Default value:

- No (disabled) if `CONFIG_BT_GATTS_ENABLE` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_GATTS_SECURITY_LEVELS_CHAR

Enable LE GATT Security Levels Characteristic

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTS_ENABLE](#)

Enable LE GATT Security Levels Characteristic

Default value:

- No (disabled) if [CONFIG_BT_GATTS_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_GATTS_KEY_MATERIAL_CHAR

Enable Encrypted Data Key Material Characteristic

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTS_ENABLE](#)

Enable the Encrypted Data Key Material characteristic in GAP service. This characteristic allows advertising data to be decrypted and authenticated using the key material (session key + IV) as defined in Bluetooth Core Specification Version 5.4. The characteristic requires encrypted link to read.

Default value:

- No (disabled) if [CONFIG_BT_GATTS_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_GATTC_ENABLE

Include GATT client module(GATTC)

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#)

This option can be close when the app work only on gatt server mode

Default value:

- Yes (enabled) if [CONFIG_BT_BLE_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

Contains:

- [CONFIG_BT_GATTC_MAX_CACHE_CHAR](#)
- [CONFIG_BT_GATTC_NOTIF_REG_MAX](#)
- [CONFIG_BT_GATTC_CACHE_NVS_FLASH](#)
- [CONFIG_BT_GATTC_CONNECT_RETRY_COUNT](#)
- [CONFIG_BT_BLE_ESTAB_LINK_CONN_TOUT](#)

CONFIG_BT_GATTC_MAX_CACHE_CHAR

Max gattc cache characteristic for discover

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTC_ENABLE](#)

Maximum GATTC cache characteristic count

Range:

- from 1 to 500 if [CONFIG_BT_GATTC_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

Default value:

- 40 if [CONFIG_BT_GATTC_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_GATTC_NOTIF_REG_MAX

Max gattc notify(indication) register number

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTC_ENABLE](#)

Maximum GATTC notify(indication) register number

Range:

- from 1 to 255 if [CONFIG_BT_GATTC_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

Default value:

- 5 if [CONFIG_BT_GATTC_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_GATTC_CACHE_NVS_FLASH

Save gattc cache data to nvs flash

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTC_ENABLE](#)

This select can save gattc cache data to nvs flash

Default value:

- No (disabled) if [CONFIG_BT_GATTC_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_GATTC_CONNECT_RETRY_COUNT

The number of attempts to reconnect if the connection establishment failed

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTC_ENABLE](#)

The number of attempts to reconnect if the connection establishment failed

Range:

- from 0 to 255 if [CONFIG_BT_GATTC_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

Default value:

- 3 if [CONFIG_BT_GATTC_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_BLE_ESTAB_LINK_CONN_TOUT

Timeout of BLE connection establishment

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_GATTC_ENABLE](#)

Bluetooth Connection establishment maximum time, if connection time exceeds this value, the connection establishment fails, ESP_GATTC_OPEN_EVT or ESP_GATTS_OPEN_EVT is triggered.

Range:

- from 1 to 60 if [CONFIG_BT_GATTC_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

Default value:

- 30 if [CONFIG_BT_GATTC_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_BLE_SMP_ENABLE

Include BLE security module(SMP)

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#)

This option can be close when the app not used the ble security connect.

Default value:

- Yes (enabled) if [CONFIG_BT_BLE_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

Contains:

- [CONFIG_BT_BLE_SMP_ID_RESET_ENABLE](#)
- [CONFIG_BT_BLE_SMP_BOND_NVS_FLASH](#)
- [CONFIG_BT_SMP_SLAVE_CON_PARAMS_UPD_ENABLE](#)

- [CONFIG_BT_BLE_RPA_SUPPORTED](#)

CONFIG_BT_SMP_SLAVE_CON_PARAMS_UPD_ENABLE

Slave enable connection parameters update during pairing

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_BLE_SMP_ENABLE](#)

In order to reduce the pairing time, slave actively initiates connection parameters update during pairing.

Default value:

- No (disabled) if [CONFIG_BT_BLE_SMP_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_BLE_SMP_ID_RESET_ENABLE

Reset device identity when all bonding records are deleted

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_BLE_SMP_ENABLE](#)

There are tracking risks associated with using a fixed or static IRK. If enabled this option, Bluedroid will assign a new randomly-generated IRK when all pairing and bonding records are deleted. This would decrease the ability of a previously paired peer to be used to determine whether a device with which it previously shared an IRK is within range.

Default value:

- No (disabled) if [CONFIG_BT_BLE_SMP_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_BLE_SMP_BOND_NVS_FLASH

Save SMP bonding keys to nvs flash

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_BLE_SMP_ENABLE](#)

This select can save SMP bonding keys to nvs flash

Default value:

- Yes (enabled) if [CONFIG_BT_BLE_SMP_ENABLE](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_BLE_RPA_SUPPORTED

Update RPA to Controller

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_ENABLED](#) > [CONFIG_BT_BLE_SMP_ENABLE](#)

This enables controller RPA list function. For ESP32, ESP32 only support network privacy mode. If this option is enabled, ESP32 will only accept advertising packets from peer devices that contain private address, HW will not receive the advertising packets contain identity address after IRK changed. If this option is disabled, address resolution will be performed in the host, so the functions that require controller to resolve address in the white list cannot be used. This option is disabled by default on ESP32, please enable or disable this option according to your own needs.

For other BLE chips, devices support network privacy mode and device privacy mode, users can switch the two modes according to their own needs. So this option is enabled by default.

Default value:

- Yes (enabled) if `CONFIG_BT_CONTROLLER_DISABLED` && `CONFIG_BT_BLE_SMP_ENABLE` && `CONFIG_BT_CONTROLLER_DISABLED` && `CONFIG_BT_BLUEDROID_ENABLED`

Bluedroid debug option Contains:

- `CONFIG_BT_BLUEDROID_MEM_DEBUG`
- `CONFIG_BT_BLUEDROID_MEM_STATS`
- `CONFIG_BT_BLUEDROID_THREAD_DEBUG`

CONFIG_BT_BLUEDROID_MEM_DEBUG

Bluedroid memory debug

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [Bluedroid debug option](#)

Bluedroid memory debug

Default value:

- No (disabled) if `CONFIG_BT_BLUEDROID_ENABLED` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLUEDROID_MEM_STATS

Bluedroid memory statistics

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [Bluedroid debug option](#)

Enable Bluedroid memory usage statistics

CONFIG_BT_BLUEDROID_THREAD_DEBUG

Bluedroid thread debug

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [Bluedroid debug option](#)

Enable Bluedroid thread debug mode. Used to debug whether the thread is blocked and dump information about the thread's related work queue.

Default value:

- No (disabled) if `CONFIG_BT_BLUEDROID_ENABLED` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLUEDROID_THREAD_BLOCK_TIME

OSI thread block time (in ms)

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [Bluedroid debug option](#) > [CONFIG_BT_BLUEDROID_THREAD_DEBUG](#)

Indicates how long it takes for the thread to execute an item before it is considered blocked.

Default value:

- 1000 if `CONFIG_BT_BLUEDROID_THREAD_DEBUG` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLUEDROID_THREAD_BLOCK_MSG

OSI thread block message count

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [Bluedroid debug option](#) > [CONFIG_BT_BLUEDROID_THREAD_DEBUG](#)

Indicates how many messages are added to the queue while the thread is executing an item before it is considered blocked.

Default value:

- 50 if `CONFIG_BT_BLUEDROID_THREAD_DEBUG` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_STACK_NO_LOG

Disable BT debug logs (minimize bin size)

Found in: *Component config > Bluetooth > Bluedroid Options*

This select can save the rodata code size

Default value:

- No (disabled) if `CONFIG_BT_BLUEDROID_ENABLED` && `CONFIG_BT_BLUEDROID_ENABLED`

BT DEBUG LOG LEVEL Contains:

- `CONFIG_BT_LOG_A2D_TRACE_LEVEL`
- `CONFIG_BT_LOG_APPL_TRACE_LEVEL`
- `CONFIG_BT_LOG_AVCT_TRACE_LEVEL`
- `CONFIG_BT_LOG_AVDT_TRACE_LEVEL`
- `CONFIG_BT_LOG_AVRC_TRACE_LEVEL`
- `CONFIG_BT_LOG_BLUFI_TRACE_LEVEL`
- `CONFIG_BT_LOG_BNEP_TRACE_LEVEL`
- `CONFIG_BT_LOG_BTC_TRACE_LEVEL`
- `CONFIG_BT_LOG_BTIF_TRACE_LEVEL`
- `CONFIG_BT_LOG_BTM_TRACE_LEVEL`
- `CONFIG_BT_LOG_GAP_TRACE_LEVEL`
- `CONFIG_BT_LOG_GATT_TRACE_LEVEL`
- `CONFIG_BT_LOG_HCI_TRACE_LEVEL`
- `CONFIG_BT_LOG_HID_TRACE_LEVEL`
- `CONFIG_BT_LOG_L2CAP_TRACE_LEVEL`
- `CONFIG_BT_LOG_MCA_TRACE_LEVEL`
- `CONFIG_BT_LOG_OSI_TRACE_LEVEL`
- `CONFIG_BT_LOG_PAN_TRACE_LEVEL`
- `CONFIG_BT_LOG_RFCOMM_TRACE_LEVEL`
- `CONFIG_BT_LOG_SDP_TRACE_LEVEL`
- `CONFIG_BT_LOG_SMP_TRACE_LEVEL`

CONFIG_BT_LOG_HCI_TRACE_LEVEL

HCI layer

Found in: *Component config > Bluetooth > Bluedroid Options > BT DEBUG LOG LEVEL*

Define BT trace level for HCI layer

Available options:

- NONE (`CONFIG_BT_LOG_HCI_TRACE_LEVEL_NONE`)
- ERROR (`CONFIG_BT_LOG_HCI_TRACE_LEVEL_ERROR`)
- WARNING (`CONFIG_BT_LOG_HCI_TRACE_LEVEL_WARNING`)
- API (`CONFIG_BT_LOG_HCI_TRACE_LEVEL_API`)
- EVENT (`CONFIG_BT_LOG_HCI_TRACE_LEVEL_EVENT`)
- DEBUG (`CONFIG_BT_LOG_HCI_TRACE_LEVEL_DEBUG`)
- VERBOSE (`CONFIG_BT_LOG_HCI_TRACE_LEVEL_VERBOSE`)

CONFIG_BT_LOG_BT_M_TRACE_LEVEL

BT_M layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for BT_M layer

Available options:

- NONE (CONFIG_BT_LOG_BT_M_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_BT_M_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_BT_M_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_BT_M_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_BT_M_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_BT_M_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_BT_M_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_L2CAP_TRACE_LEVEL

L2CAP layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for L2CAP layer

Available options:

- NONE (CONFIG_BT_LOG_L2CAP_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_L2CAP_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_L2CAP_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_L2CAP_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_L2CAP_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_L2CAP_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_L2CAP_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_RFCOMM_TRACE_LEVEL

RFCOMM layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for RFCOMM layer

Available options:

- NONE (CONFIG_BT_LOG_RFCOMM_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_RFCOMM_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_RFCOMM_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_RFCOMM_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_RFCOMM_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_RFCOMM_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_RFCOMM_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_SDP_TRACE_LEVEL

SDP layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for SDP layer

Available options:

- NONE (CONFIG_BT_LOG_SDP_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_SDP_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_SDP_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_SDP_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_SDP_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_SDP_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_SDP_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_GAP_TRACE_LEVEL

GAP layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluebird Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for GAP layer

Available options:

- NONE (CONFIG_BT_LOG_GAP_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_GAP_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_GAP_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_GAP_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_GAP_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_GAP_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_GAP_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_BNEP_TRACE_LEVEL

BNEP layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluebird Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for BNEP layer

Available options:

- NONE (CONFIG_BT_LOG_BNEP_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_BNEP_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_BNEP_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_BNEP_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_BNEP_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_BNEP_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_BNEP_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_PAN_TRACE_LEVEL

PAN layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluebird Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for PAN layer

Available options:

- NONE (CONFIG_BT_LOG_PAN_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_PAN_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_PAN_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_PAN_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_PAN_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_PAN_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_PAN_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_A2D_TRACE_LEVEL

A2D layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for A2D layer

Available options:

- NONE (CONFIG_BT_LOG_A2D_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_A2D_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_A2D_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_A2D_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_A2D_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_A2D_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_A2D_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_AVDT_TRACE_LEVEL

AVDT layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for AVDT layer

Available options:

- NONE (CONFIG_BT_LOG_AVDT_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_AVDT_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_AVDT_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_AVDT_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_AVDT_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_AVDT_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_AVDT_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_AVCT_TRACE_LEVEL

AVCT layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for AVCT layer

Available options:

- NONE (CONFIG_BT_LOG_AVCT_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_AVCT_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_AVCT_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_AVCT_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_AVCT_TRACE_LEVEL_EVENT)

- DEBUG (CONFIG_BT_LOG_AVCT_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_AVCT_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_AVRC_TRACE_LEVEL

AVRC layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for AVRC layer

Available options:

- NONE (CONFIG_BT_LOG_AVRC_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_AVRC_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_AVRC_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_AVRC_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_AVRC_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_AVRC_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_AVRC_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_MCA_TRACE_LEVEL

MCA layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for MCA layer

Available options:

- NONE (CONFIG_BT_LOG_MCA_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_MCA_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_MCA_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_MCA_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_MCA_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_MCA_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_MCA_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_HID_TRACE_LEVEL

HID layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for HID layer

Available options:

- NONE (CONFIG_BT_LOG_HID_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_HID_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_HID_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_HID_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_HID_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_HID_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_HID_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_APPL_TRACE_LEVEL

APPL layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for APPL layer

Available options:

- NONE (CONFIG_BT_LOG_APPL_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_APPL_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_APPL_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_APPL_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_APPL_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_APPL_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_APPL_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_GATT_TRACE_LEVEL

GATT layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for GATT layer

Available options:

- NONE (CONFIG_BT_LOG_GATT_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_GATT_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_GATT_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_GATT_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_GATT_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_GATT_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_GATT_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_SMP_TRACE_LEVEL

SMP layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for SMP layer

Available options:

- NONE (CONFIG_BT_LOG_SMP_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_SMP_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_SMP_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_SMP_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_SMP_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_SMP_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_SMP_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_BTIF_TRACE_LEVEL

BTIF layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for BTIF layer

Available options:

- NONE (CONFIG_BT_LOG_BTIF_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_BTIF_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_BTIF_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_BTIF_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_BTIF_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_BTIF_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_BTIF_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_BTC_TRACE_LEVEL

BTC layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for BTC layer

Available options:

- NONE (CONFIG_BT_LOG_BTC_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_BTC_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_BTC_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_BTC_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_BTC_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_BTC_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_BTC_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_OSI_TRACE_LEVEL

OSI layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for OSI layer

Available options:

- NONE (CONFIG_BT_LOG_OSI_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_OSI_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_OSI_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_OSI_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_OSI_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_OSI_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_OSI_TRACE_LEVEL_VERBOSE)

CONFIG_BT_LOG_BLUFI_TRACE_LEVEL

BLUFI layer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [BT DEBUG LOG LEVEL](#)

Define BT trace level for BLUFI layer

Available options:

- NONE (CONFIG_BT_LOG_BLUFI_TRACE_LEVEL_NONE)
- ERROR (CONFIG_BT_LOG_BLUFI_TRACE_LEVEL_ERROR)
- WARNING (CONFIG_BT_LOG_BLUFI_TRACE_LEVEL_WARNING)
- API (CONFIG_BT_LOG_BLUFI_TRACE_LEVEL_API)
- EVENT (CONFIG_BT_LOG_BLUFI_TRACE_LEVEL_EVENT)
- DEBUG (CONFIG_BT_LOG_BLUFI_TRACE_LEVEL_DEBUG)
- VERBOSE (CONFIG_BT_LOG_BLUFI_TRACE_LEVEL_VERBOSE)

CONFIG_BT_ACL_CONNECTIONS

BT/BLE MAX ACL CONNECTIONS

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Maximum BT/BLE connection count. The ESP32-C3/S3 chip supports a maximum of 10 instances, including ADV, SCAN and connections. The ESP32-C3/S3 chip can connect up to 9 devices if ADV or SCAN uses only one. If ADV and SCAN are both used, The ESP32-C3/S3 chip is connected to a maximum of 8 devices. Because Bluetooth cannot reclaim used instances once ADV or SCAN is used.

Range:

- from 1 to 50 if [CONFIG_BT_BLUEDROID_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

Default value:

- 4 if [CONFIG_BT_BLUEDROID_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_MULTI_CONNECTION_ENBALE

Enable BLE multi-connections

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Enable this option if there are multiple connections

Default value:

- Yes (enabled) if [CONFIG_BT_BLE_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_ALLOCATION_FROM_SPIRAM_FIRST

BT/BLE will first malloc the memory from the PSRAM

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

This select can save the internal RAM if there have the PSRAM

Default value:

- No (disabled) if [CONFIG_BT_BLUEDROID_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_BLE_DYNAMIC_ENV_MEMORY

Use dynamic memory allocation in BT/BLE stack

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

This select can make the allocation of memory will become more flexible

Default value:

- No (disabled) if [CONFIG_BT_BLUEDROID_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_SMP_MAX_BONDS

BT/BLE maximum bond device count

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

The number of security records for peer devices.

CONFIG_BT_BLE_ACT_SCAN_REP_ADV_SCAN

Report adv data and scan response individually when BLE active scan

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Originally, when doing BLE active scan, Bluedroid will not report adv to application layer until receive scan response. This option is used to disable the behavior. When enable this option, Bluedroid will report adv data or scan response to application layer immediately.

Memory reserved at start of DRAM for Bluetooth stack

Default value:

- No (disabled) if [CONFIG_BT_BLUEDROID_ENABLED](#) && [CONFIG_BT_BLE_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_MAX_DEVICE_NAME_LEN

length of bluetooth device name

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Bluetooth Device name length shall be no larger than 248 octets, If the broadcast data cannot contain the complete device name, then only the shortname will be displayed, the rest parts that can't fit in will be truncated.

Range:

- from 32 to 248 if [CONFIG_BT_BLUEDROID_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

Default value:

- 32 if [CONFIG_BT_BLUEDROID_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_BLE_RPA_TIMEOUT

Timeout of resolvable private address

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

This set RPA timeout of Controller and Host. Default is 900 s (15 minutes). Range is 1 s to 1 hour (3600 s).

Range:

- from 1 to 3600 if [CONFIG_BT_BLE_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

Default value:

- 900 if [CONFIG_BT_BLE_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_BLE_50_FEATURES_SUPPORTED

Enable BLE 5.0 and above features(please disable BLE 4.2 if enable BLE 5.0)

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Enabling this option activates BLE 5.0 features. This option is universally supported in chips that support BLE, except for ESP32. BLE 4.2 and BLE 5.0 cannot be used simultaneously.

Default value:

- Yes (enabled) if `CONFIG_BT_BLE_ENABLED` && `((CONFIG_BT_CONTROLLER_ENABLED && SOC_BLE_50_SUPPORTED) || CONFIG_BT_CONTROLLER_DISABLED)` && `CONFIG_BT_BLUEDROID_ENABLED`

Contains:

- `CONFIG_BT_BLE_50_DTM_TEST_EN`
- `CONFIG_BT_BLE_50_EXTEND_ADV_EN`
- `CONFIG_BT_BLE_50_EXTEND_SCAN_EN`

CONFIG_BT_BLE_50_EXTEND_ADV_EN

Enable BLE extend advertising

Found in: `Component config > Bluetooth > Bluedroid Options > CONFIG_BT_BLE_50_FEATURES_SUPPORTED`

This enables BLE extend advertising

Default value:

- Yes (enabled) if `CONFIG_BT_BLE_50_FEATURES_SUPPORTED` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_50_PERIODIC_ADV_EN

Enable BLE periodic advertising

Found in: `Component config > Bluetooth > Bluedroid Options > CONFIG_BT_BLE_50_FEATURES_SUPPORTED > CONFIG_BT_BLE_50_EXTEND_ADV_EN`

This enables BLE periodic advertising

Default value:

- Yes (enabled) if `CONFIG_BT_BLE_50_EXTEND_ADV_EN` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_50_EXTEND_SCAN_EN

Enable BLE extend scan

Found in: `Component config > Bluetooth > Bluedroid Options > CONFIG_BT_BLE_50_FEATURES_SUPPORTED`

This enables BLE extend scan

Default value:

- Yes (enabled) if `CONFIG_BT_BLE_50_FEATURES_SUPPORTED` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_50_EXTEND_SYNC_EN

Enable BLE periodic advertising sync

Found in: `Component config > Bluetooth > Bluedroid Options > CONFIG_BT_BLE_50_FEATURES_SUPPORTED > CONFIG_BT_BLE_50_EXTEND_SCAN_EN`

This enables BLE periodic advertising sync

Default value:

- Yes (enabled) if `CONFIG_BT_BLE_50_EXTEND_SCAN_EN` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_50_DTM_TEST_EN

Enable BLE 5.0 DTM test

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_50_FEATURES_SUPPORTED](#)

This enables BLE 5.0 direct test mode

Default value:

- Yes (enabled) if [CONFIG_BT_BLE_50_FEATURES_SUPPORTED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_BLE_FEAT_PERIODIC_ADV_SYNC_TRANSFER

Enable BLE periodic advertising sync transfer feature

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

This enables BLE periodic advertising sync transfer feature

Default value:

- No (disabled) if [CONFIG_BT_BLE_50_PERIODIC_ADV_EN](#) && (([CONFIG_BT_CONTROLLER_ENABLED](#) && [SOC_ESP_NIMBLE_CONTROLLER](#)) || [CONFIG_BT_CONTROLLER_DISABLED](#)) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_BLE_FEAT_PERIODIC_ADV_ENH

Enable periodic adv enhancements(adi support)

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Enable the periodic advertising enhancements

Default value:

- No (disabled) if [CONFIG_BT_BLE_50_PERIODIC_ADV_EN](#) && (([CONFIG_BT_CONTROLLER_ENABLED](#) && [SOC_ESP_NIMBLE_CONTROLLER](#)) || [CONFIG_BT_CONTROLLER_DISABLED](#)) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_BLE_FEAT_CREATE_SYNC_ENH

Enable create sync enhancements(reporting disable and duplicate filtering enable support)

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Enable the create sync enhancements

Default value:

- No (disabled) if [CONFIG_BT_BLE_50_EXTEND_SYNC_EN](#) && (([CONFIG_BT_CONTROLLER_ENABLED](#) && [SOC_ESP_NIMBLE_CONTROLLER](#)) || [CONFIG_BT_CONTROLLER_DISABLED](#)) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_BLE_FEAT_CREATE_SYNC_RETRY_MAX

Maximum retry count for periodic advertising create sync

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Set the maximum retry count when periodic advertising create sync fails with error code 0x3E (Connection Failed to be Established). Set to 0 to disable retry. Default is 3.

Range:

- from 0 to 16 if [CONFIG_BT_BLE_50_EXTEND_SYNC_EN](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

Default value:

- 3 if `CONFIG_BT_BLE_50_EXTEND_SYNC_EN` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_FEAT_ISO_EN

Enable BLE 5.2 iso feature

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Enable BLE 5.2 iso

Default value:

- No (disabled) if `CONFIG_BT_BLE_50_FEATURES_SUPPORTED` && ((`CONFIG_BT_CONTROLLER_ENABLED` && `SOC_BLE_AUDIO_SUPPORTED`) || `CONFIG_BT_CONTROLLER_DISABLED`) && `CONFIG_BT_BLUEDROID_ENABLED`

Contains:

- `CONFIG_BT_BLE_FEAT_ISO_BIG_BROADCASTER`
- `CONFIG_BT_BLE_FEAT_ISO_BIG_SYNCER`
- `CONFIG_BT_BLE_FEAT_ISO_CIG_CENTRAL`
- `CONFIG_BT_BLE_FEAT_ISO_CIG_PERIPHERAL`
- `CONFIG_BT_BLE_FEAT_ISO_60_EN`
- `CONFIG_BT_BLE_ISO_BIS_MAX_COUNT`
- `CONFIG_BT_BLE_ISO_CIS_MAX_COUNT`
- `CONFIG_BT_BLE_ISO_FLOW_CONTROL`

CONFIG_BT_BLE_FEAT_ISO_60_EN

Enable ISO v6.0 feature

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > `CONFIG_BT_BLE_FEAT_ISO_EN`

Enable iso 6.0 feature

Default value:

- No (disabled) if `CONFIG_BT_BLE_FEAT_ISO_EN` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_FEAT_ISO_BIG_BROADCASTER

Enable BLE iso BIG broadcaster

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > `CONFIG_BT_BLE_FEAT_ISO_EN`

Enable BLE 5.2 BIG broadcaster

Default value:

- Yes (enabled) if `CONFIG_BT_BLE_FEAT_ISO_EN` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_FEAT_ISO_BIG_SYNCER

Enable BLE iso BIG syncer

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > `CONFIG_BT_BLE_FEAT_ISO_EN`

Enable BLE 5.2 BIG syncer

Default value:

- Yes (enabled) if `CONFIG_BT_BLE_FEAT_ISO_EN` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_ISO_BIS_MAX_COUNT

Maximum bis count

Found in: *Component config > Bluetooth > Bluedroid Options > CONFIG_BT_BLE_FEAT_ISO_EN*

Enable BLE 5.2 CIG peripheral

Range:

- from 1 to 31 if *CONFIG_BT_BLE_FEAT_ISO_EN* && (*CONFIG_BT_BLE_FEAT_ISO_BIG_BROADCASTER* || *CONFIG_BT_BLE_FEAT_ISO_BIG_SYNCER*) && *CONFIG_BT_BLUEDROID_ENABLED*

Default value:

- 2 if *CONFIG_BT_BLE_FEAT_ISO_EN* && (*CONFIG_BT_BLE_FEAT_ISO_BIG_BROADCASTER* || *CONFIG_BT_BLE_FEAT_ISO_BIG_SYNCER*) && *CONFIG_BT_BLUEDROID_ENABLED*

CONFIG_BT_BLE_FEAT_ISO_CIG_CENTRAL

Enable BLE iso CIG central

Found in: *Component config > Bluetooth > Bluedroid Options > CONFIG_BT_BLE_FEAT_ISO_EN*

Enable BLE 5.2 CIG central

Default value:

- Yes (enabled) if *CONFIG_BT_BLE_FEAT_ISO_EN* && *CONFIG_BT_BLUEDROID_ENABLED*

CONFIG_BT_BLE_FEAT_ISO_CIG_PERIPHERAL

Enable BLE iso CIG peripheral

Found in: *Component config > Bluetooth > Bluedroid Options > CONFIG_BT_BLE_FEAT_ISO_EN*

Enable BLE 5.2 CIG peripheral

Default value:

- Yes (enabled) if *CONFIG_BT_BLE_FEAT_ISO_EN* && *CONFIG_BT_BLUEDROID_ENABLED*

CONFIG_BT_BLE_ISO_CIS_MAX_COUNT

Maximum cis count

Found in: *Component config > Bluetooth > Bluedroid Options > CONFIG_BT_BLE_FEAT_ISO_EN*

Enable BLE 5.2 CIG peripheral

Range:

- from 1 to 31 if *CONFIG_BT_BLE_FEAT_ISO_EN* && (*CONFIG_BT_BLE_FEAT_ISO_CIG_CENTRAL* || *CONFIG_BT_BLE_FEAT_ISO_CIG_PERIPHERAL*) && *CONFIG_BT_BLUEDROID_ENABLED*

Default value:

- 2 if *CONFIG_BT_BLE_FEAT_ISO_EN* && (*CONFIG_BT_BLE_FEAT_ISO_CIG_CENTRAL* || *CONFIG_BT_BLE_FEAT_ISO_CIG_PERIPHERAL*) && *CONFIG_BT_BLUEDROID_ENABLED*

CONFIG_BT_BLE_ISO_FLOW_CONTROL

Select ISO flow control type

Found in: *Component config > Bluetooth > Bluedroid Options > CONFIG_BT_BLE_FEAT_ISO_EN*

Select ISO flow control type

Available options:

- ISO standard flow control (CONFIG_BT_BLE_ISO_STD_FLOW_CTRL)
Enable ISO standard flow control
- ISO non-standard flow control (CONFIG_BT_BLE_ISO_NON_STD_FLOW_CTRL)
Enable ISO non-standard flow control

CONFIG_BT_BLE_FEAT_CTE_EN

Enable BLE CTE feature

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Enable BLE 5.1 CTE

Default value:

- No (disabled) if `CONFIG_BT_BLE_50_FEATURES_SUPPORTED` && `((CONFIG_BT_CONTROLLER_ENABLED && SOC_BLE_CTE_SUPPORTED) || CONFIG_BT_CONTROLLER_DISABLED)` && `CONFIG_BT_BLUEDROID_ENABLED`

Contains:

- [CONFIG_BT_BLE_FEAT_CTE_CONNECTION_EN](#)
- [CONFIG_BT_BLE_FEAT_CTE_CONNECTIONLESS_EN](#)

CONFIG_BT_BLE_FEAT_CTE_CONNECTIONLESS_EN

Enable BLE CTE connectionless feature

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_FEAT_CTE_EN](#)

Transmission of CTE in periodic advertising

Default value:

- Yes (enabled) if `CONFIG_BT_BLE_FEAT_CTE_EN` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_FEAT_CTE_CONNECTION_EN

Enable BLE CTE connection feature

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#) > [CONFIG_BT_BLE_FEAT_CTE_EN](#)

Transmission of CTE by ACL connection

Default value:

- Yes (enabled) if `CONFIG_BT_BLE_FEAT_CTE_EN` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_FEAT_POWER_CONTROL

Enable BLE power control feature

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Enable BLE power control feature

Default value:

- No (disabled) if `CONFIG_BT_BLE_50_FEATURES_SUPPORTED` && `((CONFIG_BT_CONTROLLER_ENABLED && SOC_BLE_POWER_CONTROL_SUPPORTED) || CONFIG_BT_CONTROLLER_DISABLED)` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_FEAT_CONN_SUBRATING

Enable BLE connection subrating feature

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Enable BLE connection subrating feature

Default value:

- No (disabled) if `CONFIG_BT_BLE_50_FEATURES_SUPPORTED` && `((CONFIG_BT_CONTROLLER_ENABLED && SOC_BLE_SUBRATE_SUPPORTED) || CONFIG_BT_CONTROLLER_DISABLED)` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_FEAT_PAWR_EN

Enable Periodic Advertisement with Response(PAwR)

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Enable BLE Periodic Advertisement with Response(PAwR) feature

Default value:

- No (disabled) if `CONFIG_BT_BLE_50_FEATURES_SUPPORTED` && `((CONFIG_BT_CONTROLLER_ENABLED && SOC_BLE_PERIODIC_ADV_WITH_RESPONSE) || CONFIG_BT_CONTROLLER_DISABLED)` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_FEAT_ADV_CODING_SELECTION

Enable Advertising Coding Selection

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Enable Advertising Coding Selection

Default value:

- No (disabled) if `CONFIG_BT_BLE_50_FEATURES_SUPPORTED` && `((CONFIG_BT_CONTROLLER_ENABLED && SOC_BLE_ADV_CODING_SELECT_SUPPORTED) || CONFIG_BT_CONTROLLER_DISABLED)` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_FEAT_CHANNEL_SOUNDING

Enable BLE channel sounding

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Enable BLE channel sounding

Default value:

- No (disabled) if `CONFIG_BT_BLE_50_FEATURES_SUPPORTED` && `((CONFIG_BT_CONTROLLER_ENABLED && SOC_BLE_CHANNEL_SOUNDING_SUPPORTED) || CONFIG_BT_CONTROLLER_DISABLED)` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_FEAT_ADV_MONITOR

Enable BLE Advertising Monitor (LE Monitor Advertisement)

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

Enable LE Monitor Advertisement feature (Bluetooth Core 6.0). Allows the host to add devices to a monitor list and receive reports when the controller detects advertising from those devices (e.g. RSSI threshold).

Default value:

- No (disabled) if `CONFIG_BT_BLE_50_FEATURES_SUPPORTED` && ((`CONFIG_BT_CONTROLLER_ENABLED` && `SOC_BLE_ADV_MONITOR_SUPPORTED`) || `CONFIG_BT_CONTROLLER_DISABLED`) && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_42_FEATURES_SUPPORTED

Enable BLE 4.2 features (please disable BLE 5.0 if enable BLE 4.2)

Found in: *Component config > Bluetooth > Bluedroid Options*

This enables BLE 4.2 features. This option is universally supported by all ESP chips with BLE capabilities. BLE 4.2 and BLE 5.0 cannot be used simultaneously.

Default value:

- No (disabled) if `CONFIG_BT_BLE_ENABLED` && `CONFIG_BT_BLUEDROID_ENABLED`

Contains:

- `CONFIG_BT_BLE_42_ADV_EN`
- `CONFIG_BT_BLE_42_DTM_TEST_EN`
- `CONFIG_BT_BLE_42_SCAN_EN`

CONFIG_BT_BLE_42_DTM_TEST_EN

Enable BLE 4.2 DTM test

Found in: *Component config > Bluetooth > Bluedroid Options > CONFIG_BT_BLE_42_FEATURES_SUPPORTED*

This enables BLE 4.2 direct test mode

Default value:

- Yes (enabled) if `CONFIG_BT_BLE_42_FEATURES_SUPPORTED` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_42_ADV_EN

Enable BLE 4.2 advertising

Found in: *Component config > Bluetooth > Bluedroid Options > CONFIG_BT_BLE_42_FEATURES_SUPPORTED*

This enables BLE v4.2 advertising

Default value:

- Yes (enabled) if `CONFIG_BT_BLE_42_FEATURES_SUPPORTED` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_42_SCAN_EN

Enable BLE 4.2 scan

Found in: *Component config > Bluetooth > Bluedroid Options > CONFIG_BT_BLE_42_FEATURES_SUPPORTED*

This enables BLE v4.2 scan

Default value:

- Yes (enabled) if `CONFIG_BT_BLE_42_FEATURES_SUPPORTED` && `CONFIG_BT_BLUEDROID_ENABLED`

CONFIG_BT_BLE_VENDOR_HCI_EN

Enable BLE Vendor HCI command and event

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

This enables BLE vendor HCI command and event

Default value:

- Yes (enabled) if [CONFIG_BT_BLE_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_BLE_HIGH_DUTY_ADV_INTERVAL

Enable BLE high duty advertising interval feature

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

This enable BLE high duty advertising interval feature

Default value:

- No (disabled) if [CONFIG_BT_BLE_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

CONFIG_BT_ABORT_WHEN_ALLOCATION_FAILS

Abort when memory allocation fails in BT/BLE stack

Found in: [Component config](#) > [Bluetooth](#) > [Bluedroid Options](#)

This enables abort when memory allocation fails

Default value:

- No (disabled) if [CONFIG_BT_BLUEDROID_ENABLED](#) && [CONFIG_BT_BLUEDROID_ENABLED](#)

NimBLE Options

 Contains:

- [BLE 5.x Features](#)
- [BLE 6.x Features](#)
- [Debugging/Testing](#)
- [Extra Features](#)
- [GAP](#)
- [GATT / ATT](#)
- [General](#)
- [Helper Utils](#)
- [Host-controller Transport](#)
- [L2CAP](#)
- [Memory Settings](#)
- [NimBLE Mesh](#)
- [Roles and Profiles](#)
- [Security \(SMP\)](#)
- [Services](#)
- [Vendor / Optimization](#)

General

 Contains:

- [CONFIG_BT_NIMBLE_MEM_ALLOC_MODE](#)
- [CONFIG_BT_NIMBLE_HOST_TASK_STACK_SIZE](#)
- [CONFIG_BT_NIMBLE_PINNED_TO_CORE_CHOICE](#)

CONFIG_BT_NIMBLE_MEM_ALLOC_MODE

Memory allocation strategy

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [General](#)

Allocation strategy for NimBLE host stack. Provides ability to allocate all required dynamic allocations from:

- Internal DRAM memory only
- External SPIRAM memory only
- Either internal or external memory based on default malloc() behavior in ESP-IDF
- Internal IRAM memory wherever applicable else internal DRAM

Available options:

- Internal memory (CONFIG_BT_NIMBLE_MEM_ALLOC_MODE_INTERNAL)
- External SPIRAM (CONFIG_BT_NIMBLE_MEM_ALLOC_MODE_EXTERNAL)
- Default alloc mode (CONFIG_BT_NIMBLE_MEM_ALLOC_MODE_DEFAULT)
- Internal IRAM (CONFIG_BT_NIMBLE_MEM_ALLOC_MODE_IRAM_8BIT)
Allows use of IRAM memory region as 8-bit accessible region.
Every unaligned (8bit or 16bit) access will result in an exception and incur penalty of certain clock cycles per unaligned read/write.

CONFIG_BT_NIMBLE_PINNED_TO_CORE_CHOICE

The CPU core on which NimBLE host will run

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [General](#)

The CPU core on which NimBLE host will run. You can choose Core 0 or Core 1. Cannot specify no-affinity

Available options:

- Core 0 (PRO CPU) (CONFIG_BT_NIMBLE_PINNED_TO_CORE_0)
- Core 1 (APP CPU) (CONFIG_BT_NIMBLE_PINNED_TO_CORE_1)

CONFIG_BT_NIMBLE_HOST_TASK_STACK_SIZE

NimBLE Host task stack size

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [General](#)

This configures stack size of NimBLE host task

Default value:

- 5120 if [CONFIG_BLE_MESH](#) && [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)
- 4096 if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

Roles and Profiles Contains:

- [CONFIG_BT_NIMBLE_ROLE_BROADCASTER](#)
- [CONFIG_BT_NIMBLE_ROLE_CENTRAL](#)
- [CONFIG_BT_NIMBLE_GATT_CLIENT](#)
- [CONFIG_BT_NIMBLE_GATT_SERVER](#)
- [CONFIG_BT_NIMBLE_ROLE_OBSERVER](#)
- [CONFIG_BT_NIMBLE_ROLE_PERIPHERAL](#)

CONFIG_BT_NIMBLE_ROLE_CENTRAL

Enable BLE Central role

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Roles and Profiles](#)

Enables BLE Central role

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_ROLE_PERIPHERAL

Enable BLE Peripheral role

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Roles and Profiles](#)

Enable BLE Peripheral role

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_ROLE_BROADCASTER

Enable BLE Broadcaster role

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Roles and Profiles](#)

Enables BLE broadcaster role

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_ROLE_OBSERVER

Enable BLE Observer role

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Roles and Profiles](#)

Enables BLE Observer role

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_GATT_CLIENT

Enable BLE GATT Client support

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Roles and Profiles](#)

Enables support for GATT Client

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_ROLE_CENTRAL](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_GATT_SERVER

Enable BLE GATT Server support

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Roles and Profiles](#)

Enables support for GATT Server

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_ROLE_PERIPHERAL](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

Security (SMP) Contains:

- [*CONFIG_BT_NIMBLE_SECURITY_ENABLE*](#)

CONFIG_BT_NIMBLE_SECURITY_ENABLE

Enable BLE SM feature

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Security \(SMP\)](#)

Enable BLE sm feature

Default value:

- Yes (enabled) if [*CONFIG_BT_NIMBLE_ENABLED*](#) && [*CONFIG_BT_NIMBLE_ENABLED*](#)

Contains:

- [*CONFIG_BT_NIMBLE_LL_CFG_FEAT_LE_ENCRYPTION*](#)
- [*CONFIG_BT_NIMBLE_SM_SC_ONLY*](#)
- [*CONFIG_BT_NIMBLE_HANDLE_REPEAT_PAIRING_DELETION*](#)
- [*CONFIG_BT_NIMBLE_MAX BONDS*](#)
- [*CONFIG_BT_NIMBLE_NVS_PERSIST*](#)
- [*CONFIG_BT_NIMBLE_SMP_ID_RESET*](#)
- [*CONFIG_BT_NIMBLE_SM_LVL*](#)
- [*CONFIG_BT_NIMBLE_SM_LEGACY*](#)
- [*CONFIG_BT_NIMBLE_SM_SC*](#)

CONFIG_BT_NIMBLE_SM_LEGACY

Security manager legacy pairing

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Security \(SMP\)](#) > [*CONFIG_BT_NIMBLE_SECURITY_ENABLE*](#)

Enable security manager legacy pairing

Default value:

- Yes (enabled) if [*CONFIG_BT_NIMBLE_SECURITY_ENABLE*](#) && [*CONFIG_BT_NIMBLE_ENABLED*](#)

CONFIG_BT_NIMBLE_SM_SC

Security manager secure connections (4.2)

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Security \(SMP\)](#) > [*CONFIG_BT_NIMBLE_SECURITY_ENABLE*](#)

Enable security manager secure connections

Default value:

- Yes (enabled) if [*CONFIG_BT_NIMBLE_SECURITY_ENABLE*](#) && [*CONFIG_BT_NIMBLE_ENABLED*](#)

CONFIG_BT_NIMBLE_SM_SC_DEBUG_KEYS

Use predefined public-private key pair

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Security \(SMP\)](#) > [*CONFIG_BT_NIMBLE_SECURITY_ENABLE*](#) > [*CONFIG_BT_NIMBLE_SM_SC*](#)

If this option is enabled, SM uses predefined DH key pair as described in Core Specification, Vol. 3, Part H, 2.3.5.6.1. This allows to decrypt air traffic easily and thus should only be used for debugging.

Default value:

- No (disabled) if `CONFIG_BT_NIMBLE_SM_SC` && `CONFIG_BT_NIMBLE_SECURITY_ENABLE` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_LL_CFG_FEAT_LE_ENCRYPTION

Enable LE encryption

Found in: `Component config > Bluetooth > NimBLE Options > Security (SMP) > CONFIG_BT_NIMBLE_SECURITY_ENABLE`

Enable encryption connection

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_SECURITY_ENABLE` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_SM_LVL

Security level

Found in: `Component config > Bluetooth > NimBLE Options > Security (SMP) > CONFIG_BT_NIMBLE_SECURITY_ENABLE`

LE Security Mode 1 Levels: 1(0). No Security 2(1). Unauthenticated pairing with encryption 3(2). Authenticated pairing with encryption 4(3). Authenticated LE Secure Connections pairing with encryption using a 128-bit strength encryption key.

Default value:

- 0 if `CONFIG_BT_NIMBLE_SECURITY_ENABLE` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_SM_SC_ONLY

Enable Secure Connections Only Mode

Found in: `Component config > Bluetooth > NimBLE Options > Security (SMP) > CONFIG_BT_NIMBLE_SECURITY_ENABLE`

Enable Secure Connections Only Mode

Default value:

- 0 if `CONFIG_BT_NIMBLE_SECURITY_ENABLE` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_SMP_ID_RESET

Reset device identity when all bonding records are deleted

Found in: `Component config > Bluetooth > NimBLE Options > Security (SMP) > CONFIG_BT_NIMBLE_SECURITY_ENABLE`

There are tracking risks associated with using a fixed or static IRK. If enabled this option, NimBLE will assign a new randomly-generated IRK when all pairing and bonding records are deleted. This would decrease the ability of a previously paired peer to be used to determine whether a device with which it previously shared an IRK is within range.

Default value:

- No (disabled) if `CONFIG_BT_NIMBLE_SECURITY_ENABLE` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_NVS_PERSIST

Persist the BLE Bonding keys in NVS

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Security \(SMP\)](#) > [CONFIG_BT_NIMBLE_SECURITY_ENABLE](#)

Enable this flag to make bonding persistent across device reboots

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_SECURITY_ENABLE](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MAX_BONDS

Maximum number of bonds to save across reboots

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Security \(SMP\)](#) > [CONFIG_BT_NIMBLE_SECURITY_ENABLE](#)

Defines maximum number of bonds to save for peer security and our security

Default value:

- 3 if [CONFIG_BT_NIMBLE_SECURITY_ENABLE](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_HANDLE_REPEAT_PAIRING_DELETION

Enable stack handling of repeat pairing

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Security \(SMP\)](#) > [CONFIG_BT_NIMBLE_SECURITY_ENABLE](#)

Use this option to let stack internally handle the request for repeat pairing. Enabling this option will delete the pairing of the device and stack will NOT post any event to application. If this option is disabled, application will get BLE_GAP_EVENT_REPEAT_PAIRING event.

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_SECURITY_ENABLE](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

GAP Contains:

- [CONFIG_BT_NIMBLE_HOST_ALLOW_CONNECT_WITH_SCAN](#)
- [CONFIG_BT_NIMBLE_HS_STOP_TIMEOUT_MS](#)
- [CONFIG_BT_NIMBLE_HOST_QUEUE_CONG_CHECK](#)
- [CONFIG_BT_NIMBLE_WHITELIST_SIZE](#)
- [CONFIG_BT_NIMBLE_ENABLE_CONN_REATTEMPT](#)
- [CONFIG_BT_NIMBLE_USE_ESP_TIMER](#)
- [CONFIG_BT_NIMBLE_HS_FLOW_CTRL](#)
- [CONFIG_BT_NIMBLE_HS_PVCY](#)
- [CONFIG_BT_NIMBLE_MAX_CCCDS](#)
- [CONFIG_BT_NIMBLE_MAX_CONNECTIONS](#)
- [CONFIG_BT_NIMBLE_CRYPTOSTACK_MBEDTLS](#)
- [CONFIG_BT_NIMBLE_RPA_TIMEOUT](#)

CONFIG_BT_NIMBLE_RPA_TIMEOUT

RPA timeout in seconds

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GAP](#)

Time interval between RPA address change.

Range:

- from 1 to 41400 if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

Default value:

- 900 if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_WHITELIST_SIZE

BLE white list size

Found in: [Component config > Bluetooth > NimBLE Options > GAP](#)

BLE Whitelist size

Range:

- from 1 to 31 if `SOC_ESP_NIMBLE_CONTROLLER` && `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`
- from 1 to 15 if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

Default value:

- 12 if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_ENABLE_CONN_REATTEMPT

Enable connection reattempts on connection establishment error

Found in: [Component config > Bluetooth > NimBLE Options > GAP](#)

Enable to make the NimBLE host to reattempt GAP connection attempt in case of connection establishment failure due to reason 0x3E.

Default value:

- Yes (enabled) if `SOC_ESP_NIMBLE_CONTROLLER` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_MAX_CONN_REATTEMPT

Maximum number connection reattempts

Found in: [Component config > Bluetooth > NimBLE Options > GAP > CONFIG_BT_NIMBLE_ENABLE_CONN_REATTEMPT](#)

Defines maximum number of connection reattempts.

Range:

- from 1 to 255 if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLE_CONN_REATTEMPT` && `CONFIG_BT_NIMBLE_ENABLED`

Default value:

- 3 if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLE_CONN_REATTEMPT` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_HS_PVCY

Enable privacy related APIs

Found in: [Component config > Bluetooth > NimBLE Options > GAP](#)

Enable this option to include host-side APIs related to BLE privacy. Disable it if your application does not require privacy features (such as RPA generation, resolving, or privacy-based pairing). Disabling this option excludes privacy-related code from the host to reduce memory usage.

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_HOST_ALLOW_CONNECT_WITH_SCAN

Allow Connections with scanning in progress

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GAP](#)

This enables support for user to initiate a new connection with scan in progress

CONFIG_BT_NIMBLE_HOST_QUEUE_CONG_CHECK

BLE queue congestion check

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GAP](#)

When scanning and scan duplicate is not enabled, if there are a lot of adv packets around or application layer handling adv packets is slow, it will cause the controller memory to run out. if enabled, adv packets will be lost when host queue is congested.

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MAX_CONNECTIONS

Maximum number of concurrent connections

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GAP](#)

Defines maximum number of concurrent BLE connections. For ESP32, user is expected to configure BTDM_CTRL_BLE_MAX_CONN from controller menu along with this option. Similarly for ESP32-C3 or ESP32-S3, user is expected to configure BT_CTRL_BLE_MAX_ACT from controller menu. For ESP32C2, ESP32C6 and ESP32H2, each connection will take about 1k DRAM.

Range:

- from 1 to 70 if [SOC_ESP_NIMBLE_CONTROLLER](#) && ([CONFIG_BT_NIMBLE_ROLE_CENTRAL](#) || [CONFIG_BT_NIMBLE_ROLE_PERIPHERAL](#)) && [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)
- from 1 to 9 if ([CONFIG_BT_NIMBLE_ROLE_CENTRAL](#) || [CONFIG_BT_NIMBLE_ROLE_PERIPHERAL](#)) && [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

Default value:

- 3 if ([CONFIG_BT_NIMBLE_ROLE_CENTRAL](#) || [CONFIG_BT_NIMBLE_ROLE_PERIPHERAL](#)) && [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MAX_CCCDS

Maximum number of CCC descriptors to save across reboots

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GAP](#)

Defines maximum number of CCC descriptors to save

Default value:

- 8 if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_CRYPTOSTACK_MBEDTLS

Override TinyCrypt with mbedTLS for crypto computations

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GAP](#)

Enable this option to choose mbedTLS instead of TinyCrypt for crypto computations.

Default value:

- No (disabled) if `CONFIG_BT_SMP_CRYPTOSTACK_TINYCRYPT && CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_HS_STOP_TIMEOUT_MS

BLE host stop timeout in msec

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GAP](#)

BLE Host stop procedure timeout in milliseconds.

Default value:

- 2000 if `CONFIG_BT_NIMBLE_ENABLED && CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_USE_ESP_TIMER

Enable Esp Timer for Nimble

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GAP](#)

Set this option to use Esp Timer which has higher priority timer instead of FreeRTOS timer

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_HS_FLOW_CTRL

Enable Host Flow control

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GAP](#)

Enable Host Flow control

CONFIG_BT_NIMBLE_HS_FLOW_CTRL_ITVL

Host Flow control interval

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GAP](#) > [CONFIG_BT_NIMBLE_HS_FLOW_CTRL](#)

Host flow control interval in msecs

Default value:

- 1000 if `CONFIG_BT_NIMBLE_HS_FLOW_CTRL && CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_HS_FLOW_CTRL_THRESH

Host Flow control threshold

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GAP](#) > [CONFIG_BT_NIMBLE_HS_FLOW_CTRL](#)

Host flow control threshold, if the number of free buffers are at or below this threshold, send an immediate number-of-completed-packets event

Default value:

- 2 if `CONFIG_BT_NIMBLE_HS_FLOW_CTRL && CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_HS_FLOW_CTRL_TX_ON_DISCONNECT

Host Flow control on disconnect

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GAP](#) > [CONFIG_BT_NIMBLE_HS_FLOW_CTRL](#)

Enable this option to send number-of-completed-packets event to controller after disconnection

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_HS_FLOW_CTRL](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

GATT / ATT Contains:

- [CONFIG_BT_NIMBLE_BLE_GATT_BLOB_TRANSFER](#)
- [CONFIG_BT_NIMBLE_GATT_CACHING_DISABLE_AUTO](#)
- [CONFIG_BT_NIMBLE_GATT_CACHING_ASSOC_ENABLE](#)
- [CONFIG_BT_NIMBLE_GATT_CACHING](#)
- [CONFIG_BT_NIMBLE_INCL_SVC_DISCOVERY](#)
- [CONFIG_BT_NIMBLE_ATT_MAX_PREP_ENTRIES](#)
- [CONFIG_BT_NIMBLE_GATT_CACHING_MAX_CONNS](#)
- [CONFIG_BT_NIMBLE_GATT_CACHING_MAX_CHRS](#)
- [CONFIG_BT_NIMBLE_GATT_CACHING_MAX_DSCLS](#)
- [CONFIG_BT_NIMBLE_GATT_MAX_PROCS](#)
- [CONFIG_BT_NIMBLE_GATT_CACHING_MAX_INCL_SVCS](#)
- [CONFIG_BT_NIMBLE_GATT_CACHING_MAX_SVCS](#)
- [CONFIG_BT_NIMBLE_ATT_PREFERRED_MTU](#)

CONFIG_BT_NIMBLE_ATT_PREFERRED_MTU

Preferred MTU size in octets

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GATT / ATT](#)

This is the default value of ATT MTU indicated by the device during an ATT MTU exchange. This value can be changed using API `ble_att_set_preferred_mtu()`

Default value:

- 256 if ([CONFIG_BT_NIMBLE_GATT_CLIENT](#) || [CONFIG_BT_NIMBLE_GATT_SERVER](#)) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_ATT_MAX_PREP_ENTRIES

Max Prepare write entries

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GATT / ATT](#)

This is the default value of ATT Maximum prepare entries

Default value:

- 64 if ([CONFIG_BT_NIMBLE_GATT_CLIENT](#) || [CONFIG_BT_NIMBLE_GATT_SERVER](#)) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_GATT_MAX_PROCS

Maximum number of GATT client procedures

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GATT / ATT](#)

Maximum number of GATT client procedures that can be executed.

Default value:

- 4 if (*CONFIG_BT_NIMBLE_GATT_CLIENT* || *CONFIG_BT_NIMBLE_GATT_SERVER*) && *CONFIG_BT_NIMBLE_ENABLED*

CONFIG_BT_NIMBLE_BLE_GATT_BLOB_TRANSFER

Blob transfer

Found in: Component config > Bluetooth > NimBLE Options > GATT / ATT

This option is used when data to be sent is more than 512 bytes. For peripheral role, BT_NIMBLE_MSYS_1_BLOCK_COUNT needs to be increased according to the need.

CONFIG_BT_NIMBLE_GATT_CACHING

Enable GATT caching

Found in: Component config > Bluetooth > NimBLE Options > GATT / ATT

Enable GATT caching

Contains:

- *CONFIG_BT_NIMBLE_GATT_CACHING_INCLUDE_SERVICES*

CONFIG_BT_NIMBLE_GATT_CACHING_INCLUDE_SERVICES

Include services in GATT caching

Found in: Component config > Bluetooth > NimBLE Options > GATT / ATT > CONFIG_BT_NIMBLE_GATT_CACHING

Enable this option to include *included services* (e.g., services referenced by other services) in the GATT database cache. Disabling this will skip caching of included service entries.

Default value:

- No (disabled) if *CONFIG_BT_NIMBLE_GATT_CACHING* && *CONFIG_BT_NIMBLE_ENABLED*

CONFIG_BT_NIMBLE_INCL_SVC_DISCOVERY

Enable Included service discovery

Found in: Component config > Bluetooth > NimBLE Options > GATT / ATT

Enable this option to start discovery for included service.

Default value:

- Yes (enabled) if *CONFIG_BT_NIMBLE_GATT_CACHING* && *CONFIG_BT_NIMBLE_GATT_CACHING_INCLUDE_SERVICES* && *CONFIG_BT_NIMBLE_ENABLED*

CONFIG_BT_NIMBLE_GATT_CACHING_MAX_CONNS

Maximum connections to be cached

Found in: Component config > Bluetooth > NimBLE Options > GATT / ATT

Set this option to set the upper limit on number of connections to be cached.

Default value:

- if *CONFIG_BT_NIMBLE_GATT_CACHING* && *CONFIG_BT_NIMBLE_ENABLED*

CONFIG_BT_NIMBLE_GATT_CACHING_MAX_SVCS

Maximum number of services per connection

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GATT / ATT](#)

Set this option to set the upper limit on number of services per connection to be cached.

Default value:

- 64 if [CONFIG_BT_NIMBLE_GATT_CACHING](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_GATT_CACHING_MAX_INCL_SVCS

Maximum number of included services per connection

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GATT / ATT](#)

Set this option to set the upper limit on number of included services per connection to be cached.

Default value:

- 64 if [CONFIG_BT_NIMBLE_GATT_CACHING](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_GATT_CACHING_MAX_CHRS

Maximum number of characteristics per connection

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GATT / ATT](#)

Set this option to set the upper limit on number of characteristics per connection to be cached.

Default value:

- 64 if [CONFIG_BT_NIMBLE_GATT_CACHING](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_GATT_CACHING_MAX_DSCLS

Maximum number of descriptors per connection

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GATT / ATT](#)

Set this option to set the upper limit on number of descriptors per connection to be cached.

Default value:

- 64 if [CONFIG_BT_NIMBLE_GATT_CACHING](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_GATT_CACHING_DISABLE_AUTO

Do not start discovery procedure automatically upon receiving Out of Sync

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GATT / ATT](#)

When client receives ATT out-of-sync error message, it will not automatically start the discovery procedure to correct the invalid cache.

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_GATT_CACHING](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_GATT_CACHING_ASSOC_ENABLE

Enable association-based GATT caching

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [GATT / ATT](#)

Enable this option to use associated address caching instead of performing service discovery.

Default value:

- No (disabled) if `CONFIG_BT_NIMBLE_GATT_CACHING` && `CONFIG_BT_NIMBLE_ENABLED`

L2CAP Contains:

- `CONFIG_BT_NIMBLE_L2CAP_ENHANCED_COC`
- `CONFIG_BT_NIMBLE_L2CAP_COC_MAX_NUM`

CONFIG_BT_NIMBLE_L2CAP_COC_MAX_NUM

Maximum number of connection oriented channels

Found in: *Component config > Bluetooth > NimBLE Options > L2CAP*

Defines maximum number of BLE Connection Oriented Channels. When set to (0), BLE COC is not compiled in

Range:

- from 0 to 9 if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

Default value:

- 0 if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_L2CAP_ENHANCED_COC

L2CAP Enhanced Connection Oriented Channel

Found in: *Component config > Bluetooth > NimBLE Options > L2CAP*

Enable Enhanced Credit Based Flow Control Mode

Default value:

- No (disabled) if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_L2CAP_COC_MAX_NUM >= 1` && `CONFIG_BT_NIMBLE_ENABLED`

Memory Settings Contains:

- `CONFIG_BT_NIMBLE_TRANSPORT_ACL_FROM_LL_COUNT`
- `CONFIG_BT_NIMBLE_TRANSPORT_EVT_DISCARD_COUNT`
- `CONFIG_BT_NIMBLE_MSYS_BUF_FROM_HEAP`
- `CONFIG_BT_NIMBLE_L2CAP_COC_SDU_BUFF_COUNT`
- `CONFIG_BT_NIMBLE_MSYS_1_BLOCK_COUNT`
- `CONFIG_BT_NIMBLE_MSYS_1_BLOCK_SIZE`
- `CONFIG_BT_NIMBLE_MSYS_2_BLOCK_COUNT`
- `CONFIG_BT_NIMBLE_MSYS_2_BLOCK_SIZE`
- `CONFIG_BT_NIMBLE_MEMPOOL_RUNTIME_ALLOC`
- `CONFIG_BT_NIMBLE_TRANSPORT_ACL_SIZE`
- `CONFIG_BT_NIMBLE_TRANSPORT_EVT_COUNT`
- `CONFIG_BT_NIMBLE_TRANSPORT_EVT_SIZE`

CONFIG_BT_NIMBLE_MSYS_1_BLOCK_COUNT

MSYS_1 Block Count

Found in: *Component config > Bluetooth > NimBLE Options > Memory Settings*

MSYS is a system level mbuf registry. For prepare write & prepare responses Mbufs are allocated out of msys_1 pool. For NIMBLE_MESH enabled cases, this block count is increased by 8 than user defined count.

Default value:

- 24 if `SOC_ESP_NIMBLE_CONTROLLER` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_MSYS_1_BLOCK_SIZE

MSYS_1 Block Size

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Memory Settings](#)

Dynamic memory size of block 1

Default value:

- 128 if SOC_ESP_NIMBLE_CONTROLLER && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MSYS_2_BLOCK_COUNT

MSYS_2 Block Count

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Memory Settings](#)

Dynamic memory count for block 2

Default value:

- 24 if [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MSYS_2_BLOCK_SIZE

MSYS_2 Block Size

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Memory Settings](#)

Dynamic memory size of block 2

Default value:

- 320 if [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MSYS_BUF_FROM_HEAP

Get Msys Mbuf from heap

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Memory Settings](#)

This option sets the source of the shared msys mbuf memory between the Host and the Controller. Allocate the memory from the heap if this option is sets, from the mempool otherwise.

Default value:

- Yes (enabled) if BT_LE_MSYS_INIT_IN_CONTROLLER && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_TRANSPORT_ACL_FROM_LL_COUNT

ACL Buffer count

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Memory Settings](#)

The number of ACL data buffers allocated for host.

Default value:

- 24 if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_TRANSPORT_ACL_SIZE

Transport ACL Buffer size

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Memory Settings](#)

This is the maximum size of the data portion of HCI ACL data packets. It does not include the HCI data header (of 4 bytes)

Default value:

- 255 if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_TRANSPORT_EVT_SIZE

Transport Event Buffer size

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Memory Settings](#)

This is the size of each HCI event buffer in bytes. In case of extended advertising, packets can be fragmented. 257 bytes is the maximum size of a packet.

Default value:

- 257 if `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`
- 70 if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_TRANSPORT_EVT_COUNT

Transport Event Buffer count

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Memory Settings](#)

This is the high priority HCI events' buffer size. High-priority event buffers are for everything except advertising reports. If there are no free high-priority event buffers then host will try to allocate a low-priority buffer instead

Default value:

- 30 if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_TRANSPORT_EVT_DISCARD_COUNT

Discardable Transport Event Buffer count

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Memory Settings](#)

This is the low priority HCI events' buffer size. Low-priority event buffers are only used for advertising reports. If there are no free low-priority event buffers, then an incoming advertising report will get dropped

Default value:

- 8 if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_L2CAP_COC_SDU_BUFF_COUNT

L2cap coc Service Data Unit Buffer count

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Memory Settings](#)

This is the service data unit buffer count for l2cap coc.

Default value:

- 1 if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_MEMPOOL_RUNTIME_ALLOC

Support on-demand runtime memory allocation for mempool

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Memory Settings](#)

When this option is enabled, mempool does not require pre-allocating memory. Instead, memory for each block will be dynamically allocated and released during mempool usage. This can significantly reduce memory consumption after mempool initialization, but may have some impact on performance.

Default value:

- No (disabled) if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_STATIC_TO_DYNAMIC` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_MEMPOOL_BLOCK_REUSED

Support block reuse for mempool runtime memory allocation

Found in: `Component config > Bluetooth > NimBLE Options > Memory Settings > CONFIG_BT_NIMBLE_MEMPOOL_RUNTIME_ALLOC`

When this option is enabled, dynamically allocated blocks will not be freed but will be reused instead. This ensures virtually no impact on performance while reducing the memory consumption of the mempool.

Default value:

- No (disabled) if `CONFIG_BT_NIMBLE_MEMPOOL_RUNTIME_ALLOC` && `CONFIG_BT_NIMBLE_ENABLED`

BLE 5.x Features Contains:

- `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT`

CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT

Enable BLE 5 feature

Found in: `Component config > Bluetooth > NimBLE Options > BLE 5.x Features`

Enable BLE 5 feature

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_ENABLED` && `SOC_BLE_50_SUPPORTED` && `CONFIG_BT_NIMBLE_ENABLED`

Contains:

- `CONFIG_BT_NIMBLE_AOA_AOD`
- `CONFIG_BT_NIMBLE_LL_CFG_FEAT_LE_2M_PHY`
- `CONFIG_BT_NIMBLE_LL_CFG_FEAT_LE_CODED_PHY`
- `CONFIG_BT_NIMBLE_EXT_ADV`
- `CONFIG_BT_NIMBLE_EXT_SCAN`
- `CONFIG_BT_NIMBLE_BLE_POWER_CONTROL`

CONFIG_BT_NIMBLE_LL_CFG_FEAT_LE_2M_PHY

Enable 2M Phy

Found in: `Component config > Bluetooth > NimBLE Options > BLE 5.x Features > CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT`

Enable 2M-PHY

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_LL_CFG_FEAT_LE_CODED_PHY

Enable coded Phy

Found in: `Component config > Bluetooth > NimBLE Options > BLE 5.x Features > CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT`

Enable coded-PHY

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_EXT_ADV

Enable extended advertising

Found in: `Component config > Bluetooth > NimBLE Options > BLE 5.x Features > CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT`

Enable this option to do extended advertising. Extended advertising will be supported from BLE 5.0 onwards.

Default value:

- No (disabled) if `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_MAX_EXT_ADV_INSTANCES

Maximum number of extended advertising instances.

Found in: `Component config > Bluetooth > NimBLE Options > BLE 5.x Features > CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT > CONFIG_BT_NIMBLE_EXT_ADV`

Change this option to set maximum number of extended advertising instances. Minimum there is always one instance of advertising. Enter how many more advertising instances you want. For ESP32C2, ESP32C6 and ESP32H2, each extended advertising instance will take about 0.5k DRAM.

Range:

- from 0 to 4 if `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`

Default value:

- 1 if `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`
- 0 if `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_EXT_ADV_MAX_SIZE

Maximum length of the advertising data.

Found in: `Component config > Bluetooth > NimBLE Options > BLE 5.x Features > CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT > CONFIG_BT_NIMBLE_EXT_ADV`

Defines the length of the extended adv data. The value should not exceed 1650.

Range:

- from 0 to 1650 if `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`

Default value:

- 1650 if `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`
- 0 if `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_ENABLE_PERIODIC_ADV

Enable periodic advertisement.

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [BLE 5.x Features](#) > [CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT](#) > [CONFIG_BT_NIMBLE_EXT_ADV](#)

Enable this option to start periodic advertisement.

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_EXT_ADV](#) && [CONFIG_BT_NIMBLE_EXT_ADV](#) && [CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_PERIODIC_ADV_ENH

Periodic adv enhancements(adi support)

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [BLE 5.x Features](#) > [CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT](#) > [CONFIG_BT_NIMBLE_EXT_ADV](#) > [CONFIG_BT_NIMBLE_ENABLE_PERIODIC_ADV](#)

Enable the periodic advertising enhancements

CONFIG_BT_NIMBLE_PERIODIC_ADV_SYNC_TRANSFER

Enable Transfer Sync Events

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [BLE 5.x Features](#) > [CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT](#) > [CONFIG_BT_NIMBLE_EXT_ADV](#) > [CONFIG_BT_NIMBLE_ENABLE_PERIODIC_ADV](#)

This enables controller transfer periodic sync events to host

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_ENABLE_PERIODIC_ADV](#) && [CONFIG_BT_NIMBLE_EXT_ADV](#) && [CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_PERIODIC_ADV_WITH_RESPONSES

Enable Periodic Advertisement with Response (EXPERIMENTAL)

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [BLE 5.x Features](#) > [CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT](#) > [CONFIG_BT_NIMBLE_EXT_ADV](#) > [CONFIG_BT_NIMBLE_ENABLE_PERIODIC_ADV](#)

This enables controller PAwR (Periodic Advertisement with Response).

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_ENABLE_PERIODIC_ADV](#) && [SOC_BLE_PERIODIC_ADV_WITH_RESPONSE](#) && [CONFIG_BT_NIMBLE_EXT_ADV](#) && [CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_EXT_SCAN

Enable extended scanning

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [BLE 5.x Features](#) > [CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT](#)

Enable this option to do extended scanning.

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_EXT_ADV` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_ENABLE_PERIODIC_SYNC

Enable periodic sync

Found in: `Component config > Bluetooth > NimBLE Options > BLE 5.x Features > CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT > CONFIG_BT_NIMBLE_EXT_SCAN`

Enable this option to receive periodic advertisement.

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_EXT_SCAN` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_MAX_PERIODIC_SYNCS

Maximum number of periodic advertising syncs

Found in: `Component config > Bluetooth > NimBLE Options > BLE 5.x Features > CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT > CONFIG_BT_NIMBLE_EXT_SCAN > CONFIG_BT_NIMBLE_ENABLE_PERIODIC_SYNC`

Set this option to set the upper limit for number of periodic sync connections. This should be less than maximum connections allowed by controller.

Range:

- from 0 to 8 if `CONFIG_BT_NIMBLE_ENABLE_PERIODIC_SYNC` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`

Default value:

- 1 if `CONFIG_BT_NIMBLE_ENABLE_PERIODIC_ADV` && `CONFIG_BT_NIMBLE_ENABLE_PERIODIC_SYNC` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`
- 0 if `CONFIG_BT_NIMBLE_ENABLE_PERIODIC_SYNC` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_MAX_PERIODIC_ADVERTISER_LIST

Maximum number of periodic advertiser list

Found in: `Component config > Bluetooth > NimBLE Options > BLE 5.x Features > CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT > CONFIG_BT_NIMBLE_EXT_SCAN > CONFIG_BT_NIMBLE_ENABLE_PERIODIC_SYNC`

Set this option to set the upper limit for number of periodic advertiser list.

Range:

- from 1 to 5 if `SOC_ESP_NIMBLE_CONTROLLER` && `CONFIG_BT_NIMBLE_ENABLE_PERIODIC_SYNC` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`

Default value:

- 5 if `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `SOC_ESP_NIMBLE_CONTROLLER` && `CONFIG_BT_NIMBLE_ENABLE_PERIODIC_SYNC` && `CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_BLE_POWER_CONTROL

Enable support for BLE Power Control

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [BLE 5.x Features](#) > [CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT](#)

Set this option to enable the Power Control feature

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT](#) && [SOC_BLE_POWER_CONTROL_SUPPORTED](#) && [CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_AOA_AOD

Direction Finding

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [BLE 5.x Features](#) > [CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT](#)

Enable support for Connectionless and Connection Oriented Direction Finding

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT](#) && [SOC_BLE_CTE_SUPPORTED](#) && [CONFIG_BT_NIMBLE_50_FEATURE_SUPPORT](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

BLE 6.x Features Contains:

- [CONFIG_BT_NIMBLE_60_FEATURE_SUPPORT](#)

CONFIG_BT_NIMBLE_60_FEATURE_SUPPORT

Enable BLE 6 feature

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [BLE 6.x Features](#)

Enable BLE 6.x feature support

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && ([SOC_BLE_60_SUPPORTED](#) || [CONFIG_BT_CONTROLLER_DISABLED](#)) && ([SOC_BLE_60_SUPPORTED](#) || [CONFIG_BT_CONTROLLER_DISABLED](#)) && [CONFIG_BT_NIMBLE_ENABLED](#)

Contains:

- [CONFIG_BT_NIMBLE_CHANNEL_SOUNDING](#)
- [CONFIG_BT_NIMBLE_MONITOR_ADV](#)

CONFIG_BT_NIMBLE_CHANNEL_SOUNDING

ble channel sounding feature

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [BLE 6.x Features](#) > [CONFIG_BT_NIMBLE_60_FEATURE_SUPPORT](#)

Used to enable/disable the channel sounding feature

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_60_FEATURE_SUPPORT](#) && ([SOC_BLE_60_SUPPORTED](#) || [CONFIG_BT_CONTROLLER_DISABLED](#)) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MONITOR_ADV

Enable Monitor Advertising

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [BLE 6.x Features](#) > [CONFIG_BT_NIMBLE_60_FEATURE_SUPPORT](#)

Enable support for Monitor Advertisers

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_60_FEATURE_SUPPORT](#) && (SOC_BLE_60_SUPPORTED || [CONFIG_BT_CONTROLLER_DISABLED](#)) && [CONFIG_BT_NIMBLE_ENABLED](#)

Services Contains:

- [CONFIG_BT_NIMBLE_ANS_SERVICE](#)
- [CONFIG_BT_NIMBLE_BAS_SERVICE](#)
- [CONFIG_BT_NIMBLE_CTS_SERVICE](#)
- [CONFIG_BT_NIMBLE_DIS_SERVICE](#)
- [CONFIG_BT_NIMBLE_GAP_SERVICE](#)
- [CONFIG_BT_NIMBLE_HTP_SERVICE](#)
- [CONFIG_BT_NIMBLE_HR_SERVICE](#)
- [CONFIG_BT_NIMBLE_HID_SERVICE](#)
- [CONFIG_BT_NIMBLE_IAS_SERVICE](#)
- [CONFIG_BT_NIMBLE_IPSS_SERVICE](#)
- [CONFIG_BT_NIMBLE_LLS_SERVICE](#)
- [CONFIG_BT_NIMBLE_PROX_SERVICE](#)
- [CONFIG_BT_NIMBLE_SPS_SERVICE](#)
- [CONFIG_BT_NIMBLE_TPS_SERVICE](#)

CONFIG_BT_NIMBLE_PROX_SERVICE

Proximity service

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#)

Enable Proximity Service support

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_ANS_SERVICE

Alert Notification service

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#)

Enable Alert Notification Service support

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_CTS_SERVICE

Current Time Service

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#)

Enable Current Time Service support

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_GATT_SERVER` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_HTP_SERVICE

Health Thermometer service

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#)

Enable Health Thermometer Service support

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_GATT_SERVER` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_IPSS_SERVICE

Internet Protocol Support service

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#)

Enable Internet Protocol Service support

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_GATT_SERVER` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_TPS_SERVICE

Tx Power service

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#)

Enable Tx Power Service support

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_GATT_SERVER` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_IAS_SERVICE

Immediate Alert service

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#)

Enable Immediate Alert Service support

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_GATT_SERVER` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_LLS_SERVICE

Link Loss service

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#)

Enable Link Loss Service support

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_GATT_SERVER` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_SPS_SERVICE

Serial Port service

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#)

Enable Serial Port Service support

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_HR_SERVICE

Heart Rate service

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#)

Enable HeartRate Service support

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_HID_SERVICE

Human Interface Device service

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#)

Enable HID service support

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

Contains:

- [CONFIG_BT_NIMBLE_SVC_HID_MAX_RPTS](#)
- [CONFIG_BT_NIMBLE_SVC_HID_MAX_INSTANCES](#)

CONFIG_BT_NIMBLE_SVC_HID_MAX_INSTANCES

Maximum HID service instances

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_HID_SERVICE](#)

Defines maximum number of HID service instances

Default value:

- 2 if [CONFIG_BT_NIMBLE_HID_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_HID_MAX_RPTS

Maximum HID Report characteristics per service instance

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_HID_SERVICE](#)

Defines maximum number of report characteristics per service instance

Default value:

- 3 if [CONFIG_BT_NIMBLE_HID_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_BAS_SERVICE

Battery service

Found in: *Component config > Bluetooth > NimBLE Options > Services*

Enable Battery service support

Default value:

- Yes (enabled) if *CONFIG_BT_NIMBLE_GATT_SERVER* && *CONFIG_BT_NIMBLE_ENABLED*

Contains:

- *CONFIG_BT_NIMBLE_SVC_BAS_BATTERY_LEVEL_NOTIFY*

CONFIG_BT_NIMBLE_SVC_BAS_BATTERY_LEVEL_NOTIFY

BAS Battery Level NOTIFY permission

Found in: *Component config > Bluetooth > NimBLE Options > Services > CONFIG_BT_NIMBLE_BAS_SERVICE*

Enable/Disable notifications on BAS Battery Level Characteristic

CONFIG_BT_NIMBLE_DIS_SERVICE

Device Information Service

Found in: *Component config > Bluetooth > NimBLE Options > Services*

Enable Device Information service support

Default value:

- Yes (enabled) if *CONFIG_BT_NIMBLE_GATT_SERVER* && *CONFIG_BT_NIMBLE_ENABLED*

Contains:

- *CONFIG_BT_NIMBLE_SVC_DIS_INCLUDED*
- *CONFIG_BT_NIMBLE_SVC_DIS_FIRMWARE_REVISION*
- *CONFIG_BT_NIMBLE_SVC_DIS_HARDWARE_REVISION*
- *CONFIG_BT_NIMBLE_SVC_DIS_MANUFACTURER_NAME*
- *CONFIG_BT_NIMBLE_SVC_DIS_PNP_ID*
- *CONFIG_BT_NIMBLE_SVC_DIS_SERIAL_NUMBER*
- *CONFIG_BT_NIMBLE_SVC_DIS_SOFTWARE_REVISION*
- *CONFIG_BT_NIMBLE_SVC_DIS_SYSTEM_ID*

CONFIG_BT_NIMBLE_SVC_DIS_MANUFACTURER_NAME

Manufacturer Name

Found in: *Component config > Bluetooth > NimBLE Options > Services > CONFIG_BT_NIMBLE_DIS_SERVICE*

Enable the DIS characteristic Manufacturer Name String characteristic

Default value:

- No (disabled) if *CONFIG_BT_NIMBLE_DIS_SERVICE* && *CONFIG_BT_NIMBLE_GATT_SERVER* && *CONFIG_BT_NIMBLE_ENABLED*

CONFIG_BT_NIMBLE_SVC_DIS_SERIAL_NUMBER

Serial Number

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_DIS_SERVICE](#)

Enable the DIS Serial Number characteristic

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_DIS_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_DIS_HARDWARE_REVISION

Hardware Revision

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_DIS_SERVICE](#)

Enable the DIS Hardware Revision characteristic

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_DIS_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_DIS_FIRMWARE_REVISION

Firmware Revision

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_DIS_SERVICE](#)

Enable the DIS Firmware Revision characteristic

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_DIS_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_DIS_SOFTWARE_REVISION

Software Revision

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_DIS_SERVICE](#)

Enable the DIS Software Revision characteristic

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_DIS_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_DIS_SYSTEM_ID

System ID

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_DIS_SERVICE](#)

Enable the DIS System ID characteristic

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_DIS_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_DIS_PNP_ID

PnP ID

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_DIS_SERVICE](#)

Enable the DIS PnP ID characteristic

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_DIS_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_DIS_INCLUDED

DIS as an Included Service

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_DIS_SERVICE](#)

Use DIS as an included service

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_DIS_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_GAP_SERVICE

GAP Service

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#)

Enable GAP Service support

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

Contains:

- [CONFIG_BT_NIMBLE_SVC_GAP_DEVICE_NAME](#)
- [CONFIG_BT_NIMBLE_SVC_GAP_APPEARANCE](#)
- [GAP Appearance write permissions](#)
- [CONFIG_BT_NIMBLE_SVC_GAP_CENT_ADDR_RESOLUTION](#)
- [GAP device name write permissions](#)
- [CONFIG_BT_NIMBLE_SVC_GAP_GATT_SECURITY_LEVEL](#)
- [CONFIG_BT_NIMBLE_GAP_DEVICE_NAME_MAX_LEN](#)
- [Peripheral Preferred Connection Parameters \(PPCP\) settings](#)
- [CONFIG_BT_NIMBLE_SVC_GAP_RPA_ONLY](#)

CONFIG_BT_NIMBLE_SVC_GAP_DEVICE_NAME

BLE GAP default device name

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_GAP_SERVICE](#)

The Device Name characteristic shall contain the name of the device as an UTF-8 string. This name can be changed by using API `ble_svc_gap_device_name_set()`

Default value:

- "nimble" if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_GAP_DEVICE_NAME_MAX_LEN

Maximum length of BLE device name in octets

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_GAP_SERVICE](#)

Device Name characteristic value shall be 0 to 248 octets in length

Default value:

- 31 if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_GAP_APPEARANCE

External appearance of the device

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_GAP_SERVICE](#)

Standard BLE GAP Appearance value in HEX format e.g. 0x02C0

Default value:

- 0 if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_GAP_GATT_SECURITY_LEVEL

LE GATT Security Level Characteristic

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_GAP_SERVICE](#)

Enable the LE GATT Security Level Characteristic

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_GAP_RPA_ONLY

Resolvable Private Address Only characteristic

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_GAP_SERVICE](#)

Enable the Resolvable Private Address Only characteristic

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_GAP_CENT_ADDR_RESOLUTION

GAP Characteristic - Central Address Resolution

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_GAP_SERVICE](#)

Whether or not Central Address Resolution characteristic is supported on the device, and if supported, whether or not Central Address Resolution is supported.

- Central Address Resolution characteristic not supported
- Central Address Resolution not supported
- Central Address Resolution supported

Available options:

- Characteristic not supported (CONFIG_BT_NIMBLE_SVC_GAP_CAR_CHAR_NOT_SUPP)
- Central Address Resolution not supported (CONFIG_BT_NIMBLE_SVC_GAP_CAR_NOT_SUPP)
- Central Address Resolution supported (CONFIG_BT_NIMBLE_SVC_GAP_CAR_SUPP)

GAP Appearance write permissions Contains:

- *CONFIG_BT_NIMBLE_SVC_GAP_APPEAR_WRITE*

CONFIG_BT_NIMBLE_SVC_GAP_APPEAR_WRITE

Write

Found in: Component config > Bluetooth > NimBLE Options > Services > CONFIG_BT_NIMBLE_GAP_SERVICE > GAP Appearance write permissions

Enable write permission (BLE_GATT_CHR_F_WRITE)

Default value:

- No (disabled) if *CONFIG_BT_NIMBLE_GAP_SERVICE* && *CONFIG_BT_NIMBLE_GAP_SERVICE* && *CONFIG_BT_NIMBLE_GATT_SERVER* && *CONFIG_BT_NIMBLE_ENABLED*

CONFIG_BT_NIMBLE_SVC_GAP_APPEAR_WRITE_ENC

Write with encryption

Found in: Component config > Bluetooth > NimBLE Options > Services > CONFIG_BT_NIMBLE_GAP_SERVICE > GAP Appearance write permissions > CONFIG_BT_NIMBLE_SVC_GAP_APPEAR_WRITE

Enable write with encryption permission (BLE_GATT_CHR_F_WRITE_ENC)

CONFIG_BT_NIMBLE_SVC_GAP_APPEAR_WRITE_AUTHEN

Write with authentication

Found in: Component config > Bluetooth > NimBLE Options > Services > CONFIG_BT_NIMBLE_GAP_SERVICE > GAP Appearance write permissions > CONFIG_BT_NIMBLE_SVC_GAP_APPEAR_WRITE

Enable write with authentication permission (BLE_GATT_CHR_F_WRITE_AUTHEN)

CONFIG_BT_NIMBLE_SVC_GAP_APPEAR_WRITE_AUTHOR

Write with authorisation

Found in: Component config > Bluetooth > NimBLE Options > Services > CONFIG_BT_NIMBLE_GAP_SERVICE > GAP Appearance write permissions > CONFIG_BT_NIMBLE_SVC_GAP_APPEAR_WRITE

Enable write with authorisation permission (BLE_GATT_CHR_F_WRITE_AUTHOR)

Default value:

- No (disabled) if *CONFIG_BT_NIMBLE_SVC_GAP_APPEAR_WRITE* && *CONFIG_BT_NIMBLE_GAP_SERVICE* && *CONFIG_BT_NIMBLE_GAP_SERVICE* && *CONFIG_BT_NIMBLE_GATT_SERVER* && *CONFIG_BT_NIMBLE_ENABLED*

GAP device name write permissions Contains:

- [CONFIG_BT_NIMBLE_SVC_GAP_NAME_WRITE](#)

CONFIG_BT_NIMBLE_SVC_GAP_NAME_WRITE

Write

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_GAP_SERVICE](#) > [GAP device name write permissions](#)

Enable write permission (BLE_GATT_CHR_F_WRITE)

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_GAP_NAME_WRITE_ENC

Write with encryption

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_GAP_SERVICE](#) > [GAP device name write permissions](#) > [CONFIG_BT_NIMBLE_SVC_GAP_NAME_WRITE](#)

Enable write with encryption permission (BLE_GATT_CHR_F_WRITE_ENC)

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_SVC_GAP_NAME_WRITE](#) && [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_GAP_NAME_WRITE_AUTHEN

Write with authentication

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_GAP_SERVICE](#) > [GAP device name write permissions](#) > [CONFIG_BT_NIMBLE_SVC_GAP_NAME_WRITE](#)

Enable write with authentication permission (BLE_GATT_CHR_F_WRITE_AUTHEN)

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_SVC_GAP_NAME_WRITE](#) && [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_GAP_NAME_WRITE_AUTHOR

Write with authorisation

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_GAP_SERVICE](#) > [GAP device name write permissions](#) > [CONFIG_BT_NIMBLE_SVC_GAP_NAME_WRITE](#)

Enable write with authorisation permission (BLE_GATT_CHR_F_WRITE_AUTHOR)

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_SVC_GAP_NAME_WRITE](#) && [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

Peripheral Preferred Connection Parameters (PPCP) settings Contains:

- [CONFIG_BT_NIMBLE_SVC_GAP_PPCP_MAX_CONN_INTERVAL](#)
- [CONFIG_BT_NIMBLE_SVC_GAP_PPCP_MIN_CONN_INTERVAL](#)
- [CONFIG_BT_NIMBLE_SVC_GAP_PPCP_SLAVE_LATENCY](#)
- [CONFIG_BT_NIMBLE_SVC_GAP_PPCP_SUPERVISION_TMO](#)

CONFIG_BT_NIMBLE_SVC_GAP_PPCP_MAX_CONN_INTERVAL

PPCP Connection Interval Max (Unit: 1.25 ms)

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_GAP_SERVICE](#) > [Peripheral Preferred Connection Parameters \(PPCP\) settings](#)

Peripheral Preferred Connection Parameter: Connection Interval maximum value Interval Max = value * 1.25 ms

Default value:

- 0 if [CONFIG_BT_NIMBLE_ROLE_PERIPHERAL](#) && [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_GAP_PPCP_MIN_CONN_INTERVAL

PPCP Connection Interval Min (Unit: 1.25 ms)

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_GAP_SERVICE](#) > [Peripheral Preferred Connection Parameters \(PPCP\) settings](#)

Peripheral Preferred Connection Parameter: Connection Interval minimum value Interval Min = value * 1.25 ms

Default value:

- 0 if [CONFIG_BT_NIMBLE_ROLE_PERIPHERAL](#) && [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_GAP_PPCP_SLAVE_LATENCY

PPCP Slave Latency

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_GAP_SERVICE](#) > [Peripheral Preferred Connection Parameters \(PPCP\) settings](#)

Peripheral Preferred Connection Parameter: Slave Latency

Default value:

- 0 if [CONFIG_BT_NIMBLE_GAP_SERVICE](#) && [CONFIG_BT_NIMBLE_GATT_SERVER](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SVC_GAP_PPCP_SUPERVISION_TMO

PPCP Supervision Timeout (Unit: 10 ms)

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Services](#) > [CONFIG_BT_NIMBLE_GAP_SERVICE](#) > [Peripheral Preferred Connection Parameters \(PPCP\) settings](#)

Peripheral Preferred Connection Parameter: Supervision Timeout Timeout = Value * 10 ms

Default value:

- 0 if `CONFIG_BT_NIMBLE_GAP_SERVICE` && `CONFIG_BT_NIMBLE_GAP_SERVICE` && `CONFIG_BT_NIMBLE_GAP_SERVICE` && `CONFIG_BT_NIMBLE_GATT_SERVER` && `CONFIG_BT_NIMBLE_ENABLED`

Extra Features Contains:

- `CONFIG_BT_NIMBLE_GATTC_AUTO_PAIR`
- `CONFIG_BT_NIMBLE_ADV_UUID_CONCAT`
- `CONFIG_BT_NIMBLE_BLUFI_ENABLE`
- `CONFIG_BT_NIMBLE_CPFD_CAFD`
- `CONFIG_BT_NIMBLE_DTM_MODE_TEST`
- `CONFIG_BT_NIMBLE_DYNAMIC_SERVICE`
- `CONFIG_BT_NIMBLE_MEM_OPTIMIZATION`
- `CONFIG_BT_NIMBLE_SM_SIGN_CNT`
- `CONFIG_BT_NIMBLE_SUBRATE`
- `CONFIG_BT_NIMBLE_STATIC_PASSKEY`
- `CONFIG_BT_NIMBLE_ENC_ADV_DATA`
- `CONFIG_BT_NIMBLE_GATTC_PROC_PREEMPTION_PROTECT`
- `CONFIG_BT_NIMBLE_MAX_EADS`
- `CONFIG_BT_NIMBLE_EATT_CHAN_NUM`
- `CONFIG_BT_NIMBLE_RECONFIG_MTU`
- `CONFIG_BT_NIMBLE_LOW_SPEED_MODE`
- `CONFIG_BT_NIMBLE_STATIC_TO_DYNAMIC`

CONFIG_BT_NIMBLE_DYNAMIC_SERVICE

Enable dynamic services

Found in: *Component config > Bluetooth > NimBLE Options > Extra Features*

This enables user to add/remove Gatt services at runtime

CONFIG_BT_NIMBLE_BLUFI_ENABLE

Enable blufi functionality

Found in: *Component config > Bluetooth > NimBLE Options > Extra Features*

Set this option to enable blufi functionality.

Default value:

- No (disabled) if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_ENC_ADV_DATA

Encrypted Advertising Data

Found in: *Component config > Bluetooth > NimBLE Options > Extra Features*

This option is used to enable encrypted advertising data.

CONFIG_BT_NIMBLE_ADV_UUID_CONCAT

concatenate uuids while parsing advertising data

Found in: *Component config > Bluetooth > NimBLE Options > Extra Features*

Enables concatenation of multiple UUIDs of the same type while parsing advertising data on the central device. When disabled, only the last parsed UUID of a given type is retained.

Default value:

- No (disabled) if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_MAX_EADS

Maximum number of EAD devices to save across reboots

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Extra Features](#)

Defines maximum number of encrypted advertising data key material to save

Default value:

- 10 if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENC_ADV_DATA](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_GATTC_PROC_PREEMPTION_PROTECT

Gatt-proc preemption protect check

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Extra Features](#)

When BLE and Wireless protocol/IEEE 802.15.4 operate in coexistence, BLE preemption can disrupt the GATT context, causing the service discovery callback to not be invoked. A temporary list is maintained to preserve the GATT context and use it in case of preemption.

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_GATTC_AUTO_PAIR

Automatically pair upon receiving service request failure

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Extra Features](#)

If enabled, when a service request (e.g. read, write) to a server fails, and the ATT error suggests insufficient security, then the central will initiate pairing and retry the service request.

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_EATT_CHAN_NUM

Maximum number of EATT channels

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Extra Features](#)

Defines the number of channels EATT bearers can use

Default value:

- 0 if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_SUBRATE

Enable Subrate Change

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Extra Features](#)

Enable connection subrate change feature

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_STATIC_PASSKEY

Enable support for Static Passkey

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Extra Features](#)

Enable support for user defined static passkey for SMP

Default value:

- No (disabled) if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_SECURITY_ENABLE` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_DTM_MODE_TEST

Enable DTM related test procedure support

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Extra Features](#)

Enable DTM related test procedure support

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_MEM_OPTIMIZATION

Enable memory optimization

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Extra Features](#)

This option enables memory optimizations in the NimBLE stack.

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_STATIC_TO_DYNAMIC

Use Dynamic Memory allocations instead of static

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Extra Features](#)

This option enables use of dynamic memory allocation instead of static allocation for variable global / local variable. This helps in improving memory footprint

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_SM_SIGN_CNT

Enable Sign counter operations

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Extra Features](#)

Enable Sign counter operations

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_CPFDF_CAFD

Enable CPFDF/CAPD support

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Extra Features](#)

Enable CPFDF and CAPD gatt descriptors support

Default value:

- Yes (enabled) if `CONFIG_BT_NIMBLE_ENABLED` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_RECONFIG_MTU

Reconfig L2CAP MTU support

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Extra Features](#)

Enable Support to reconfigure L2CAP MTU / MPS

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_LOW_SPEED_MODE

Reduce DRAM/IRAM usage by moving GATT/GAP/NPL functions to FLASH

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Extra Features](#)

Disabling IRAM-optimized paths will place functions in FLASH (.irom.text) instead of DRAM (.iram.text/.text in internal RAM). Suitable for low-speed GAP/GATT operations to reduce RAM usage.

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_ENABLED](#)

NimBLE Mesh Contains:

- [CONFIG_BT_NIMBLE_MESH](#)

CONFIG_BT_NIMBLE_MESH

Enable BLE mesh functionality

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [NimBLE Mesh](#)

Enable BLE Mesh example present in upstream mynewt-nimble and not maintained by Espressif.

IDF maintains ESP-BLE-MESH as the official Mesh solution. Please refer to ESP-BLE-MESH guide at: `./doc/esp32/api-guides/esp-ble-mesh/ble-mesh-index`

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

Contains:

- [CONFIG_BT_NIMBLE_MESH_PROVISIONER](#)
- [CONFIG_BT_NIMBLE_MESH_PROV](#)
- [CONFIG_BT_NIMBLE_MESH_GATT_PROXY](#)
- [CONFIG_BT_NIMBLE_MESH_FRIEND](#)
- [CONFIG_BT_NIMBLE_MESH_LOW_POWER](#)
- [CONFIG_BT_NIMBLE_MESH_PROXY](#)
- [CONFIG_BT_NIMBLE_MESH_RELAY](#)
- [CONFIG_BT_NIMBLE_MESH_DEVICE_NAME](#)
- [CONFIG_BT_NIMBLE_MESH_NODE_COUNT](#)

CONFIG_BT_NIMBLE_MESH_PROXY

Enable mesh proxy functionality

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [NimBLE Mesh](#) > [CONFIG_BT_NIMBLE_MESH](#)

Enable proxy. This is automatically set whenever NIMBLE_MESH_PB_GATT or NIMBLE_MESH_GATT_PROXY is set

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_MESH](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MESH_PROV

Enable BLE mesh provisioning

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [NimBLE Mesh](#) > [CONFIG_BT_NIMBLE_MESH](#)

Enable mesh provisioning

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_MESH](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MESH_PB_ADV

Enable mesh provisioning over advertising bearer

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [NimBLE Mesh](#) > [CONFIG_BT_NIMBLE_MESH](#) > [CONFIG_BT_NIMBLE_MESH_PROV](#)

Enable this option to allow the device to be provisioned over the advertising bearer

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_MESH_PROV](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MESH_PB_GATT

Enable mesh provisioning over GATT bearer

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [NimBLE Mesh](#) > [CONFIG_BT_NIMBLE_MESH](#) > [CONFIG_BT_NIMBLE_MESH_PROV](#)

Enable this option to allow the device to be provisioned over the GATT bearer

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_MESH_PROV](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MESH_GATT_PROXY

Enable GATT Proxy functionality

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [NimBLE Mesh](#) > [CONFIG_BT_NIMBLE_MESH](#)

This option enables support for the Mesh GATT Proxy Service, i.e. the ability to act as a proxy between a Mesh GATT Client and a Mesh network

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_MESH](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MESH_RELAY

Enable mesh relay functionality

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [NimBLE Mesh](#) > [CONFIG_BT_NIMBLE_MESH](#)

Support for acting as a Mesh Relay Node

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_MESH](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MESH_LOW_POWER

Enable mesh low power mode

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [NimBLE Mesh](#) > [CONFIG_BT_NIMBLE_MESH](#)

Enable this option to be able to act as a Low Power Node

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_MESH](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MESH_FRIEND

Enable mesh friend functionality

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [NimBLE Mesh](#) > [CONFIG_BT_NIMBLE_MESH](#)

Enable this option to be able to act as a Friend Node

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_MESH](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MESH_DEVICE_NAME

Set mesh device name

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [NimBLE Mesh](#) > [CONFIG_BT_NIMBLE_MESH](#)

This value defines Bluetooth Mesh device/node name

Default value:

- "nimble-mesh-node" if [CONFIG_BT_NIMBLE_MESH](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MESH_NODE_COUNT

Set mesh node count

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [NimBLE Mesh](#) > [CONFIG_BT_NIMBLE_MESH](#)

Defines mesh node count.

Default value:

- 1 if [CONFIG_BT_NIMBLE_MESH](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_MESH_PROVISIONER

Enable BLE mesh provisioner

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [NimBLE Mesh](#) > [CONFIG_BT_NIMBLE_MESH](#)

Enable mesh provisioner.

Default value:

- 0 if [CONFIG_BT_NIMBLE_MESH](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

Host-controller Transport Contains:

- [CONFIG_BT_NIMBLE_TRANSPORT_UART](#)
- [CONFIG_BT_NIMBLE_HCI_UART_CTS_PIN](#)
- [CONFIG_BT_NIMBLE_USE_HCI_UART_FLOW_CTRL](#)
- [CONFIG_BT_NIMBLE_HCI_UART_RTS_PIN](#)

CONFIG_BT_NIMBLE_TRANSPORT_UART

Enable Uart Transport

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Host-controller Transport](#)

Use UART transport

Default value:

- Yes (enabled) if [CONFIG_BT_CONTROLLER_DISABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_TRANSPORT_UART_PORT

Uart port

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Host-controller Transport](#) > [CONFIG_BT_NIMBLE_TRANSPORT_UART](#)

Uart port

Default value:

- 1 if [CONFIG_BT_CONTROLLER_DISABLED](#) && [CONFIG_BT_NIMBLE_TRANSPORT_UART](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_HCI_USE_UART_BAUDRATE

Uart Hci Baud Rate

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Host-controller Transport](#) > [CONFIG_BT_NIMBLE_TRANSPORT_UART](#)

Uart Baud Rate

Available options:

- 115200 ([CONFIG_UART_BAUDRATE_115200](#))
- 230400 ([CONFIG_UART_BAUDRATE_230400](#))
- 460800 ([CONFIG_UART_BAUDRATE_460800](#))
- 921600 ([CONFIG_UART_BAUDRATE_921600](#))

CONFIG_BT_NIMBLE_USE_HCI_UART_PARITY

Uart PARITY

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Host-controller Transport](#) > [CONFIG_BT_NIMBLE_TRANSPORT_UART](#)

Uart Parity

Available options:

- None ([CONFIG_UART_PARITY_NONE](#))
- Odd ([CONFIG_UART_PARITY_ODD](#))
- Even ([CONFIG_UART_PARITY_EVEN](#))

CONFIG_BT_NIMBLE_UART_RX_PIN

UART Rx pin

Found in: Component config > Bluetooth > NimBLE Options > Host-controller Transport > CONFIG_BT_NIMBLE_TRANSPORT_UART

Rx pin for Nimble Transport

Default value:

- 5 if `CONFIG_BT_CONTROLLER_DISABLED` && `CONFIG_BT_NIMBLE_TRANSPORT_UART` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_UART_TX_PIN

UART Tx pin

Found in: Component config > Bluetooth > NimBLE Options > Host-controller Transport > CONFIG_BT_NIMBLE_TRANSPORT_UART

Tx pin for Nimble Transport

Default value:

- 4 if `CONFIG_BT_CONTROLLER_DISABLED` && `CONFIG_BT_NIMBLE_TRANSPORT_UART` && `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_USE_HCI_UART_FLOW_CTRL

Uart Flow Control

Found in: Component config > Bluetooth > NimBLE Options > Host-controller Transport

Uart Flow Control

Available options:

- Disable (`CONFIG_UART_HW_FLOWCTRL_DISABLE`)
- Enable hardware flow control (`CONFIG_UART_HW_FLOWCTRL_CTS_RTS`)

CONFIG_BT_NIMBLE_HCI_UART_RTS_PIN

UART Rts Pin

Found in: Component config > Bluetooth > NimBLE Options > Host-controller Transport

UART HCI RTS pin

Default value:

- 19 if `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_NIMBLE_HCI_UART_CTS_PIN

UART Cts Pin

Found in: Component config > Bluetooth > NimBLE Options > Host-controller Transport

UART HCI CTS pin

Default value:

- 23 if `CONFIG_BT_NIMBLE_ENABLED`

Debugging/Testing Contains:

- [CONFIG_BT_NIMBLE_DEBUG](#)
- [CONFIG_BT_NIMBLE_PRINT_ERR_NAME](#)
- [CONFIG_BT_NIMBLE_LOG_LEVEL](#)
- [CONFIG_BT_NIMBLE_MEM_DEBUG](#)
- [CONFIG_BT_NIMBLE_TEST_THROUGHPUT_TEST](#)

CONFIG_BT_NIMBLE_MEM_DEBUG

NimBLE memory debug

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Debugging/Testing](#)

Enable NimBLE memory debug to track memory allocations and deallocations. This feature records the function name, line number, and size of each allocation, and provides APIs to display memory usage statistics, find memory leaks, and monitor peak memory usage.

CONFIG_BT_NIMBLE_LOG_LEVEL

NimBLE Host log verbosity

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Debugging/Testing](#)

Select NimBLE log level. Note that the selected NimBLE log verbosity can not exceed the level set in "Component config --> Log output --> Default log verbosity".

Available options:

- No logs ([CONFIG_BT_NIMBLE_LOG_LEVEL_NONE](#))
- Error logs ([CONFIG_BT_NIMBLE_LOG_LEVEL_ERROR](#))
- Warning logs ([CONFIG_BT_NIMBLE_LOG_LEVEL_WARNING](#))
- Info logs ([CONFIG_BT_NIMBLE_LOG_LEVEL_INFO](#))
- Debug logs ([CONFIG_BT_NIMBLE_LOG_LEVEL_DEBUG](#))

CONFIG_BT_NIMBLE_PRINT_ERR_NAME

Enable feature to print Error description

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Debugging/Testing](#)

Enable feature to give useful explanation for HCI errors

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_DEBUG

Enable extra runtime asserts and host debugging

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Debugging/Testing](#)

This enables extra runtime asserts and host debugging

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_TEST_THROUGHPUT_TEST

Throughput Test Mode enable

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Debugging/Testing](#)

Enable the throughput test mode

Default value:

- No (disabled) if [CONFIG_BT_NIMBLE_ENABLED](#)

Vendor / Optimization Contains:

- [CONFIG_BT_NIMBLE_ADV_SEND_CONSTANT_DID](#)
- [CONFIG_BT_NIMBLE_SCAN_ALLOW_ENH_ADI_FILTER](#)
- [CONFIG_BT_NIMBLE_HIGH_DUTY_ADV_ITVL](#)
- [CONFIG_BT_NIMBLE_VS_SUPPORT](#)
- [CONFIG_BT_NIMBLE_OPTIMIZE_MULTI_CONN](#)

CONFIG_BT_NIMBLE_VS_SUPPORT

Enable support for VSC and VSE

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Vendor / Optimization](#)

This option is used to enable support for sending Vendor Specific HCI commands and handling Vendor Specific HCI Events.

CONFIG_BT_NIMBLE_OPTIMIZE_MULTI_CONN

Enable the optimization of multi-connection

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Vendor / Optimization](#)

This option enables the use of vendor-specific APIs for multi-connections, which can greatly enhance the stability of coexistence between numerous central and peripheral devices. It will prohibit the usage of standard APIs.

Default value:

- No (disabled) if `SOC_BLE_MULTI_CONN_OPTIMIZATION` && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_HIGH_DUTY_ADV_ITVL

Enable BLE high duty advertising interval feature

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Vendor / Optimization](#)

This enable BLE high duty advertising interval feature

CONFIG_BT_NIMBLE_ADV_SEND_CONSTANT_DID

Enable BLE Ext Adv to Send with Specific DID

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Vendor / Optimization](#)

This enable vendor-specific APIs to send specific DID for different advertising data in extended advertising.

CONFIG_BT_NIMBLE_SCAN_ALLOW_ENH_ADI_FILTER

Enable BLE Ext Scan to Receive Packet with Specific ADI

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Vendor / Optimization](#)

This enables vendor-specific APIs to only receive specific groups of ADIs, and implements DID's duplicate filtering for extended scanning. Enabling this can reduce the unnecessary reception of secondary channel PDUs carrying non-targeted ADIs.

Helper Utils Contains:

- [CONFIG_BT_NIMBLE_EXTRA_ADV_FIELDS](#)
- [CONFIG_BT_NIMBLE_UTIL_API](#)
- [CONFIG_BT_NIMBLE_CHK_HOST_STATUS](#)

CONFIG_BT_NIMBLE_CHK_HOST_STATUS

Enable host status query support

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Helper Utils](#)

Enable support to query host status

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_UTIL_API

Enable helper APIs

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Helper Utils](#)

Enable support for misc APIs like set data addr change, read local resolv addr and set host feature

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_NIMBLE_EXTRA_ADV_FIELDS

Enable extra Adv report fields parsing

Found in: [Component config](#) > [Bluetooth](#) > [NimBLE Options](#) > [Helper Utils](#)

Enable support for parsing of more adv fields as per spec

Default value:

- Yes (enabled) if [CONFIG_BT_NIMBLE_ENABLED](#) && [CONFIG_BT_NIMBLE_ENABLED](#)

Controller Options

CONFIG_BT_RELEASE_IRAM

Release Bluetooth text (READ DOCS FIRST)

Found in: [Component config](#) > [Bluetooth](#)

This option release Bluetooth text section and merge Bluetooth data, bss & text into a large free heap region when esp_bt_mem_release is called, total saving ~21kB or more of IRAM. ESP32-C2 only 3 configurable PMP entries available, rest of them are hard-coded. We cannot split the memory into 3 different regions (IRAM, BLE-IRAM, DRAM). So this option will disable the PMP (ESP_SYSTEM_PMP_IDRAM_SPLIT)

Default value:

- No (disabled) if [CONFIG_BT_ENABLED](#) && BT_LE_RELEASE_IRAM_SUPPORTED

Common Options Contains:

- [BT Logs](#)
- [CONFIG_BT_LE_USED_MEM_STATISTICS_ENABLED](#)
- [CONFIG_BT_ALARM_MAX_NUM](#)
- [CONFIG_BT_SMP_CRYPTO_STACK](#)

CONFIG_BT_ALARM_MAX_NUM

Maximum number of Bluetooth alarms

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#)

This option decides the maximum number of alarms which could be used by Bluetooth host.

Default value:

- 50 if [CONFIG_BT_BLUEDROID_ENABLED](#) || [CONFIG_BT_NIMBLE_ENABLED](#)

CONFIG_BT_SMP_CRYPTO_STACK

SMP cryptographic stack

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#)

Select the cryptographic library to use for SMP operations (AES, AES-CMAC, ECDH P-256).

Available options:

- TinyCrypt ([CONFIG_BT_SMP_CRYPTO_STACK_TINYCRYPT](#))
Use TinyCrypt library for cryptographic operations. TinyCrypt is a lightweight cryptographic library designed for constrained devices. This can reduce code size compared to the native implementation. This is the default option.
- mbedTLS ([CONFIG_BT_SMP_CRYPTO_STACK_MBEDTLS](#))
Use mbedTLS library for cryptographic operations. This can provide hardware acceleration on supported platforms and reduce code size by sharing crypto implementations with other components.

BT Logs Contains:

- [CONFIG_BT_LOG_CRITICAL_ONLY](#)
- [CONFIG_BLE_LOG_ENABLED](#)
- [Legacy SPI Log Output \(Deprecated - use BT Log Async Output instead\)](#)
- [Log Sources](#)

Log Sources Contains:

- [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

CONFIG_BT_LE_CONTROLLER_LOG_ENABLED

Enable Controller logs

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#)

Default value:

- No (disabled) if SOC_ESP_NIMBLE_CONTROLLER

Contains:

- [CONFIG_BT_LE_CONTROLLER_LOG_DUMP_ONLY](#)
- [CONFIG_BT_LE_CONTROLLER_LOG_MOD_OUTPUT_SWITCH](#)
- [CONFIG_BT_LE_CONTROLLER_LOG_OUTPUT_LEVEL](#)
- [CONFIG_BT_LE_CONTROLLER_LOG_OUTPUT_MODE](#)
- [CONFIG_BT_LE_CONTROLLER_LOG_HCI_ENABLED](#)
- [CONFIG_BT_LE_CONTROLLER_LOG_CTRL_ENABLED](#)
- [CONFIG_BT_LE_CONTROLLER_LOG_TASK_WDT_USER_HANDLER_ENABLE](#)
- [CONFIG_BT_LE_CONTROLLER_LOG_WRAP_PANIC_HANDLER_ENABLE](#)
- [CONFIG_BT_LE_LOG_CTRL_BUF1_SIZE](#)
- [CONFIG_BT_LE_LOG_HCI_BUF_SIZE](#)
- [CONFIG_BT_LE_CONTROLLER_LOG_SPI_OUT_ENABLED](#)
- [CONFIG_BT_LE_LOG_CTRL_BUF2_SIZE](#)
- [CONFIG_BT_LE_CONTROLLER_LOG_STORAGE_ENABLE](#)

CONFIG_BT_LE_CONTROLLER_LOG_OUTPUT_MODE

Controller log output mode

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#) > [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

Available options:

- BLE Log v2 mode ([CONFIG_BT_LE_CONTROLLER_LOG_MODE_BLE_LOG_V2](#))
Utilize BLE Log v2 for controller log
- Legacy log mode (Deprecated) ([CONFIG_BT_LE_CTRL_LEGACY_LOG_MODE_ENABLED](#))

CONFIG_BT_LE_CONTROLLER_LOG_CTRL_ENABLED

Enable controller link layer logs

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#) > [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

Default value:

- Yes (enabled) if [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#) && [CONFIG_BT_LE_CTRL_LEGACY_LOG_MODE_ENABLED](#)

CONFIG_BT_LE_CONTROLLER_LOG_HCI_ENABLED

Enable controller HCI logs

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#) > [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

Default value:

- Yes (enabled) if [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#) && [CONFIG_BT_LE_CTRL_LEGACY_LOG_MODE_ENABLED](#)

CONFIG_BT_LE_CONTROLLER_LOG_DUMP_ONLY

Controller log dump mode only

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#) > [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

Only operate in dump mode. Logs are cached internally only, not output asynchronously.

Default value:

- Yes (enabled) if [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#) && [CONFIG_BT_LE_CTRL_LEGACY_LOG_MODE_ENABLED](#)

CONFIG_BT_LE_CONTROLLER_LOG_SPI_OUT_ENABLED

Output controller logs to SPI bus (Experimental)

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#) > [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

Output ble controller logs to SPI bus

CONFIG_BT_LE_CONTROLLER_LOG_STORAGE_ENABLE

Store controller logs to flash (Experimental)

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#) > [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

Store ble controller logs to flash memory.

CONFIG_BT_LE_CONTROLLER_LOG_PARTITION_SIZE

Controller log partition size (multiples of 4K)

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#) > [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#) > [CONFIG_BT_LE_CONTROLLER_LOG_STORAGE_ENABLE](#)

The size of ble controller log partition shall be a multiples of 4K. The name of log partition shall be "bt_ctrl_log". The partition type shall be ESP_PARTITION_TYPE_DATA. The partition sub_type shall be ESP_PARTITION_SUBTYPE_ANY.

Default value:

- 65536 if [CONFIG_BT_LE_CONTROLLER_LOG_STORAGE_ENABLE](#) && [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#) && [CONFIG_BT_LE_CTRL_LEGACY_LOG_MODE_ENABLED](#)

CONFIG_BT_LE_LOG_CTRL_BUF1_SIZE

First controller log buffer size

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#) > [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

Configure the size of the first BLE controller LOG buffer.

Default value:

- 4096 if [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#) && [CONFIG_BT_LE_CTRL_LEGACY_LOG_MODE_ENABLED](#)

CONFIG_BT_LE_LOG_CTRL_BUF2_SIZE

Second controller log buffer size

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#) > [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

Configure the size of the second BLE controller LOG buffer.

Default value:

- 1024 if [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#) && [CONFIG_BT_LE_CTRL_LEGACY_LOG_MODE_ENABLED](#)

CONFIG_BT_LE_LOG_HCI_BUF_SIZE

HCI log buffer size

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#) > [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

Configure the size of the BLE HCI LOG buffer.

Default value:

- 4096 if [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#) && [CONFIG_BT_LE_CTRL_LEGACY_LOG_MODE_ENABLED](#)

CONFIG_BT_LE_CONTROLLER_LOG_OUTPUT_LEVEL

Controller log output level

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#) > [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

The output level of controller log.

Range:

- from 0 to 5 if [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

Default value:

- 2 if [CONFIG_BT_LOG_CRITICAL_ONLY_CTRL](#) && [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)
- 1 if [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

CONFIG_BT_LE_CONTROLLER_LOG_MOD_OUTPUT_SWITCH

Controller log module output switch

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#) > [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

Bitmask to enable/disable logging for individual controller modules. 0xFFFFFFFF enables all modules.

Range:

- from 0 to 0xFFFFFFFF if [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

Default value:

- "0xFFFFFFFF" if [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

CONFIG_BT_LE_CONTROLLER_LOG_WRAP_PANIC_HANDLER_ENABLE

Enable wrap panic handler

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#) > [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

Wrap esp_panic_handler to get controller logs when PC pointer exception crashes.

Default value:

- No (disabled) if [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

CONFIG_BT_LE_CONTROLLER_LOG_TASK_WDT_USER_HANDLER_ENABLE

Enable esp_task_wdt_isr_user_handler implementation

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Log Sources](#) > [CONFIG_BT_LE_CONTROLLER_LOG_ENABLED](#)

Implement esp_task_wdt_isr_user_handler to get controller logs when task wdt issue is triggered.

Default value:

- No (disabled) if `CONFIG_BT_LE_CONTROLLER_LOG_ENABLED`

CONFIG_BT_LOG_CRITICAL_ONLY

Enable bandwidth-optimized log mode (critical logs only)

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#)

Enable bandwidth-optimized logging for the BLE Log Async Output system. When enabled, only high-severity logs are captured and log encoding is applied to reduce UART/SPI DMA bandwidth usage.

Each stack component below can be independently enabled. Requires a DMA-backed output peripheral —configure in: BLE Log Module → BLE Log peripheral choice.

Contains:

- [CONFIG_BT_LOG_CRITICAL_ONLY_CTRL](#)
- [CONFIG_BT_LOG_CRITICAL_ONLY_HOST](#)
- [CONFIG_BT_LOG_CRITICAL_ONLY_MESH](#)

CONFIG_BT_LOG_CRITICAL_ONLY_CTRL

Controller: bandwidth-optimized logging

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BT_LOG_CRITICAL_ONLY](#)

Enable controller log output via the async transport with a reduced output level for bandwidth optimization. The controller log level defaults to 2 when active.

Default value:

- Yes (enabled) if `(SOC_ESP_NIMBLE_CONTROLLER_BT_CTRL_RUN_IN_FLASH_ONLY) && CONFIG_BT_LOG_CRITICAL_ONLY` ||

CONFIG_BT_LOG_CRITICAL_ONLY_HOST

Host: bandwidth-optimized logging

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BT_LOG_CRITICAL_ONLY](#)

Enable host stack log encoding via the async transport. For Bluedroid, the per-level compression options control which severity levels are encoded —configure in: Settings of BLE Log Compression → BLE Host log compression.

Default value:

- Yes (enabled) if `CONFIG_BT_LOG_CRITICAL_ONLY`

CONFIG_BT_LOG_CRITICAL_ONLY_MESH

Mesh: bandwidth-optimized logging

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BT_LOG_CRITICAL_ONLY](#)

Enable mesh log encoding via the async transport. The per-level compression options are configured in: Settings of BLE Log Compression → BLE Mesh log compression.

Default value:

- Yes (enabled) if `SOC_BLE_MESH_SUPPORTED && CONFIG_BLE_MESH && CONFIG_BT_LOG_CRITICAL_ONLY`

CONFIG_BLE_LOG_ENABLED

Enable BT Log Async Output (Dev Only)

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#)

Enable BT Log Async Output

Default value:

- No (disabled)

CONFIG_BLE_LOG_TASK_STACK_SIZE

Stack size for BLE Log Task

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

Stack size for BLE Log Task

Default value:

- 1024 if [CONFIG_BLE_LOG_ENABLED](#)
- 1024 if [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_LBM_TRANS_BUF_SIZE

Total buffer memory per common LBM (bytes)

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

Total buffer memory allocated for each common pool log buffer manager (LBM). This memory is divided equally among internal transport buffers. Must be a multiple of BLE_LOG_TRANS_BUF_CNT (currently 4).

The common pool contains: - BLE_LOG_LBM_ATOMIC_LOCK_TASK_CNT atomic LBMs (task context) - BLE_LOG_LBM_ATOMIC_LOCK_ISR_CNT atomic LBMs (ISR context) - 2 spinlock-protected LBMs (one for task, one for ISR fallback)

Total common pool memory: (ATOMIC_TASK_CNT + ATOMIC_ISR_CNT + 2) * BLE_LOG_LBM_TRANS_BUF_SIZE

Default value:

- 2048 if [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_LBM_ATOMIC_LOCK_TASK_CNT

Count of log buffer managers with atomic lock protection for task context

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

BLE Log module will search for an LBM with atomic lock protection first; if all LBMs with atomic lock protection are unavailable, BLE Log module will try to use the LBM with spin lock protection. So the more LBMs with atomic lock protection are created, the better the logging performance will be.

Default value:

- 2 if [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_LBM_ATOMIC_LOCK_ISR_CNT

Count of log buffer managers with atomic lock protection for ISR context

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

BLE Log module will search for an LBM with atomic lock protection first; if all LBMs with atomic lock protection are unavailable, BLE Log module will try to use the LBM with spin lock protection. So the more LBMs with atomic lock protection are created, the more ISRs can nest.

Default value:

- 1 if [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_IS_ESP_CONTROLLER

Current BLE Controller is ESP BLE Controller

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

Current BLE Controller is ESP BLE Controller

Default value:

- Yes (enabled) if [CONFIG_BT_CONTROLLER_ENABLED](#) && [SOC_ESP_NIMBLE_CONTROLLER](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_IS_ESP_LEGACY_CONTROLLER

Current BLE Controller is ESP BLE Legacy Controller

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

Current BLE Controller is ESP BLE Legacy Controller

CONFIG_BLE_LOG_LL_ENABLED

Enable BLE Log for Link Layer

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

Enable BLE Log for Link Layer

Default value:

- No (disabled) if [CONFIG_BT_CONTROLLER_ENABLED](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_LBM_LL_TRANS_BUF_SIZE

Total buffer memory per Link Layer LBM (bytes)

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [CONFIG_BLE_LOG_LL_ENABLED](#)

Total buffer memory allocated for each Link Layer dedicated log buffer manager (LBM). This memory is divided equally among internal transport buffers. Must be a multiple of BLE_LOG_TRANS_BUF_CNT (currently 4).

There are 2 Link Layer LBMs without lock protection (each is accessed from a single context only): - LL task LBM (Link Layer task context logs) - LL HCI LBM (Link Layer HCI context logs)

Total LL pool memory: 2 * BLE_LOG_LBM_LL_TRANS_BUF_SIZE

Default value:

- 2048 if [CONFIG_BLE_LOG_LL_ENABLED](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_LL_HCI_LOG_PAYLOAD_LEN_LIMIT_ENABLED

Enable LL HCI Log Payload Length Limit

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [CONFIG_BLE_LOG_LL_ENABLED](#)

Enable length limit for LL HCI Log payload (addr_append). When enabled, if len_append exceeds the configured limit, it will be truncated to the maximum length.

Default value:

- No (disabled) if [CONFIG_BLE_LOG_LL_ENABLED](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_LL_HCI_LOG_PAYLOAD_LEN_LIMIT

LL HCI Log Payload Length Limit (bytes)

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [CONFIG_BLE_LOG_LL_ENABLED](#) > [CONFIG_BLE_LOG_LL_HCI_LOG_PAYLOAD_LEN_LIMIT_ENABLED](#)

Maximum length for LL HCI Log payload (len_append). When the feature is enabled and len_append exceeds this value, it will be truncated.

Default value:

- 32 if [CONFIG_BLE_LOG_LL_HCI_LOG_PAYLOAD_LEN_LIMIT_ENABLED](#) && [CONFIG_BLE_LOG_LL_ENABLED](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_HOST_SIDE_HCI_LOG_ENABLED

Enable BLE Host side HCI Logging

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

Enable BLE Host side HCI Logging

Default value:

- Yes (enabled) if [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_TS_ENABLED

Enable BLE Log Timestamp Synchronization (TS)

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

Enable BLE Log TS with external logging module

Default value:

- No (disabled) if [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_SYNC_IO_NUM

GPIO number for Timestamp Synchronization (TS) toggle output

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [CONFIG_BLE_LOG_TS_ENABLED](#)

GPIO number for TS toggle output

Default value:

- 0 if [CONFIG_BLE_LOG_TS_ENABLED](#) && [CONFIG_BLE_LOG_TS_ENABLED](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_TS_TRIGGER_TIMEOUT_MS

Timeout (ms) for Timestamp Synchronization toggle

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [CONFIG_BLE_LOG_TS_ENABLED](#)

Timeout (ms) for Timestamp Synchronization toggle

Default value:

- 1000 if [CONFIG_BLE_LOG_TS_ENABLED](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_TS_TRIGGER_CHOICE

BLE Log Timestamp Synchronization trigger choice

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [CONFIG_BLE_LOG_TS_ENABLED](#)

Choose BLE Log Timestamp Synchronization trigger

Available options:

- BLE Log Timestamp Synchronization trigger - ESP Timer ([CONFIG_BLE_LOG_TS_TRIGGER_ESP_TIMER](#))
ESP Timer based periodic TS trigger
- BLE Log Timestamp Synchronization trigger - Task Event ([CONFIG_BLE_LOG_TS_TRIGGER_TASK_EVENT](#))
Task Event based TS trigger (Light Sleep Test Compatibility)

CONFIG_BLE_LOG_TS_TRIGGER_ESP_TIMER_ISR_DISPATCH_METHOD

Utilize ISR dispatch method for ESP Timer as Timestamp Synchronization trigger

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [CONFIG_BLE_LOG_TS_ENABLED](#)

Utilize ISR dispatch method for ESP Timer as Timestamp Synchronization trigger

Default value:

- No (disabled) if [CONFIG_BLE_LOG_TS_TRIGGER_ESP_TIMER](#) && [CONFIG_BLE_LOG_TS_ENABLED](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_PRPH_CHOICE

BLE Log peripheral choice

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

Choose BLE Log peripheral

Available options:

- Dummy transport ([CONFIG_BLE_LOG_PRPH_DUMMY](#))
Dummy transport (dump only)
- Utilize SPI master DMA driver as transport ([CONFIG_BLE_LOG_PRPH_SPI_MASTER_DMA](#))
Utilize SPI master DMA driver as transport
- Utilize UART DMA driver as transport ([CONFIG_BLE_LOG_PRPH_UART_DMA](#))
Utilize UART DMA driver as transport

CONFIG_BLE_LOG_PRPH_SPI_MASTER_DMA_MOSI_IO_NUM

GPIO number of MOSI port for SPI master DMA transport

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

GPIO number of MOSI port for SPI master DMA transport

Default value:

- 0 if [CONFIG_BLE_LOG_PRPH_SPI_MASTER_DMA](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_PRPH_SPI_MASTER_DMA_SCLK_IO_NUM

GPIO number of SCLK port for SPI master DMA transport

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

GPIO number of SCLK port for SPI master DMA transport

Default value:

- 0 if [CONFIG_BLE_LOG_PRPH_SPI_MASTER_DMA](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_PRPH_SPI_MASTER_DMA_CS_IO_NUM

GPIO number of CS port for SPI master DMA transport

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

GPIO number of CS port for SPI master DMA transport

Default value:

- 0 if [CONFIG_BLE_LOG_PRPH_SPI_MASTER_DMA](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_PRPH_UART_DMA_PORT

UART port number for UART DMA transport

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

UART port number for UART DMA

Default value:

- 0 if [CONFIG_BLE_LOG_PRPH_UART_DMA](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_PRPH_UART_DMA_BAUD_RATE

Baud rate of UART port for UART DMA transport

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

Determine the baud rate of UART port

Default value:

- 3000000 if [CONFIG_BLE_LOG_PRPH_UART_DMA](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_LOG_PRPH_UART_DMA_TX_IO_NUM

GPIO number for UART TX

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#)

GPIO number for UART TX

Default value:

- 0 if [CONFIG_BLE_LOG_PRPH_UART_DMA](#) && [CONFIG_BLE_LOG_ENABLED](#)

Settings of BLE Log Compression Contains:

- [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#)

CONFIG_BLE_COMPRESSED_LOG_ENABLE

Enable BLE log compression(Preview, Please read help information)

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#)

Compress BLE logs during application build to reduce flash usage and improve output speed. When enabled, log data from Bluetooth Low Energy components will be compressed before storage, optimizing both memory footprint and transmission efficiency.

Note: This library depends on additional Python packages. It will function correctly only after these dependencies are installed; refer to: "[components/bt/common/ble_log/log_compression/README.en.md](#)" for installation instructions.

Default value:

- No (disabled) if [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE

Enable BLE Mesh log compression(Preview)

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#)

Apply compression to ESP-BLE-MESH protocol stack logs. Requires base BLE compression to be enabled. Specifically optimizes log storage and transmission for ble mesh.

Note: This library depends on additional Python packages. It will function correctly only after these dependencies are installed; refer to: "[components/bt/common/ble_log/log_compression/README.en.md](#)" for installation instructions.

If the required packages are not installed, the log-compression mechanism will remain disabled even when this Config is enabled.

Default value:

- No (disabled) if [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_MESH](#) && [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_LOG_ENABLED](#)

Contains:

- [CONFIG_BLE_MESH_COMPRESSED_LOG_BUFFER_LEN](#)
- *Select the net buf log tag to be compressed*
- *Select the stack log tag to be compressed*

CONFIG_BLE_MESH_COMPRESSED_LOG_BUFFER_LEN

BLE Mesh log buffer length

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) > [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#)

Maximum output length for a single log

Default value:

- 400 if [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_LOG_ENABLED](#)

Select the stack log tag to be compressed Contains:

- [CONFIG_BLE_MESH_STACK_DEBUG_LOG_COMPRESSION](#)
- [CONFIG_BLE_MESH_STACK_ERR_LOG_COMPRESSION](#)
- [CONFIG_BLE_MESH_STACK_INFO_LOG_COMPRESSION](#)
- [CONFIG_BLE_MESH_STACK_WARN_LOG_COMPRESSION](#)

CONFIG_BLE_MESH_STACK_ERR_LOG_COMPRESSION

Compress ERROR log of ESP-BLE-MESH

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) > [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) > [Select the stack log tag to be compressed](#)

The error log in the BLE-MESH component will be compressed

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_MESH_STACK_ERR_LOG_PRESERVE

Keep the original error log statement

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) > [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) > [Select the stack log tag to be compressed](#) > [CONFIG_BLE_MESH_STACK_ERR_LOG_COMPRESSION](#)

When this option is enabled, the log data will be output through both the compressed log interface and the original UART interface at the same time, meaning that the log statements will appear on both paths. However, please note that this dual-output approach introduces additional code and string constants, which will increase the size of the firmware binary file. When this option is disabled, the logs will no longer be printed through the original UART output path; instead, they will only be output through the compressed log interface. As the code and strings related to the original UART output are omitted, the size of the firmware binary file can be effectively reduced.

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH_STACK_ERR_LOG_COMPRESSION](#) && [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_MESH_STACK_WARN_LOG_COMPRESSION

Compress warn log of ESP-BLE-MESH

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) > [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) > [Select the stack log tag to be compressed](#)

The warn log in the BLE-MESH component will be compressed

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_MESH_STACK_WARN_LOG_PRESERVE

Keep the original warn log statement

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) > [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) > [Select the stack log tag to be compressed](#) > [CONFIG_BLE_MESH_STACK_WARN_LOG_COMPRESSION](#)

Please refer to the help information in [BLE_MESH_STACK_ERR_LOG_PRESERVE](#)

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH_STACK_WARN_LOG_COMPRESSION](#) && [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_MESH_STACK_INFO_LOG_COMPRESSION

Compress info log of ESP-BLE-MESH

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) > [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) > [Select the stack log tag to be compressed](#)

The info log in the BLE-MESH component will be compressed

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_MESH_STACK_INFO_LOG_PRESERVE

Keep the original info log statement

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) > [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) > [Select the stack log tag to be compressed](#) > [CONFIG_BLE_MESH_STACK_INFO_LOG_COMPRESSION](#)

Please refer to the help information in [BLE_MESH_STACK_ERR_LOG_PRESERVE](#)

Default value:

- No (disabled) if [CONFIG_BLE_MESH_STACK_INFO_LOG_COMPRESSION](#) && [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_MESH_STACK_DEBUG_LOG_COMPRESSION

Compress debug log of ESP-BLE-MESH

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) > [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) > [Select the stack log tag to be compressed](#)

The debug log in the BLE-MESH component will be compressed

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_MESH_STACK_DEBUG_LOG_PRESERVE

Keep the original debug log statement

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) > [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) > [Select the stack log tag to be compressed](#) > [CONFIG_BLE_MESH_STACK_DEBUG_LOG_COMPRESSION](#)

Please refer to the help information in [BLE_MESH_STACK_ERR_LOG_PRESERVE](#)

Default value:

- No (disabled) if [CONFIG_BLE_MESH_STACK_DEBUG_LOG_COMPRESSION](#) && [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_LOG_ENABLED](#)

Select the net buf log tag to be compressed Contains:

- [CONFIG_BLE_MESH_NET_BUF_DEBUG_LOG_COMPRESSION](#)
- [CONFIG_BLE_MESH_NET_BUF_ERR_LOG_COMPRESSION](#)
- [CONFIG_BLE_MESH_NET_BUF_INFO_LOG_COMPRESSION](#)
- [CONFIG_BLE_MESH_NET_BUF_WARN_LOG_COMPRESSION](#)

CONFIG_BLE_MESH_NET_BUF_ERR_LOG_COMPRESSION

Compress ERROR log of ESP-BLE-MESH

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) > [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) > [Select the net buf log tag to be compressed](#)

The error log in the BLE-MESH component will be compressed

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_MESH_NET_BUF_ERR_LOG_PRESERVE

Keep the original error log statement

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) > [CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE](#) > [Select the net buf log tag to be compressed](#) > [CONFIG_BLE_MESH_NET_BUF_ERR_LOG_COMPRESSION](#)

When this option is enabled, the log data will be output through both the compressed log interface and the original UART interface at the same time, meaning that the log statements will appear on both paths. However, please note that this dual-output approach introduces additional code and string constants, which will increase the size of the firmware binary file. When this option is disabled, the logs will no longer be printed through the original UART output path; instead, they will only be output through the compressed log interface. As the code and strings related to the original UART output are omitted, the size of the firmware binary file can be effectively reduced.

Default value:

- Yes (enabled) if `CONFIG_BLE_MESH_NET_BUF_ERR_LOG_COMPRESSION` && `CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_MESH_NET_BUF_WARN_LOG_COMPRESSION

Compress warn log of ESP-BLE-MESH

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE > Select the net buf log tag to be compressed

The warn log in the BLE-MESH component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_MESH_NET_BUF_WARN_LOG_PRESERVE

Keep the original warn log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE > Select the net buf log tag to be compressed > CONFIG_BLE_MESH_NET_BUF_WARN_LOG_COMPRESSION

Please refer to the help information in `BLE_MESH_NET_BUF_ERR_LOG_PRESERVE`

Default value:

- Yes (enabled) if `CONFIG_BLE_MESH_NET_BUF_WARN_LOG_COMPRESSION` && `CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_MESH_NET_BUF_INFO_LOG_COMPRESSION

Compress info log of ESP-BLE-MESH

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE > Select the net buf log tag to be compressed

The info log in the BLE-MESH component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_MESH_NET_BUF_INFO_LOG_PRESERVE

Keep the original info log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE > Select the net buf log tag to be compressed > CONFIG_BLE_MESH_NET_BUF_INFO_LOG_COMPRESSION

Please refer to the help information in BLE_MESH_NET_BUF_ERR_LOG_PRESERVE

Default value:

- No (disabled) if `CONFIG_BLE_MESH_NET_BUF_INFO_LOG_COMPRESSION` && `CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_MESH_NET_BUF_DEBUG_LOG_COMPRESSION

Compress debug log of ESP-BLE-MESH

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE > Select the net buf log tag to be compressed

The debug log in the BLE-MESH component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_MESH_NET_BUF_DEBUG_LOG_PRESERVE

Keep the original debug log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE > Select the net buf log tag to be compressed > CONFIG_BLE_MESH_NET_BUF_DEBUG_LOG_COMPRESSION

Please refer to the help information in BLE_MESH_NET_BUF_ERR_LOG_PRESERVE

Default value:

- No (disabled) if `CONFIG_BLE_MESH_NET_BUF_DEBUG_LOG_COMPRESSION` && `CONFIG_BLE_MESH_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE

Enable BLE Host log compression(Preview, only Bluedroid Host for now)

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE

Apply compression to host logs. Requires base BLE compression to be enabled. Specifically optimizes log storage and transmission.

Note: This library depends on additional Python packages. It will function correctly only after these dependencies are installed; refer to: "components/bt/common/ble_log/log_compression/README.en.md" for installation instructions.

Default value:

- No (disabled) if `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BT_BLUEDROID_ENABLED` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

Contains:

- `CONFIG_BLE_HOST_COMPRESSED_LOG_BUFFER_LEN`
- Select the APPL layer log tag to be compressed
- Select the BTM layer log tag to be compressed
- Select the GAP layer log tag to be compressed
- Select the GATT layer log tag to be compressed
- Select the LA2CAP layer log tag to be compressed
- Select the SMP layer log tag to be compressed

CONFIG_BLE_HOST_COMPRESSED_LOG_BUFFER_LEN

Host log buffer length

Found in: `Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE`

Maximum output length for a single log

Default value:

- 300 if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

Select the BTM layer log tag to be compressed **Contains:**

- `CONFIG_BLE_BLUEDROID_BTM_API_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_BTM_DEBUG_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_BTM_ERROR_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_BTM_EVENT_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_BTM_VERBOSE_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_BTM_WARNING_LOG_COMPRESSION`

CONFIG_BLE_BLUEDROID_BTM_ERROR_LOG_COMPRESSION

Compress error log of Bluedroid host

Found in: `Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the BTM layer log tag to be compressed`

The error log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_BTM_ERROR_LOG_PRESERVE

Keep the original error log statement

Found in: `Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the BTM layer log tag to be compressed > CONFIG_BLE_BLUEDROID_BTM_ERROR_LOG_COMPRESSION`

When this option is enabled, the log data will be output through both the compressed log interface and the original UART interface at the same time, meaning that the log statements will appear on both paths. However, please note that this dual-output approach introduces additional code and string constants, which will increase the size of the firmware binary file. When this option is disabled, the logs will no longer be printed through the original UART output path; instead, they will only be output through the compressed log interface. As the code and strings related to the original UART output are omitted, the size of the firmware binary file can be effectively reduced.

Default value:

- Yes (enabled) if `CONFIG_BLE_BLUEDROID_BTМ_ERROR_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_BTМ_WARNING_LOG_COMPRESSION

Compress warning log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the BTМ layer log tag to be compressed

The warning log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_BTМ_WARNING_LOG_PRESERVE

Keep the original warning log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the BTМ layer log tag to be compressed > CONFIG_BLE_BLUEDROID_BTМ_WARNING_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_BTМ_ERROR_LOG_PRESERVE

Default value:

- Yes (enabled) if `CONFIG_BLE_BLUEDROID_BTМ_WARNING_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_BTМ_API_LOG_COMPRESSION

Compress api log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the BTМ layer log tag to be compressed

The api log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_BTAPI_LOG_PRESERVE

Keep the original api log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the BTM layer log tag to be compressed > CONFIG_BLE_BLUEDROID_BTAPI_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_BTAPI_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if `CONFIG_BLE_BLUEDROID_BTAPI_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_BTAPI_EVENT_LOG_COMPRESSION

Compress event log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the BTM layer log tag to be compressed

The event log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_BTAPI_EVENT_LOG_PRESERVE

Keep the original event log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the BTM layer log tag to be compressed > CONFIG_BLE_BLUEDROID_BTAPI_EVENT_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_BTAPI_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if `CONFIG_BLE_BLUEDROID_BTAPI_EVENT_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_BTAPI_DEBUG_LOG_COMPRESSION

Compress debug log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the BTM layer log tag to be compressed

The debug log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_BTМ_DEBUG_LOG_PRESERVE

Keep the original debug log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the BTМ layer log tag to be compressed > CONFIG_BLE_BLUEDROID_BTМ_DEBUG_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_BTМ_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if *CONFIG_BLE_BLUEDROID_BTМ_DEBUG_LOG_COMPRESSION* && *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_BTМ_VERBOSE_LOG_COMPRESSION

Compress verbose log of Bluedroid host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the BTМ layer log tag to be compressed

The verbose log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_BTМ_VERBOSE_LOG_PRESERVE

Keep the original verbose log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the BTМ layer log tag to be compressed > CONFIG_BLE_BLUEDROID_BTМ_VERBOSE_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_BTМ_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if *CONFIG_BLE_BLUEDROID_BTМ_VERBOSE_LOG_COMPRESSION* && *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

Select the LA2CAP layer log tag to be compressed Contains:

- *CONFIG_BLE_BLUEDROID_L2CAP_API_LOG_COMPRESSION*
- *CONFIG_BLE_BLUEDROID_L2CAP_DEBUG_LOG_COMPRESSION*
- *CONFIG_BLE_BLUEDROID_L2CAP_ERROR_LOG_COMPRESSION*
- *CONFIG_BLE_BLUEDROID_L2CAP_EVENT_LOG_COMPRESSION*
- *CONFIG_BLE_BLUEDROID_L2CAP_VERBOSE_LOG_COMPRESSION*
- *CONFIG_BLE_BLUEDROID_L2CAP_WARNING_LOG_COMPRESSION*

CONFIG_BLE_BLUEDROID_L2CAP_ERROR_LOG_COMPRESSION

Compress error log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the LA2CAP layer log tag to be compressed

The error log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_L2CAP_ERROR_LOG_PRESERVE

Keep the original error log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the LA2CAP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_L2CAP_ERROR_LOG_COMPRESSION

When this option is enabled, the log data will be output through both the compressed log interface and the original UART interface at the same time, meaning that the log statements will appear on both paths. However, please note that this dual-output approach introduces additional code and string constants, which will increase the size of the firmware binary file. When this option is disabled, the logs will no longer be printed through the original UART output path; instead, they will only be output through the compressed log interface. As the code and strings related to the original UART output are omitted, the size of the firmware binary file can be effectively reduced.

Default value:

- Yes (enabled) if `CONFIG_BLE_BLUEDROID_L2CAP_ERROR_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_L2CAP_WARNING_LOG_COMPRESSION

Compress warning log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the LA2CAP layer log tag to be compressed

The warning log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_L2CAP_WARNING_LOG_PRESERVE

Keep the original warning log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the LA2CAP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_L2CAP_WARNING_LOG_COMPRESSION

Please refer to the help information in `BLE_BLUEDROID_L2CAP_ERROR_LOG_PRESERVE`

Default value:

- Yes (enabled) if `CONFIG_BLE_BLUEDROID_L2CAP_WARNING_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_L2CAP_API_LOG_COMPRESSION

Compress api log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the LA2CAP layer log tag to be compressed

The api log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_L2CAP_API_LOG_PRESERVE

Keep the original api log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the LA2CAP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_L2CAP_API_LOG_COMPRESSION

Please refer to the help information in `BLE_BLUEDROID_L2CAP_ERROR_LOG_PRESERVE`

Default value:

- No (disabled) if `CONFIG_BLE_BLUEDROID_L2CAP_API_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_L2CAP_EVENT_LOG_COMPRESSION

Compress event log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the LA2CAP layer log tag to be compressed

The event log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_L2CAP_EVENT_LOG_PRESERVE

Keep the original event log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the LA2CAP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_L2CAP_EVENT_LOG_COMPRESSION

Please refer to the help information in `BLE_BLUEDROID_L2CAP_ERROR_LOG_PRESERVE`

Default value:

- No (disabled) if `CONFIG_BLE_BLUEDROID_L2CAP_EVENT_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_L2CAP_DEBUG_LOG_COMPRESSION

Compress debug log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the LA2CAP layer log tag to be compressed

The debug log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_L2CAP_DEBUG_LOG_PRESERVE

Keep the original debug log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the LA2CAP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_L2CAP_DEBUG_LOG_COMPRESSION

Please refer to the help information in `BLE_BLUEDROID_L2CAP_ERROR_LOG_PRESERVE`

Default value:

- No (disabled) if `CONFIG_BLE_BLUEDROID_L2CAP_DEBUG_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_L2CAP_VERBOSE_LOG_COMPRESSION

Compress verbose log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the LA2CAP layer log tag to be compressed

The verbose log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_L2CAP_VERBOSE_LOG_PRESERVE

Keep the original verbose log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the LA2CAP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_L2CAP_VERBOSE_LOG_COMPRESSION

Please refer to the help information in `BLE_BLUEDROID_L2CAP_ERROR_LOG_PRESERVE`

Default value:

- No (disabled) if `CONFIG_BLE_BLUEDROID_L2CAP_VERBOSE_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

Select the GAP layer log tag to be compressed Contains:

- `CONFIG_BLE_BLUEDROID_GAP_API_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_GAP_DEBUG_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_GAP_ERROR_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_GAP_EVENT_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_GAP_VERBOSE_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_GAP_WARNING_LOG_COMPRESSION`

CONFIG_BLE_BLUEDROID_GAP_ERROR_LOG_COMPRESSION

Compress error log of Bluetooth host

Found in: `Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GAP layer log tag to be compressed`

The error log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GAP_ERROR_LOG_PRESERVE

Keep the original error log statement

Found in: `Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GAP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_GAP_ERROR_LOG_COMPRESSION`

When this option is enabled, the log data will be output through both the compressed log interface and the original UART interface at the same time, meaning that the log statements will appear on both paths. However, please note that this dual-output approach introduces additional code and string constants, which will increase the size of the firmware binary file. When this option is disabled, the logs will no longer be printed through the original UART output path; instead, they will only be output through the compressed log interface. As the code and strings related to the original UART output are omitted, the size of the firmware binary file can be effectively reduced.

Default value:

- Yes (enabled) if `CONFIG_BLE_BLUEDROID_GAP_ERROR_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GAP_WARNING_LOG_COMPRESSION

Compress warning log of Bluetooth host

Found in: `Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GAP layer log tag to be compressed`

The warning log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GAP_WARNING_LOG_PRESERVE

Keep the original warning log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GAP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_GAP_WARNING_LOG_COMPRESSION

Please refer to the help information in `BLE_BLUEDROID_GAP_ERROR_LOG_PRESERVE`

Default value:

- Yes (enabled) if `CONFIG_BLE_BLUEDROID_GAP_WARNING_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GAP_API_LOG_COMPRESSION

Compress api log of Bluedroid host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GAP layer log tag to be compressed

The api log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GAP_API_LOG_PRESERVE

Keep the original api log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GAP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_GAP_API_LOG_COMPRESSION

Please refer to the help information in `BLE_BLUEDROID_GAP_ERROR_LOG_PRESERVE`

Default value:

- No (disabled) if `CONFIG_BLE_BLUEDROID_GAP_API_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GAP_EVENT_LOG_COMPRESSION

Compress event log of Bluedroid host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GAP layer log tag to be compressed

The event log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GAP_EVENT_LOG_PRESERVE

Keep the original event log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GAP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_GAP_EVENT_LOG_COMPRESSION

Please refer to the help information in `BLE_BLUEDROID_GAP_ERROR_LOG_PRESERVE`

Default value:

- No (disabled) if `CONFIG_BLE_BLUEDROID_GAP_EVENT_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GAP_DEBUG_LOG_COMPRESSION

Compress debug log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GAP layer log tag to be compressed

The debug log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GAP_DEBUG_LOG_PRESERVE

Keep the original debug log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GAP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_GAP_DEBUG_LOG_COMPRESSION

Please refer to the help information in `BLE_BLUEDROID_GAP_ERROR_LOG_PRESERVE`

Default value:

- No (disabled) if `CONFIG_BLE_BLUEDROID_GAP_DEBUG_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GAP_VERBOSE_LOG_COMPRESSION

Compress verbose log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GAP layer log tag to be compressed

The verbose log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GAP_VERBOSE_LOG_PRESERVE

Keep the original verbose log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GAP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_GAP_VERBOSE_LOG_COMPRESSION

Please refer to the help information in `BLE_BLUEDROID_GAP_ERROR_LOG_PRESERVE`

Default value:

- No (disabled) if `CONFIG_BLE_BLUEDROID_GAP_VERBOSE_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

Select the GATT layer log tag to be compressed Contains:

- `CONFIG_BLE_BLUEDROID_GATT_API_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_GATT_DEBUG_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_GATT_ERROR_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_GATT_EVENT_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_GATT_VERBOSE_LOG_COMPRESSION`
- `CONFIG_BLE_BLUEDROID_GATT_WARNING_LOG_COMPRESSION`

CONFIG_BLE_BLUEDROID_GATT_ERROR_LOG_COMPRESSION

Compress error log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GATT layer log tag to be compressed

The error log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GATT_ERROR_LOG_PRESERVE

Keep the original error log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GATT layer log tag to be compressed > CONFIG_BLE_BLUEDROID_GATT_ERROR_LOG_COMPRESSION

When this option is enabled, the log data will be output through both the compressed log interface and the original UART interface at the same time, meaning that the log statements will appear on both paths. However, please note that this dual-output approach introduces additional code and string constants, which will increase the size of the firmware binary file. When this option is disabled, the logs will no longer be printed through the original UART output path; instead, they will only be output through the

compressed log interface. As the code and strings related to the original UART output are omitted, the size of the firmware binary file can be effectively reduced.

Default value:

- Yes (enabled) if `CONFIG_BLE_BLUEDROID_GATT_ERROR_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GATT_WARNING_LOG_COMPRESSION

Compress warning log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GATT layer log tag to be compressed

The warning log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GATT_WARNING_LOG_PRESERVE

Keep the original warning log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GATT layer log tag to be compressed > CONFIG_BLE_BLUEDROID_GATT_WARNING_LOG_COMPRESSION

Please refer to the help information in `BLE_BLUEDROID_GATT_ERROR_LOG_PRESERVE`

Default value:

- Yes (enabled) if `CONFIG_BLE_BLUEDROID_GATT_WARNING_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GATT_API_LOG_COMPRESSION

Compress api log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GATT layer log tag to be compressed

The api log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_GATT_API_LOG_PRESERVE

Keep the original api log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE

> Select the GATT layer log tag to be compressed > CONFIG_BLE_BLUEDROID_GATT_API_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_GATT_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if *CONFIG_BLE_BLUEDROID_GATT_API_LOG_COMPRESSION* && *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_GATT_EVENT_LOG_COMPRESSION

Compress event log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GATT layer log tag to be compressed

The event log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_GATT_EVENT_LOG_PRESERVE

Keep the original event log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GATT layer log tag to be compressed > CONFIG_BLE_BLUEDROID_GATT_EVENT_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_GATT_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if *CONFIG_BLE_BLUEDROID_GATT_EVENT_LOG_COMPRESSION* && *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_GATT_DEBUG_LOG_COMPRESSION

Compress debug log of Bluetooth host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GATT layer log tag to be compressed

The debug log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_GATT_DEBUG_LOG_PRESERVE

Keep the original debug log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GATT layer log tag to be compressed > CONFIG_BLE_BLUEDROID_GATT_DEBUG_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_GATT_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if *CONFIG_BLE_BLUEDROID_GATT_DEBUG_LOG_COMPRESSION* && *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_GATT_VERBOSE_LOG_COMPRESSION

Compress verbose log of Bluedroid host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GATT layer log tag to be compressed

The verbose log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_GATT_VERBOSE_LOG_PRESERVE

Keep the original verbose log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the GATT layer log tag to be compressed > CONFIG_BLE_BLUEDROID_GATT_VERBOSE_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_GATT_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if *CONFIG_BLE_BLUEDROID_GATT_VERBOSE_LOG_COMPRESSION* && *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

Select the SMP layer log tag to be compressed Contains:

- *CONFIG_BLE_BLUEDROID_SMP_API_LOG_COMPRESSION*
- *CONFIG_BLE_BLUEDROID_SMP_DEBUG_LOG_COMPRESSION*
- *CONFIG_BLE_BLUEDROID_SMP_ERROR_LOG_COMPRESSION*
- *CONFIG_BLE_BLUEDROID_SMP_EVENT_LOG_COMPRESSION*
- *CONFIG_BLE_BLUEDROID_SMP_VERBOSE_LOG_COMPRESSION*
- *CONFIG_BLE_BLUEDROID_SMP_WARNING_LOG_COMPRESSION*

CONFIG_BLE_BLUEDROID_SMP_ERROR_LOG_COMPRESSION

Compress error log of Bluedroid host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the SMP layer log tag to be compressed

The error log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_SMP_ERROR_LOG_PRESERVE

Keep the original error log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the SMP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_SMP_ERROR_LOG_COMPRESSION

When this option is enabled, the log data will be output through both the compressed log interface and the original UART interface at the same time, meaning that the log statements will appear on both paths. However, please note that this dual-output approach introduces additional code and string constants, which will increase the size of the firmware binary file. When this option is disabled, the logs will no longer be printed through the original UART output path; instead, they will only be output through the compressed log interface. As the code and strings related to the original UART output are omitted, the size of the firmware binary file can be effectively reduced.

Default value:

- Yes (enabled) if `CONFIG_BLE_BLUEDROID_SMP_ERROR_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_SMP_WARNING_LOG_COMPRESSION

Compress warning log of Bluedroid host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the SMP layer log tag to be compressed

The warning log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_SMP_WARNING_LOG_PRESERVE

Keep the original warning log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the SMP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_SMP_WARNING_LOG_COMPRESSION

Please refer to the help information in `BLE_BLUEDROID_SMP_ERROR_LOG_PRESERVE`

Default value:

- Yes (enabled) if `CONFIG_BLE_BLUEDROID_SMP_WARNING_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_SMP_API_LOG_COMPRESSION

Compress api log of Bluedroid host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the SMP layer log tag to be compressed

The api log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_SMP_API_LOG_PRESERVE

Keep the original api log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the SMP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_SMP_API_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_SMP_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if `CONFIG_BLE_BLUEDROID_SMP_API_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_SMP_EVENT_LOG_COMPRESSION

Compress event log of Bluedroid host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the SMP layer log tag to be compressed

The event log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_SMP_EVENT_LOG_PRESERVE

Keep the original event log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the SMP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_SMP_EVENT_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_SMP_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if `CONFIG_BLE_BLUEDROID_SMP_EVENT_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_SMP_DEBUG_LOG_COMPRESSION

Compress debug log of Bluedroid host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the SMP layer log tag to be compressed

The debug log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_SMP_DEBUG_LOG_PRESERVE

Keep the original debug log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the SMP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_SMP_DEBUG_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_SMP_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if `CONFIG_BLE_BLUEDROID_SMP_DEBUG_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_SMP_VERBOSE_LOG_COMPRESSION

Compress verbose log of Bluedroid host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the SMP layer log tag to be compressed

The verbose log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

CONFIG_BLE_BLUEDROID_SMP_VERBOSE_LOG_PRESERVE

Keep the original verbose log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the SMP layer log tag to be compressed > CONFIG_BLE_BLUEDROID_SMP_VERBOSE_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_SMP_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if `CONFIG_BLE_BLUEDROID_SMP_VERBOSE_LOG_COMPRESSION` && `CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_COMPRESSED_LOG_ENABLE` && `CONFIG_BLE_LOG_ENABLED`

Select the APPL layer log tag to be compressed Contains:

- [CONFIG_BLE_BLUEDROID_APPL_API_LOG_COMPRESSION](#)
- [CONFIG_BLE_BLUEDROID_APPL_DEBUG_LOG_COMPRESSION](#)
- [CONFIG_BLE_BLUEDROID_APPL_ERROR_LOG_COMPRESSION](#)
- [CONFIG_BLE_BLUEDROID_APPL_EVENT_LOG_COMPRESSION](#)
- [CONFIG_BLE_BLUEDROID_APPL_VERBOSE_LOG_COMPRESSION](#)
- [CONFIG_BLE_BLUEDROID_APPL_WARNING_LOG_COMPRESSION](#)

CONFIG_BLE_BLUEDROID_APPL_ERROR_LOG_COMPRESSION

Compress error log of Bluetooth host

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) > [CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE](#) > [Select the APPL layer log tag to be compressed](#)

The error log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if [CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_BLUEDROID_APPL_ERROR_LOG_PRESERVE

Keep the original error log statement

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) > [CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE](#) > [Select the APPL layer log tag to be compressed](#) > [CONFIG_BLE_BLUEDROID_APPL_ERROR_LOG_COMPRESSION](#)

When this option is enabled, the log data will be output through both the compressed log interface and the original UART interface at the same time, meaning that the log statements will appear on both paths. However, please note that this dual-output approach introduces additional code and string constants, which will increase the size of the firmware binary file. When this option is disabled, the logs will no longer be printed through the original UART output path; instead, they will only be output through the compressed log interface. As the code and strings related to the original UART output are omitted, the size of the firmware binary file can be effectively reduced.

Default value:

- Yes (enabled) if [CONFIG_BLE_BLUEDROID_APPL_ERROR_LOG_COMPRESSION](#) && [CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_BLUEDROID_APPL_WARNING_LOG_COMPRESSION

Compress warning log of Bluetooth host

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [CONFIG_BLE_LOG_ENABLED](#) > [Settings of BLE Log Compression](#) > [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) > [CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE](#) > [Select the APPL layer log tag to be compressed](#)

The warning log in the Bluetooth host component will be compressed

Default value:

- Yes (enabled) if [CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_COMPRESSED_LOG_ENABLE](#) && [CONFIG_BLE_LOG_ENABLED](#)

CONFIG_BLE_BLUEDROID_APPL_WARNING_LOG_PRESERVE

Keep the original warning log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the APPL layer log tag to be compressed > CONFIG_BLE_BLUEDROID_APPL_WARNING_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_APPL_ERROR_LOG_PRESERVE

Default value:

- Yes (enabled) if *CONFIG_BLE_BLUEDROID_APPL_WARNING_LOG_COMPRESSION* && *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_APPL_API_LOG_COMPRESSION

Compress api log of Bluedroid host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the APPL layer log tag to be compressed

The api log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_APPL_API_LOG_PRESERVE

Keep the original api log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the APPL layer log tag to be compressed > CONFIG_BLE_BLUEDROID_APPL_API_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_APPL_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if *CONFIG_BLE_BLUEDROID_APPL_API_LOG_COMPRESSION* && *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_APPL_EVENT_LOG_COMPRESSION

Compress event log of Bluedroid host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the APPL layer log tag to be compressed

The event log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_APPL_EVENT_LOG_PRESERVE

Keep the original event log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the APPL layer log tag to be compressed > CONFIG_BLE_BLUEDROID_APPL_EVENT_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_APPL_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if *CONFIG_BLE_BLUEDROID_APPL_EVENT_LOG_COMPRESSION* && *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_APPL_DEBUG_LOG_COMPRESSION

Compress debug log of Bluedroid host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the APPL layer log tag to be compressed

The debug log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_APPL_DEBUG_LOG_PRESERVE

Keep the original debug log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the APPL layer log tag to be compressed > CONFIG_BLE_BLUEDROID_APPL_DEBUG_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_APPL_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if *CONFIG_BLE_BLUEDROID_APPL_DEBUG_LOG_COMPRESSION* && *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_APPL_VERBOSE_LOG_COMPRESSION

Compress verbose log of Bluedroid host

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the APPL layer log tag to be compressed

The verbose log in the Bluedroid host component will be compressed

Default value:

- Yes (enabled) if *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

CONFIG_BLE_BLUEDROID_APPL_VERBOSE_LOG_PRESERVE

Keep the original verbose log statement

Found in: Component config > Bluetooth > Common Options > BT Logs > CONFIG_BLE_LOG_ENABLED > Settings of BLE Log Compression > CONFIG_BLE_COMPRESSED_LOG_ENABLE > CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE > Select the APPL layer log tag to be compressed > CONFIG_BLE_BLUEDROID_APPL_VERBOSE_LOG_COMPRESSION

Please refer to the help information in BLE_BLUEDROID_APPL_ERROR_LOG_PRESERVE

Default value:

- No (disabled) if *CONFIG_BLE_BLUEDROID_APPL_VERBOSE_LOG_COMPRESSION* && *CONFIG_BLE_HOST_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_COMPRESSED_LOG_ENABLE* && *CONFIG_BLE_LOG_ENABLED*

Legacy SPI Log Output (Deprecated - use BT Log Async Output instead) Contains:

- *CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED*

CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED

Output ble logs to SPI bus (Experimental)

Found in: Component config > Bluetooth > Common Options > BT Logs > Legacy SPI Log Output (Deprecated - use BT Log Async Output instead)

Output ble logs to SPI bus

CONFIG_BT_BLE_LOG_SPI_OUT_UL_TASK_BUF_SIZE

SPI transaction buffer size for upper layer task logs

Found in: Component config > Bluetooth > Common Options > BT Logs > Legacy SPI Log Output (Deprecated - use BT Log Async Output instead) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED

SPI transaction buffer size for upper layer task logs. There will be 2 SPI DMA buffers with the same size.

Default value:

- 512 if *CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED*

CONFIG_BT_BLE_LOG_SPI_OUT_HCI_ENABLED

Enable HCI log output to SPI

Found in: Component config > Bluetooth > Common Options > BT Logs > Legacy SPI Log Output (Deprecated - use BT Log Async Output instead) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED

Enable logging of HCI packets to the SPI bus when BLE SPI log output is enabled.

Default value:

- No (disabled) if *CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED*

CONFIG_BT_BLE_LOG_SPI_OUT_HCI_BUF_SIZE

SPI transaction buffer size for HCI logs

Found in: Component config > Bluetooth > Common Options > BT Logs > Legacy SPI Log Output (Deprecated - use BT Log Async Output instead) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED > CONFIG_BT_BLE_LOG_SPI_OUT_HCI_ENABLED

SPI transaction buffer size for HCI logs. There will be 2 SPI DMA buffers with the same size.

Default value:

- 1024 if `CONFIG_BT_BLE_LOG_SPI_OUT_HCI_ENABLED` && `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_HCI_TASK_CNT

HCI task count

Found in: Component config > Bluetooth > Common Options > BT Logs > Legacy SPI Log Output (Deprecated - use BT Log Async Output instead) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED > CONFIG_BT_BLE_LOG_SPI_OUT_HCI_ENABLED

HCI task count

Default value:

- 1 if `CONFIG_BT_BLE_LOG_SPI_OUT_HCI_ENABLED` && `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_HOST_ENABLED

Enable Host log output to SPI

Found in: Component config > Bluetooth > Common Options > BT Logs > Legacy SPI Log Output (Deprecated - use BT Log Async Output instead) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED

This configuration applies to the logs of both Bluedroid Host and NimBLE Host. When BLE SPI log output is enabled, this option allows host logs to be transmitted via SPI.

Default value:

- No (disabled) if `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_HOST_BUF_SIZE

SPI transaction buffer size for host logs

Found in: Component config > Bluetooth > Common Options > BT Logs > Legacy SPI Log Output (Deprecated - use BT Log Async Output instead) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED > CONFIG_BT_BLE_LOG_SPI_OUT_HOST_ENABLED

SPI transaction buffer size for host logs. There will be 2 SPI DMA buffers with the same size.

Default value:

- 1024 if `CONFIG_BT_BLE_LOG_SPI_OUT_HOST_ENABLED` && `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_HOST_TASK_CNT

Host task count

Found in: Component config > Bluetooth > Common Options > BT Logs > Legacy SPI Log Output (Deprecated - use BT Log Async Output instead) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED > CONFIG_BT_BLE_LOG_SPI_OUT_HOST_ENABLED

Host task count.

Default value:

- 2 if `CONFIG_BT_BLE_LOG_SPI_OUT_HOST_ENABLED` && `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_LL_ENABLED

Enable Controller log output to SPI

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Legacy SPI Log Output \(Deprecated - use BT Log Async Output instead\)](#) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED

Enable controller log output to SPI bus.

Default value:

- No (disabled) if `CONFIG_BT_LE_CONTROLLER_LOG_SPI_OUT_ENABLED` && `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_LL_TASK_BUF_SIZE

SPI transaction buffer size for lower layer task logs

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Legacy SPI Log Output \(Deprecated - use BT Log Async Output instead\)](#) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED > CONFIG_BT_BLE_LOG_SPI_OUT_LL_ENABLED

SPI transaction buffer size for lower layer task logs. There will be 2 SPI DMA buffers with the same size.

Default value:

- 1024 if `CONFIG_BT_BLE_LOG_SPI_OUT_LL_ENABLED` && `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_LL_ISR_BUF_SIZE

SPI transaction buffer size for lower layer ISR logs

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Legacy SPI Log Output \(Deprecated - use BT Log Async Output instead\)](#) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED > CONFIG_BT_BLE_LOG_SPI_OUT_LL_ENABLED

SPI transaction buffer size for lower layer ISR logs. There will be 2 SPI DMA buffers with the same size.

Default value:

- 512 if `CONFIG_BT_BLE_LOG_SPI_OUT_LL_ENABLED` && `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_LL_HCI_BUF_SIZE

SPI transaction buffer size for lower layer HCI logs

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Legacy SPI Log Output \(Deprecated - use BT Log Async Output instead\)](#) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED > CONFIG_BT_BLE_LOG_SPI_OUT_LL_ENABLED

SPI transaction buffer size for upper layer HCI logs. There will be 2 SPI DMA buffers with the same size

Default value:

- 512 if `CONFIG_BT_BLE_LOG_SPI_OUT_LL_ENABLED` && `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_MOSI_IO_NUM

GPIO number of SPI MOSI

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Legacy SPI Log Output \(Deprecated - use BT Log Async Output instead\)](#) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED

GPIO number of SPI MOSI

Default value:

- 0 if [CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED](#)

CONFIG_BT_BLE_LOG_SPI_OUT_SCLK_IO_NUM

GPIO number of SPI SCLK

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Legacy SPI Log Output \(Deprecated - use BT Log Async Output instead\)](#) > [CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED](#)

GPIO number of SPI SCLK

Default value:

- 1 if [CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED](#)

CONFIG_BT_BLE_LOG_SPI_OUT_CS_IO_NUM

GPIO number of SPI CS

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Legacy SPI Log Output \(Deprecated - use BT Log Async Output instead\)](#) > [CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED](#)

GPIO number of SPI CS

Default value:

- 2 if [CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED](#)

CONFIG_BT_BLE_LOG_SPI_OUT_TS_SYNC_ENABLED

Enable ble log & logic analyzer log time sync

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Legacy SPI Log Output \(Deprecated - use BT Log Async Output instead\)](#) > [CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED](#)

Enable ble log & logic analyzer log time sync

Default value:

- Yes (enabled) if [CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED](#)

CONFIG_BT_BLE_LOG_SPI_OUT_SYNC_IO_NUM

GPIO number of SYNC IO

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Legacy SPI Log Output \(Deprecated - use BT Log Async Output instead\)](#) > [CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED](#) > [CONFIG_BT_BLE_LOG_SPI_OUT_TS_SYNC_ENABLED](#)

GPIO number of SYNC IO

Default value:

- 3 if [CONFIG_BT_BLE_LOG_SPI_OUT_TS_SYNC_ENABLED](#) && [CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED](#)

CONFIG_BT_BLE_LOG_SPI_OUT_FLUSH_TIMER_ENABLED

Enable periodic buffer flush out

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Legacy SPI Log Output \(Deprecated - use BT Log Async Output instead\)](#) > [CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED](#)

Enable periodic buffer flush out Not recommended when SPI receiver is unavailable

Default value:

- No (disabled) if `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_FLUSH_TIMEOUT

Buffer flush out period in unit of ms

Found in: Component config > Bluetooth > Common Options > BT Logs > Legacy SPI Log Output (Deprecated - use BT Log Async Output instead) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED > CONFIG_BT_BLE_LOG_SPI_OUT_FLUSH_TIMER_ENABLED

Buffer flush out period in unit of ms

Default value:

- 1000 if `CONFIG_BT_BLE_LOG_SPI_OUT_FLUSH_TIMER_ENABLED` && `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_LE_AUDIO_ENABLED

Enable LE Audio log output to SPI

Found in: Component config > Bluetooth > Common Options > BT Logs > Legacy SPI Log Output (Deprecated - use BT Log Async Output instead) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED

Enable LE Audio log output to SPI

Default value:

- No (disabled) if `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_LE_AUDIO_BUF_SIZE

SPI transaction buffer size for LE Audio logs

Found in: Component config > Bluetooth > Common Options > BT Logs > Legacy SPI Log Output (Deprecated - use BT Log Async Output instead) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED > CONFIG_BT_BLE_LOG_SPI_OUT_LE_AUDIO_ENABLED

SPI transaction buffer size for LE Audio logs. There will be 2 SPI DMA buffers with the same size.

Default value:

- 1024 if `CONFIG_BT_BLE_LOG_SPI_OUT_LE_AUDIO_ENABLED` && `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_LE_AUDIO_TASK_CNT

LE audio task count

Found in: Component config > Bluetooth > Common Options > BT Logs > Legacy SPI Log Output (Deprecated - use BT Log Async Output instead) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED > CONFIG_BT_BLE_LOG_SPI_OUT_LE_AUDIO_ENABLED

LE audio task count

Default value:

- 1 if `CONFIG_BT_BLE_LOG_SPI_OUT_LE_AUDIO_ENABLED` && `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_MESH_ENABLED

Enable BLE mesh log output to SPI

Found in: Component config > Bluetooth > Common Options > BT Logs > Legacy SPI Log Output (Deprecated - use BT Log Async Output instead) > CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED

Enable BLE mesh log output to SPI

Default value:

- No (disabled) if `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_MESH_BUF_SIZE

SPI transaction buffer size for BLE mesh logs

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Legacy SPI Log Output \(Deprecated - use BT Log Async Output instead\)](#) > [CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED](#) > [CONFIG_BT_BLE_LOG_SPI_OUT_MESH_ENABLED](#)

SPI transaction buffer size for BLE mesh logs. There will be 2 SPI DMA buffers with the same size.

Default value:

- 1024 if `CONFIG_BT_BLE_LOG_SPI_OUT_MESH_ENABLED` && `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_BLE_LOG_SPI_OUT_MESH_TASK_CNT

Mesh task count

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#) > [BT Logs](#) > [Legacy SPI Log Output \(Deprecated - use BT Log Async Output instead\)](#) > [CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED](#) > [CONFIG_BT_BLE_LOG_SPI_OUT_MESH_ENABLED](#)

Mesh task count

Default value:

- 3 if `CONFIG_BT_BLE_LOG_SPI_OUT_MESH_ENABLED` && `CONFIG_BT_BLE_LOG_SPI_OUT_ENABLED`

CONFIG_BT_LE_USED_MEM_STATISTICS_ENABLED

Enable used memory statistics

Found in: [Component config](#) > [Bluetooth](#) > [Common Options](#)

Used in internal tests only. Enable used memory statistics.

Default value:

- No (disabled)

CONFIG_BT_HCI_LOG_DEBUG_EN

Enable Bluetooth HCI debug mode

Found in: [Component config](#) > [Bluetooth](#)

This option is used to enable bluetooth debug mode, which saves the hci layer data stream.

Default value:

- No (disabled) if `CONFIG_BT_BLUEDROID_ENABLED` || `CONFIG_BT_NIMBLE_ENABLED`

CONFIG_BT_HCI_LOG_DATA_BUFFER_SIZE

Size of the cache used for HCI data in Bluetooth HCI debug mode (N*1024 bytes)

Found in: [Component config](#) > [Bluetooth](#) > [CONFIG_BT_HCI_LOG_DEBUG_EN](#)

This option is to configure the buffer size of the hci data steam cache in hci debug mode. This is a ring buffer, the new data will overwrite the oldest data if the buffer is full.

Range:

- from 1 to 100 if `CONFIG_BT_HCI_LOG_DEBUG_EN`

Default value:

- 5 if `CONFIG_BT_HCI_LOG_DEBUG_EN`

CONFIG_BT_HCI_LOG_ADV_BUFFER_SIZE

Size of the cache used for adv report in Bluetooth HCI debug mode (N*1024 bytes)

Found in: Component config > Bluetooth > CONFIG_BT_HCI_LOG_DEBUG_EN

This option is to configure the buffer size of the hci adv report cache in hci debug mode. This is a ring buffer, the new data will overwrite the oldest data if the buffer is full.

Range:

- from 1 to 100 if `CONFIG_BT_HCI_LOG_DEBUG_EN`

Default value:

- 8 if `CONFIG_BT_HCI_LOG_DEBUG_EN`

CONFIG_BLE_MESH

ESP BLE Mesh Support

Found in: Component config

This option enables ESP BLE Mesh support. The specific features that are available may depend on other features that have been enabled in the stack, such as Bluetooth Support, Bluedroid Support & GATT support.

Contains:

- *BLE Mesh and BLE coexistence support*
- `CONFIG_BLE_MESH_GATT_PROXY_CLIENT`
- `CONFIG_BLE_MESH_GATT_PROXY_SERVER`
- *BLE Mesh NET BUF DEBUG LOG LEVEL*
- `CONFIG_BLE_MESH_PROV`
- `CONFIG_BLE_MESH_PROXY`
- *BLE Mesh specific test option*
- *BLE Mesh STACK DEBUG LOG LEVEL*
- `CONFIG_BLE_MESH_NO_LOG`
- `CONFIG_BLE_MESH_IVU_DIVIDER`
- `CONFIG_BLE_MESH_FAST_PROV`
- `CONFIG_BLE_MESH_FREERTOS_STATIC_ALLOC`
- `CONFIG_BLE_MESH_EXPERIMENTAL`
- `CONFIG_BLE_MESH_CRPL`
- `CONFIG_BLE_MESH_RX_SDU_MAX`
- `CONFIG_BLE_MESH_MODEL_KEY_COUNT`
- `CONFIG_BLE_MESH_APP_KEY_COUNT`
- `CONFIG_BLE_MESH_MODEL_GROUP_COUNT`
- `CONFIG_BLE_MESH_LABEL_COUNT`
- `CONFIG_BLE_MESH_SUBNET_COUNT`
- `CONFIG_BLE_MESH_TX_SEG_MAX`
- `CONFIG_BLE_MESH_RX_SEG_MSG_COUNT`
- `CONFIG_BLE_MESH_TX_SEG_MSG_COUNT`
- `CONFIG_BLE_MESH_MEM_ALLOC_MODE`
- `CONFIG_BLE_MESH_MSG_CACHE_SIZE`
- `CONFIG_BLE_MESH_NOT_RELAY_REPLAY_MSG`
- `CONFIG_BLE_MESH_ADV_BUF_COUNT`
- `CONFIG_BLE_MESH_PB_GATT`
- `CONFIG_BLE_MESH_PB_ADV`
- `CONFIG_BLE_MESH_IVU_RECOVERY_IVI`
- `CONFIG_BLE_MESH_RELAY`
- `CONFIG_BLE_MESH_SAR_ENHANCEMENT`
- `CONFIG_BLE_MESH_SETTINGS`

- [CONFIG_BLE_MESH_ACTIVE_SCAN](#)
- [CONFIG_BLE_MESH_DEINIT](#)
- [CONFIG_BLE_MESH_USE_DUPLICATE_SCAN](#)
- [CONFIG_BLE_MESH_V11_SUPPORT](#)
- [Support for BLE Mesh Client/Server models](#)
- [Support for BLE Mesh Foundation models](#)
- [CONFIG_BLE_MESH_NODE](#)
- [CONFIG_BLE_MESH_PROVISIONER](#)
- [CONFIG_BLE_MESH_FRIEND](#)
- [CONFIG_BLE_MESH_LOW_POWER](#)
- [CONFIG_BLE_MESH_PROXY_CLI_SRV_COEXIST](#)
- [CONFIG_BLE_MESH_HCI_5_0](#)
- [CONFIG_BLE_MESH_USE_BLE_50](#)
- [CONFIG_BLE_MESH_RANDOM_ADV_INTERVAL](#)
- [CONFIG_BLE_MESH_IV_UPDATE_TEST](#)
- [CONFIG_BLE_MESH_CLIENT_MSG_TIMEOUT](#)
- [CONFIG_BLE_MESH_USE_UNIFIED_CRYPTO](#)

CONFIG_BLE_MESH_HCI_5_0

Support sending 20ms non-connectable adv packets

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

It is a temporary solution and needs further modifications.

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_V11_SUPPORT

Support ESP BLE Mesh v1.1 features (Preview)

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Support BLE Mesh v1.1 features

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_RANDOM_ADV_INTERVAL

Support using random adv interval for mesh packets

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Enable this option to allow using random advertising interval for mesh packets. And this could help avoid collision of advertising packets.

Default value:

- No (disabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_USE_UNIFIED_CRYPTO

Use the unified BLE tinycrypt implementation(Deprecated)

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

This configuration path is deprecated and will be removed in a future version. Please use the corresponding Kconfig options under \$IDF_PATH/components/bt/common/Kconfig.

Default value:

- No (disabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_USE_BLE_50

Support using BLE 5.0 APIs for BLE Mesh

Found in: *Component config* > *CONFIG_BLE_MESH*

This option to enable BLE Mesh using some BLE 5.0 APIs.

Default value:

- No (disabled) if *CONFIG_BLE_MESH_EXPERIMENTAL* && *CONFIG_BLE_MESH*

Contains:

- *CONFIG_BLE_MESH_EXT_ADV*
- *CONFIG_BLE_MESH_ADV_INST_ID*
- *CONFIG_BLE_MESH_SUPPORT_MULTI_ADV*

CONFIG_BLE_MESH_EXT_ADV

Enable extended advertising for BLE Mesh

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_USE_BLE_50*

Enable broadcasting of BLE Mesh messages using BLE 5.0 extended advertising. This allows control of extended advertising parameters (e.g., PHY selection) while maintaining standard BLE Mesh packet length.

Default value:

- No (disabled) if *CONFIG_BLE_MESH_USE_BLE_50* && *CONFIG_BLE_MESH*

Contains:

- *CONFIG_BLE_MESH_LONG_PACKET*
- *CONFIG_BLE_MESH_EXT_ADV_BUF_COUNT*
- *CONFIG_BLE_MESH_EXT_RELAY_ADV_BUF_COUNT*

CONFIG_BLE_MESH_EXT_ADV_BUF_COUNT

Number of extended advertising buffers

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_USE_BLE_50* > *CONFIG_BLE_MESH_EXT_ADV*

Number of buffer slots allocated for extended advertising packets.

Range:

- from 6 to 256 if *CONFIG_BLE_MESH_EXT_ADV* && *CONFIG_BLE_MESH*

Default value:

- 60 if *CONFIG_BLE_MESH_EXT_ADV* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_EXT_RELAY_ADV_BUF_COUNT

Number of extended relay advertising buffers

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_USE_BLE_50* > *CONFIG_BLE_MESH_EXT_ADV*

Number of buffer slots allocated for extended advertising packets used in message relay.

Range:

- from 0 to 256 if *CONFIG_BLE_MESH_EXT_ADV* && *CONFIG_BLE_MESH_RELAY* && *CONFIG_BLE_MESH*

Default value:

- 60 if *CONFIG_BLE_MESH_EXT_ADV* && *CONFIG_BLE_MESH_RELAY* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_LONG_PACKET

Enable non-standard long packet mode for BLE Mesh

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_USE_BLE_50](#) > [CONFIG_BLE_MESH_EXT_ADV](#)

Enable extended-length advertising packets for BLE Mesh using BLE 5.0 extended advertising. This overrides the standard BLE Mesh packet length limitations.

Default value:

- No (disabled) if [CONFIG_BLE_MESH_EXT_ADV](#) && [CONFIG_BLE_MESH](#)

Contains:

- [CONFIG_BLE_MESH_LONG_PACKET_ADV_BUF_COUNT](#)
- [CONFIG_BLE_MESH_LONG_PACKET_RELAY_ADV_BUF_COUNT](#)
- [CONFIG_BLE_MESH_LONG_PACKET_ADV_LEN](#)
- [CONFIG_BLE_MESH_LONG_PACKET_RX_SEG_CNT](#)
- [CONFIG_BLE_MESH_LONG_PACKET_TX_SEG_CNT](#)

CONFIG_BLE_MESH_LONG_PACKET_ADV_LEN

Maximum advertising payload length

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_USE_BLE_50](#) > [CONFIG_BLE_MESH_EXT_ADV](#) > [CONFIG_BLE_MESH_LONG_PACKET](#)

Maximum payload length for extended advertising packets (bytes). Range: 30-249 bytes. Default: 105. Note: Maximum access payload = (LENGTH - 17) × SEGMENT_COUNT - Transmission: Uses [BLE_MESH_LONG_PACKET_TX_SEG_CNT](#) - Reception: Uses [BLE_MESH_LONG_PACKET_RX_SEG_CNT](#)

Range:

- from 30 to 249 if [CONFIG_BLE_MESH_LONG_PACKET](#) && [CONFIG_BLE_MESH](#)

Default value:

- 105 if [CONFIG_BLE_MESH_LONG_PACKET](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_LONG_PACKET_ADV_BUF_COUNT

Long packet advertising buffer count

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_USE_BLE_50](#) > [CONFIG_BLE_MESH_EXT_ADV](#) > [CONFIG_BLE_MESH_LONG_PACKET](#)

Number of advertising buffers allocated for long packet transmissions.

Default value:

- 20 if [CONFIG_BLE_MESH_LONG_PACKET](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_LONG_PACKET_RELAY_ADV_BUF_COUNT

Long packet relay buffer count

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_USE_BLE_50](#) > [CONFIG_BLE_MESH_EXT_ADV](#) > [CONFIG_BLE_MESH_LONG_PACKET](#)

Number of advertising buffers allocated for relay long packets.

Default value:

- 20 if [CONFIG_BLE_MESH_LONG_PACKET](#) && [CONFIG_BLE_MESH_RELAY](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_LONG_PACKET_TX_SEG_CNT

Maximum transmission segments per message

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_USE_BLE_50* > *CONFIG_BLE_MESH_EXT_ADV* > *CONFIG_BLE_MESH_LONG_PACKET*

Maximum number of segments for outgoing long packet messages. Upper bound: min(BLE_MESH_TX_SEG_MAX, BLE_MESH_LONG_PACKET_ADV_BUF_COUNT)

Range:

- from 1 to if *CONFIG_BLE_MESH_TX_SEG_MAX* < *CONFIG_BLE_MESH_LONG_PACKET_ADV_BUF_COUNT* && *CONFIG_BLE_MESH_LONG_PACKET* && *CONFIG_BLE_MESH*
- from 1 to if *CONFIG_BLE_MESH_TX_SEG_MAX* >= *CONFIG_BLE_MESH_LONG_PACKET_ADV_BUF_COUNT* && *CONFIG_BLE_MESH_LONG_PACKET* && *CONFIG_BLE_MESH*

Default value:

- 20 if *CONFIG_BLE_MESH_LONG_PACKET* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_LONG_PACKET_RX_SEG_CNT

Maximum reception segments per message

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_USE_BLE_50* > *CONFIG_BLE_MESH_EXT_ADV* > *CONFIG_BLE_MESH_LONG_PACKET*

Maximum number of segments supported for receiving long packet messages.

Range:

- from 1 to if *CONFIG_BLE_MESH_LONG_PACKET* && *CONFIG_BLE_MESH*

Default value:

- 20 if *CONFIG_BLE_MESH_LONG_PACKET* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_ADV_INST_ID

Extended adv instance for Mesh normal packets

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_USE_BLE_50*

Extended ADV instance used by Mesh normal advertising packets.

Range:

- from 0 to 3 if *CONFIG_BLE_MESH_USE_BLE_50* && *CONFIG_BLE_MESH*

Default value:

- 0 if *CONFIG_BLE_MESH_USE_BLE_50* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_SUPPORT_MULTI_ADV

Support using multiple adv instance for BLE Mesh

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_USE_BLE_50*

Enable this option to support using multiple adv instance while running BLE Mesh.

Default value:

- No (disabled) if *CONFIG_BLE_MESH_USE_BLE_50* && *CONFIG_BLE_MESH*

Contains:

- *CONFIG_BLE_MESH_PROXY_ADV_INST_ID*
- *CONFIG_BLE_MESH_SEPARATE_BLE_ADV_INSTANCE*
- *CONFIG_BLE_MESH_SEPARATE_RELAY_ADV_INSTANCE*

CONFIG_BLE_MESH_PROXY_ADV_INST_ID

Extended adv instance for Mesh proxy packets

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_USE_BLE_50* > *CONFIG_BLE_MESH_SUPPORT_MULTI_ADV*

Extended ADV instance used by Mesh proxy advertising packets.

Range:

- from 0 to 3 if *CONFIG_BLE_MESH_PROXY* && (*CONFIG_BLE_MESH_PB_GATT* || *CONFIG_BLE_MESH_GATT_PROXY_SERVER*) && *CONFIG_BLE_MESH_SUPPORT_MULTI_ADV* && *CONFIG_BLE_MESH*

Default value:

- 1 if *CONFIG_BLE_MESH_PROXY* && (*CONFIG_BLE_MESH_PB_GATT* || *CONFIG_BLE_MESH_GATT_PROXY_SERVER*) && *CONFIG_BLE_MESH_SUPPORT_MULTI_ADV* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_SEPARATE_RELAY_ADV_INSTANCE

Use separate extended adv instance for Mesh relay packets

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_USE_BLE_50* > *CONFIG_BLE_MESH_SUPPORT_MULTI_ADV*

Enable this option to support using a separate extended ADV instance for Mesh relay packets.

Default value:

- No (disabled) if *CONFIG_BLE_MESH_SUPPORT_MULTI_ADV* && *CONFIG_BLE_MESH_RELAY_ADV_BUF* && *CONFIG_BLE_MESH*

Contains:

- *CONFIG_BLE_MESH_RELAY_ADV_INST_ID*

CONFIG_BLE_MESH_RELAY_ADV_INST_ID

Extended adv instance for Mesh relay packets

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_USE_BLE_50* > *CONFIG_BLE_MESH_SUPPORT_MULTI_ADV* > *CONFIG_BLE_MESH_SEPARATE_RELAY_ADV_INSTANCE*

Extended ADV instance used by Mesh relay advertising packets.

Range:

- from 0 to 3 if *CONFIG_BLE_MESH_SEPARATE_RELAY_ADV_INSTANCE* && *CONFIG_BLE_MESH*

Default value:

- 2 if *CONFIG_BLE_MESH_SEPARATE_RELAY_ADV_INSTANCE* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_SEPARATE_BLE_ADV_INSTANCE

Use separate extended adv instance for BLE normal packets

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_USE_BLE_50* > *CONFIG_BLE_MESH_SUPPORT_MULTI_ADV*

Enable this option to support using a separate extended ADV instance for normal BLE advertising packets.

Default value:

- No (disabled) if *CONFIG_BLE_MESH_SUPPORT_MULTI_ADV* && *CONFIG_BLE_MESH_SUPPORT_BLE_ADV* && *CONFIG_BLE_MESH*

Contains:

- [CONFIG_BLE_MESH_BLE_ADV_INST_ID](#)

CONFIG_BLE_MESH_BLE_ADV_INST_ID

Extended adv instance for normal BLE packets

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_USE_BLE_50](#) > [CONFIG_BLE_MESH_SUPPORT_MULTI_ADV](#) > [CONFIG_BLE_MESH_SEPARATE_BLE_ADV_INSTANCE](#)

Extended ADV instance used by normal BLE advertising packets.

Range:

- from 0 to 3 if [CONFIG_BLE_MESH_SEPARATE_BLE_ADV_INSTANCE](#) && [CONFIG_BLE_MESH](#)

Default value:

- 3 if [CONFIG_BLE_MESH_SEPARATE_BLE_ADV_INSTANCE](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_USE_DUPLICATE_SCAN

Support Duplicate Scan in BLE Mesh

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Enable this option to allow using specific duplicate scan filter in BLE Mesh, and Scan Duplicate Type must be set by choosing the option in the Bluetooth Controller section in menuconfig, which is "Scan Duplicate By Device Address and Advertising Data".

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_ACTIVE_SCAN

Support Active Scan in BLE Mesh

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Enable this option to allow using BLE Active Scan for BLE Mesh.

CONFIG_BLE_MESH_MEM_ALLOC_MODE

Memory allocation strategy

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Allocation strategy for BLE Mesh stack, essentially provides ability to allocate all required dynamic allocations from,

- Internal DRAM memory only
- External SPIRAM memory only
- Either internal or external memory based on default malloc() behavior in ESP-IDF
- Internal IRAM memory wherever applicable else internal DRAM

Recommended mode here is always internal (*), since that is most preferred from security perspective. But if application requirement does not allow sufficient free internal memory then alternate mode can be selected.

(*) In case of ESP32-S2/ESP32-S3, hardware allows encryption of external SPIRAM contents provided hardware flash encryption feature is enabled. In that case, using external SPIRAM allocation strategy is also safe choice from security perspective.

Available options:

- Internal DRAM (CONFIG_BLE_MESH_MEM_ALLOC_MODE_INTERNAL)
- External SPIRAM (CONFIG_BLE_MESH_MEM_ALLOC_MODE_EXTERNAL)
- Default alloc mode (CONFIG_BLE_MESH_MEM_ALLOC_MODE_DEFAULT)
Enable this option to use the default memory allocation strategy when external SPIRAM is enabled. See the SPIRAM options for more details.
- Internal IRAM (CONFIG_BLE_MESH_MEM_ALLOC_MODE_IRAM_8BIT)
Allows to use IRAM memory region as 8bit accessible region. Every unaligned (8bit or 16bit) access will result in an exception and incur penalty of certain clock cycles per unaligned read/write.

CONFIG_BLE_MESH_FREERTOS_STATIC_ALLOC

Enable FreeRTOS static allocation

Found in: *Component config* > *CONFIG_BLE_MESH*

Enable this option to use FreeRTOS static allocation APIs for BLE Mesh, which provides the ability to use different dynamic memory (i.e. SPIRAM or IRAM) for FreeRTOS objects. If this option is disabled, the FreeRTOS static allocation APIs will not be used, and internal DRAM will be allocated for FreeRTOS objects.

Default value:

- No (disabled) if ESP32_IRAM_AS_8BIT_ACCESSIBLE_MEMORY && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_FREERTOS_STATIC_ALLOC_MODE

Memory allocation for FreeRTOS objects

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_FREERTOS_STATIC_ALLOC*

Choose the memory to be used for FreeRTOS objects.

Available options:

- External SPIRAM (CONFIG_BLE_MESH_FREERTOS_STATIC_ALLOC_EXTERNAL)
If enabled, BLE Mesh allocates dynamic memory from external SPIRAM for FreeRTOS objects, i.e. mutex, queue, and task stack. External SPIRAM can only be used for task stack when SPIRAM_ALLOW_STACK_EXTERNAL_MEMORY is enabled. See the SPIRAM options for more details.
- Internal IRAM (CONFIG_BLE_MESH_FREERTOS_STATIC_ALLOC_IRAM_8BIT)
If enabled, BLE Mesh allocates dynamic memory from internal IRAM for FreeRTOS objects, i.e. mutex, queue. Note: IRAM region cannot be used as task stack.

CONFIG_BLE_MESH_DEINIT

Support de-initialize BLE Mesh stack

Found in: *Component config* > *CONFIG_BLE_MESH*

If enabled, users can use the function `esp_ble_mesh_deinit()` to de-initialize the whole BLE Mesh stack.

Default value:

- Yes (enabled) if *CONFIG_BLE_MESH*

BLE Mesh and BLE coexistence support Contains:

- [CONFIG_BLE_MESH_SUPPORT_BLE_SCAN](#)
- [CONFIG_BLE_MESH_SUPPORT_BLE_ADV](#)

CONFIG_BLE_MESH_SUPPORT_BLE_ADV

Support sending normal BLE advertising packets

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh and BLE coexistence support](#)

When selected, users can send normal BLE advertising packets with specific API.

Default value:

- No (disabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_BLE_ADV_BUF_COUNT

Number of advertising buffers for BLE advertising packets

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh and BLE coexistence support](#) > [CONFIG_BLE_MESH_SUPPORT_BLE_ADV](#)

Number of advertising buffers for BLE packets available.

Range:

- from 1 to 255 if [CONFIG_BLE_MESH_SUPPORT_BLE_ADV](#) && [CONFIG_BLE_MESH](#)

Default value:

- 3 if [CONFIG_BLE_MESH_SUPPORT_BLE_ADV](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_SUPPORT_BLE_SCAN

Support scanning normal BLE advertising packets

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh and BLE coexistence support](#)

When selected, users can register a callback and receive normal BLE advertising packets in the application layer.

Default value:

- No (disabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_FAST_PROV

Enable BLE Mesh Fast Provisioning

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Enable this option to allow BLE Mesh fast provisioning solution to be used. When there are multiple unprovisioned devices around, fast provisioning can greatly reduce the time consumption of the whole provisioning process. When this option is enabled, and after an unprovisioned device is provisioned into a node successfully, it can be changed to a temporary Provisioner.

Default value:

- No (disabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_NODE

Support for BLE Mesh Node

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Enable the device to be provisioned into a node. This option should be enabled when an unprovisioned device is going to be provisioned into a node and communicate with other nodes in the BLE Mesh network.

CONFIG_BLE_MESH_PROVISIONER

Support for BLE Mesh Provisioner

Found in: *Component config* > *CONFIG_BLE_MESH*

Enable the device to be a Provisioner. The option should be enabled when a device is going to act as a Provisioner and provision unprovisioned devices into the BLE Mesh network.

CONFIG_BLE_MESH_WAIT_FOR_PROV_MAX_DEV_NUM

Maximum number of unprovisioned devices that can be added to device queue

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_PROVISIONER*

This option specifies how many unprovisioned devices can be added to device queue for provisioning. Users can use this option to define the size of the queue in the bottom layer which is used to store unprovisioned device information (e.g. Device UUID, address).

Range:

- from 1 to 100 if *CONFIG_BLE_MESH_PROVISIONER* && *CONFIG_BLE_MESH*

Default value:

- 10 if *CONFIG_BLE_MESH_PROVISIONER* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_MAX_PROV_NODES

Maximum number of devices that can be provisioned by Provisioner

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_PROVISIONER*

This option specifies how many devices can be provisioned by a Provisioner. This value indicates the maximum number of unprovisioned devices which can be provisioned by a Provisioner. For instance, if the value is 6, it means the Provisioner can provision up to 6 unprovisioned devices. Theoretically a Provisioner without the limitation of its memory can provision up to 32766 unprovisioned devices, here we limit the maximum number to 100 just to limit the memory used by a Provisioner. The bigger the value is, the more memory it will cost by a Provisioner to store the information of nodes.

Range:

- from 1 to 1000 if *CONFIG_BLE_MESH_PROVISIONER* && *CONFIG_BLE_MESH*

Default value:

- 10 if *CONFIG_BLE_MESH_PROVISIONER* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_PBA_SAME_TIME

Maximum number of PB-ADV running at the same time by Provisioner

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_PROVISIONER*

This option specifies how many devices can be provisioned at the same time using PB-ADV. For examples, if the value is 2, it means a Provisioner can provision two unprovisioned devices with PB-ADV at the same time.

Range:

- from 1 to 10 if *CONFIG_BLE_MESH_PB_ADV* && *CONFIG_BLE_MESH_PROVISIONER* && *CONFIG_BLE_MESH*

Default value:

- 2 if *CONFIG_BLE_MESH_PB_ADV* && *CONFIG_BLE_MESH_PROVISIONER* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_PBG_SAME_TIME

Maximum number of PB-GATT running at the same time by Provisioner

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_PROVISIONER*

This option specifies how many devices can be provisioned at the same time using PB-GATT. For example, if the value is 2, it means a Provisioner can provision two unprovisioned devices with PB-GATT at the same time.

Range:

- from 1 to 5 if *CONFIG_BLE_MESH_PB_GATT* && *CONFIG_BLE_MESH_PROVISIONER* && *CONFIG_BLE_MESH*

Default value:

- 1 if *CONFIG_BLE_MESH_PB_GATT* && *CONFIG_BLE_MESH_PROVISIONER* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_PROVISIONER_SUBNET_COUNT

Maximum number of mesh subnets that can be created by Provisioner

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_PROVISIONER*

This option specifies how many subnets per network a Provisioner can create. Indeed, this value decides the number of network keys which can be added by a Provisioner.

Range:

- from 1 to 4096 if *CONFIG_BLE_MESH_PROVISIONER* && *CONFIG_BLE_MESH*

Default value:

- 3 if *CONFIG_BLE_MESH_PROVISIONER* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_PROVISIONER_APP_KEY_COUNT

Maximum number of application keys that can be owned by Provisioner

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_PROVISIONER*

This option specifies how many application keys the Provisioner can have. Indeed, this value decides the number of the application keys which can be added by a Provisioner.

Range:

- from 1 to 4096 if *CONFIG_BLE_MESH_PROVISIONER* && *CONFIG_BLE_MESH*

Default value:

- 3 if *CONFIG_BLE_MESH_PROVISIONER* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_PROVISIONER_RECV_HB

Support receiving Heartbeat messages

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_PROVISIONER*

When this option is enabled, Provisioner can call specific functions to enable or disable receiving Heartbeat messages and notify them to the application layer.

Default value:

- No (disabled) if *CONFIG_BLE_MESH_PROVISIONER* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_PROVISIONER_RECV_HB_FILTER_SIZE

Maximum number of filter entries for receiving Heartbeat messages

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_PROVISIONER* > *CONFIG_BLE_MESH_PROVISIONER_RECV_HB*

This option specifies how many heartbeat filter entries Provisioner supports. The heartbeat filter (acceptlist or rejectlist) entries are used to store a list of SRC and DST which can be used to decide if a heartbeat message will be processed and notified to the application layer by Provisioner. Note: The filter is an empty rejectlist by default.

Range:

- from 1 to 1000 if `CONFIG_BLE_MESH_PROVISIONER_RECV_HB` && `CONFIG_BLE_MESH_PROVISIONER` && `CONFIG_BLE_MESH`

Default value:

- 3 if `CONFIG_BLE_MESH_PROVISIONER_RECV_HB` && `CONFIG_BLE_MESH_PROVISIONER` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_PROV

BLE Mesh Provisioning support

Found in: `Component config` > `CONFIG_BLE_MESH`

Enable this option to support BLE Mesh Provisioning functionality. For BLE Mesh, this option should be always enabled.

Default value:

- Yes (enabled) if `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_PROV_EPA

BLE Mesh enhanced provisioning authentication

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_PROV`

Enable this option to support BLE Mesh enhanced provisioning authentication functionality. This option can increase the security level of provisioning. It is recommended to enable this option.

Default value:

- Yes (enabled) if `CONFIG_BLE_MESH_PROV` && `CONFIG_BLE_MESH_V11_SUPPORT` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_CERT_BASED_PROV

Support Certificate-based provisioning

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_PROV`

Enable this option to support BLE Mesh Certificate-Based Provisioning.

Default value:

- No (disabled) if `CONFIG_BLE_MESH_PROV` && `CONFIG_BLE_MESH_V11_SUPPORT` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_RECORD_FRAG_MAX_SIZE

Maximum size of the provisioning record fragment that Provisioner can receive

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_PROV` > `CONFIG_BLE_MESH_CERT_BASED_PROV`

This option sets the maximum size of the provisioning record fragment that the Provisioner can receive. The range depends on provisioning bearer.

Range:

- from 1 to 57 if `CONFIG_BLE_MESH_CERT_BASED_PROV` && `CONFIG_BLE_MESH`

Default value:

- 56 if `CONFIG_BLE_MESH_CERT_BASED_PROV` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_PB_ADV

Provisioning support using the advertising bearer (PB-ADV)

Found in: *Component config* > *CONFIG_BLE_MESH*

Enable this option to allow the device to be provisioned over the advertising bearer. This option should be enabled if PB-ADV is going to be used during provisioning procedure.

Default value:

- Yes (enabled) if *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_UNPROVISIONED_BEACON_INTERVAL

Interval between two consecutive Unprovisioned Device Beacon

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_PB_ADV*

This option specifies the interval of sending two consecutive unprovisioned device beacon, users can use this option to change the frequency of sending unprovisioned device beacon. For example, if the value is 5, it means the unprovisioned device beacon will send every 5 seconds. When the option of BLE_MESH_FAST_PROV is selected, the value is better to be 3 seconds, or less.

Range:

- from 1 to 100 if *CONFIG_BLE_MESH_NODE* && *CONFIG_BLE_MESH_PB_ADV* && *CONFIG_BLE_MESH*

Default value:

- 5 if *CONFIG_BLE_MESH_NODE* && *CONFIG_BLE_MESH_PB_ADV* && *CONFIG_BLE_MESH*
- 3 if *CONFIG_BLE_MESH_FAST_PROV* && *CONFIG_BLE_MESH_NODE* && *CONFIG_BLE_MESH_PB_ADV* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_PB_GATT

Provisioning support using GATT (PB-GATT)

Found in: *Component config* > *CONFIG_BLE_MESH*

Enable this option to allow the device to be provisioned over GATT. This option should be enabled if PB-GATT is going to be used during provisioning procedure.

Virtual option enabled whenever any Proxy protocol is needed

CONFIG_BLE_MESH_PROXY

BLE Mesh Proxy protocol support

Found in: *Component config* > *CONFIG_BLE_MESH*

Enable this option to support BLE Mesh Proxy protocol used by PB-GATT and other proxy pdu transmission.

Default value:

- Yes (enabled) if *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_GATT_PROXY_SERVER

BLE Mesh GATT Proxy Server

Found in: *Component config* > *CONFIG_BLE_MESH*

This option enables support for Mesh GATT Proxy Service, i.e. the ability to act as a proxy between a Mesh GATT Client and a Mesh network. This option should be enabled if a node is going to be a Proxy Server.

Default value:

- Yes (enabled) if `CONFIG_BLE_MESH_NODE` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_NODE_ID_TIMEOUT

Node Identity advertising timeout

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_GATT_PROXY_SERVER`

This option determines for how long the local node advertises using Node Identity. The given value is in seconds. The specification limits this to 60 seconds and lists it as the recommended value as well. So leaving the default value is the safest option. When an unprovisioned device is provisioned successfully and becomes a node, it will start to advertise using Node Identity during the time set by this option. And after that, Network ID will be advertised.

Range:

- from 1 to 60 if `CONFIG_BLE_MESH_GATT_PROXY_SERVER` && `CONFIG_BLE_MESH`

Default value:

- 60 if `CONFIG_BLE_MESH_GATT_PROXY_SERVER` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_PROXY_FILTER_SIZE

Maximum number of filter entries per Proxy Client

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_GATT_PROXY_SERVER`

This option specifies how many Proxy Filter entries the local node supports. The entries of Proxy filter (whitelist or blacklist) are used to store a list of addresses which can be used to decide which messages will be forwarded to the Proxy Client by the Proxy Server.

Range:

- from 1 to 32767 if `CONFIG_BLE_MESH_GATT_PROXY_SERVER` && `CONFIG_BLE_MESH`

Default value:

- 4 if `CONFIG_BLE_MESH_GATT_PROXY_SERVER` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_PROXY_PRIVACY

Support Proxy Privacy

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_GATT_PROXY_SERVER`

The Proxy Privacy parameter controls the privacy of the Proxy Server over the connection. The value of the Proxy Privacy parameter is controlled by the type of proxy connection, which is dependent on the bearer used by the proxy connection.

Default value:

- Yes (enabled) if `CONFIG_BLE_MESH_PRB_SRV` && `CONFIG_BLE_MESH_GATT_PROXY_SERVER` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_PROXY_SOLIC_PDU_RX

Support receiving Proxy Solicitation PDU

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_GATT_PROXY_SERVER`

Enable this option to support receiving Proxy Solicitation PDU.

CONFIG_BLE_MESH_PROXY_SOLIC_RX_CRPL

Maximum capacity of solicitation replay protection list

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_GATT_PROXY_SERVER* > *CONFIG_BLE_MESH_PROXY_SOLIC_PDU_RX*

This option specifies the maximum capacity of the solicitation replay protection list. The solicitation replay protection list is used to reject Solicitation PDUs that were already processed by a node, which will store the solicitation src and solicitation sequence number of the received Solicitation PDU message.

Range:

- from 1 to 65536 if *CONFIG_BLE_MESH_PROXY_SOLIC_PDU_RX* && *CONFIG_BLE_MESH*

Default value:

- 2 if *CONFIG_BLE_MESH_PROXY_SOLIC_PDU_RX* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_PROXY_CLI_SRV_COEXIST

Support Proxy Client and Proxy Server coexistence

Found in: *Component config* > *CONFIG_BLE_MESH*

Enable this option to support the coexistence of proxy client and proxy server.

Default value:

- No (disabled) if *CONFIG_BLE_MESH_EXPERIMENTAL* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_GATT_PROXY_CLIENT

BLE Mesh GATT Proxy Client

Found in: *Component config* > *CONFIG_BLE_MESH*

This option enables support for Mesh GATT Proxy Client. The Proxy Client can use the GATT bearer to send mesh messages to a node that supports the advertising bearer.

Default value:

- No (disabled) if *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_PROXY_SOLIC_PDU_TX

Support sending Proxy Solicitation PDU

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_GATT_PROXY_CLIENT*

Enable this option to support sending Proxy Solicitation PDU.

CONFIG_BLE_MESH_PROXY_SOLIC_TX_SRC_COUNT

Maximum number of SSRC that can be used by Proxy Client

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_GATT_PROXY_CLIENT* > *CONFIG_BLE_MESH_PROXY_SOLIC_PDU_TX*

This option specifies the maximum number of Solicitation Source (SSRC) that can be used by Proxy Client for sending a Solicitation PDU. A Proxy Client may use the primary address or any of the secondary addresses as the SSRC for a Solicitation PDU. So for a Proxy Client, it's better to choose the value based on its own element count.

Range:

- from 1 to 16 if *CONFIG_BLE_MESH_PROXY_SOLIC_PDU_TX* && *CONFIG_BLE_MESH*

Default value:

- 2 if *CONFIG_BLE_MESH_PROXY_SOLIC_PDU_TX* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_SETTINGS

Store BLE Mesh configuration persistently

Found in: *Component config* > *CONFIG_BLE_MESH*

When selected, the BLE Mesh stack will take care of storing/restoring the BLE Mesh configuration persistently in flash. If the device is a BLE Mesh node, when this option is enabled, the configuration of the device will be stored persistently, including unicast address, NetKey, AppKey, etc. And if the device is a BLE Mesh Provisioner, the information of the device will be stored persistently, including the information of provisioned nodes, NetKey, AppKey, etc.

Default value:

- No (disabled) if *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_STORE_TIMEOUT

Delay (in seconds) before storing anything persistently

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_SETTINGS*

This value defines in seconds how soon any pending changes are actually written into persistent storage (flash) after a change occurs. The option allows nodes to delay a certain period of time to save proper information to flash. The default value is 0, which means information will be stored immediately once there are updates.

Range:

- from 0 to 1000000 if *CONFIG_BLE_MESH_SETTINGS* && *CONFIG_BLE_MESH*

Default value:

- 0 if *CONFIG_BLE_MESH_SETTINGS* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_SEQ_STORE_RATE

How often the sequence number gets updated in storage

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_SETTINGS*

This value defines how often the local sequence number gets updated in persistent storage (i.e. flash). e.g. a value of 100 means that the sequence number will be stored to flash on every 100th increment. If the node sends messages very frequently a higher value makes more sense, whereas if the node sends infrequently a value as low as 0 (update storage for every increment) can make sense. When the stack gets initialized it will add sequence number to the last stored one, so that it starts off with a value that's guaranteed to be larger than the last one used before power off.

Range:

- from 0 to 1000000 if *CONFIG_BLE_MESH_SETTINGS* && *CONFIG_BLE_MESH*

Default value:

- 0 if *CONFIG_BLE_MESH_SETTINGS* && *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_RPL_STORE_TIMEOUT

Minimum frequency that the RPL gets updated in storage

Found in: *Component config* > *CONFIG_BLE_MESH* > *CONFIG_BLE_MESH_SETTINGS*

This value defines in seconds how soon the RPL (Replay Protection List) gets written to persistent storage after a change occurs. If the node receives messages frequently, then a large value is recommended. If the node receives messages rarely, then the value can be as low as 0 (which means the RPL is written into the storage immediately). Note that if the node operates in a security-sensitive case, and there is a risk of sudden power-off, then a value of 0 is strongly recommended. Otherwise, a power loss before RPL being written into the storage may introduce message replay attacks and system security will be in a vulnerable state.

Range:

- from 0 to 1000000 if `CONFIG_BLE_MESH_SETTINGS` && `CONFIG_BLE_MESH`

Default value:

- 0 if `CONFIG_BLE_MESH_SETTINGS` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_SETTINGS_BACKWARD_COMPATIBILITY

A specific option for settings backward compatibility

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_SETTINGS`

This option is created to solve the issue of failure in recovering node information after mesh stack updates. In the old version mesh stack, there is no key of "mesh/role" in nvs. In the new version mesh stack, key of "mesh/role" is added in nvs, recovering node information needs to check "mesh/role" key in nvs and implements selective recovery of mesh node information. Therefore, there may be failure in recovering node information during node restarting after OTA.

The new version mesh stack adds the option of "mesh/role" because we have added the support of storing Provisioner information, while the old version only supports storing node information.

If users are updating their nodes from old version to new version, we recommend enabling this option, so that system could set the flag in advance before recovering node information and make sure the node information recovering could work as expected.

Default value:

- No (disabled) if `CONFIG_BLE_MESH_NODE` && `CONFIG_BLE_MESH_SETTINGS` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_SPECIFIC_PARTITION

Use a specific NVS partition for BLE Mesh

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_SETTINGS`

When selected, the mesh stack will use a specified NVS partition instead of default NVS partition. Note that the specified partition must be registered with NVS using `nvs_flash_init_partition()` API, and the partition must exist in the csv file. When Provisioner needs to store a large amount of nodes' information in the flash (e.g. more than 20), this option is recommended to be enabled.

Default value:

- No (disabled) if `CONFIG_BLE_MESH_SETTINGS` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_PARTITION_NAME

Name of the NVS partition for BLE Mesh

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_SETTINGS` > `CONFIG_BLE_MESH_SPECIFIC_PARTITION`

This value defines the name of the specified NVS partition used by the mesh stack.

Default value:

- "ble_mesh" if `CONFIG_BLE_MESH_SPECIFIC_PARTITION` && `CONFIG_BLE_MESH_SETTINGS` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_USE_MULTIPLE_NAMESPACE

Support using multiple NVS namespaces by Provisioner

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_SETTINGS`

When selected, Provisioner can use different NVS namespaces to store different instances of mesh information. For example, if in the first room, Provisioner uses NetKey A, AppKey A and provisions three devices, these information will be treated as mesh information instance A. When the Provisioner

moves to the second room, it uses NetKey B, AppKey B and provisions two devices, then the information will be treated as mesh information instance B. Here instance A and instance B will be stored in different namespaces. With this option enabled, Provisioner needs to use specific functions to open the corresponding NVS namespace, restore the mesh information, release the mesh information or erase the mesh information.

Default value:

- No (disabled) if `CONFIG_BLE_MESH_PROVISIONER` && `CONFIG_BLE_MESH_SETTINGS` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_MAX_NVS_NAMESPACE

Maximum number of NVS namespaces

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_SETTINGS` > `CONFIG_BLE_MESH_USE_MULTIPLE_NAMESPACE`

This option specifies the maximum NVS namespaces supported by Provisioner.

Range:

- from 1 to 255 if `CONFIG_BLE_MESH_USE_MULTIPLE_NAMESPACE` && `CONFIG_BLE_MESH_SETTINGS` && `CONFIG_BLE_MESH`

Default value:

- 2 if `CONFIG_BLE_MESH_USE_MULTIPLE_NAMESPACE` && `CONFIG_BLE_MESH_SETTINGS` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_SUBNET_COUNT

Maximum number of mesh subnets per network

Found in: `Component config` > `CONFIG_BLE_MESH`

This option specifies how many subnets a Mesh network can have at the same time. Indeed, this value decides the number of the network keys which can be owned by a node.

Range:

- from 1 to 4096 if `CONFIG_BLE_MESH`

Default value:

- 3 if `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_APP_KEY_COUNT

Maximum number of application keys per network

Found in: `Component config` > `CONFIG_BLE_MESH`

This option specifies how many application keys the device can store per network. Indeed, this value decides the number of the application keys which can be owned by a node.

Range:

- from 1 to 4096 if `CONFIG_BLE_MESH`

Default value:

- 3 if `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_MODEL_KEY_COUNT

Maximum number of application keys per model

Found in: `Component config` > `CONFIG_BLE_MESH`

This option specifies the maximum number of application keys to which each model can be bound.

Range:

- from 1 to 4096 if `CONFIG_BLE_MESH`

Default value:

- 3 if `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_MODEL_GROUP_COUNT

Maximum number of group address subscriptions per model

Found in: `Component config` > `CONFIG_BLE_MESH`

This option specifies the maximum number of addresses to which each model can be subscribed.

Range:

- from 1 to 4096 if `CONFIG_BLE_MESH`

Default value:

- 3 if `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_LABEL_COUNT

Maximum number of Label UUIDs used for Virtual Addresses

Found in: `Component config` > `CONFIG_BLE_MESH`

This option specifies how many Label UUIDs can be stored. Indeed, this value decides the number of the Virtual Addresses can be supported by a node.

Range:

- from 0 to 4096 if `CONFIG_BLE_MESH`

Default value:

- 3 if `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_CRPL

Maximum capacity of the replay protection list

Found in: `Component config` > `CONFIG_BLE_MESH`

This option specifies the maximum capacity of the replay protection list. It is similar to Network message cache size, but has a different purpose. The replay protection list is used to prevent a node from replay attack, which will store the source address and sequence number of the received mesh messages. For Provisioner, the replay protection list size should not be smaller than the maximum number of nodes whose information can be stored. And the element number of each node should also be taken into consideration. For example, if Provisioner can provision up to 20 nodes and each node contains two elements, then the replay protection list size of Provisioner should be at least 40.

Range:

- from 2 to 65535 if `CONFIG_BLE_MESH`

Default value:

- 10 if `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_NOT_RELAY_REPLAY_MSG

Not relay replayed messages in a mesh network

Found in: `Component config` > `CONFIG_BLE_MESH`

There may be many expired messages in a complex mesh network that would be considered replayed messages. Enable this option will refuse to relay such messages, which could help to reduce invalid packets in the mesh network. However, it should be noted that enabling this option may result in packet loss in certain environments. Therefore, users need to decide whether to enable this option according to the actual usage situation.

Default value:

- No (disabled) if `CONFIG_BLE_MESH_EXPERIMENTAL` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_MSG_CACHE_SIZE

Network message cache size

Found in: *Component config* > *CONFIG_BLE_MESH*

Number of messages that are cached for the network. This helps prevent unnecessary decryption operations and unnecessary relays. This option is similar to Replay protection list, but has a different purpose. A node is not required to cache the entire Network PDU and may cache only part of it for tracking, such as values for SRC/SEQ or others.

Range:

- from 2 to 65535 if *CONFIG_BLE_MESH*

Default value:

- 10 if *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_ADV_BUF_COUNT

Number of advertising buffers

Found in: *Component config* > *CONFIG_BLE_MESH*

Number of advertising buffers available. The transport layer reserves ADV_BUF_COUNT - 3 buffers for outgoing segments. The maximum outgoing SDU size is 12 times this value (out of which 4 or 8 bytes are used for the Transport Layer MIC). For example, 5 segments means the maximum SDU size is 60 bytes, which leaves 56 bytes for application layer data using a 4-byte MIC, or 52 bytes using an 8-byte MIC.

Range:

- from 6 to 256 if *CONFIG_BLE_MESH*

Default value:

- 60 if *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_IVU_DIVIDER

Divider for IV Update state refresh timer

Found in: *Component config* > *CONFIG_BLE_MESH*

When the IV Update state enters Normal operation or IV Update in Progress, we need to keep track of how many hours has passed in the state, since the specification requires us to remain in the state at least for 96 hours (Update in Progress has an additional upper limit of 144 hours).

In order to fulfill the above requirement, even if the node might be powered off once in a while, we need to store persistently how many hours the node has been in the state. This doesn't necessarily need to happen every hour (thanks to the flexible duration range). The exact cadence will depend a lot on the ways that the node will be used and what kind of power source it has.

Since there is no single optimal answer, this configuration option allows specifying a divider, i.e. how many intervals the 96 hour minimum gets split into. After each interval the duration that the node has been in the current state gets stored to flash. E.g. the default value of 4 means that the state is saved every 24 hours (96 / 4).

Range:

- from 2 to 96 if *CONFIG_BLE_MESH*

Default value:

- 4 if *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_IVU_RECOVERY_IVI

Recovery the IV index when the latest whole IV update procedure is missed

Found in: *Component config* > *CONFIG_BLE_MESH*

According to Section 3.10.5 of Mesh Specification v1.0.1. If a node in Normal Operation receives a Secure Network beacon with an IV index equal to the last known IV index+1 and the IV Update Flag set to 0, the node may update its IV without going to the IV Update in Progress state, or it may initiate an IV Index Recovery procedure (Section 3.10.6), or it may ignore the Secure Network beacon. The node makes the choice depending on the time since last IV update and the likelihood that the node has missed the Secure Network beacons with the IV update Flag. When the above situation is encountered, this option can be used to decide whether to perform the IV index recovery procedure.

Default value:

- No (disabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_SAR_ENHANCEMENT

Segmentation and reassembly enhancement

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Enable this option to use the enhanced segmentation and reassembly mechanism introduced in Bluetooth Mesh Protocol 1.1.

Default value:

- No (disabled) if [CONFIG_BLE_MESH_V11_SUPPORT](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_TX_SEG_MSG_COUNT

Maximum number of simultaneous outgoing segmented messages

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Maximum number of simultaneous outgoing multi-segment and/or reliable messages. The default value is 1, which means the device can only send one segmented message at a time. And if another segmented message is going to be sent, it should wait for the completion of the previous one. If users are going to send multiple segmented messages at the same time, this value should be configured properly.

Range:

- from 1 to if [CONFIG_BLE_MESH](#)

Default value:

- 1 if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_RX_SEG_MSG_COUNT

Maximum number of simultaneous incoming segmented messages

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Maximum number of simultaneous incoming multi-segment and/or reliable messages. The default value is 1, which means the device can only receive one segmented message at a time. And if another segmented message is going to be received, it should wait for the completion of the previous one. If users are going to receive multiple segmented messages at the same time, this value should be configured properly.

Range:

- from 1 to 255 if [CONFIG_BLE_MESH](#)

Default value:

- 1 if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_RX_SDU_MAX

Maximum incoming Upper Transport Access PDU length

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Maximum incoming Upper Transport Access PDU length. Leave this to the default value, unless you really need to optimize memory usage.

Range:

- from 36 to 384 if [CONFIG_BLE_MESH](#)

Default value:

- 384 if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_TX_SEG_MAX

Maximum number of segments in outgoing messages

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Maximum number of segments supported for outgoing messages. This value should typically be fine-tuned based on what models the local node supports, i.e. what's the largest message payload that the node needs to be able to send. This value affects memory and call stack consumption, which is why the default is lower than the maximum that the specification would allow (32 segments).

The maximum outgoing SDU size is 12 times this number (out of which 4 or 8 bytes is used for the Transport Layer MIC). For example, 5 segments means the maximum SDU size is 60 bytes, which leaves 56 bytes for application layer data using a 4-byte MIC and 52 bytes using an 8-byte MIC.

Be sure to specify a sufficient number of advertising buffers when setting this option to a higher value. There must be at least three more advertising buffers ([BLE_MESH_ADV_BUF_COUNT](#)) as there are outgoing segments.

Range:

- from 2 to 32 if [CONFIG_BLE_MESH](#)

Default value:

- 32 if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_RELAY

Relay support

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Support for acting as a Mesh Relay Node. Enabling this option will allow a node to support the Relay feature, and the Relay feature can still be enabled or disabled by proper configuration messages. Disabling this option will let a node not support the Relay feature.

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH_NODE](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_RELAY_ADV_BUF

Use separate advertising buffers for relay packets

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_RELAY](#)

When selected, self-send packets will be put in a high-priority queue and relay packets will be put in a low-priority queue.

Default value:

- No (disabled) if [CONFIG_BLE_MESH_RELAY](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_RELAY_ADV_BUF_COUNT

Number of advertising buffers for relay packets

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_RELAY](#) > [CONFIG_BLE_MESH_RELAY_ADV_BUF](#)

Number of advertising buffers for relay packets available.

Range:

- from 6 to 256 if `CONFIG_BLE_MESH_RELAY_ADV_BUF` && `CONFIG_BLE_MESH_RELAY` && `CONFIG_BLE_MESH`

Default value:

- 60 if `CONFIG_BLE_MESH_RELAY_ADV_BUF` && `CONFIG_BLE_MESH_RELAY` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_LOW_POWER

Support for Low Power features

Found in: `Component config` > `CONFIG_BLE_MESH`

Enable this option to operate as a Low Power Node. If low power consumption is required by a node, this option should be enabled. And once the node enters the mesh network, it will try to find a Friend node and establish a friendship.

CONFIG_BLE_MESH_LPN_ESTABLISHMENT

Perform Friendship establishment using low power

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_LOW_POWER`

Perform the Friendship establishment using low power with the help of a reduced scan duty cycle. The downside of this is that the node may miss out on messages intended for it until it has successfully set up Friendship with a Friend node. When this option is enabled, the node will stop scanning for a period of time after a Friend Request or Friend Poll is sent, so as to reduce more power consumption.

Default value:

- No (disabled) if `CONFIG_BLE_MESH_LOW_POWER` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_LPN_AUTO

Automatically start looking for Friend nodes once provisioned

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_LOW_POWER`

Once provisioned, automatically enable LPN functionality and start looking for Friend nodes. If this option is disabled LPN mode needs to be manually enabled by calling `bt_mesh_lpn_set(true)`. When an unprovisioned device is provisioned successfully and becomes a node, enabling this option will trigger the node starts to send Friend Request at a certain period until it finds a proper Friend node.

Default value:

- No (disabled) if `CONFIG_BLE_MESH_LOW_POWER` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_LPN_AUTO_TIMEOUT

Time from last received message before going to LPN mode

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_LOW_POWER` > `CONFIG_BLE_MESH_LPN_AUTO`

Time in seconds from the last received message, that the node waits out before starting to look for Friend nodes.

Range:

- from 0 to 3600 if `CONFIG_BLE_MESH_LPN_AUTO` && `CONFIG_BLE_MESH_LOW_POWER` && `CONFIG_BLE_MESH`

Default value:

- 15 if `CONFIG_BLE_MESH_LPN_AUTO` && `CONFIG_BLE_MESH_LOW_POWER` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_LPN_RETRY_TIMEOUT

Retry timeout for Friend requests

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_LOW_POWER](#)

Time in seconds between Friend Requests, if a previous Friend Request did not yield any acceptable Friend Offers.

Range:

- from 1 to 3600 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

Default value:

- 6 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_LPN_RSSI_FACTOR

RSSIFactor, used in Friend Offer Delay calculation

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_LOW_POWER](#)

The contribution of the RSSI, measured by the Friend node, used in Friend Offer Delay calculations. 0 = 1, 1 = 1.5, 2 = 2, 3 = 2.5. RSSIFactor, one of the parameters carried by Friend Request sent by Low Power node, which is used to calculate the Friend Offer Delay.

Range:

- from 0 to 3 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

Default value:

- 0 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_LPN_RECV_WIN_FACTOR

ReceiveWindowFactor, used in Friend Offer Delay calculation

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_LOW_POWER](#)

The contribution of the supported Receive Window used in Friend Offer Delay calculations. 0 = 1, 1 = 1.5, 2 = 2, 3 = 2.5. ReceiveWindowFactor, one of the parameters carried by Friend Request sent by Low Power node, which is used to calculate the Friend Offer Delay.

Range:

- from 0 to 3 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

Default value:

- 0 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_LPN_MIN_QUEUE_SIZE

Minimum size of the acceptable friend queue (MinQueueSizeLog)

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_LOW_POWER](#)

The MinQueueSizeLog field is defined as $\log_2(N)$, where N is the minimum number of maximum size Lower Transport PDUs that the Friend node can store in its Friend Queue. As an example, MinQueueSizeLog value 1 gives N = 2, and value 7 gives N = 128.

Range:

- from 1 to 7 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

Default value:

- 1 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_LPN_RECV_DELAY

Receive delay requested by the local node

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_LOW_POWER](#)

The ReceiveDelay is the time between the Low Power node sending a request and listening for a response. This delay allows the Friend node time to prepare the response. The value is in units of milliseconds.

Range:

- from 10 to 255 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

Default value:

- 100 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_LPN_POLL_TIMEOUT

The value of the PollTimeout timer

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_LOW_POWER](#)

PollTimeout timer is used to measure time between two consecutive requests sent by a Low Power node. If no requests are received the Friend node before the PollTimeout timer expires, then the friendship is considered terminated. The value is in units of 100 milliseconds, so e.g. a value of 300 means 30 seconds. The smaller the value, the faster the Low Power node tries to get messages from corresponding Friend node and vice versa.

Range:

- from 10 to 244735 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

Default value:

- 300 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_LPN_INIT_POLL_TIMEOUT

The starting value of the PollTimeout timer

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_LOW_POWER](#)

The initial value of the PollTimeout timer when Friendship is to be established for the first time. After this, the timeout gradually grows toward the actual PollTimeout, doubling in value for each iteration. The value is in units of 100 milliseconds, so e.g. a value of 300 means 30 seconds.

Range:

- from 10 to if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

Default value:

- if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_LPN_SCAN_LATENCY

Latency for enabling scanning

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_LOW_POWER](#)

Latency (in milliseconds) is the time it takes to enable scanning. In practice, it means how much time in advance of the Receive Window, the request to enable scanning is made.

Range:

- from 0 to 50 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

Default value:

- 10 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_LPN_GROUPS

Number of groups the LPN can subscribe to

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_LOW_POWER](#)

Maximum number of groups to which the LPN can subscribe.

Range:

- from 0 to 16384 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

Default value:

- 8 if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_LPN_SUB_ALL_NODES_ADDR

Automatically subscribe all nodes address

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_LOW_POWER](#)

Automatically subscribe all nodes address when friendship established.

Default value:

- No (disabled) if [CONFIG_BLE_MESH_LOW_POWER](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_FRIEND

Support for Friend feature

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Enable this option to be able to act as a Friend Node.

CONFIG_BLE_MESH_FRIEND_RECV_WIN

Friend Receive Window

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_FRIEND](#)

Receive Window in milliseconds supported by the Friend node.

Range:

- from 1 to 255 if [CONFIG_BLE_MESH_FRIEND](#) && [CONFIG_BLE_MESH](#)

Default value:

- 255 if [CONFIG_BLE_MESH_FRIEND](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_FRIEND_QUEUE_SIZE

Minimum number of buffers supported per Friend Queue

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_FRIEND](#)

Minimum number of buffers available to be stored for each local Friend Queue. This option decides the size of each buffer which can be used by a Friend node to store messages for each Low Power node.

Range:

- from 2 to 65536 if [CONFIG_BLE_MESH_FRIEND](#) && [CONFIG_BLE_MESH](#)

Default value:

- 16 if [CONFIG_BLE_MESH_FRIEND](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_FRIEND_SUB_LIST_SIZE

Friend Subscription List Size

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [CONFIG_BLE_MESH_FRIEND](#)

Size of the Subscription List that can be supported by a Friend node for a Low Power node. And Low Power node can send Friend Subscription List Add or Friend Subscription List Remove messages to the Friend node to add or remove subscription addresses.

Range:

- from 0 to 1023 if [CONFIG_BLE_MESH_FRIEND](#) && [CONFIG_BLE_MESH](#)

Default value:

- 3 if `CONFIG_BLE_MESH_FRIEND` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_FRIEND_LPN_COUNT

Number of supported LPN nodes

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_FRIEND`

Number of Low Power Nodes with which a Friend can have Friendship simultaneously. A Friend node can have friendship with multiple Low Power nodes at the same time, while a Low Power node can only establish friendship with only one Friend node at the same time.

Range:

- from 1 to 1000 if `CONFIG_BLE_MESH_FRIEND` && `CONFIG_BLE_MESH`

Default value:

- 2 if `CONFIG_BLE_MESH_FRIEND` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_FRIEND_SEG_RX

Number of incomplete segment lists per LPN

Found in: `Component config` > `CONFIG_BLE_MESH` > `CONFIG_BLE_MESH_FRIEND`

Number of incomplete segment lists tracked for each Friends' LPN. In other words, this determines from how many elements can segmented messages destined for the Friend queue be received simultaneously.

Range:

- from 1 to 1000 if `CONFIG_BLE_MESH_FRIEND` && `CONFIG_BLE_MESH`

Default value:

- 1 if `CONFIG_BLE_MESH_FRIEND` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_NO_LOG

Disable BLE Mesh debug logs (minimize bin size)

Found in: `Component config` > `CONFIG_BLE_MESH`

Select this to save the BLE Mesh related rodata code size. Enabling this option will disable the output of BLE Mesh debug log.

Default value:

- No (disabled) if `CONFIG_BLE_MESH` && `CONFIG_BLE_MESH`

BLE Mesh STACK DEBUG LOG LEVEL Contains:

- `CONFIG_BLE_MESH_STACK_TRACE_LEVEL`

CONFIG_BLE_MESH_STACK_TRACE_LEVEL

BLE_MESH_STACK

Found in: `Component config` > `CONFIG_BLE_MESH` > `BLE Mesh STACK DEBUG LOG LEVEL`

Define BLE Mesh trace level for BLE Mesh stack.

Available options:

- NONE (`CONFIG_BLE_MESH_TRACE_LEVEL_NONE`)
- ERROR (`CONFIG_BLE_MESH_TRACE_LEVEL_ERROR`)
- WARNING (`CONFIG_BLE_MESH_TRACE_LEVEL_WARNING`)
- INFO (`CONFIG_BLE_MESH_TRACE_LEVEL_INFO`)

- `DEBUG (CONFIG_BLE_MESH_TRACE_LEVEL_DEBUG)`
- `VERBOSE (CONFIG_BLE_MESH_TRACE_LEVEL_VERBOSE)`

BLE Mesh NET BUF DEBUG LOG LEVEL Contains:

- [`CONFIG\_BLE\_MESH\_NET\_BUF\_TRACE\_LEVEL`](#)

CONFIG_BLE_MESH_NET_BUF_TRACE_LEVEL

`BLE_MESH_NET_BUF`

Found in: *Component config* > [`CONFIG\_BLE\_MESH`](#) > [`BLE Mesh NET BUF DEBUG LOG LEVEL`](#)

Define BLE Mesh trace level for BLE Mesh net buffer.

Available options:

- `NONE (CONFIG_BLE_MESH_NET_BUF_TRACE_LEVEL_NONE)`
- `ERROR (CONFIG_BLE_MESH_NET_BUF_TRACE_LEVEL_ERROR)`
- `WARNING (CONFIG_BLE_MESH_NET_BUF_TRACE_LEVEL_WARNING)`
- `INFO (CONFIG_BLE_MESH_NET_BUF_TRACE_LEVEL_INFO)`
- `DEBUG (CONFIG_BLE_MESH_NET_BUF_TRACE_LEVEL_DEBUG)`
- `VERBOSE (CONFIG_BLE_MESH_NET_BUF_TRACE_LEVEL_VERBOSE)`

CONFIG_BLE_MESH_CLIENT_MSG_TIMEOUT

Timeout(ms) for client message response

Found in: *Component config* > [`CONFIG\_BLE\_MESH`](#)

Timeout value used by the node to get response of the acknowledged message which is sent by the client model. This value indicates the maximum time that a client model waits for the response of the sent acknowledged messages. If a client model uses 0 as the timeout value when sending acknowledged messages, then the default value will be used which is four seconds.

Range:

- from 100 to 1200000 if [`CONFIG\_BLE\_MESH`](#)

Default value:

- 4000 if [`CONFIG\_BLE\_MESH`](#)

Support for BLE Mesh Foundation models Contains:

- [`CONFIG\_BLE\_MESH\_BRC\_CLI`](#)
- [`CONFIG\_BLE\_MESH\_BRC\_SRV`](#)
- [`CONFIG\_BLE\_MESH\_CFG\_CLI`](#)
- [`CONFIG\_BLE\_MESH\_DF\_CLI`](#)
- [`CONFIG\_BLE\_MESH\_DF\_SRV`](#)
- [`CONFIG\_BLE\_MESH\_HEALTH\_CLI`](#)
- [`CONFIG\_BLE\_MESH\_HEALTH\_SRV`](#)
- [`CONFIG\_BLE\_MESH\_LCD\_CLI`](#)
- [`CONFIG\_BLE\_MESH\_LCD\_SRV`](#)
- [`CONFIG\_BLE\_MESH\_PRB\_CLI`](#)
- [`CONFIG\_BLE\_MESH\_PRB\_SRV`](#)
- [`CONFIG\_BLE\_MESH\_ODP\_CLI`](#)
- [`CONFIG\_BLE\_MESH\_ODP\_SRV`](#)
- [`CONFIG\_BLE\_MESH\_AGG\_CLI`](#)
- [`CONFIG\_BLE\_MESH\_AGG\_SRV`](#)
- [`CONFIG\_BLE\_MESH\_RPR\_CLI`](#)
- [`CONFIG\_BLE\_MESH\_RPR\_SRV`](#)

- [*CONFIG_BLE_MESH_SAR_CLI*](#)
- [*CONFIG_BLE_MESH_SAR_SRV*](#)
- [*CONFIG_BLE_MESH_SRPL_CLI*](#)
- [*CONFIG_BLE_MESH_SRPL_SRV*](#)
- [*CONFIG_BLE_MESH_COMP_DATA_1*](#)
- [*CONFIG_BLE_MESH_COMP_DATA_128*](#)
- [*CONFIG_BLE_MESH_MODELS_METADATA_0*](#)

CONFIG_BLE_MESH_CFG_CLI

Configuration Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Configuration Client model.

CONFIG_BLE_MESH_HEALTH_CLI

Health Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Health Client model.

CONFIG_BLE_MESH_HEALTH_SRV

Health Server model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Health Server model.

Default value:

- Yes (enabled) if [*CONFIG_BLE_MESH*](#)

CONFIG_BLE_MESH_BRC_CLI

Bridge Configuration Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Bridge Configuration Client model.

CONFIG_BLE_MESH_BRC_SRV

Bridge Configuration Server model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Bridge Configuration Server model.

Default value:

- No (disabled) if [*CONFIG_BLE_MESH_V11_SUPPORT*](#) && [*CONFIG_BLE_MESH*](#)

CONFIG_BLE_MESH_MAX_BRIDGING_TABLE_ENTRY_COUNT

Maximum number of Bridging Table entries

*Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#) > [*CONFIG_BLE_MESH_BRC_SRV*](#)*

Maximum number of Bridging Table entries that the Bridge Configuration Server can support.

Range:

- from 16 to 65535 if `CONFIG_BLE_MESH_BRC_SRV` && `CONFIG_BLE_MESH_V11_SUPPORT` && `CONFIG_BLE_MESH`

Default value:

- 16 if `CONFIG_BLE_MESH_BRC_SRV` && `CONFIG_BLE_MESH_V11_SUPPORT` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_BRIDGE_CRPL

Maximum capacity of bridge replay protection list

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#) > [CONFIG_BLE_MESH_BRC_SRV](#)

This option specifies the maximum capacity of the bridge replay protection list. The bridge replay protection list is used to prevent a bridged subnet from replay attack, which will store the source address and sequence number of the received bridge messages.

Range:

- from 1 to 255 if `CONFIG_BLE_MESH_BRC_SRV` && `CONFIG_BLE_MESH_V11_SUPPORT` && `CONFIG_BLE_MESH`

Default value:

- 5 if `CONFIG_BLE_MESH_BRC_SRV` && `CONFIG_BLE_MESH_V11_SUPPORT` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_PRB_CLI

Mesh Private Beacon Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Mesh Private Beacon Client model.

CONFIG_BLE_MESH_PRB_SRV

Mesh Private Beacon Server model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Mesh Private Beacon Server model.

CONFIG_BLE_MESH_ODP_CLI

On-Demand Private Proxy Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for On-Demand Private Proxy Client model.

CONFIG_BLE_MESH_ODP_SRV

On-Demand Private Proxy Server model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for On-Demand Private Proxy Server model.

CONFIG_BLE_MESH_SRPL_CLI

Solicitation PDU RPL Configuration Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Solicitation PDU RPL Configuration Client model.

CONFIG_BLE_MESH_SRPL_SRV

Solicitation PDU RPL Configuration Server model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Solicitation PDU RPL Configuration Server model. Note: This option depends on the functionality of receiving Solicitation PDU. If the device doesn't support receiving Solicitation PDU, then there is no need to enable this server model.

CONFIG_BLE_MESH_AGG_CLI

Opcodes Aggregator Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Opcodes Aggregator Client model.

CONFIG_BLE_MESH_AGG_SRV

Opcodes Aggregator Server model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Opcodes Aggregator Server model.

CONFIG_BLE_MESH_SAR_CLI

SAR Configuration Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for SAR Configuration Client model.

CONFIG_BLE_MESH_SAR_SRV

SAR Configuration Server model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for SAR Configuration Server model.

CONFIG_BLE_MESH_COMP_DATA_1

Support Composition Data Page 1

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Composition Data Page 1 contains information about the relationships among models. Each model either can be a root model or can extend other models.

CONFIG_BLE_MESH_COMP_DATA_128

Support Composition Data Page 128

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Composition Data Page 128 is used to indicate the structure of elements, features, and models of a node after the successful execution of the Node Address Refresh procedure or the Node Composition Refresh procedure, or after the execution of the Node Removal procedure followed by the provisioning process. Composition Data Page 128 shall be present if the node supports the Remote Provisioning Server model; otherwise it is optional.

CONFIG_BLE_MESH_MODELS_METADATA_0

Support Models Metadata Page 0

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

The Models Metadata state contains metadata of a node's models. The Models Metadata state is composed of a number of pages of information. Models Metadata Page 0 shall be present if the node supports the Large Composition Data Server model.

CONFIG_BLE_MESH_MODELS_METADATA_128

Support Models Metadata Page 128

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#) > [CONFIG_BLE_MESH_MODELS_METADATA_0](#)

The Models Metadata state contains metadata of a node's models. The Models Metadata state is composed of a number of pages of information. Models Metadata Page 128 contains metadata for the node's models after the successful execution of the Node Address Refresh procedure or the Node Composition Refresh procedure, or after the execution of the Node Removal procedure followed by the provisioning process. Models Metadata Page 128 shall be present if the node supports the Remote Provisioning Server model and the node supports the Large Composition Data Server model.

CONFIG_BLE_MESH_LCD_CLI

Large Composition Data Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Large Composition Data Client model.

CONFIG_BLE_MESH_LCD_SRV

Large Composition Data Server model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Large Composition Data Server model.

CONFIG_BLE_MESH_RPR_CLI

Remote Provisioning Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Remote Provisioning Client model

CONFIG_BLE_MESH_RPR_CLI_PROV_SAME_TIME

Maximum number of PB-Remote running at the same time by Provisioner

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#) > [CONFIG_BLE_MESH_RPR_CLI](#)

This option specifies how many devices can be provisioned at the same time using PB-REMOTE. For example, if the value is 2, it means a Provisioner can provision two unprovisioned devices with PB-REMOTE at the same time.

Range:

- from 1 to 5 if [CONFIG_BLE_MESH_RPR_CLI](#) && [CONFIG_BLE_MESH_V11_SUPPORT](#) && [CONFIG_BLE_MESH](#)

Default value:

- 2 if `CONFIG_BLE_MESH_RPR_CLI` && `CONFIG_BLE_MESH_V11_SUPPORT` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_RPR_SRV

Remote Provisioning Server model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Remote Provisioning Server model

CONFIG_BLE_MESH_RPR_SRV_MAX_SCANNED_ITEMS

Maximum number of device information can be scanned

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#) > [CONFIG_BLE_MESH_RPR_SRV](#)

This option specifies how many device information can a Remote Provisioning Server store each time while scanning.

Range:

- from 4 to 255 if `CONFIG_BLE_MESH_RPR_SRV` && `CONFIG_BLE_MESH_V11_SUPPORT` && `CONFIG_BLE_MESH`

Default value:

- 10 if `CONFIG_BLE_MESH_RPR_SRV` && `CONFIG_BLE_MESH_V11_SUPPORT` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_RPR_SRV_ACTIVE_SCAN

Support Active Scan for remote provisioning

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#) > [CONFIG_BLE_MESH_RPR_SRV](#)

Enable this option to support Active Scan for remote provisioning.

CONFIG_BLE_MESH_RPR_SRV_MAX_EXT_SCAN

Maximum number of extended scan procedures

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#) > [CONFIG_BLE_MESH_RPR_SRV](#)

This option specifies how many extended scan procedures can be started by the Remote Provisioning Server.

Range:

- from 1 to 10 if `CONFIG_BLE_MESH_RPR_SRV` && `CONFIG_BLE_MESH_V11_SUPPORT` && `CONFIG_BLE_MESH`

Default value:

- 1 if `CONFIG_BLE_MESH_RPR_SRV` && `CONFIG_BLE_MESH_V11_SUPPORT` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_DF_CLI

Directed Forwarding Configuration Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Directed Forwarding Configuration Client model.

CONFIG_BLE_MESH_DF_SRV

Directed Forwarding Configuration Server model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#)

Enable support for Directed Forwarding Configuration Server model.

CONFIG_BLE_MESH_MAX_DISC_TABLE_ENTRY_COUNT

Maximum number of discovery table entries in a given subnet

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#) > [CONFIG_BLE_MESH_DF_SRV](#)

Maximum number of Discovery Table entries supported by the node in a given subnet.

Range:

- from 2 to 255 if [CONFIG_BLE_MESH_DF_SRV](#) && [CONFIG_BLE_MESH_V11_SUPPORT](#) && [CONFIG_BLE_MESH](#)

Default value:

- 2 if [CONFIG_BLE_MESH_DF_SRV](#) && [CONFIG_BLE_MESH_V11_SUPPORT](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_MAX_FORWARD_TABLE_ENTRY_COUNT

Maximum number of forward table entries in a given subnet

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#) > [CONFIG_BLE_MESH_DF_SRV](#)

Maximum number of Forward Table entries supported by the node in a given subnet.

Range:

- from 2 to 64 if [CONFIG_BLE_MESH_DF_SRV](#) && [CONFIG_BLE_MESH_V11_SUPPORT](#) && [CONFIG_BLE_MESH](#)

Default value:

- 2 if [CONFIG_BLE_MESH_DF_SRV](#) && [CONFIG_BLE_MESH_V11_SUPPORT](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_MAX_DEPS_NODES_PER_PATH

Maximum number of dependent nodes per path

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#) > [CONFIG_BLE_MESH_DF_SRV](#)

Maximum size of dependent nodes list supported by each forward table entry.

Range:

- from 2 to 64 if [CONFIG_BLE_MESH_DF_SRV](#) && [CONFIG_BLE_MESH_V11_SUPPORT](#) && [CONFIG_BLE_MESH](#)

Default value:

- 2 if [CONFIG_BLE_MESH_DF_SRV](#) && [CONFIG_BLE_MESH_V11_SUPPORT](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_PATH_MONITOR_TEST

Enable Path Monitoring test mode

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Foundation models](#) > [CONFIG_BLE_MESH_DF_SRV](#)

The option only removes the Path Use timer; all other behavior of the device is not changed. If Path Monitoring test mode is going to be used, this option should be enabled.

Default value:

- No (disabled) if `CONFIG_BLE_MESH_DF_SRV` && `CONFIG_BLE_MESH_V11_SUPPORT` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_SUPPORT_DIRECTED_PROXY

Enable Directed Proxy functionality

Found in: Component config > CONFIG_BLE_MESH > Support for BLE Mesh Foundation models > CONFIG_BLE_MESH_DF_SRV

Support Directed Proxy functionality.

Default value:

- Yes (enabled) if `CONFIG_BLE_MESH_GATT_PROXY_SERVER` && `CONFIG_BLE_MESH_DF_SRV` && `CONFIG_BLE_MESH_V11_SUPPORT` && `CONFIG_BLE_MESH`

Support for BLE Mesh Client/Server models Contains:

- *Binary Larger Object Transfer model*
- `CONFIG_BLE_MESH_MBT_CLI`
- `CONFIG_BLE_MESH_MBT_SRV`
- *Device Firmware Update model*
- `CONFIG_BLE_MESH_GENERIC_BATTERY_CLI`
- `CONFIG_BLE_MESH_GENERIC_DEF_TRANS_TIME_CLI`
- `CONFIG_BLE_MESH_GENERIC_LEVEL_CLI`
- `CONFIG_BLE_MESH_GENERIC_LOCATION_CLI`
- `CONFIG_BLE_MESH_GENERIC_ONOFF_CLI`
- `CONFIG_BLE_MESH_GENERIC_POWER_LEVEL_CLI`
- `CONFIG_BLE_MESH_GENERIC_POWER_ONOFF_CLI`
- `CONFIG_BLE_MESH_GENERIC_PROPERTY_CLI`
- `CONFIG_BLE_MESH_GENERIC_SERVER`
- `CONFIG_BLE_MESH_LIGHT_CTL_CLI`
- `CONFIG_BLE_MESH_LIGHT_HSL_CLI`
- `CONFIG_BLE_MESH_LIGHT_LC_CLI`
- `CONFIG_BLE_MESH_LIGHT_LIGHTNESS_CLI`
- `CONFIG_BLE_MESH_LIGHT_XYL_CLI`
- `CONFIG_BLE_MESH_LIGHTING_SERVER`
- `CONFIG_BLE_MESH_SCENE_CLI`
- `CONFIG_BLE_MESH_SCHEDULER_CLI`
- `CONFIG_BLE_MESH_SENSOR_CLI`
- `CONFIG_BLE_MESH_SENSOR_SERVER`
- `CONFIG_BLE_MESH_TIME_SCENE_SERVER`
- `CONFIG_BLE_MESH_TIME_CLI`

CONFIG_BLE_MESH_GENERIC_ONOFF_CLI

Generic OnOff Client model

Found in: Component config > CONFIG_BLE_MESH > Support for BLE Mesh Client/Server models

Enable support for Generic OnOff Client model.

CONFIG_BLE_MESH_GENERIC_LEVEL_CLI

Generic Level Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Generic Level Client model.

CONFIG_BLE_MESH_GENERIC_DEF_TRANS_TIME_CLI

Generic Default Transition Time Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Generic Default Transition Time Client model.

CONFIG_BLE_MESH_GENERIC_POWER_ONOFF_CLI

Generic Power OnOff Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Generic Power OnOff Client model.

CONFIG_BLE_MESH_GENERIC_POWER_LEVEL_CLI

Generic Power Level Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Generic Power Level Client model.

CONFIG_BLE_MESH_GENERIC_BATTERY_CLI

Generic Battery Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Generic Battery Client model.

CONFIG_BLE_MESH_GENERIC_LOCATION_CLI

Generic Location Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Generic Location Client model.

CONFIG_BLE_MESH_GENERIC_PROPERTY_CLI

Generic Property Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Generic Property Client model.

CONFIG_BLE_MESH_SENSOR_CLI

Sensor Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Sensor Client model.

CONFIG_BLE_MESH_TIME_CLI

Time Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Time Client model.

CONFIG_BLE_MESH_SCENE_CLI

Scene Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Scene Client model.

CONFIG_BLE_MESH_SCHEDULER_CLI

Scheduler Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Scheduler Client model.

CONFIG_BLE_MESH_LIGHT_LIGHTNESS_CLI

Light Lightness Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Light Lightness Client model.

CONFIG_BLE_MESH_LIGHT_CTL_CLI

Light CTL Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Light CTL Client model.

CONFIG_BLE_MESH_LIGHT_HSL_CLI

Light HSL Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Light HSL Client model.

CONFIG_BLE_MESH_LIGHT_XYL_CLI

Light XYL Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Light XYL Client model.

CONFIG_BLE_MESH_LIGHT_LC_CLI

Light LC Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Light LC Client model.

CONFIG_BLE_MESH_GENERIC_SERVER

Generic server models

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Generic server models.

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_SENSOR_SERVER

Sensor server models

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Sensor server models.

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_TIME_SCENE_SERVER

Time and Scenes server models

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Time and Scenes server models.

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_LIGHTING_SERVER

Lighting server models

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for Lighting server models.

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_MBT_CLI

BLOB Transfer Client model(Deprecated)

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for BLOB Transfer Client model. Warn: This version of the Mesh Binary Large Object Transfer Model will be deprecated, and a new version will be released in the future.

Default value:

- No (disabled) if [CONFIG_BLE_MESH_V11_SUPPORT](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_MAX_BLOB_RECEIVERS

Maximum number of simultaneous blob receivers

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [CONFIG_BLE_MESH_MBT_CLI](#)

Maximum number of BLOB Transfer Server models that can participating in the BLOB transfer with a BLOB Transfer Client model.

Range:

- from 1 to 255 if [CONFIG_BLE_MESH_MBT_CLI](#) && [CONFIG_BLE_MESH](#)

Default value:

- 2 if [CONFIG_BLE_MESH_MBT_CLI](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_MBT_SRV

BLOB Transfer Server model(Deprecated)

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#)

Enable support for BLOB Transfer Server model. Warn: This version of the Mesh Binary Large Object Transfer Model will be deprecated, and a new version will be released in the future.

Default value:

- No (disabled) if [CONFIG_BLE_MESH_V11_SUPPORT](#) && [CONFIG_BLE_MESH](#)

Binary Larger Object Transfer model Contains:

- [BLOB models common configuration](#)
- [CONFIG_BLE_MESH_BLOB_CLI](#)
- [CONFIG_BLE_MESH_BLOB_SRV](#)

CONFIG_BLE_MESH_BLOB_SRV

Support for BLOB Transfer Server model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Binary Larger Object Transfer model](#)

Enable the Binary Large Object (BLOB) Transfer Server model.

CONFIG_BLE_MESH_BLOB_SRV_PULL_REQ_COUNT

Number of chunks to request for each pull

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Binary Larger Object Transfer model](#) > [CONFIG_BLE_MESH_BLOB_SRV](#)

In Pull mode (Pull BLOB Transfer Mode), the BLOB Transfer Server requests a fixed number of chunks from the Client. Use this option to control the chunk count in the request. If the BLOB Transfer Server is instantiated on a Low Power node, the pull request count will be trimmed to not overflow the Friend queue.

Range:

- from 1 to 16 if [CONFIG_BLE_MESH_BLOB_SRV](#) && [CONFIG_BLE_MESH](#)

Default value:

- 4 if [CONFIG_BLE_MESH_BLOB_SRV](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_BLOB_SIZE_MAX

Largest BLOB size in bytes

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Binary Larger Object Transfer model](#) > [CONFIG_BLE_MESH_BLOB_SRV](#)

The maximum object size a BLOB Transfer Server can receive.

Range:

- from 1 to 3257617792 if [CONFIG_BLE_MESH_BLOB_SRV](#) && [CONFIG_BLE_MESH](#)

Default value:

- 524288 if [CONFIG_BLE_MESH_BLOB_SRV](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_BLOB_BLOCK_SIZE_MIN

Minimum block size

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Binary Larger Object Transfer model](#) > [CONFIG_BLE_MESH_BLOB_SRV](#)

Minimum acceptable block size in a BLOB transfer. The transfer block size will be some number that is a power of two, and is between block size min and block size max. If no such number exists, a compile time warning will be issued.

Range:

- from 64 to 1048576 if [CONFIG_BLE_MESH_BLOB_SRV](#) && [CONFIG_BLE_MESH](#)

Default value:

- 4096 if [CONFIG_BLE_MESH_BLOB_SRV](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_BLOB_BLOCK_SIZE_MAX

Maximum block size

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Binary Larger Object Transfer model](#) > [CONFIG_BLE_MESH_BLOB_SRV](#)

Maximum acceptable block size in a BLOB transfer. The transfer block size will be some number that is a power of two, and is between block size min and block size max. If no such number exists, a compile time warning will be issued.

Range:

- from 1 to 1048576 if [CONFIG_BLE_MESH_BLOB_SRV](#) && [CONFIG_BLE_MESH](#)

Default value:

- 4096 if [CONFIG_BLE_MESH_BLOB_SRV](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_BLOB_REPORT_TIMEOUT

Partial Block Report interval in Pull mode

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Binary Larger Object Transfer model](#) > [CONFIG_BLE_MESH_BLOB_SRV](#)

The timer value that Pull BLOB Transfer Server uses to report missed chunks.

Range:

- from 1 to 31 if [CONFIG_BLE_MESH_BLOB_SRV](#) && [CONFIG_BLE_MESH](#)

Default value:

- 10 if [CONFIG_BLE_MESH_BLOB_SRV](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_RX_BLOB_CHUNK_SIZE

BLOB Server chunk size

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Binary Larger Object Transfer model](#) > [CONFIG_BLE_MESH_BLOB_SRV](#)

Set the chunk size for the BLOB Server. The actual maximum chunk size depends on how many segments are possible and will be clamped to the max possible if set above. see also: BLE_MESH_RX_SEG_MAX, and the maximum SDU a node can receive.

CONFIG_BLE_MESH_BLOB_CLI

Support for BLOB Transfer Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Binary Larger Object Transfer model](#)

Enable the Binary Large Object (BLOB) Transfer Client model.

CONFIG_BLE_MESH_BLOB_CLI_BLOCK_RETRIES

Number of retries per block

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Binary Larger Object Transfer model](#) > [CONFIG_BLE_MESH_BLOB_CLI](#)

Controls the number of times the client will attempt to resend missing chunks to the BLOB receivers for every block.

Default value:

- 5 if [CONFIG_BLE_MESH_BLOB_CLI](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_TX_BLOB_CHUNK_SIZE

BLOB Client chunk size

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Binary Larger Object Transfer model](#) > [CONFIG_BLE_MESH_BLOB_CLI](#)

Set the chunk size for the BLOB Client. The actual maximum chunk size depends on how many segments are possible and will be clamped to the max possible if set above. see also: [BLE_MESH_TX_SEG_MAX](#), and the maximum SDU a node can receive.

CONFIG_BLE_MESH_TX_BLOB_CHUNK_SEND_INTERVAL

BLOB Client chunk send interval

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Binary Larger Object Transfer model](#) > [CONFIG_BLE_MESH_BLOB_CLI](#)

Set the interval in milliseconds in which chunks are sent during the BLOB transfer. Note: Without a delay between each sent chunk, the network might become too busy with the BLOB transfer for other communications to succeed. Note: Timing restrictions, like the timeout base, should be considered or changed accordingly when setting this interval. Otherwise, the interval might be too big for the timeout settings and cause timeouts.

Range:

- from 0 to 2147483647 if [CONFIG_BLE_MESH_BLOB_CLI](#) && [CONFIG_BLE_MESH](#)

Default value:

- 0 if [CONFIG_BLE_MESH_BLOB_CLI](#) && [CONFIG_BLE_MESH](#)

BLOB models common configuration

 Contains:

- [CONFIG_BLE_MESH_ALIGN_CHUNK_SIZE_TO_MAX_SEGMENT](#)
- [CONFIG_BLE_MESH_BLOB_CHUNK_COUNT_MAX](#)

CONFIG_BLE_MESH_BLOB_CHUNK_COUNT_MAX

Maximum chunk count per block

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Binary Larger Object Transfer model](#) > [BLOB models common configuration](#)

A BLOB transfer contains several blocks. Each block is made up of several chunks. This option controls the maximum chunk count per block.

Range:

- from 1 to 2992 if [CONFIG_BLE_MESH](#)

Default value:

- 256 if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_ALIGN_CHUNK_SIZE_TO_MAX_SEGMENT

Align chunk size to max segmented message size

Found in: Component config > CONFIG_BLE_MESH > Support for BLE Mesh Client/Server models > Binary Larger Object Transfer model > BLOB models common configuration

Default value:

- Yes (enabled) if *CONFIG_BLE_MESH*

Device Firmware Update model Contains:

- *CONFIG_BLE_MESH_DFU_SLOTS*
- *Firmware Update model configuration*
- *CONFIG_BLE_MESH_DFD_CLI*
- *CONFIG_BLE_MESH_DFD_SRV*
- *CONFIG_BLE_MESH_DFU_CLI*
- *CONFIG_BLE_MESH_DFU_SRV*

CONFIG_BLE_MESH_DFU_SRV

Support for Firmware Update Server model

Found in: Component config > CONFIG_BLE_MESH > Support for BLE Mesh Client/Server models > Device Firmware Update model

Enable the Firmware Update Server model.

CONFIG_BLE_MESH_DFU_CLI

Support for Firmware Update Client model

Found in: Component config > CONFIG_BLE_MESH > Support for BLE Mesh Client/Server models > Device Firmware Update model

Enable the Firmware Update Client model.

Firmware Update model configuration Contains:

- *CONFIG_BLE_MESH_DFU_FWID_MAXLEN*
- *CONFIG_BLE_MESH_DFU_METADATA_MAXLEN*
- *CONFIG_BLE_MESH_DFU_URI_MAXLEN*
- *CONFIG_BLE_MESH_DFU_METADATA*

CONFIG_BLE_MESH_DFU_FWID_MAXLEN

DFU FWID max length

Found in: Component config > CONFIG_BLE_MESH > Support for BLE Mesh Client/Server models > Device Firmware Update model > Firmware Update model configuration

This value defines the maximum length of an image's firmware ID.

Range:

- from 0 to 106 if *CONFIG_BLE_MESH*

Default value:

- 16 if *CONFIG_BLE_MESH*

CONFIG_BLE_MESH_DFU_METADATA_MAXLEN

DFU metadata max length

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Device Firmware Update model](#) > [Firmware Update model configuration](#)

This value defines the maximum length of an image's metadata.

Range:

- from 18 to 255 if [CONFIG_BLE_MESH_DFU_METADATA](#) && [CONFIG_BLE_MESH](#)
- from 0 to 255 if [CONFIG_BLE_MESH](#)

Default value:

- 32 if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_DFU_METADATA

Support for the default metadata format

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Device Firmware Update model](#) > [Firmware Update model configuration](#)

This option adds a set of functions that can be used to encode and decode a firmware metadata using the format defined in the Bluetooth mesh DFU subsystem.

CONFIG_BLE_MESH_DFU_URI_MAXLEN

DFU URI max length

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Device Firmware Update model](#) > [Firmware Update model configuration](#)

This value defines the maximum length of an image's URI, not including a string terminator.

Range:

- from 0 to 255 if [CONFIG_BLE_MESH](#)

Default value:

- 32 if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_DFU_SLOTS

Firmware image slot manager

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Device Firmware Update model](#)

Enable the DFU image slot manager, for managing firmware distribution slots for the Firmware Update Client model.

Default value:

- Yes (enabled) if [CONFIG_BLE_MESH_DFU_CLI](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_DFU_SLOT_CNT

Number of firmware image slots

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Device Firmware Update model](#) > [CONFIG_BLE_MESH_DFU_SLOTS](#)

This value defines the number of firmware slots the DFU image slot manager can keep simultaneously.

Range:

- from 1 to 32767 if [CONFIG_BLE_MESH_DFU_SLOTS](#) && [CONFIG_BLE_MESH](#)

Default value:

- 1 if [CONFIG_BLE_MESH_DFU_SLOTS](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_DFD_CLI

Support for Device Distribution Client model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Device Firmware Update model](#)

Enable the Device Distribution Client model

CONFIG_BLE_MESH_DFD_SRV

Support for Firmware Distribution Server model

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Device Firmware Update model](#)

Enable the Firmware Distribution Server model.

CONFIG_BLE_MESH_DFD_SRV_SLOT_MAX_SIZE

Largest DFU image that can be stored

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Device Firmware Update model](#) > [CONFIG_BLE_MESH_DFD_SRV](#)

This value defines the largest DFU image a single slot can store.

Range:

- from 0 to if [CONFIG_BLE_MESH_DFD_SRV](#) && [CONFIG_BLE_MESH](#)

Default value:

- if [CONFIG_BLE_MESH_DFD_SRV](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_DFD_SRV_SLOT_SPACE

Total DFU image storage space

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Device Firmware Update model](#) > [CONFIG_BLE_MESH_DFD_SRV](#)

This value defines the total storage space dedicated to storing DFU images on the Firmware Distribution Server.

Range:

- from 0 to 4294967295 if [CONFIG_BLE_MESH_DFD_SRV](#) && [CONFIG_BLE_MESH](#)

Default value:

- if [CONFIG_BLE_MESH_DFD_SRV](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_DFD_SRV_TARGETS_MAX

Maximum Target node count

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Device Firmware Update model](#) > [CONFIG_BLE_MESH_DFD_SRV](#)

This value defines the maximum number of Target nodes the Firmware Distribution Server can target simultaneously.

Range:

- from 1 to 65535 if [CONFIG_BLE_MESH_DFD_SRV](#) && [CONFIG_BLE_MESH](#)

Default value:

- 8 if [CONFIG_BLE_MESH_DFD_SRV](#) && [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_DFD_SRV_OOB_UPLOAD

Support for DFU image OOB upload

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [Support for BLE Mesh Client/Server models](#) > [Device Firmware Update model](#) > [CONFIG_BLE_MESH_DFD_SRV](#)

This enables support for OOB upload of firmware images for distribution. This makes several callbacks and use of the init macro `BLE_MESH_DFD_SRV_INIT_OOB` mandatory. See the API documentation for `bt_mesh_dfd_srv_cb` for details about the mandatory callbacks.

CONFIG_BLE_MESH_IV_UPDATE_TEST

Test the IV Update Procedure

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

This option removes the 96 hour limit of the IV Update Procedure and lets the state to be changed at any time. If IV Update test mode is going to be used, this option should be enabled.

Default value:

- No (disabled) if [CONFIG_BLE_MESH](#)

BLE Mesh specific test option Contains:

- [CONFIG_BLE_MESH_DEBUG](#)
- [CONFIG_BLE_MESH_SHELL](#)
- [CONFIG_BLE_MESH_BQB_TEST](#)
- [CONFIG_BLE_MESH_SELF_TEST](#)
- [CONFIG_BLE_MESH_TEST_AUTO_ENTER_NETWORK](#)
- [CONFIG_BLE_MESH_TEST_USE_WHITE_LIST](#)

CONFIG_BLE_MESH_SELF_TEST

Perform BLE Mesh self-tests

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#)

This option adds extra self-tests which are run every time BLE Mesh networking is initialized.

Default value:

- No (disabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_BQB_TEST

Enable BLE Mesh specific internal test

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#)

This option is used to enable some internal functions for auto-pts test.

Default value:

- No (disabled) if [CONFIG_BLE_MESH](#)

CONFIG_BLE_MESH_TEST_AUTO_ENTER_NETWORK

Unprovisioned device enters mesh network automatically

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#)

With this option enabled, an unprovisioned device can automatically enters mesh network using a specific test function without the provisioning procedure. And on the Provisioner side, a test function needs to be invoked to add the node information into the mesh stack.

Default value:

- Yes (enabled) if `CONFIG_BLE_MESH_SELF_TEST` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_TEST_USE_WHITE_LIST

Use white list to filter mesh advertising packets

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#)

With this option enabled, users can use white list to filter mesh advertising packets while scanning.

Default value:

- No (disabled) if `CONFIG_BLE_MESH_SELF_TEST` && `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_SHELL

Enable BLE Mesh shell

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#)

Activate shell module that provides BLE Mesh commands to the console.

Default value:

- No (disabled) if `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_DEBUG

Enable BLE Mesh debug logs

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#)

Enable debug logs for the BLE Mesh functionality.

Default value:

- No (disabled) if `CONFIG_BLE_MESH`

CONFIG_BLE_MESH_DEBUG_NET

Network layer debug

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#) > [CONFIG_BLE_MESH_DEBUG](#)

Enable Network layer debug logs for the BLE Mesh functionality.

CONFIG_BLE_MESH_DEBUG_TRANS

Transport layer debug

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#) > [CONFIG_BLE_MESH_DEBUG](#)

Enable Transport layer debug logs for the BLE Mesh functionality.

CONFIG_BLE_MESH_DEBUG_BEACON

Beacon debug

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#) > [CONFIG_BLE_MESH_DEBUG](#)

Enable Beacon-related debug logs for the BLE Mesh functionality.

CONFIG_BLE_MESH_DEBUG_CRYPTO

Crypto debug

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#) > [CONFIG_BLE_MESH_DEBUG](#)

Enable cryptographic debug logs for the BLE Mesh functionality.

CONFIG_BLE_MESH_DEBUG_PROV

Provisioning debug

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#) > [CONFIG_BLE_MESH_DEBUG](#)

Enable Provisioning debug logs for the BLE Mesh functionality.

CONFIG_BLE_MESH_DEBUG_ACCESS

Access layer debug

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#) > [CONFIG_BLE_MESH_DEBUG](#)

Enable Access layer debug logs for the BLE Mesh functionality.

CONFIG_BLE_MESH_DEBUG_MODEL

Foundation model debug

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#) > [CONFIG_BLE_MESH_DEBUG](#)

Enable Foundation Models debug logs for the BLE Mesh functionality.

CONFIG_BLE_MESH_DEBUG_ADV

Advertising debug

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#) > [CONFIG_BLE_MESH_DEBUG](#)

Enable advertising debug logs for the BLE Mesh functionality.

CONFIG_BLE_MESH_DEBUG_LOW_POWER

Low Power debug

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#) > [CONFIG_BLE_MESH_DEBUG](#)

Enable Low Power debug logs for the BLE Mesh functionality.

CONFIG_BLE_MESH_DEBUG_FRIEND

Friend debug

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#) > [CONFIG_BLE_MESH_DEBUG](#)

Enable Friend debug logs for the BLE Mesh functionality.

CONFIG_BLE_MESH_DEBUG_PROXY

Proxy debug

Found in: [Component config](#) > [CONFIG_BLE_MESH](#) > [BLE Mesh specific test option](#) > [CONFIG_BLE_MESH_DEBUG](#)

Enable Proxy protocol debug logs for the BLE Mesh functionality.

CONFIG_BLE_MESH_EXPERIMENTAL

Make BLE Mesh experimental features visible

Found in: [Component config](#) > [CONFIG_BLE_MESH](#)

Make BLE Mesh Experimental features visible. Experimental features list:

- CONFIG_BLE_MESH_NOT_RELAY_REPLAY_MSG
- CONFIG_BLE_MESH_USE_BLE_50
- CONFIG_BLE_MESH_PROXY_CLI_SRV_COEXIST

Default value:

- No (disabled) if [CONFIG_BLE_MESH](#)

Driver Configurations Contains:

- [Analog Comparator Configuration](#)
- [DAC Configuration](#)
- [GPIO Configuration](#)
- [GPTimer Configuration](#)
- [I2C Configuration](#)
- [I2S Configuration](#)
- [LEDC Configuration](#)
- [Legacy ADC Configuration](#)
- [MCPWM Configuration](#)
- [Parallel IO Configuration](#)
- [PCNT Configuration](#)
- [RMT Configuration](#)
- [Sigma Delta Modulator Configuration](#)
- [SPI Configuration](#)
- [Temperature sensor Configuration](#)
- [TWAI Configuration](#)
- [UART Configuration](#)
- [USB Serial/JTAG Configuration](#)

Legacy ADC Configuration Contains:

- [CONFIG_ADC_DISABLE_DAC](#)
- [Legacy ADC Calibration Configuration](#)
- [CONFIG_ADC_SKIP_LEGACY_CONFLICT_CHECK](#)
- [CONFIG_ADC_SUPPRESS_DEPRECATED_WARN](#)

CONFIG_ADC_DISABLE_DAC

Disable DAC when ADC2 is used on GPIO 25 and 26

Found in: [Component config](#) > [Driver Configurations](#) > [Legacy ADC Configuration](#)

If this is set, the ADC2 driver will disable the output of the DAC corresponding to the specified channel. This is the default value.

For testing, disable this option so that we can measure the output of DAC by internal ADC.

Default value:

- Yes (enabled) if SOC_DAC_SUPPORTED

CONFIG_ADC_SUPPRESS_DEPRECATED_WARN

Suppress legacy driver deprecated warning

Found in: [Component config](#) > [Driver Configurations](#) > [Legacy ADC Configuration](#)

Whether to suppress the deprecation warnings when using legacy adc driver (driver/adc.h). If you want to continue using the legacy driver, and don't want to see related deprecation warnings, you can enable this option.

Default value:

- No (disabled)

CONFIG_ADC_SKIP_LEGACY_CONFLICT_CHECK

Skip legacy conflict check

Found in: [Component config](#) > [Driver Configurations](#) > [Legacy ADC Configuration](#)

This configuration option allows the user to bypass the conflict check mechanism with legacy code.

Default value:

- No (disabled)

Legacy ADC Calibration Configuration

 Contains:

- [CONFIG_ADC_CALI_SUPPRESS_DEPRECATED_WARN](#)

CONFIG_ADC_CALI_SUPPRESS_DEPRECATED_WARN

Suppress legacy driver deprecated warning

Found in: [Component config](#) > [Driver Configurations](#) > [Legacy ADC Configuration](#) > [Legacy ADC Calibration Configuration](#)

Whether to suppress the deprecation warnings when using legacy adc calibration driver (esp_adc_cal.h). If you want to continue using the legacy driver, and don't want to see related deprecation warnings, you can enable this option.

Default value:

- No (disabled)

SPI Configuration

 Contains:

- [CONFIG_SPI_MASTER_ISR_IN_IRAM](#)
- [CONFIG_SPI_SLAVE_ISR_IN_IRAM](#)
- [CONFIG_SPI_MASTER_IN_IRAM](#)
- [CONFIG_SPI_SLAVE_IN_IRAM](#)

CONFIG_SPI_MASTER_IN_IRAM

Place transmitting functions of SPI master into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [SPI Configuration](#)

Normally only the ISR of SPI master is placed in the IRAM, so that it can work without the flash when interrupt is triggered. For other functions, there's some possibility that the flash cache miss when running inside and out of SPI functions, which may increase the interval of SPI transactions. Enable this to put `queue\_trans`, `get\_trans\_result` and `transmit` functions into the IRAM to avoid possible cache miss.

This configuration won't be available if `CONFIG_FREERTOS_PLACE_FUNCTIONS_INTO_FLASH` is enabled.

During unit test, this is enabled to measure the ideal case of api.

CONFIG_SPI_MASTER_ISR_IN_IRAM

Place SPI master ISR function into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [SPI Configuration](#)

Place the SPI master ISR in to IRAM to avoid possible cache miss.

Enabling this configuration is possible only when `HEAP_PLACE_FUNCTION_INTO_FLASH` is disabled since the spi master uses can allocate transactions buffers into DMA memory section using the heap component API that ipso facto has to be placed in IRAM.

Also you can forbid the ISR being disabled during flash writing access, by add `ESP_INTR_FLAG_IRAM` when initializing the driver.

CONFIG_SPI_SLAVE_IN_IRAM

Place transmitting functions of SPI slave into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [SPI Configuration](#)

Normally only the ISR of SPI slave is placed in the IRAM, so that it can work without the flash when interrupt is triggered. For other functions, there's some possibility that the flash cache miss when running inside and out of SPI functions, which may increase the interval of SPI transactions. Enable this to put `put_queue\_trans`, `get\_trans\_result` and `transmit` functions into the IRAM to avoid possible cache miss.

Default value:

- No (disabled)

CONFIG_SPI_SLAVE_ISR_IN_IRAM

Place SPI slave ISR function into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [SPI Configuration](#)

Place the SPI slave ISR in to IRAM to avoid possible cache miss.

Also you can forbid the ISR being disabled during flash writing access, by add `ESP_INTR_FLAG_IRAM` when initializing the driver.

Default value:

- Yes (enabled)

TWAI Configuration Contains:

- [CONFIG_TWAI_ISR_IN_IRAM](#)

CONFIG_TWAI_ISR_IN_IRAM

Place TWAI ISR function into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [TWAI Configuration](#)

Place the TWAI ISR in to IRAM. This will allow the ISR to avoid cache misses, and also be able to run whilst the cache is disabled (such as when writing to SPI Flash). Note that if this option is enabled: - Users should also set the `ESP_INTR_FLAG_IRAM` in the driver configuration structure when installing the driver (see docs for specifics). - Alert logging (i.e., setting of the `TWAI_ALERT_AND_LOG` flag) will have no effect.

Default value:

- No (disabled) if SOC_TWAI_SUPPORTED

Temperature sensor Configuration Contains:

- [CONFIG_TEMP_SENSOR_ENABLE_DEBUG_LOG](#)
- [CONFIG_TEMP_SENSOR_SKIP_LEGACY_CONFLICT_CHECK](#)
- [CONFIG_TEMP_SENSOR_SUPPRESS_DEPRECATED_WARN](#)
- [CONFIG_TEMP_SENSOR_ISR_IRAM_SAFE](#)

CONFIG_TEMP_SENSOR_SUPPRESS_DEPRECATED_WARN

Suppress legacy driver deprecated warning

Found in: [Component config](#) > [Driver Configurations](#) > [Temperature sensor Configuration](#)

Whether to suppress the deprecation warnings when using legacy temperature sensor driver (driver/temp_sensor.h). If you want to continue using the legacy driver, and don't want to see related deprecation warnings, you can enable this option.

Default value:

- No (disabled) if SOC_TEMP_SENSOR_SUPPORTED

CONFIG_TEMP_SENSOR_SKIP_LEGACY_CONFLICT_CHECK

Skip legacy conflict check

Found in: [Component config](#) > [Driver Configurations](#) > [Temperature sensor Configuration](#)

This configuration option allows the user to bypass the conflict check mechanism with legacy code.

Default value:

- No (disabled) if SOC_TEMP_SENSOR_SUPPORTED

CONFIG_TEMP_SENSOR_ENABLE_DEBUG_LOG

Enable debug log

Found in: [Component config](#) > [Driver Configurations](#) > [Temperature sensor Configuration](#)

Whether to enable the debug log message for temperature sensor driver. Note that, this option only controls the temperature sensor driver log, won't affect other drivers.

Default value:

- No (disabled) if SOC_TEMP_SENSOR_SUPPORTED

CONFIG_TEMP_SENSOR_ISR_IRAM_SAFE

Temperature sensor ISR Iram-Safe

Found in: [Component config](#) > [Driver Configurations](#) > [Temperature sensor Configuration](#)

Ensure the Temperature Sensor interrupt is Iram-Safe by allowing the interrupt handler to be executable when the cache is disabled (e.g. SPI Flash write).

Default value:

- No (disabled) if SOC_TEMPERATURE_SENSOR_INTR_SUPPORT && SOC_TEMP_SENSOR_SUPPORTED

UART Configuration Contains:

- [CONFIG_UART_ISR_IN_IRAM](#)

CONFIG_UART_ISR_IN_IRAM

Place UART ISR function into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [UART Configuration](#)

If this option is not selected, UART interrupt will be disabled for a long time and may cause data lost when doing spi flash operation.

GPIO Configuration

 Contains:

- [CONFIG_GPIO_CTRL_FUNC_IN_IRAM](#)

CONFIG_GPIO_CTRL_FUNC_IN_IRAM

Place GPIO control functions into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [GPIO Configuration](#)

Place GPIO control functions (like `intr_disable/set_level`) into IRAM, so that these functions can be IRAM-safe and able to be called in the other IRAM interrupt context.

Default value:

- No (disabled)

Sigma Delta Modulator Configuration

 Contains:

- [CONFIG_SDM_ENABLE_DEBUG_LOG](#)
- [CONFIG_SDM_CTRL_FUNC_IN_IRAM](#)
- [CONFIG_SDM_SKIP_LEGACY_CONFLICT_CHECK](#)
- [CONFIG_SDM_SUPPRESS_DEPRECATED_WARN](#)

CONFIG_SDM_CTRL_FUNC_IN_IRAM

Place SDM control functions into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [Sigma Delta Modulator Configuration](#)

Place SDM control functions (like `set_duty`) into IRAM, so that these functions can be IRAM-safe and able to be called in the other IRAM interrupt context. Enabling this option can improve driver performance as well.

Default value:

- No (disabled) if `SOC_SDM_SUPPORTED`

CONFIG_SDM_SUPPRESS_DEPRECATED_WARN

Suppress legacy driver deprecated warning

Found in: [Component config](#) > [Driver Configurations](#) > [Sigma Delta Modulator Configuration](#)

Whether to suppress the deprecation warnings when using legacy sigma delta driver. If you want to continue using the legacy driver, and don't want to see related deprecation warnings, you can enable this option.

Default value:

- No (disabled) if `SOC_SDM_SUPPORTED`

CONFIG_SDM_SKIP_LEGACY_CONFLICT_CHECK

Skip legacy conflict check

Found in: [Component config](#) > [Driver Configurations](#) > [Sigma Delta Modulator Configuration](#)

This configuration option allows the user to bypass the conflict check mechanism with legacy code.

Default value:

- No (disabled) if SOC_SDM_SUPPORTED

CONFIG_SDM_ENABLE_DEBUG_LOG

Enable debug log

Found in: [Component config](#) > [Driver Configurations](#) > [Sigma Delta Modulator Configuration](#)

Whether to enable the debug log message for SDM driver. Note that, this option only controls the SDM driver log, won't affect other drivers.

Default value:

- No (disabled) if SOC_SDM_SUPPORTED

Analog Comparator Configuration Contains:

- [CONFIG_ANA_CMPR_ISR_IRAM_SAFE](#)
- [CONFIG_ANA_CMPR_ENABLE_DEBUG_LOG](#)
- [CONFIG_ANA_CMPR_CTRL_FUNC_IN_IRAM](#)

CONFIG_ANA_CMPR_ISR_IRAM_SAFE

Analog comparator ISR IRAM-Safe

Found in: [Component config](#) > [Driver Configurations](#) > [Analog Comparator Configuration](#)

Ensure the Analog Comparator interrupt is IRAM-Safe by allowing the interrupt handler to be executable when the cache is disabled (e.g. SPI Flash write).

Default value:

- No (disabled)

CONFIG_ANA_CMPR_CTRL_FUNC_IN_IRAM

Place Analog Comparator control functions into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [Analog Comparator Configuration](#)

Place Analog Comparator control functions (like `ana_cmpr_set_internal_reference`) into IRAM, so that these functions can be IRAM-safe and able to be called in an IRAM interrupt context. Enabling this option can improve driver performance as well.

Default value:

- No (disabled)

CONFIG_ANA_CMPR_ENABLE_DEBUG_LOG

Enable debug log

Found in: [Component config](#) > [Driver Configurations](#) > [Analog Comparator Configuration](#)

Whether to enable the debug log message for Analog Comparator driver. Note that, this option only controls the Analog Comparator driver log, won't affect other drivers.

Default value:

- No (disabled)

GPTimer Configuration Contains:

- [*CONFIG_GPTIMER_ENABLE_DEBUG_LOG*](#)
- [*CONFIG_GPTIMER_ISR_IRAM_SAFE*](#)
- [*CONFIG_GPTIMER_CTRL_FUNC_IN_IRAM*](#)
- [*CONFIG_GPTIMER_ISR_HANDLER_IN_IRAM*](#)
- [*CONFIG_GPTIMER_SKIP_LEGACY_CONFLICT_CHECK*](#)
- [*CONFIG_GPTIMER_SUPPRESS_DEPRECATED_WARN*](#)

CONFIG_GPTIMER_ISR_HANDLER_IN_IRAM

Place GPTimer ISR handler into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [GPTimer Configuration](#)

Place GPTimer ISR handler into IRAM for better performance and fewer cache misses.

Default value:

- Yes (enabled)

CONFIG_GPTIMER_CTRL_FUNC_IN_IRAM

Place GPTimer control functions into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [GPTimer Configuration](#)

Place GPTimer control functions (like start/stop) into IRAM, so that these functions can be IRAM-safe and able to be called in the other IRAM interrupt context. Enabling this option can improve driver performance as well.

Default value:

- No (disabled)

CONFIG_GPTIMER_ISR_IRAM_SAFE

GPTimer ISR IRAM-Safe

Found in: [Component config](#) > [Driver Configurations](#) > [GPTimer Configuration](#)

Ensure the GPTimer interrupt is IRAM-Safe by allowing the interrupt handler to be executable when the cache is disabled (e.g. SPI Flash write).

Default value:

- No (disabled)

CONFIG_GPTIMER_SUPPRESS_DEPRECATED_WARN

Suppress legacy driver deprecated warning

Found in: [Component config](#) > [Driver Configurations](#) > [GPTimer Configuration](#)

Whether to suppress the deprecation warnings when using legacy timer group driver (driver/timer.h). If you want to continue using the legacy driver, and don't want to see related deprecation warnings, you can enable this option.

Default value:

- No (disabled)

CONFIG_GPTIMER_SKIP_LEGACY_CONFLICT_CHECK

Skip legacy conflict check

Found in: [Component config](#) > [Driver Configurations](#) > [GPTimer Configuration](#)

This configuration option allows the user to bypass the conflict check mechanism with legacy code.

Default value:

- No (disabled)

CONFIG_GPTIMER_ENABLE_DEBUG_LOG

Enable debug log

Found in: [Component config](#) > [Driver Configurations](#) > [GPTimer Configuration](#)

Whether to enable the debug log message for GPTimer driver. Note that, this option only controls the GPTimer driver log, won't affect other drivers.

Default value:

- No (disabled)

PCNT Configuration Contains:

- [CONFIG_PCNT_ENABLE_DEBUG_LOG](#)
- [CONFIG_PCNT_ISR_IRAM_SAFE](#)
- [CONFIG_PCNT_CTRL_FUNC_IN_IRAM](#)
- [CONFIG_PCNT_SKIP_LEGACY_CONFLICT_CHECK](#)
- [CONFIG_PCNT_SUPPRESS_DEPRECATED_WARN](#)

CONFIG_PCNT_CTRL_FUNC_IN_IRAM

Place PCNT control functions into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [PCNT Configuration](#)

Place PCNT control functions (like start/stop) into IRAM, so that these functions can be IRAM-safe and able to be called in the other IRAM interrupt context. Enabling this option can improve driver performance as well.

Default value:

- No (disabled)

CONFIG_PCNT_ISR_IRAM_SAFE

PCNT ISR IRAM-Safe

Found in: [Component config](#) > [Driver Configurations](#) > [PCNT Configuration](#)

Ensure the PCNT interrupt is IRAM-Safe by allowing the interrupt handler to be executable when the cache is disabled (e.g. SPI Flash write).

Default value:

- No (disabled)

CONFIG_PCNT_SUPPRESS_DEPRECATED_WARN

Suppress legacy driver deprecated warning

Found in: [Component config](#) > [Driver Configurations](#) > [PCNT Configuration](#)

Whether to suppress the deprecation warnings when using legacy PCNT driver (driver/pcnt.h). If you want to continue using the legacy driver, and don't want to see related deprecation warnings, you can enable this option.

Default value:

- No (disabled)

CONFIG_PCNT_SKIP_LEGACY_CONFLICT_CHECK

Skip legacy conflict check

Found in: [Component config](#) > [Driver Configurations](#) > [PCNT Configuration](#)

This configuration option allows the user to bypass the conflict check mechanism with legacy code.

Default value:

- No (disabled)

CONFIG_PCNT_ENABLE_DEBUG_LOG

Enable debug log

Found in: [Component config](#) > [Driver Configurations](#) > [PCNT Configuration](#)

Whether to enable the debug log message for PCNT driver. Note that, this option only controls the PCNT driver log, won't affect other drivers.

Default value:

- No (disabled)

RMT Configuration

Contains:

- [CONFIG_RMT_ENABLE_DEBUG_LOG](#)
- [CONFIG_RMT_RECV_FUNC_IN_IRAM](#)
- [CONFIG_RMT_ISR_IRAM_SAFE](#)
- [CONFIG_RMT_SKIP_LEGACY_CONFLICT_CHECK](#)
- [CONFIG_RMT_SUPPRESS_DEPRECATED_WARN](#)

CONFIG_RMT_ISR_IRAM_SAFE

RMT ISR IRAM-Safe

Found in: [Component config](#) > [Driver Configurations](#) > [RMT Configuration](#)

Ensure the RMT interrupt is IRAM-Safe by allowing the interrupt handler to be executable when the cache is disabled (e.g. SPI Flash write).

Default value:

- No (disabled)

CONFIG_RMT_RECV_FUNC_IN_IRAM

Place RMT receive function into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [RMT Configuration](#)

Place RMT receive function into IRAM, so that the receive function can be IRAM-safe and able to be called when the flash cache is disabled. Enabling this option can improve driver performance as well.

Default value:

- No (disabled)

CONFIG_RMT_SUPPRESS_DEPRECATED_WARN

Suppress legacy driver deprecated warning

Found in: [Component config](#) > [Driver Configurations](#) > [RMT Configuration](#)

Whether to suppress the deprecation warnings when using legacy rmt driver (driver/rmt.h). If you want to continue using the legacy driver, and don't want to see related deprecation warnings, you can enable this option.

Default value:

- No (disabled)

CONFIG_RMT_SKIP_LEGACY_CONFLICT_CHECK

Skip legacy conflict check

Found in: [Component config](#) > [Driver Configurations](#) > [RMT Configuration](#)

This configuration option allows the user to bypass the conflict check mechanism with legacy code.

Default value:

- No (disabled)

CONFIG_RMT_ENABLE_DEBUG_LOG

Enable debug log

Found in: [Component config](#) > [Driver Configurations](#) > [RMT Configuration](#)

Whether to enable the debug log message for RMT driver. Note that, this option only controls the RMT driver log, won't affect other drivers.

Default value:

- No (disabled)

MCPWM Configuration Contains:

- [CONFIG_MCPWM_ENABLE_DEBUG_LOG](#)
- [CONFIG_MCPWM_CTRL_FUNC_IN_IRAM](#)
- [CONFIG_MCPWM_ISR_IRAM_SAFE](#)
- [CONFIG_MCPWM_SKIP_LEGACY_CONFLICT_CHECK](#)
- [CONFIG_MCPWM_SUPPRESS_DEPRECATED_WARN](#)

CONFIG_MCPWM_ISR_IRAM_SAFE

Place MCPWM ISR function into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [MCPWM Configuration](#)

This will ensure the MCPWM interrupt handle is IRAM-Safe, allow to avoid flash cache misses, and also be able to run whilst the cache is disabled. (e.g. SPI Flash write)

Default value:

- No (disabled)

CONFIG_MCPWM_CTRL_FUNC_IN_IRAM

Place MCPWM control functions into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [MCPWM Configuration](#)

Place MCPWM control functions (like `set_compare_value`) into IRAM, so that these functions can be IRAM-safe and able to be called in the other IRAM interrupt context. Enabling this option can improve driver performance as well.

Default value:

- No (disabled)

CONFIG_MCPWM_SUPPRESS_DEPRECATED_WARN

Suppress legacy driver deprecated warning

Found in: [Component config](#) > [Driver Configurations](#) > [MCPWM Configuration](#)

Whether to suppress the deprecation warnings when using legacy MCPWM driver (driver/mcpwm.h). If you want to continue using the legacy driver, and don't want to see related deprecation warnings, you can enable this option.

Default value:

- No (disabled)

CONFIG_MCPWM_ENABLE_DEBUG_LOG

Enable debug log

Found in: [Component config](#) > [Driver Configurations](#) > [MCPWM Configuration](#)

Whether to enable the debug log message for MCPWM driver. Note that, this option only controls the MCPWM driver log, won't affect other drivers.

Default value:

- No (disabled)

CONFIG_MCPWM_SKIP_LEGACY_CONFLICT_CHECK

Skip legacy conflict check

Found in: [Component config](#) > [Driver Configurations](#) > [MCPWM Configuration](#)

This configuration option allows the user to bypass the conflict check mechanism with legacy code.

Default value:

- No (disabled)

I2S Configuration

 Contains:

- [CONFIG_I2S_ENABLE_DEBUG_LOG](#)
- [CONFIG_I2S_ISR_IRAM_SAFE](#)
- [CONFIG_I2S_SKIP_LEGACY_CONFLICT_CHECK](#)
- [CONFIG_I2S_SUPPRESS_DEPRECATED_WARN](#)

CONFIG_I2S_ISR_IRAM_SAFE

I2S ISR IRAM-Safe

Found in: [Component config](#) > [Driver Configurations](#) > [I2S Configuration](#)

Ensure the I2S interrupt is IRAM-Safe by allowing the interrupt handler to be executable when the cache is disabled (e.g. SPI Flash write).

Default value:

- No (disabled)

CONFIG_I2S_SUPPRESS_DEPRECATED_WARN

Suppress legacy driver deprecated warning

Found in: [Component config](#) > [Driver Configurations](#) > [I2S Configuration](#)

Enable this option will suppress the deprecation warnings of using APIs in legacy I2S driver.

Default value:

- No (disabled)

CONFIG_I2S_SKIP_LEGACY_CONFLICT_CHECK

Skip legacy conflict check

Found in: [Component config](#) > [Driver Configurations](#) > [I2S Configuration](#)

This configuration option allows the user to bypass the conflict check mechanism with legacy code.

Default value:

- No (disabled)

CONFIG_I2S_ENABLE_DEBUG_LOG

Enable I2S debug log

Found in: [Component config](#) > [Driver Configurations](#) > [I2S Configuration](#)

Whether to enable the debug log message for I2S driver. Note that, this option only controls the I2S driver log, will not affect other drivers.

Default value:

- No (disabled)

DAC Configuration

 Contains:

- [CONFIG_DAC_DMA_AUTO_16BIT_ALIGN](#)
- [CONFIG_DAC_ISR_IRAM_SAFE](#)
- [CONFIG_DAC_ENABLE_DEBUG_LOG](#)
- [CONFIG_DAC_CTRL_FUNC_IN_IRAM](#)
- [CONFIG_DAC_SKIP_LEGACY_CONFLICT_CHECK](#)
- [CONFIG_DAC_SUPPRESS_DEPRECATED_WARN](#)

CONFIG_DAC_CTRL_FUNC_IN_IRAM

Place DAC control functions into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [DAC Configuration](#)

Place DAC control functions (e.g. 'dac_one_shot_output_voltage') into IRAM, so that this function can be IRAM-safe and able to be called in the other IRAM interrupt context. Enabling this option can improve driver performance as well.

Default value:

- No (disabled) if SOC_DAC_SUPPORTED

CONFIG_DAC_ISR_IRAM_SAFE

DAC ISR IRAM-Safe

Found in: [Component config](#) > [Driver Configurations](#) > [DAC Configuration](#)

Ensure the DAC interrupt is IRAM-Safe by allowing the interrupt handler to be executable when the cache is disabled (e.g. SPI Flash write).

Default value:

- No (disabled) if SOC_DAC_SUPPORTED

CONFIG_DAC_SUPPRESS_DEPRECATED_WARN

Suppress legacy driver deprecated warning

Found in: [Component config](#) > [Driver Configurations](#) > [DAC Configuration](#)

Whether to suppress the deprecation warnings when using legacy DAC driver (driver/dac.h). If you want to continue using the legacy driver, and don't want to see related deprecation warnings, you can enable this option.

Default value:

- No (disabled) if SOC_DAC_SUPPORTED

CONFIG_DAC_ENABLE_DEBUG_LOG

Enable debug log

Found in: [Component config](#) > [Driver Configurations](#) > [DAC Configuration](#)

Whether to enable the debug log message for DAC driver. Note that, this option only controls the DAC driver log, won't affect other drivers.

Default value:

- No (disabled) if SOC_DAC_SUPPORTED

CONFIG_DAC_DMA_AUTO_16BIT_ALIGN

Align the continuous data to 16 bit automatically

Found in: [Component config](#) > [Driver Configurations](#) > [DAC Configuration](#)

Whether to left shift the continuous data to align every bytes to 16 bits in the driver. On ESP32, although the DAC resolution is only 8 bits, the hardware requires 16 bits data in continuous mode. By enabling this option, the driver will left shift 8 bits for the input data automatically. Only disable this option when you decide to do this step by yourself. Note that the driver will allocate a new piece of memory to save the converted data.

Default value:

- Yes (enabled) if SOC_DAC_DMA_16BIT_ALIGN && SOC_DAC_SUPPORTED

CONFIG_DAC_SKIP_LEGACY_CONFLICT_CHECK

Skip legacy conflict check

Found in: [Component config](#) > [Driver Configurations](#) > [DAC Configuration](#)

This configuration option allows the user to bypass the conflict check mechanism with legacy code.

Default value:

- No (disabled) if SOC_DAC_SUPPORTED

USB Serial/JTAG Configuration

 Contains:

- [CONFIG_USJ_NO_AUTO_LS_ON_CONNECTION](#)

CONFIG_USJ_NO_AUTO_LS_ON_CONNECTION

Don't enter the automatic light sleep when USB Serial/JTAG port is connected

Found in: [Component config](#) > [Driver Configurations](#) > [USB Serial/JTAG Configuration](#)

If enabled, the chip will constantly monitor the connection status of the USB Serial/JTAG port. As long as the USB Serial/JTAG is connected, a ESP_PM_NO_LIGHT_SLEEP power management lock will be acquired to prevent the system from entering light sleep. This option can be useful if serial monitoring is needed via USB Serial/JTAG while power management is enabled, as the USB Serial/JTAG cannot work under light sleep and after waking up from light sleep. Note. This option can only control the automatic Light-Sleep behavior. If esp_light_sleep_start() is called manually from the program, enabling this option will not prevent light sleep entry even if the USB Serial/JTAG is in use.

Parallel IO Configuration Contains:

- [CONFIG_PARLIO_ENABLE_DEBUG_LOG](#)
- [CONFIG_PARLIO_ISR_IRAM_SAFE](#)

CONFIG_PARLIO_ENABLE_DEBUG_LOG

Enable debug log

Found in: [Component config](#) > [Driver Configurations](#) > [Parallel IO Configuration](#)

Whether to enable the debug log message for parallel IO driver. Note that, this option only controls the parallel IO driver log, won't affect other drivers.

Default value:

- No (disabled)

CONFIG_PARLIO_ISR_IRAM_SAFE

Parallel IO ISR IRAM-Safe

Found in: [Component config](#) > [Driver Configurations](#) > [Parallel IO Configuration](#)

Ensure the Parallel IO interrupt is IRAM-Safe by allowing the interrupt handler to be executable when the cache is disabled (e.g. SPI Flash write).

Default value:

- No (disabled)

LEDC Configuration Contains:

- [CONFIG_LEDC_CTRL_FUNC_IN_IRAM](#)

CONFIG_LEDC_CTRL_FUNC_IN_IRAM

Place LEDC control functions into IRAM

Found in: [Component config](#) > [Driver Configurations](#) > [LEDC Configuration](#)

Place LEDC control functions (ledc_update_duty and ledc_stop) into IRAM, so that these functions can be IRAM-safe and able to be called in an IRAM context. Enabling this option can improve driver performance as well.

Default value:

- No (disabled)

I2C Configuration Contains:

- [CONFIG_I2C_ENABLE_DEBUG_LOG](#)
- [CONFIG_I2C_ISR_IRAM_SAFE](#)
- [CONFIG_I2C_SKIP_LEGACY_CONFLICT_CHECK](#)

CONFIG_I2C_ISR_IRAM_SAFE

I2C ISR IRAM-Safe

Found in: [Component config](#) > [Driver Configurations](#) > [I2C Configuration](#)

Ensure the I2C interrupt is IRAM-Safe by allowing the interrupt handler to be executable when the cache is disabled (e.g. SPI Flash write). note: This cannot be used in the I2C legacy driver.

Default value:

- No (disabled)

CONFIG_I2C_ENABLE_DEBUG_LOG

Enable I2C debug log

Found in: [Component config](#) > [Driver Configurations](#) > [I2C Configuration](#)

Whether to enable the debug log message for I2C driver. Note that this option only controls the I2C driver log, will not affect other drivers.

note: This cannot be used in the I2C legacy driver.

Default value:

- No (disabled)

CONFIG_I2C_SKIP_LEGACY_CONFLICT_CHECK

Skip legacy conflict check

Found in: [Component config](#) > [Driver Configurations](#) > [I2C Configuration](#)

This configuration option allows the user to bypass the conflict check mechanism with legacy code.

Default value:

- No (disabled)

eFuse Bit Manager Contains:

- [CONFIG_EFUSE_VIRTUAL](#)
- [CONFIG_EFUSE_CUSTOM_TABLE](#)

CONFIG_EFUSE_CUSTOM_TABLE

Use custom eFuse table

Found in: [Component config](#) > [eFuse Bit Manager](#)

Allows to generate a structure for eFuse from the CSV file.

Default value:

- No (disabled)

CONFIG_EFUSE_CUSTOM_TABLE_FILENAME

Custom eFuse CSV file

Found in: [Component config](#) > [eFuse Bit Manager](#) > [CONFIG_EFUSE_CUSTOM_TABLE](#)

Name of the custom eFuse CSV filename. This path is evaluated relative to the project root directory.

Default value:

- "main/esp_efuse_custom_table.csv" if [CONFIG_EFUSE_CUSTOM_TABLE](#)

CONFIG_EFUSE_VIRTUAL

Simulate eFuse operations in RAM

Found in: [Component config](#) > [eFuse Bit Manager](#)

If "n" - No virtual mode. All eFuse operations are real and use eFuse registers. If "y" - The virtual mode is enabled and all eFuse operations (read and write) are redirected to RAM instead of eFuse registers, all permanent changes (via eFuse) are disabled. Log output will state changes that would be applied, but they will not be.

If it is "y", then SECURE_FLASH_ENCRYPTION_MODE_RELEASE cannot be used. Because the EFUSE VIRT mode is for testing only.

During startup, the eFuses are copied into RAM. This mode is useful for fast tests.

Default value:

- No (disabled)

CONFIG_EFUSE_VIRTUAL_KEEP_IN_FLASH

Keep eFuses in flash

Found in: [Component config](#) > [eFuse Bit Manager](#) > [CONFIG_EFUSE_VIRTUAL](#)

In addition to the "Simulate eFuse operations in RAM" option, this option just adds a feature to keep eFuses after reboots in flash memory. To use this mode the partition_table should have the *efuse* partition. partition.csv: "efuse_em, data, efuse, , 0x2000,"

During startup, the eFuses are copied from flash or, in case if flash is empty, from real eFuse to RAM and then update flash. This mode is useful when need to keep changes after reboot (testing secure_boot and flash_encryption).

CONFIG_EFUSE_VIRTUAL_LOG_ALL_WRITES

Log all virtual writes

Found in: [Component config](#) > [eFuse Bit Manager](#) > [CONFIG_EFUSE_VIRTUAL](#)

If enabled, log efuse burns. This shows changes that would be made.

ESP-TLS Contains:

- [CONFIG_ESP_TLS_INSECURE](#)
- [CONFIG_ESP_TLS_LIBRARY_CHOOSE](#)
- [CONFIG_ESP_TLS_CLIENT_SESSION_TICKETS](#)
- [CONFIG_ESP_DEBUG_WOLFSSL](#)
- [CONFIG_ESP_TLS_SERVER](#)
- [CONFIG_ESP_TLS_PSK_VERIFICATION](#)
- [CONFIG_ESP_WOLFSSL_SMALL_CERT_VERIFY](#)
- [CONFIG_ESP_TLS_USE_DS_PERIPHERAL](#)

CONFIG_ESP_TLS_LIBRARY_CHOOSE

Choose SSL/TLS library for ESP-TLS (See help for more Info)

Found in: [Component config](#) > [ESP-TLS](#)

The ESP-TLS APIs support multiple backend TLS libraries. Currently mbedTLS and WolfSSL are supported. Different TLS libraries may support different features and have different resource usage. Consult the ESP-TLS documentation in ESP-IDF Programming guide for more details.

Available options:

- mbedTLS ([CONFIG_ESP_TLS_USING_MBEDTLS](#))
- wolfSSL (License info in wolfSSL directory [README](#)) ([CONFIG_ESP_TLS_USING_WOLFSSL](#))

CONFIG_ESP_TLS_USE_DS_PERIPHERAL

Use Digital Signature (DS) Peripheral with ESP-TLS

Found in: [Component config](#) > [ESP-TLS](#)

Enable use of the Digital Signature Peripheral for ESP-TLS. The DS peripheral can only be used when it is appropriately configured for TLS. Consult the ESP-TLS documentation in ESP-IDF Programming Guide for more details.

Default value:

- Yes (enabled)

CONFIG_ESP_TLS_CLIENT_SESSION_TICKETS

Enable client session tickets

Found in: [Component config](#) > [ESP-TLS](#)

Enable session ticket support as specified in RFC5077.

CONFIG_ESP_TLS_SERVER

Enable ESP-TLS Server

Found in: [Component config](#) > [ESP-TLS](#)

Enable support for creating server side SSL/TLS session, available for mbedTLS as well as wolfSSL TLS library.

CONFIG_ESP_TLS_SERVER_SESSION_TICKETS

Enable server session tickets

Found in: [Component config](#) > [ESP-TLS](#) > [CONFIG_ESP_TLS_SERVER](#)

Enable session ticket support as specified in RFC5077

CONFIG_ESP_TLS_SERVER_SESSION_TICKET_TIMEOUT

Server session ticket timeout in seconds

Found in: [Component config](#) > [ESP-TLS](#) > [CONFIG_ESP_TLS_SERVER](#) > [CONFIG_ESP_TLS_SERVER_SESSION_TICKETS](#)

Sets the session ticket timeout used in the tls server.

Default value:

- 86400 if [CONFIG_ESP_TLS_SERVER_SESSION_TICKETS](#)

CONFIG_ESP_TLS_SERVER_CERT_SELECT_HOOK

Certificate selection hook

Found in: [Component config](#) > [ESP-TLS](#) > [CONFIG_ESP_TLS_SERVER](#)

Ability to configure and use a certificate selection callback during server handshake, to select a certificate to present to the client based on the TLS extensions supplied in the client hello (alpn, sni, etc).

CONFIG_ESP_TLS_SERVER_MIN_AUTH_MODE_OPTIONAL

ESP-TLS Server: Set minimum Certificate Verification mode to Optional

Found in: [Component config](#) > [ESP-TLS](#) > [CONFIG_ESP_TLS_SERVER](#)

When this option is enabled, the peer (here, the client) certificate is checked by the server, however the handshake continues even if verification failed. By default, the peer certificate is not checked and ignored by the server.

`mbedtls_ssl_get_verify_result()` can be called after the handshake is complete to retrieve status of verification.

CONFIG_ESP_TLS_PSK_VERIFICATION

Enable PSK verification

Found in: [Component config](#) > [ESP-TLS](#)

Enable support for pre shared key ciphers, supported for both mbedTLS as well as wolfSSL TLS library.

CONFIG_ESP_TLS_INSECURE

Allow potentially insecure options

Found in: [Component config](#) > [ESP-TLS](#)

You can enable some potentially insecure options. These options should only be used for testing purposes. Only enable these options if you are very sure.

CONFIG_ESP_TLS_SKIP_SERVER_CERT_VERIFY

Skip server certificate verification by default (WARNING: ONLY FOR TESTING PURPOSE, READ HELP)

Found in: [Component config](#) > [ESP-TLS](#) > [CONFIG_ESP_TLS_INSECURE](#)

After enabling this option the esp-tls client will skip the server certificate verification by default. Note that this option will only modify the default behaviour of esp-tls client regarding server cert verification. The default behaviour should only be applicable when no other option regarding the server cert verification is opted in the esp-tls config (e.g. `crt_bundle_attach`, `use_global_ca_store` etc.). WARNING : Enabling this option comes with a potential risk of establishing a TLS connection with a server which has a fake identity, provided that the server certificate is not provided either through API or other mechanism like `ca_store` etc.

CONFIG_ESP_WOLFSSL_SMALL_CERT_VERIFY

Enable SMALL_CERT_VERIFY

Found in: [Component config](#) > [ESP-TLS](#)

Enables server verification with Intermediate CA cert, does not authenticate full chain of trust upto the root CA cert (After Enabling this option client only needs to have Intermediate CA certificate of the server to authenticate server, root CA cert is not necessary).

Default value:

- Yes (enabled) if [CONFIG_ESP_TLS_USING_WOLFSSL](#)

CONFIG_ESP_DEBUG_WOLFSSL

Enable debug logs for wolfSSL

Found in: [Component config](#) > [ESP-TLS](#)

Enable detailed debug prints for wolfSSL SSL library.

ADC and ADC Calibration

 Contains:

- [ADC Calibration Configurations](#)
- [CONFIG_ADC_CONTINUOUS_ISR_IRAM_SAFE](#)
- [CONFIG_ADC_DISABLE_DAC_OUTPUT](#)
- [CONFIG_ADC_ONESHOT_CTRL_FUNC_IN_IRAM](#)

CONFIG_ADC_ONESHOT_CTRL_FUNC_IN_IRAM

Place ISR version ADC oneshot mode read function into IRAM

Found in: [Component config](#) > [ADC and ADC Calibration](#)

Place ISR version ADC oneshot mode read function into IRAM.

Default value:

- No (disabled)

CONFIG_ADC_CONTINUOUS_ISR_IRAM_SAFE

ADC continuous mode driver ISR IRAM-Safe

Found in: [Component config](#) > [ADC and ADC Calibration](#)

Ensure the ADC continuous mode ISR is IRAM-Safe. When enabled, the ISR handler will be available when the cache is disabled.

Default value:

- No (disabled) if SOC_ADC_DMA_SUPPORTED

ADC Calibration Configurations

CONFIG_ADC_DISABLE_DAC_OUTPUT

Disable DAC when ADC2 is in use

Found in: [Component config](#) > [ADC and ADC Calibration](#)

By default, this is set. The ADC oneshot driver will disable the output of the corresponding DAC channels: ESP32: IO25 and IO26 ESP32S2: IO17 and IO18

Disable this option so as to measure the output of DAC by internal ADC, for test usage.

Default value:

- Yes (enabled) if SOC_DAC_SUPPORTED

Wireless Coexistence Contains:

- [CONFIG_ESP_COEX_EXTERNAL_COEXIST_ENABLE](#)
- [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- [CONFIG_ESP_COEX_SW_COEXIST_ENABLE](#)
- [CONFIG_ESP_COEX_POWER_MANAGEMENT](#)

CONFIG_ESP_COEX_SW_COEXIST_ENABLE

Software controls WiFi/Bluetooth coexistence

Found in: [Component config](#) > [Wireless Coexistence](#)

If enabled, WiFi & Bluetooth coexistence is controlled by software rather than hardware. Recommended for heavy traffic scenarios. Both coexistence configuration options are automatically managed, no user intervention is required. If only Bluetooth is used, it is recommended to disable this option to reduce binary file size.

Default value:

- Yes (enabled) if [CONFIG_IEEE802154_ENABLED](#) && [CONFIG_BT_ENABLED](#)

CONFIG_ESP_COEX_EXTERNAL_COEXIST_ENABLE

External Coexistence

Found in: [Component config](#) > [Wireless Coexistence](#)

If enabled, HW External coexistence arbitration is managed by GPIO pins. It can support three types of wired combinations so far which are 1-wired/2-wired/3-wired. User can select GPIO pins in application code with configure interfaces.

This function depends on BT-off because currently we do not support external coex and internal coex simultaneously.

CONFIG_ESP_COEX_POWER_MANAGEMENT

Support power management under coexistence

Found in: [Component config](#) > [Wireless Coexistence](#)

If enabled, coexist power management will be enabled.

Default value:

- No (disabled) if [CONFIG_ESP_COEX_SW_COEXIST_ENABLE](#)

CONFIG_ESP_COEX_GPIO_DEBUG

GPIO debugging for coexistence

Found in: [Component config](#) > [Wireless Coexistence](#)

Support coexistence GPIO debugging

CONFIG_ESP_COEX_GPIO_DEBUG_DIAG

Debugging Diagram

Found in: [Component config](#) > [Wireless Coexistence](#) > [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Select type of debugging diagram

Available options:

- General (CONFIG_ESP_COEX_GPIO_DEBUG_DIAG_GENERAL)
- Wi-Fi (CONFIG_ESP_COEX_GPIO_DEBUG_DIAG_WIFI)

CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT

Max number of debugging GPIOs

Found in: [Component config](#) > [Wireless Coexistence](#) > [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Range:

- from 0 to 12 if [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Default value:

- 12 if [CONFIG_ESP_COEX_GPIO_DEBUG](#)

CONFIG_ESP_COEX_GPIO_DEBUG_IO_IDX0

Actual IO num for Debug IO ID0

Found in: [Component config](#) > [Wireless Coexistence](#) > [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Range:

- from 0 to 56 if `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 0` && `CONFIG_ESP_COEX_GPIO_DEBUG`

Default value:

- 2 if `IDF_TARGET_ESP32C5` && `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 0` && `CONFIG_ESP_COEX_GPIO_DEBUG`
- 4 if `IDF_TARGET_ESP32C61` && `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 0` && `CONFIG_ESP_COEX_GPIO_DEBUG`
- 1 if `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 0` && `CONFIG_ESP_COEX_GPIO_DEBUG`

CONFIG_ESP_COEX_GPIO_DEBUG_IO_IDX1

Actual IO num for Debug IO ID1

Found in: *Component config > Wireless Coexistence > CONFIG_ESP_COEX_GPIO_DEBUG*

Range:

- from 0 to 56 if `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 1` && `CONFIG_ESP_COEX_GPIO_DEBUG`

Default value:

- 3 if `IDF_TARGET_ESP32C5` && `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 1` && `CONFIG_ESP_COEX_GPIO_DEBUG`
- 5 if `IDF_TARGET_ESP32C61` && `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 1` && `CONFIG_ESP_COEX_GPIO_DEBUG`
- 2 if `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 1` && `CONFIG_ESP_COEX_GPIO_DEBUG`

CONFIG_ESP_COEX_GPIO_DEBUG_IO_IDX2

Actual IO num for Debug IO ID2

Found in: *Component config > Wireless Coexistence > CONFIG_ESP_COEX_GPIO_DEBUG*

Range:

- from 0 to 56 if `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 2` && `CONFIG_ESP_COEX_GPIO_DEBUG`

Default value:

- 4 if `IDF_TARGET_ESP32C5` && `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 2` && `CONFIG_ESP_COEX_GPIO_DEBUG`
- 6 if `IDF_TARGET_ESP32C61` && `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 2` && `CONFIG_ESP_COEX_GPIO_DEBUG`
- 3 if `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 2` && `CONFIG_ESP_COEX_GPIO_DEBUG`

CONFIG_ESP_COEX_GPIO_DEBUG_IO_IDX3

Actual IO num for Debug IO ID3

Found in: *Component config > Wireless Coexistence > CONFIG_ESP_COEX_GPIO_DEBUG*

Range:

- from 0 to 56 if `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 3` && `CONFIG_ESP_COEX_GPIO_DEBUG`

Default value:

- 5 if `IDF_TARGET_ESP32C5` && `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 3` && `CONFIG_ESP_COEX_GPIO_DEBUG`
- 7 if `IDF_TARGET_ESP32C61` && `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 3` && `CONFIG_ESP_COEX_GPIO_DEBUG`
- 4 if `CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT > 3` && `CONFIG_ESP_COEX_GPIO_DEBUG`

CONFIG_ESP_COEX_GPIO_DEBUG_IO_IDX4

Actual IO num for Debug IO ID4

Found in: [Component config](#) > [Wireless Coexistence](#) > [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Range:

- from 0 to 56 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 4 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Default value:

- 27 if IDF_TARGET_ESP32C5 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 4 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 0 if IDF_TARGET_ESP32C61 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 4 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 5 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 4 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

CONFIG_ESP_COEX_GPIO_DEBUG_IO_IDX5

Actual IO num for Debug IO ID5

Found in: [Component config](#) > [Wireless Coexistence](#) > [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Range:

- from 0 to 56 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 5 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Default value:

- 6 if IDF_TARGET_ESP32C5 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 5 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 1 if IDF_TARGET_ESP32C61 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 5 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 6 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 5 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

CONFIG_ESP_COEX_GPIO_DEBUG_IO_IDX6

Actual IO num for Debug IO ID6

Found in: [Component config](#) > [Wireless Coexistence](#) > [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Range:

- from 0 to 56 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 6 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Default value:

- 7 if IDF_TARGET_ESP32C5 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 6 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 8 if IDF_TARGET_ESP32C61 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 6 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 7 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 6 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

CONFIG_ESP_COEX_GPIO_DEBUG_IO_IDX7

Actual IO num for Debug IO ID7

Found in: [Component config](#) > [Wireless Coexistence](#) > [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Range:

- from 0 to 56 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 7 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Default value:

- 26 if IDF_TARGET_ESP32C5 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 7 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 2 if IDF_TARGET_ESP32C61 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 7 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 8 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 7 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

CONFIG_ESP_COEX_GPIO_DEBUG_IO_IDX8

Actual IO num for Debug IO ID8

Found in: [Component config](#) > [Wireless Coexistence](#) > [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Range:

- from 0 to 56 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 8 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Default value:

- 24 if IDF_TARGET_ESP32C5 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 8 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 3 if IDF_TARGET_ESP32C61 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 8 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 9 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 8 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

CONFIG_ESP_COEX_GPIO_DEBUG_IO_IDX9

Actual IO num for Debug IO ID9

Found in: [Component config](#) > [Wireless Coexistence](#) > [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Range:

- from 0 to 56 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 9 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Default value:

- 23 if IDF_TARGET_ESP32C5 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 9 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 9 if IDF_TARGET_ESP32C61 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 9 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 10 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 9 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

CONFIG_ESP_COEX_GPIO_DEBUG_IO_IDX10

Actual IO num for Debug IO ID10

Found in: [Component config](#) > [Wireless Coexistence](#) > [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Range:

- from 0 to 56 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 10 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Default value:

- 10 if IDF_TARGET_ESP32C5 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 10 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 13 if IDF_TARGET_ESP32C61 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 10 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 11 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 10 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

CONFIG_ESP_COEX_GPIO_DEBUG_IO_IDX11

Actual IO num for Debug IO ID11

Found in: [Component config](#) > [Wireless Coexistence](#) > [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Range:

- from 0 to 56 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 11 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Default value:

- 9 if IDF_TARGET_ESP32C5 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 11 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 12 if IDF_TARGET_ESP32C61 && [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 11 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)
- 12 if [CONFIG_ESP_COEX_GPIO_DEBUG_IO_COUNT](#) > 11 && [CONFIG_ESP_COEX_GPIO_DEBUG](#)

Common ESP-related Contains:

- [CONFIG_ESP_ERR_TO_NAME_LOOKUP](#)

CONFIG_ESP_ERR_TO_NAME_LOOKUP

Enable lookup of error code strings

Found in: [Component config](#) > [Common ESP-related](#)

Functions `esp_err_to_name()` and `esp_err_to_name_r()` return string representations of error codes from a pre-generated lookup table. This option can be used to turn off the use of the look-up table in order to save memory but this comes at the price of sacrificing distinguishable (meaningful) output string representations.

Default value:

- Yes (enabled)

Ethernet Contains:

- [CONFIG_ETH_TRANSMIT_MUTEX](#)
- [CONFIG_ETH_USE_OPENETH](#)
- [CONFIG_ETH_USE_SPI_ETHERNET](#)

CONFIG_ETH_USE_SPI_ETHERNET

Support SPI to Ethernet Module

Found in: [Component config](#) > [Ethernet](#)

ESP-IDF can also support some SPI-Ethernet modules.

Default value:

- Yes (enabled)

Contains:

- [CONFIG_ETH_SPI_ETHERNET_DM9051](#)
- [CONFIG_ETH_SPI_ETHERNET_KSZ8851SNL](#)
- [CONFIG_ETH_SPI_ETHERNET_W5500](#)

CONFIG_ETH_SPI_ETHERNET_DM9051

Use DM9051

Found in: [Component config](#) > [Ethernet](#) > [CONFIG_ETH_USE_SPI_ETHERNET](#)

DM9051 is a fast Ethernet controller with an SPI interface. It's also integrated with a 10/100M PHY and MAC. Select this to enable DM9051 driver.

CONFIG_ETH_SPI_ETHERNET_W5500

Use W5500 (MAC RAW)

Found in: [Component config](#) > [Ethernet](#) > [CONFIG_ETH_USE_SPI_ETHERNET](#)

W5500 is a HW TCP/IP embedded Ethernet controller. TCP/IP stack, 10/100 Ethernet MAC and PHY are embedded in a single chip. However the driver in ESP-IDF only enables the RAW MAC mode, making it compatible with the software TCP/IP stack. Say yes to enable W5500 driver.

CONFIG_ETH_SPI_ETHERNET_KSZ8851SNL

Use KSZ8851SNL

Found in: [Component config](#) > [Ethernet](#) > [CONFIG_ETH_USE_SPI_ETHERNET](#)

The KSZ8851SNL is a single-chip Fast Ethernet controller consisting of a 10/100 physical layer transceiver (PHY), a MAC, and a Serial Peripheral Interface (SPI). Select this to enable KSZ8851SNL driver.

CONFIG_ETH_USE_OPENETH

Support OpenCores Ethernet MAC (for use with QEMU)

Found in: [Component config](#) > [Ethernet](#)

OpenCores Ethernet MAC driver can be used when an ESP-IDF application is executed in QEMU. This driver is not supported when running on a real chip.

Default value:

- No (disabled)

Contains:

- [CONFIG_ETH_OPENETH_DMA_RX_BUFFER_NUM](#)
- [CONFIG_ETH_OPENETH_DMA_TX_BUFFER_NUM](#)

CONFIG_ETH_OPENETH_DMA_RX_BUFFER_NUM

Number of Ethernet DMA Rx buffers

Found in: [Component config](#) > [Ethernet](#) > [CONFIG_ETH_USE_OPENETH](#)

Number of DMA receive buffers, each buffer is 1600 bytes.

Range:

- from 1 to 64 if [CONFIG_ETH_USE_OPENETH](#)

Default value:

- 4 if [CONFIG_ETH_USE_OPENETH](#)

CONFIG_ETH_OPENETH_DMA_TX_BUFFER_NUM

Number of Ethernet DMA Tx buffers

Found in: [Component config](#) > [Ethernet](#) > [CONFIG_ETH_USE_OPENETH](#)

Number of DMA transmit buffers, each buffer is 1600 bytes.

Range:

- from 1 to 64 if [CONFIG_ETH_USE_OPENETH](#)

Default value:

- 1 if [CONFIG_ETH_USE_OPENETH](#)

CONFIG_ETH_TRANSMIT_MUTEX

Enable Transmit Mutex

Found in: [Component config](#) > [Ethernet](#)

Prevents multiple accesses when Ethernet interface is used as shared resource and multiple functionalities might try to access it at a time.

Default value:

- No (disabled)

Event Loop Library Contains:

- [CONFIG_ESP_EVENT_LOOP_PROFILING](#)
- [CONFIG_ESP_EVENT_POST_FROM_ISR](#)

CONFIG_ESP_EVENT_LOOP_PROFILING

Enable event loop profiling

Found in: [Component config](#) > [Event Loop Library](#)

Enables collections of statistics in the event loop library such as the number of events posted to/received by an event loop, number of callbacks involved, number of events dropped to a full event loop queue, run time of event handlers, and number of times/run time of each event handler.

Default value:

- No (disabled)

CONFIG_ESP_EVENT_POST_FROM_ISR

Support posting events from ISRs

Found in: [Component config](#) > [Event Loop Library](#)

Enable posting events from interrupt handlers.

Default value:

- Yes (enabled)

CONFIG_ESP_EVENT_POST_FROM_IRAM_ISR

Support posting events from ISRs placed in IRAM

Found in: [Component config](#) > [Event Loop Library](#) > [CONFIG_ESP_EVENT_POST_FROM_ISR](#)

Enable posting events from interrupt handlers placed in IRAM. Enabling this option places API functions `esp_event_post` and `esp_event_post_to` in IRAM.

Default value:

- Yes (enabled)

GDB Stub Contains:

- [*CONFIG_ESP_GDBSTUB_SUPPORT_TASKS*](#)
- [*CONFIG_ESP_SYSTEM_GDBSTUB_RUNTIME*](#)

CONFIG_ESP_SYSTEM_GDBSTUB_RUNTIME

GDBStub at runtime

Found in: [*Component config*](#) > [*GDB Stub*](#)

Enable builtin GDBStub. This allows to debug the target device using serial port: - Run 'idf.py monitor'. - Wait for the device to initialize. - Press Ctrl+C to interrupt the execution and enter GDB attached to your device for debugging. NOTE: all UART input will be handled by GDBStub.

CONFIG_ESP_GDBSTUB_SUPPORT_TASKS

Enable listing FreeRTOS tasks through GDB Stub

Found in: [*Component config*](#) > [*GDB Stub*](#)

If enabled, GDBStub can supply the list of FreeRTOS tasks to GDB. Thread list can be queried from GDB using 'info threads' command. Note that if GDB task lists were corrupted, this feature may not work. If GDBStub fails, try disabling this feature.

Default value:

- Yes (enabled)

CONFIG_ESP_GDBSTUB_MAX_TASKS

Maximum number of tasks supported by GDB Stub

Found in: [*Component config*](#) > [*GDB Stub*](#) > [*CONFIG_ESP_GDBSTUB_SUPPORT_TASKS*](#)

Set the number of tasks which GDB Stub will support.

Default value:

- 32

ESP HID Contains:

- [*CONFIG_ESPHID_TASK_SIZE_BLE*](#)
- [*CONFIG_ESPHID_TASK_SIZE_BT*](#)

CONFIG_ESPHID_TASK_SIZE_BT

Task stack size for ESP HID BR/EDR

Found in: [*Component config*](#) > [*ESP HID*](#)

This is the stack size for the BT HID task. Default is 2048 bytes.

Range:

- from 2048 to 10240

Default value:

- 2048

CONFIG_ESPHID_TASK_SIZE_BLE

Task stack size for ESP HID BLE

Found in: [Component config](#) > [ESP HID](#)

This is the stack size for the BLE HID task. Default is 4096 bytes.

Range:

- from 2048 to 10240

Default value:

- 4096

ESP HTTP client Contains:

- [CONFIG_ESP_HTTP_CLIENT_ENABLE_BASIC_AUTH](#)
- [CONFIG_ESP_HTTP_CLIENT_ENABLE_DIGEST_AUTH](#)
- [CONFIG_ESP_HTTP_CLIENT_ENABLE_HTTPS](#)

CONFIG_ESP_HTTP_CLIENT_ENABLE_HTTPS

Enable https

Found in: [Component config](#) > [ESP HTTP client](#)

This option will enable https protocol by linking esp-tls library and initializing SSL transport

Default value:

- Yes (enabled)

CONFIG_ESP_HTTP_CLIENT_ENABLE_BASIC_AUTH

Enable HTTP Basic Authentication

Found in: [Component config](#) > [ESP HTTP client](#)

This option will enable HTTP Basic Authentication. It is disabled by default as Basic auth uses unencrypted encoding, so it introduces a vulnerability when not using TLS

Default value:

- No (disabled)

CONFIG_ESP_HTTP_CLIENT_ENABLE_DIGEST_AUTH

Enable HTTP Digest Authentication

Found in: [Component config](#) > [ESP HTTP client](#)

This option will enable HTTP Digest Authentication. It is enabled by default, but use of this configuration is not recommended as the password can be derived from the exchange, so it introduces a vulnerability when not using TLS

Default value:

- No (disabled)

HTTP Server Contains:

- [CONFIG_HTTPD_QUEUE_WORK_BLOCKING](#)
- [CONFIG_HTTPD_PURGE_BUF_LEN](#)
- [CONFIG_HTTPD_LOG_PURGE_DATA](#)
- [CONFIG_HTTPD_MAX_REQ_HDR_LEN](#)
- [CONFIG_HTTPD_MAX_URI_LEN](#)
- [CONFIG_HTTPD_ERR_RESP_NO_DELAY](#)
- [CONFIG_HTTPD_WS_SUPPORT](#)

CONFIG_HTTPD_MAX_REQ_HDR_LEN

Max HTTP Request Header Length

Found in: [Component config](#) > [HTTP Server](#)

This sets the maximum supported size of headers section in HTTP request packet to be processed by the server

Default value:

- 512

CONFIG_HTTPD_MAX_URI_LEN

Max HTTP URI Length

Found in: [Component config](#) > [HTTP Server](#)

This sets the maximum supported size of HTTP request URI to be processed by the server

Default value:

- 512

CONFIG_HTTPD_ERR_RESP_NO_DELAY

Use TCP_NODELAY socket option when sending HTTP error responses

Found in: [Component config](#) > [HTTP Server](#)

Using TCP_NODELAY socket option ensures that HTTP error response reaches the client before the underlying socket is closed. Please note that turning this off may cause multiple test failures

Default value:

- Yes (enabled)

CONFIG_HTTPD_PURGE_BUF_LEN

Length of temporary buffer for purging data

Found in: [Component config](#) > [HTTP Server](#)

This sets the size of the temporary buffer used to receive and discard any remaining data that is received from the HTTP client in the request, but not processed as part of the server HTTP request handler.

If the remaining data is larger than the available buffer size, the buffer will be filled in multiple iterations. The buffer should be small enough to fit on the stack, but large enough to avoid excessive iterations.

Default value:

- 32

CONFIG_HTTPD_LOG_PURGE_DATA

Log purged content data at Debug level

Found in: [Component config](#) > [HTTP Server](#)

Enabling this will log discarded binary HTTP request data at Debug level. For large content data this may not be desirable as it will clutter the log.

Default value:

- No (disabled)

CONFIG_HTTPD_WS_SUPPORT

WebSocket server support

Found in: [Component config](#) > [HTTP Server](#)

This sets the WebSocket server support.

Default value:

- No (disabled)

CONFIG_HTTPD_QUEUE_WORK_BLOCKING

httpd_queue_work as blocking API

Found in: [Component config](#) > [HTTP Server](#)

This makes httpd_queue_work() API to wait until a message space is available on UDP control socket. It internally uses a counting semaphore with count set to `LWIP_UDP_RECVMBOX_SIZE` to achieve this. This config will slightly change API behavior to block until message gets delivered on control socket.

ESP HTTPS OTA Contains:

- [CONFIG_ESP_HTTPS_OTA_ALLOW_HTTP](#)
- [CONFIG_ESP_HTTPS_OTA_DECRYPT_CB](#)

CONFIG_ESP_HTTPS_OTA_DECRYPT_CB

Provide decryption callback

Found in: [Component config](#) > [ESP HTTPS OTA](#)

Exposes an additional callback whereby firmware data could be decrypted before being processed by OTA update component. This can help to integrate external encryption related format and removal of such encapsulation layer from firmware image.

Default value:

- No (disabled)

CONFIG_ESP_HTTPS_OTA_ALLOW_HTTP

Allow HTTP for OTA (WARNING: ONLY FOR TESTING PURPOSE, READ HELP)

Found in: [Component config](#) > [ESP HTTPS OTA](#)

It is highly recommended to keep HTTPS (along with server certificate validation) enabled. Enabling this option comes with potential risk of: - Non-encrypted communication channel with server - Accepting firmware upgrade image from server with fake identity

Default value:

- No (disabled)

ESP HTTPS server Contains:

- [CONFIG_ESP_HTTPS_SERVER_ENABLE](#)

CONFIG_ESP_HTTPS_SERVER_ENABLE

Enable ESP_HTTPS_SERVER component

Found in: [Component config](#) > [ESP HTTPS server](#)

Enable ESP HTTPS server component

Hardware Settings Contains:

- [CONFIG_ESP_ENABLE_PVT](#)
- [Chip revision](#)
- [Crypto DPA Protection](#)
- [CONFIG_ESP_ECDSA_ENABLE_P192_CURVE](#)
- [ETM Configuration](#)
- [CONFIG_ESP_CRYPTOP_FORCE_ECC_CONSTANT_TIME_POINT_MUL](#)
- [GDMA Configuration](#)
- [MAC Config](#)
- [Main XTAL Config](#)
- [Peripheral Control](#)
- [RTC Clock Config](#)
- [Sleep Config](#)

Chip revision Contains:

- [CONFIG_ESP_REV_NEW_CHIP_TEST](#)
- [CONFIG_ESP_EFUSE_BLOCK_REV_MIN_FULL](#)
- [CONFIG_ESP32P4_REV_MIN](#)

CONFIG_ESP32P4_REV_MIN

Minimum Supported ESP32-P4 Revision

Found in: [Component config](#) > [Hardware Settings](#) > [Chip revision](#)

Required minimum chip revision. ESP-IDF will check for it and reject to boot if the chip revision fails the check. This ensures the chip used will have some modifications (features, or bugfixes).

The complied binary will only support chips above this revision, this will also help to reduce binary size.

Available options:

- Rev v0.0 (CONFIG_ESP32P4_REV_MIN_0)

CONFIG_ESP_EFUSE_BLOCK_REV_MIN_FULL

Minimum Supported ESP32-P4 eFuse Block Revision

Found in: [Component config](#) > [Hardware Settings](#) > [Chip revision](#)

Required minimum eFuse Block revision. ESP-IDF will check it at the 2nd bootloader stage whether the current image can work correctly for this eFuse Block revision. So that to avoid running an incompatible image on a SoC that contains breaking change in the eFuse Block. If you want to update this value to run the image that not compatible with the current eFuse Block revision, please contact to Espressif's business team for details: <https://www.espressif.com.cn/en/contact-us/sales-questions>

Default value:

- 0

CONFIG_ESP_REV_NEW_CHIP_TEST

Internal test mode

Found in: [Component config](#) > [Hardware Settings](#) > [Chip revision](#)

For internal chip testing, a small number of new versions chips didn't update the version field in eFuse, you can enable this option to force the software recognize the chip version based on the rev selected in menuconfig.

Default value:

- No (disabled)

MAC Config Contains:

- [*CONFIG_ESP_MAC_USE_CUSTOM_MAC_AS_BASE_MAC*](#)
- [*CONFIG_ESP32P4_UNIVERSAL_MAC_ADDRESSES*](#)

CONFIG_ESP32P4_UNIVERSAL_MAC_ADDRESSES

Number of universally administered (by IEEE) MAC address

Found in: [Component config](#) > [Hardware Settings](#) > [MAC Config](#)

TODO IDF-6514

Available options:

- Two (CONFIG_ESP32P4_UNIVERSAL_MAC_ADDRESSES_TWO)

CONFIG_ESP_MAC_USE_CUSTOM_MAC_AS_BASE_MAC

Enable using custom mac as base mac

Found in: [Component config](#) > [Hardware Settings](#) > [MAC Config](#)

When this configuration is enabled, the user can invoke *esp_read_mac* to obtain the desired type of MAC using a custom MAC as the base MAC.

Default value:

- No (disabled)

Sleep Config Contains:

- [*CONFIG_ESP_SLEEP_GPIO_ENABLE_INTERNAL_RESISTORS*](#)
- [*CONFIG_ESP_SLEEP_CACHE_SAFE_ASSERTION*](#)
- [*CONFIG_ESP_SLEEP_EVENT_CALLBACKS*](#)
- [*CONFIG_ESP_SLEEP_DEBUG*](#)
- [*CONFIG_ESP_SLEEP_WAIT_FLASH_READY_EXTRA_DELAY*](#)
- [*CONFIG_ESP_SLEEP_GPIO_RESET_WORKAROUND*](#)
- [*CONFIG_ESP_SLEEP_POWER_DOWN_FLASH*](#)
- [*CONFIG_ESP_SLEEP_MSPI_NEED_ALL_IO_PU*](#)
- [*CONFIG_ESP_SLEEP_FLASH_LEAKAGE_WORKAROUND*](#)
- [*CONFIG_ESP_SLEEP_PSRAM_LEAKAGE_WORKAROUND*](#)

CONFIG_ESP_SLEEP_POWER_DOWN_FLASH

Power down flash in light sleep when there is no SPIRAM

Found in: [Component config](#) > [Hardware Settings](#) > [Sleep Config](#)

If enabled, chip will try to power down flash as part of *esp_light_sleep_start()*, which costs more time when chip wakes up. Can only be enabled if there is no SPIRAM configured.

This option will power down flash under a strict but relatively safe condition. Also, it is possible to power down flash under a relaxed condition by using *esp_sleep_pd_config()* to set *ESP_PD_DOMAIN_VDDSDIO* to *ESP_PD_OPTION_OFF*. It should be noted that there is a risk in powering down flash, you can refer *ESP-IDF Programming Guide/API Reference/System API/Sleep Modes/Power-down of Flash* for more details.

CONFIG_ESP_SLEEP_FLASH_LEAKAGE_WORKAROUND

Pull-up Flash CS pin in light sleep

Found in: [Component config](#) > [Hardware Settings](#) > [Sleep Config](#)

All IOs will be set to isolate(floating) state by default during sleep. Since the power supply of SPI Flash is not lost during lightsleep, if its CS pin is recognized as low level(selected state) in the floating state, there will be a large current leakage, and the data in Flash may be corrupted by random signals on other SPI pins. Select this option will set the CS pin of Flash to PULL-UP state during sleep, but this will increase the sleep current about 10 uA. If you are developing with esp32xx modules, you must select this option, but if you are developing with chips, you can also pull up the CS pin of SPI Flash in the external circuit to save power consumption caused by internal pull-up during sleep. (!!! Don't deselect this option if you don't have external SPI Flash CS pin pullups.)

CONFIG_ESP_SLEEP_PSRAM_LEAKAGE_WORKAROUND

Pull-up PSRAM CS pin in light sleep

Found in: [Component config](#) > [Hardware Settings](#) > [Sleep Config](#)

All IOs will be set to isolate(floating) state by default during sleep. Since the power supply of PSRAM is not lost during lightsleep, if its CS pin is recognized as low level(selected state) in the floating state, there will be a large current leakage, and the data in PSRAM may be corrupted by random signals on other SPI pins. Select this option will set the CS pin of PSRAM to PULL-UP state during sleep, but this will increase the sleep current about 10 uA. If you are developing with esp32xx modules, you must select this option, but if you are developing with chips, you can also pull up the CS pin of PSRAM in the external circuit to save power consumption caused by internal pull-up during sleep. (!!! Don't deselect this option if you don't have external PSRAM CS pin pullups.)

Default value:

- Yes (enabled) if [CONFIG_SPIRAM](#)

CONFIG_ESP_SLEEP_MSPI_NEED_ALL_IO_PU

Pull-up all SPI pins in light sleep

Found in: [Component config](#) > [Hardware Settings](#) > [Sleep Config](#)

To reduce leakage current, some types of SPI Flash/RAM only need to pull up the CS pin during light sleep. But there are also some kinds of SPI Flash/RAM that need to pull up all pins. It depends on the SPI Flash/RAM chip used.

CONFIG_ESP_SLEEP_GPIO_RESET_WORKAROUND

light sleep GPIO reset workaround

Found in: [Component config](#) > [Hardware Settings](#) > [Sleep Config](#)

All existing chips except esp32 and esp32s2 will reset on wake-up if a GPIO receives a small electrostatic pulse during light sleep, with specific conditions.

- GPIO needs to be configured as input-mode only
- The pin receives a small electrostatic pulse, and reset occurs when the pulse voltage is higher than 6 V

For GPIO set to input mode only, it is not a good practice to leave it open/floating, The hardware design needs to controlled it with determined supply or ground voltage is necessary.

This option provides a software workaround for this issue. Configure to isolate all GPIO pins in sleep state.

Default value:

- Yes (enabled)

CONFIG_ESP_SLEEP_WAIT_FLASH_READY_EXTRA_DELAY

Extra delay (in us) after flash powerdown sleep wakeup to wait flash ready

Found in: [Component config](#) > [Hardware Settings](#) > [Sleep Config](#)

When the chip exits sleep, the CPU and the flash chip are powered on at the same time. CPU will run rom code (deepsleep) or ram code (lightsleep) first, and then load or execute code from flash.

Some flash chips need sufficient time to pass between power on and first read operation. By default, without any extra delay, this time is approximately 900us, although some flash chip types need more than that.

(!!! Please adjust this value according to the Data Sheet of SPI Flash used in your project.) In Flash Data Sheet, the parameters that define the Flash ready timing after power-up (minimum time from Vcc(min) to CS active) usually named tVSL in ELECTRICAL CHARACTERISTICS chapter, and the configuration value here should be: ESP_SLEEP_WAIT_FLASH_READY_EXTRA_DELAY = tVSL - 900

For esp32 and esp32s3, the default extra delay is set to 2000us. When optimizing startup time for applications which require it, this value may be reduced.

If you are seeing "flash read err, 1000" message printed to the console after deep sleep reset on esp32, or triggered RTC_WDT/LP_WDT after lightsleep wakeup, try increasing this value. (For esp32, the delay will be executed in both deep sleep and light sleep wake up flow. For chips after esp32, the delay will be executed only in light sleep flow, the delay controlled by the EFUSE_FLASH_TPUW in ROM will be executed in deepsleep wake up flow.)

Range:

- from 0 to 5000

Default value:

- 0

CONFIG_ESP_SLEEP_CACHE_SAFE_ASSERTION

Check the cache safety of the sleep wakeup code in sleep process

Found in: [Component config](#) > [Hardware Settings](#) > [Sleep Config](#)

Enabling it will check the cache safety of the code before the flash power is ready after light sleep wakeup, and check PM_SLP_IRAM_OPT related code cache safety. This option is only for code quality inspection. Enabling it will increase the time overhead of entering and exiting sleep. It is not recommended to enable it in the release version.

Default value:

- No (disabled)

CONFIG_ESP_SLEEP_DEBUG

esp sleep debug

Found in: [Component config](#) > [Hardware Settings](#) > [Sleep Config](#)

Enable esp sleep debug.

Default value:

- No (disabled)

CONFIG_ESP_SLEEP_GPIO_ENABLE_INTERNAL_RESISTORS

Allow to enable internal pull-up/downs for the Deep-Sleep wakeup IOs

Found in: [Component config](#) > [Hardware Settings](#) > [Sleep Config](#)

When using rtc gpio wakeup source during deepsleep without external pull-up/downs, you may want to make use of the internal ones.

Default value:

- Yes (enabled)

CONFIG_ESP_SLEEP_EVENT_CALLBACKS

Enable registration of sleep event callbacks

Found in: [Component config](#) > [Hardware Settings](#) > [Sleep Config](#)

If enabled, it allows user to register sleep event callbacks. It is primarily designed for internal developers and customers can use PM_LIGHT_SLEEP_CALLBACKS as an alternative.

NOTE: These callbacks are executed from the IDLE task context hence you cannot have any blocking calls in your callbacks.

NOTE: Enabling these callbacks may change sleep duration calculations based on time spent in callback and hence it is highly recommended to keep them as short as possible.

Default value:

- No (disabled) if [CONFIG_FREERTOS_USE_TICKLESS_IDLE](#)

RTC Clock Config Contains:

- [CONFIG_RTC_CLK_CAL_CYCLES](#)
- [CONFIG_RTC_CLK_SRC](#)

CONFIG_RTC_CLK_SRC

RTC clock source

Found in: [Component config](#) > [Hardware Settings](#) > [RTC Clock Config](#)

Choose which clock is used as RTC clock source.

Available options:

- Internal 150 kHz RC oscillator ([CONFIG_RTC_CLK_SRC_INT_RC](#))
- External 32kHz crystal ([CONFIG_RTC_CLK_SRC_EXT_CRYS](#))
- External 32kHz oscillator at 32K_XP pin ([CONFIG_RTC_CLK_SRC_EXT_OSC](#))

CONFIG_RTC_CLK_CAL_CYCLES

Number of cycles for RTC_SLOW_CLK calibration

Found in: [Component config](#) > [Hardware Settings](#) > [RTC Clock Config](#)

When the startup code initializes RTC_SLOW_CLK, it can perform calibration by comparing the RTC_SLOW_CLK frequency with main XTAL frequency. This option sets the number of RTC_SLOW_CLK cycles measured by the calibration routine. Higher numbers increase calibration precision, which may be important for applications which spend a lot of time in deep sleep. Lower numbers reduce startup time.

When this option is set to 0, clock calibration will not be performed at startup, and approximate clock frequencies will be assumed:

- 150000 Hz if internal RC oscillator is used as clock source. For this use value 1024.
- **32768 Hz if the 32k crystal oscillator is used. For this use value 3000 or more.** In case more value will help improve the definition of the launch of the crystal. If the crystal could not start, it will be switched to internal RC.

Range:

- from 0 to 8190 if [CONFIG_RTC_CLK_SRC_EXT_CRYS](#)
- from 0 to 32766

Default value:

- 3000 if [CONFIG_RTC_CLK_SRC_EXT_CRYST](#) || [CONFIG_RTC_CLK_SRC_EXT_OSC](#)
- 1024

Peripheral Control Contains:

- [CONFIG_PERIPH_CTRL_FUNC_IN_IRAM](#)

CONFIG_PERIPH_CTRL_FUNC_IN_IRAM

Place peripheral control functions into IRAM

Found in: [Component config](#) > [Hardware Settings](#) > [Peripheral Control](#)

Place peripheral control functions (e.g. `periph_module_reset`) into IRAM, so that these functions can be IRAM-safe and able to be called in the other IRAM interrupt context.

Default value:

- No (disabled)

ETM Configuration Contains:

- [CONFIG_ETM_ENABLE_DEBUG_LOG](#)

CONFIG_ETM_ENABLE_DEBUG_LOG

Enable debug log

Found in: [Component config](#) > [Hardware Settings](#) > [ETM Configuration](#)

Whether to enable the debug log message for ETM core driver. Note that, this option only controls the ETM related driver log, won't affect other drivers.

Default value:

- No (disabled)

GDMA Configuration Contains:

- [CONFIG_GDMA_ENABLE_DEBUG_LOG](#)
- [CONFIG_GDMA_ISR_IRAM_SAFE](#)
- [CONFIG_GDMA_CTRL_FUNC_IN_IRAM](#)

CONFIG_GDMA_CTRL_FUNC_IN_IRAM

Place GDMA control functions into IRAM

Found in: [Component config](#) > [Hardware Settings](#) > [GDMA Configuration](#)

Place GDMA control functions (like start/stop/append/reset) into IRAM, so that these functions can be IRAM-safe and able to be called in the other IRAM interrupt context. Enabling this option can improve driver performance as well.

Default value:

- No (disabled)

CONFIG_GDMA_ISR_IRAM_SAFE

GDMA ISR IRAM-Safe

Found in: [Component config](#) > [Hardware Settings](#) > [GDMA Configuration](#)

This will ensure the GDMA interrupt handler is IRAM-Safe, allow to avoid flash cache misses, and also be able to run whilst the cache is disabled. (e.g. SPI Flash write).

Default value:

- No (disabled)

CONFIG_GDMA_ENABLE_DEBUG_LOG

Enable debug log

Found in: [Component config](#) > [Hardware Settings](#) > [GDMA Configuration](#)

Whether to enable the debug log message for GDMA driver. Note that, this option only controls the GDMA driver log, won't affect other drivers.

Default value:

- No (disabled)

Main XTAL Config Contains:

- [CONFIG_XTAL_FREQ_SEL](#)

CONFIG_XTAL_FREQ_SEL

Main XTAL frequency

Found in: [Component config](#) > [Hardware Settings](#) > [Main XTAL Config](#)

This option selects the operating frequency of the XTAL (crystal) clock used to drive the ESP target. The selected value MUST reflect the frequency of the given hardware.

Note: The XTAL_FREQ_AUTO option allows the ESP target to automatically estimating XTAL clock's operating frequency. However, this feature is only supported on the ESP32. The ESP32 uses the internal 8MHz as a reference when estimating. Due to the internal oscillator's frequency being temperature dependent, usage of the XTAL_FREQ_AUTO is not recommended in applications that operate in high ambient temperatures or use high-temperature qualified chips and modules.

Available options:

- 24 MHz (CONFIG_XTAL_FREQ_24)
- 26 MHz (CONFIG_XTAL_FREQ_26)
- 32 MHz (CONFIG_XTAL_FREQ_32)
- 40 MHz (CONFIG_XTAL_FREQ_40)
- Autodetect (CONFIG_XTAL_FREQ_AUTO)

Crypto DPA Protection Contains:

- [CONFIG_ESP_CRYPTO_DPA_PROTECTION_AT_STARTUP](#)

CONFIG_ESP_CRYPTO_DPA_PROTECTION_AT_STARTUP

Enable crypto DPA protection at startup

Found in: [Component config](#) > [Hardware Settings](#) > [Crypto DPA Protection](#)

This config controls the DPA (Differential Power Analysis) protection knob for the crypto peripherals. DPA protection dynamically adjusts the clock frequency of the crypto peripheral. DPA protection helps

to make it difficult to perform SCA attacks on the crypto peripherals. However, there is also associated performance impact based on the security level set. Please refer to the TRM for more details.

Default value:

- Yes (enabled) if SOC_CRYPTODPA_PROTECTION_SUPPORTED

CONFIG_ESP_CRYPTODPA_PROTECTION_LEVEL

DPA protection level

Found in: [Component config](#) > [Hardware Settings](#) > [Crypto DPA Protection](#) > [CONFIG_ESP_CRYPTODPA_PROTECTION_AT_STARTUP](#)

Configure the DPA protection security level

Available options:

- Security level low (CONFIG_ESP_CRYPTODPA_PROTECTION_LEVEL_LOW)
- Security level medium (CONFIG_ESP_CRYPTODPA_PROTECTION_LEVEL_MEDIUM)
- Security level high (CONFIG_ESP_CRYPTODPA_PROTECTION_LEVEL_HIGH)

CONFIG_ESP_CRYPTOFORCE_ECC_CONSTANT_TIME_POINT_MUL

Forcefully enable ECC constant time point multiplication operations

Found in: [Component config](#) > [Hardware Settings](#)

If enabled, the app startup code will burn the ECC_FORCE_CONST_TIME efuse bit to force the ECC peripheral to always perform constant time point multiplication operations, irrespective of the ECC_MULT_SECURITY_MODE status bit that is present in the ECC_MULT_CONF_REG register. By default, ESP-IDF configures the ECC peripheral to perform constant time point multiplication operations, so enabling this config would provide security enhancement only in the cases when trusted boot is not enabled and the attacker tries carrying out non-constant time point multiplication operations by changing the default ESP-IDF configurations. Performing constant time operations protect the ECC multiplication operations from timing attacks.

Default value:

- No (disabled) if SOC_ECC_CONSTANT_TIME_POINT_MUL

CONFIG_ESP_ECDSA_ENABLE_P192_CURVE

Enable ECDSA 192-curve operations

Found in: [Component config](#) > [Hardware Settings](#)

By default, only the 256-bit curve operations are allowed. If this configuration is enabled, it will set the eFuse to allow ECDSA operations using both the 192-bit and 256-bit curves

Default value:

- No (disabled) if SOC_ECDSA_P192_CURVE_DEFAULT_DISABLED

CONFIG_ESP_ENABLE_PVT

Auto adjust hp & lp voltage using pvt function (MUST ENABLE FOR MP)

Found in: [Component config](#) > [Hardware Settings](#)

If enabled, hp & lp voltage can be auto adjust by PVT characteristic. Otherwise, internal voltage will be set to fix dbias. This is a must for stable mass production. Disable for debugging only.

Default value:

- Yes (enabled) if SOC_PMU_PVT_SUPPORTED

LCD and Touch Panel Contains:

- [LCD Peripheral Configuration](#)

LCD Peripheral Configuration Contains:

- [CONFIG_LCD_ENABLE_DEBUG_LOG](#)
- [CONFIG_LCD_PANEL_IO_FORMAT_BUF_SIZE](#)
- [CONFIG_LCD_RGB_RESTART_IN_VSYNC](#)
- [CONFIG_LCD_RGB_ISR_IRAM_SAFE](#)

CONFIG_LCD_PANEL_IO_FORMAT_BUF_SIZE

LCD panel io format buffer size

Found in: [Component config](#) > [LCD and Touch Panel](#) > [LCD Peripheral Configuration](#)

LCD driver allocates an internal buffer to transform the data into a proper format, because of the endian order mismatch. This option is to set the size of the buffer, in bytes.

Default value:

- 32

CONFIG_LCD_ENABLE_DEBUG_LOG

Enable debug log

Found in: [Component config](#) > [LCD and Touch Panel](#) > [LCD Peripheral Configuration](#)

Whether to enable the debug log message for LCD driver. Note that, this option only controls the LCD driver log, won't affect other drivers.

Default value:

- No (disabled)

CONFIG_LCD_RGB_ISR_IRAM_SAFE

RGB LCD ISR IRAM-Safe

Found in: [Component config](#) > [LCD and Touch Panel](#) > [LCD Peripheral Configuration](#)

Ensure the LCD interrupt is IRAM-Safe by allowing the interrupt handler to be executable when the cache is disabled (e.g. SPI Flash write). If you want the LCD driver to keep flushing the screen even when cache ops disabled, you can enable this option. Note, this will also increase the IRAM usage.

Default value:

- No (disabled) if SOC_LCD_RGB_SUPPORTED

CONFIG_LCD_RGB_RESTART_IN_VSYNC

Restart transmission in VSYNC

Found in: [Component config](#) > [LCD and Touch Panel](#) > [LCD Peripheral Configuration](#)

Reset the GDMA channel every VBlank to stop permanent desyncs from happening. Only need to enable it when in your application, the DMA can't deliver data as fast as the LCD consumes it.

Default value:

- No (disabled) if SOC_LCD_RGB_SUPPORTED

ESP NETIF Adapter Contains:

- [CONFIG_ESP_NETIF_SET_DNS_PER_DEFAULT_NETIF](#)
- [CONFIG_ESP_NETIF_BRIDGE_EN](#)
- [CONFIG_ESP_NETIF_L2_TAP](#)
- [CONFIG_ESP_NETIF_IP_LOST_TIMER_INTERVAL](#)
- [CONFIG_ESP_NETIF_USE_TCPIP_STACK_LIB](#)
- [CONFIG_ESP_NETIF_RECEIVE_REPORT_ERRORS](#)

CONFIG_ESP_NETIF_IP_LOST_TIMER_INTERVAL

IP Address lost timer interval (seconds)

Found in: [Component config](#) > [ESP NETIF Adapter](#)

The value of 0 indicates the IP lost timer is disabled, otherwise the timer is enabled.

The IP address may be lost because of some reasons, e.g. when the station disconnects from soft-AP, or when DHCP IP renew fails etc. If the IP lost timer is enabled, it will be started everytime the IP is lost. Event SYSTEM_EVENT_STA_LOST_IP will be raised if the timer expires. The IP lost timer is stopped if the station get the IP again before the timer expires.

Range:

- from 0 to 65535

Default value:

- 120

CONFIG_ESP_NETIF_USE_TCPIP_STACK_LIB

TCP/IP Stack Library

Found in: [Component config](#) > [ESP NETIF Adapter](#)

Choose the TCP/IP Stack to work, for example, LwIP, uIP, etc.

Available options:

- LwIP (CONFIG_ESP_NETIF_TCPIP_LWIP)
lwIP is a small independent implementation of the TCP/IP protocol suite.
- Loopback (CONFIG_ESP_NETIF_LOOPBACK)
Dummy implementation of esp-netif functionality which connects driver transmit to receive function. This option is for testing purpose only

CONFIG_ESP_NETIF_RECEIVE_REPORT_ERRORS

Use esp_err_t to report errors from esp_netif_receive

Found in: [Component config](#) > [ESP NETIF Adapter](#)

Enable if esp_netif_receive() should return error code. This is useful to inform upper layers that packet input to TCP/IP stack failed, so the upper layers could implement flow control. This option is disabled by default due to backward compatibility and will be enabled in v6.0 (IDF-7194)

Default value:

- No (disabled)

CONFIG_ESP_NETIF_L2_TAP

Enable netif L2 TAP support

Found in: [Component config](#) > [ESP NETIF Adapter](#)

A user program can read/write link layer (L2) frames from/to ESP TAP device. The ESP TAP device can be currently associated only with Ethernet physical interfaces.

CONFIG_ESP_NETIF_L2_TAP_MAX_FDS

Maximum number of opened L2 TAP File descriptors

Found in: [Component config](#) > [ESP NETIF Adapter](#) > [CONFIG_ESP_NETIF_L2_TAP](#)

Maximum number of opened File descriptors (FD's) associated with ESP TAP device. ESP TAP FD's take up a certain amount of memory, and allowing fewer FD's to be opened at the same time conserves memory.

Range:

- from 1 to 10 if [CONFIG_ESP_NETIF_L2_TAP](#)

Default value:

- 5 if [CONFIG_ESP_NETIF_L2_TAP](#)

CONFIG_ESP_NETIF_L2_TAP_RX_QUEUE_SIZE

Size of L2 TAP Rx queue

Found in: [Component config](#) > [ESP NETIF Adapter](#) > [CONFIG_ESP_NETIF_L2_TAP](#)

Maximum number of frames queued in opened File descriptor. Once the queue is full, the newly arriving frames are dropped until the queue has enough room to accept incoming traffic (Tail Drop queue management).

Range:

- from 1 to 100 if [CONFIG_ESP_NETIF_L2_TAP](#)

Default value:

- 20 if [CONFIG_ESP_NETIF_L2_TAP](#)

CONFIG_ESP_NETIF_BRIDGE_EN

Enable LwIP IEEE 802.1D bridge

Found in: [Component config](#) > [ESP NETIF Adapter](#)

Enable LwIP IEEE 802.1D bridge support in ESP-NETIF. Note that "Number of clients store data in netif" (LWIP_NUM_NETIF_CLIENT_DATA) option needs to be properly configured to be LwIP bridge available!

Default value:

- No (disabled)

CONFIG_ESP_NETIF_SET_DNS_PER_DEFAULT_NETIF

Enable DNS server per interface

Found in: [Component config](#) > [ESP NETIF Adapter](#)

Enable this option to use the DNS server which belongs to the selected default network interface. This feature collects DNS server and netif information from LWIP core modules. Whenever a new default netif is selected, global DNS servers in LWIP are updated with the netif related servers.

Default value:

- No (disabled)

Partition API Configuration

Power Management Contains:

- [CONFIG_PM_LIGHTSLEEP_RTC_OSC_CAL_INTERVAL](#)
- [CONFIG_PM_SLP_DISABLE_GPIO](#)
- [CONFIG_PM_LIGHT_SLEEP_CALLBACKS](#)
- [CONFIG_PM_POWER_DOWN_CPU_IN_LIGHT_SLEEP](#)
- [CONFIG_PM_POWER_DOWN_PERIPHERAL_IN_LIGHT_SLEEP](#)
- [CONFIG_PM_SLP_IRAM_OPT](#)
- [CONFIG_PM_RTOS_IDLE_OPT](#)
- [CONFIG_PM_ENABLE](#)

CONFIG_PM_ENABLE

Support for power management

Found in: [Component config](#) > [Power Management](#)

If enabled, application is compiled with support for power management. This option has run-time overhead (increased interrupt latency, longer time to enter idle state), and it also reduces accuracy of RTOS ticks and timers used for timekeeping. Enable this option if application uses power management APIs.

Default value:

- No (disabled) if `__DOXYGEN__`

CONFIG_PM_DFS_INIT_AUTO

Enable dynamic frequency scaling (DFS) at startup

Found in: [Component config](#) > [Power Management](#) > [CONFIG_PM_ENABLE](#)

If enabled, startup code configures dynamic frequency scaling. Max CPU frequency is set to `DEFAULT_CPU_FREQ_MHZ` setting, min frequency is set to XTAL frequency. If disabled, DFS will not be active until the application configures it using `esp_pm_configure` function.

Default value:

- No (disabled) if [CONFIG_PM_ENABLE](#)

CONFIG_PM_PROFILING

Enable profiling counters for PM locks

Found in: [Component config](#) > [Power Management](#) > [CONFIG_PM_ENABLE](#)

If enabled, `esp_pm_*` functions will keep track of the amount of time each of the power management locks has been held, and `esp_pm_dump_locks` function will print this information. This feature can be used to analyze which locks are preventing the chip from going into a lower power state, and see what time the chip spends in each power saving mode. This feature does incur some run-time overhead, so should typically be disabled in production builds.

Default value:

- No (disabled) if [CONFIG_PM_ENABLE](#)

CONFIG_PM_TRACE

Enable debug tracing of PM using GPIOs

Found in: [Component config](#) > [Power Management](#) > [CONFIG_PM_ENABLE](#)

If enabled, some GPIOs will be used to signal events such as RTOS ticks, frequency switching, entry/exit from idle state. Refer to `pm_trace.c` file for the list of GPIOs. This feature is intended to be used when analyzing/debugging behavior of power management implementation, and should be kept disabled in applications.

Default value:

- No (disabled) if [CONFIG_PM_ENABLE](#)

CONFIG_PM_SLP_IRAM_OPT

Put lightsleep related codes in internal RAM

Found in: [Component config](#) > [Power Management](#)

If enabled, about 2.1KB of lightsleep related source code would be in IRAM and chip would sleep longer for 310us at 160MHz CPU frequency most each time. This feature is intended to be used when lower power consumption is needed while there is enough place in IRAM to place source code.

Default value:

- No (disabled)

CONFIG_PM_RTOS_IDLE_OPT

Put RTOS IDLE related codes in internal RAM

Found in: [Component config](#) > [Power Management](#)

If enabled, about 180Bytes of RTOS_IDLE related source code would be in IRAM and chip would sleep longer for 20us at 160MHz CPU frequency most each time. This feature is intended to be used when lower power consumption is needed while there is enough place in IRAM to place source code.

CONFIG_PM_SLP_DISABLE_GPIO

Disable all GPIO when chip at sleep

Found in: [Component config](#) > [Power Management](#)

This feature is intended to disable all GPIO pins at automatic sleep to get a lower power mode. If enabled, chips will disable all GPIO pins at automatic sleep to reduce about 200~300 uA current. If you want to specifically use some pins normally as chip wakes when chip sleeps, you can call 'gpio_sleep_sel_dis' to disable this feature on those pins. You can also keep this feature on and call 'gpio_sleep_set_direction' and 'gpio_sleep_set_pull_mode' to have a different GPIO configuration at sleep. Warning: If you want to enable this option on ESP32, you should enable [GPIO_ESP32_SUPPORT_SWITCH_SLP_PULL](#) at first, otherwise you will not be able to switch pullup/pulldown mode.

CONFIG_PM_LIGHTSLEEP_RTC_OSC_CAL_INTERVAL

Calibrate the RTC_FAST/SLOW clock every N times of light sleep

Found in: [Component config](#) > [Power Management](#)

The value of this option determines the calibration interval of the RTC_FAST/SLOW clock during sleep when power management is enabled. When it is configured as N, the RTC_FAST/SLOW clock will be calibrated every N times of lightsleep. Decreasing this value will increase the time the chip is in the active state, thereby increasing the average power consumption of the chip. Increasing this value can reduce the average power consumption, but when the external environment changes drastically and the chip RTC_FAST/SLOW oscillator frequency drifts, it may cause system instability.

Range:

- from 1 to 128 if [CONFIG_PM_ENABLE](#)

Default value:

- 1 if [CONFIG_PM_ENABLE](#)

CONFIG_PM_POWER_DOWN_CPU_IN_LIGHT_SLEEP

Power down CPU in light sleep

Found in: [Component config](#) > [Power Management](#)

If enabled, the CPU will be powered down in light sleep, ESP chips supports saving and restoring CPU's running context before and after light sleep, the feature provides applications with seamless CPU powerdowned lightsleep without user awareness. But this will takes up some internal memory. On esp32c3 soc, enabling this option will consume 1.68 KB of internal RAM and will reduce sleep current consumption by about 100 uA. On esp32s3 soc, enabling this option will consume 8.58 KB of internal RAM and will reduce sleep current consumption by about 650 uA.

Default value:

- Yes (enabled)

CONFIG_PM_POWER_DOWN_PERIPHERAL_IN_LIGHT_SLEEP

Power down Digital Peripheral in light sleep (EXPERIMENTAL)

Found in: [Component config](#) > [Power Management](#)

If enabled, digital peripherals will try to powered down in light sleep, then all related peripherals will not be available during sleep, including wake-up sources from the peripherals (For detailed availability information, see the note of the corresponding wakeup source enable function). The chip will automatically save/restore register context during sleep/wakeup to make the upper layer user unaware of the peripheral powerdown during sleep. Enabling this option will increase static RAM and heap usage but will also significantly reduce power. consumption during lightsleep, the actual memory cost depends on the peripherals you have initialized, for specific power consumption data in this mode, please refer to Electrical Characteristics section in the chip datasheet. (In order to save/restore the context of the necessary hardware for FreeRTOS to run, it will need at least 4.55 KB free heap at sleep time. Otherwise sleep will not power down the peripherals.)

Note1: Please use this option with caution, the current IDF does not support the retention of all peripherals. When the digital peripherals are powered off and a sleep and wake-up is completed, the peripherals that have not saved the running context are equivalent to performing a reset. !!! Please confirm the peripherals used in your application and their sleep retention support status before enabling this option, peripherals sleep retention driver support status is tracked in power_management.rst

Note2: When this option is enabled simultaneously with FREERTOS_USE_TICKLESS_IDLE, since the UART will be powered down, the uart FIFO will be flushed before sleep to avoid data loss, however, this has the potential to block the sleep process and cause the wakeup time to be skipped, which will cause the tick of freertos to not be compensated correctly when returning from sleep and cause the system to crash. To avoid this, you can increase FREERTOS_IDLE_TIME_BEFORE_SLEEP threshold in menuconfig.

Note3: Enabling this option does not necessarily mean that the peripheral power domain will be turned down during sleep. The control priority of `esp_sleep_pd_config` is higher than this option, user code can still prevent the peripheral power domain from powering down during sleep by `esp_sleep_pd_config(ESP_PD_DOMAIN_TOP, ESP_PD_OPTION_ON)`. In addition, whether the peripheral power domain is powered down during sleep also depends on the sleep working strategy selected by the driver. If any module belonging to the peripheral power domain chooses not to be powered down during sleep, then the peripheral power domain will not be powered off either.

Default value:

- No (disabled) if SOC_PAU_SUPPORTED

CONFIG_PM_LIGHT_SLEEP_CALLBACKS

Enable registration of pm light sleep callbacks

Found in: [Component config](#) > [Power Management](#)

If enabled, it allows user to register entry and exit callbacks which are called before and after entering auto light sleep.

NOTE: These callbacks are executed from the IDLE task context hence you cannot have any blocking calls in your callbacks.

NOTE: Enabling these callbacks may change sleep duration calculations based on time spent in callback and hence it is highly recommended to keep them as short as possible

Default value:

- No (disabled) if [CONFIG_FREERTOS_USE_TICKLESS_IDLE](#)

ESP PSRAM Contains:

- [CONFIG_SPIRAM](#)

CONFIG_SPIRAM

Support for external PSRAM

Found in: [Component config](#) > [ESP PSRAM](#)

This enables support for an external PSRAM chip, connected in parallel with the main SPI flash chip.

PSRAM config Contains:

- [CONFIG_SPIRAM_ALLOW_BSS_SEG_EXTERNAL_MEMORY](#)
- [CONFIG_SPIRAM_ALLOW_NOINIT_SEG_EXTERNAL_MEMORY](#)
- [CONFIG_SPIRAM_ALLOW_STACK_EXTERNAL_MEMORY](#)
- [CONFIG_SPIRAM_ECC_ENABLE](#)
- [CONFIG_SPIRAM_BOOT_INIT](#)
- [CONFIG_SPIRAM_MODE](#)
- [CONFIG_SPIRAM_MALLOC_ALWAYSINTERNAL](#)
- [CONFIG_SPIRAM_MALLOC_RESERVE_INTERNAL](#)
- [CONFIG_SPIRAM_MEMTEST](#)
- [CONFIG_SPIRAM_SPEED](#)
- [CONFIG_SPIRAM_USE](#)
- [CONFIG_SPIRAM_TRY_ALLOCATE_WIFI_LWIP](#)

CONFIG_SPIRAM_MODE

Line Mode of PSRAM chip in use

Found in: [Component config](#) > [ESP PSRAM](#) > [CONFIG_SPIRAM](#) > [PSRAM config](#)

Available options:

- 16-Line-Mode PSRAM ([CONFIG_SPIRAM_MODE_HEX](#))

CONFIG_SPIRAM_SPEED

Set PSRAM clock speed

Found in: [Component config](#) > [ESP PSRAM](#) > [CONFIG_SPIRAM](#) > [PSRAM config](#)

Select the speed for the PSRAM chip.

Available options:

- 20MHz clock speed ([CONFIG_SPIRAM_SPEED_20M](#))

CONFIG_SPIRAM_ECC_ENABLE

Enable PSRAM ECC

Found in: [Component config](#) > [ESP PSRAM](#) > [CONFIG_SPIRAM](#) > [PSRAM config](#)

Enable Error-Correcting Code function when accessing PSRAM.

If enabled, 1/8 of the PSRAM total size will be reserved for error-correcting code.

CONFIG_SPIRAM_ALLOW_STACK_EXTERNAL_MEMORY

Allow external memory as an argument to xTaskCreateStatic

Found in: [Component config](#) > [ESP PSRAM](#) > [CONFIG_SPIRAM](#) > [PSRAM config](#)

Accessing memory in PSRAM has certain restrictions, so task stacks allocated by xTaskCreate are by default allocated from internal RAM.

This option allows for passing memory allocated from PSRAM to be passed to xTaskCreateStatic. This should only be used for tasks where the stack is never accessed while the L2Cache is disabled, e.g. during SPI Flash operations

CONFIG_SPIRAM_BOOT_INIT

Initialize SPI RAM during startup

Found in: [Component config](#) > [ESP PSRAM](#) > [CONFIG_SPIRAM](#) > [PSRAM config](#)

If this is enabled, the SPI RAM will be enabled during initial boot. Unless you have specific requirements, you'll want to leave this enabled so memory allocated during boot-up can also be placed in SPI RAM.

CONFIG_SPIRAM_IGNORE_NOTFOUND

Ignore PSRAM when not found

Found in: [Component config](#) > [ESP PSRAM](#) > [CONFIG_SPIRAM](#) > [PSRAM config](#) > [CONFIG_SPIRAM_BOOT_INIT](#)

Normally, if psram initialization is enabled during compile time but not found at runtime, it is seen as an error making the CPU panic. If this is enabled, booting will complete but no PSRAM will be available. In particular, it is important to note that when SPIRAM_IGNORE_NOTFOUND is enabled, some WIFI related configs will be set to the default value used when SPIRAM is disabled. It can be assumed that by enabling this config, WIFI and LWIP will assume that the current chip does not have SPIRAM. SPIRAM_TRY_ALLOCATE_WIFI_LWIP will not work, buffers in WIFI and LWIP will be set to smaller ranges and default values. WIFI_CACHE_TX_BUFFER_NUM and ESP_WIFI_AMSDU_TX_ENABLED will be disabled, ESP_WIFI_TX_BUFFER will be set to ESP_WIFI_DYNAMIC_TX_BUFFER.

CONFIG_SPIRAM_USE

SPI RAM access method

Found in: [Component config](#) > [ESP PSRAM](#) > [CONFIG_SPIRAM](#) > [PSRAM config](#)

The SPI RAM can be accessed in multiple methods: by just having it available as an unmanaged memory region in the CPU's memory map, by integrating it in the heap as 'special' memory needing heap_caps_malloc to allocate, or by fully integrating it making malloc() also able to return SPI RAM pointers.

Available options:

- Integrate RAM into memory map (CONFIG_SPIRAM_USE_MEMMAP)

- Make RAM allocatable using `heap_caps_malloc(..., MALLOC_CAP_SPIRAM)` (`CONFIG_SPIRAM_USE_CAPS_ALLOC`)
- Make RAM allocatable using `malloc()` as well (`CONFIG_SPIRAM_USE_MALLOC`)

CONFIG_SPIRAM_MEMTEST

Run memory test on SPI RAM initialization

Found in: [Component config](#) > [ESP PSRAM](#) > [CONFIG_SPIRAM](#) > [PSRAM config](#)

Runs a rudimentary memory test on initialization. Aborts when memory test fails. Disable this for slightly faster startup.

CONFIG_SPIRAM_MALLOC_ALWAYSINTERNAL

Maximum `malloc()` size, in bytes, to always put in internal memory

Found in: [Component config](#) > [ESP PSRAM](#) > [CONFIG_SPIRAM](#) > [PSRAM config](#)

If `malloc()` is capable of also allocating SPI-connected ram, its allocation strategy will prefer to allocate chunks less than this size in internal memory, while allocations larger than this will be done from external RAM. If allocation from the preferred region fails, an attempt is made to allocate from the non-preferred region instead, so `malloc()` will not suddenly fail when either internal or external memory is full.

CONFIG_SPIRAM_TRY_ALLOCATE_WIFI_LWIP

Try to allocate memories of WiFi and LWIP in SPIRAM firstly. If failed, allocate internal memory

Found in: [Component config](#) > [ESP PSRAM](#) > [CONFIG_SPIRAM](#) > [PSRAM config](#)

Try to allocate memories of WiFi and LWIP in SPIRAM firstly. If failed, try to allocate internal memory then.

CONFIG_SPIRAM_MALLOC_RESERVE_INTERNAL

Reserve this amount of bytes for data that specifically needs to be in DMA or internal memory

Found in: [Component config](#) > [ESP PSRAM](#) > [CONFIG_SPIRAM](#) > [PSRAM config](#)

Because the external/internal RAM allocation strategy is not always perfect, it sometimes may happen that the internal memory is entirely filled up. This causes allocations that are specifically done in internal memory, for example the stack for new tasks or memory to service DMA or have memory that's also available when SPI cache is down, to fail. This option reserves a pool specifically for requests like that; the memory in this pool is not given out when a normal `malloc()` is called.

Set this to 0 to disable this feature.

Note that because FreeRTOS stacks are forced to internal memory, they will also use this memory pool; be sure to keep this in mind when adjusting this value.

Note also that the DMA reserved pool may not be one single contiguous memory region, depending on the configured size and the static memory usage of the app.

CONFIG_SPIRAM_ALLOW_BSS_SEG_EXTERNAL_MEMORY

Allow .bss segment placed in external memory

Found in: [Component config](#) > [ESP PSRAM](#) > [CONFIG_SPIRAM](#) > [PSRAM config](#)

If enabled, variables with `EXT_RAM_BSS_ATTR` attribute will be placed in SPIRAM instead of internal DRAM. BSS section of `lwip`, `net80211`, `pp`, `bt` libraries will be automatically placed in SPIRAM. BSS sections from other object files and libraries can also be placed in SPIRAM through linker fragment scheme `extram_bss`.

Note that the variables placed in SPIRAM using EXT_RAM_BSS_ATTR will be zero initialized.

CONFIG_SPIRAM_ALLOW_NOINIT_SEG_EXTERNAL_MEMORY

Allow .noinit segment placed in external memory

Found in: *Component config > ESP PSRAM > CONFIG_SPIRAM > PSRAM config*

If enabled, noinit variables can be placed in PSRAM using EXT_RAM_NOINIT_ATTR.

Note the values placed into this section will not be initialized at startup and should keep its value after software restart.

ESP Ringbuf Contains:

- *CONFIG_RINGBUF_PLACE_FUNCTIONS_INTO_FLASH*

CONFIG_RINGBUF_PLACE_FUNCTIONS_INTO_FLASH

Place non-ISR ringbuf functions into flash

Found in: *Component config > ESP Ringbuf*

Place non-ISR ringbuf functions (like xRingbufferCreate/xRingbufferSend) into flash. This frees up IRAM, but the functions can no longer be called when the cache is disabled.

Default value:

- No (disabled)

CONFIG_RINGBUF_PLACE_ISR_FUNCTIONS_INTO_FLASH

Place ISR ringbuf functions into flash

Found in: *Component config > ESP Ringbuf > CONFIG_RINGBUF_PLACE_FUNCTIONS_INTO_FLASH*

Place ISR ringbuf functions (like xRingbufferSendFromISR/xRingbufferReceiveFromISR) into flash. This frees up IRAM, but the functions can no longer be called when the cache is disabled or from an IRAM interrupt context.

This option is not compatible with ESP-IDF drivers which are configured to run the ISR from an IRAM context, e.g. CONFIG_UART_ISR_IN_IRAM.

Default value:

- No (disabled) if *CONFIG_RINGBUF_PLACE_FUNCTIONS_INTO_FLASH*

ESP System Settings Contains:

- *CONFIG_ESP_SYSTEM_RTC_EXT_XTAL_BOOTSTRAP_CYCLES*
- *Cache config*
- *CONFIG_ESP_CONSOLE_UART*
- *CONFIG_ESP_CONSOLE_SECONDARY*
- *CONFIG_ESP_DEFAULT_CPU_FREQ_MHZ*
- *CONFIG_ESP_SYSTEM_ALLOW_RTC_FAST_MEM_AS_HEAP*
- *CONFIG_ESP_TASK_WDT_EN*
- *CONFIG_ESP_SYSTEM_EVENT_TASK_STACK_SIZE*
- *CONFIG_ESP_SYSTEM_USE_EH_FRAME*
- *CONFIG_ESP_SYSTEM_HW_STACK_GUARD*
- *CONFIG_ESP_XT_WDT*
- *CONFIG_ESP_SYSTEM_CHECK_INT_LEVEL*
- *CONFIG_ESP_INT_WDT*
- *CONFIG_ESP_MAIN_TASK_AFFINITY*
- *CONFIG_ESP_MAIN_TASK_STACK_SIZE*

- [CONFIG_ESP_DEBUG_OCDAWARE](#)
- [Memory protection](#)
- [CONFIG_ESP_MINIMAL_SHARED_STACK_SIZE](#)
- [CONFIG_ESP_DEBUG_STUBS_ENABLE](#)
- [CONFIG_ESP_SYSTEM_PANIC](#)
- [CONFIG_ESP_SYSTEM_PANIC_REBOOT_DELAY_SECONDS](#)
- [CONFIG_ESP_PANIC_HANDLER_IRAM](#)
- [CONFIG_ESP_SYSTEM_EVENT_QUEUE_SIZE](#)
- [CONFIG_ESP_CONSOLE_UART_BAUDRATE](#)
- [CONFIG_ESP_CONSOLE_UART_NUM](#)
- [CONFIG_ESP_CONSOLE_UART_RX_GPIO](#)
- [CONFIG_ESP_CONSOLE_UART_TX_GPIO](#)

CONFIG_ESP_DEFAULT_CPU_FREQ_MHZ

CPU frequency

Found in: [Component config](#) > [ESP System Settings](#)

CPU frequency to be set on application startup.

Available options:

- 40 MHz (CONFIG_ESP_DEFAULT_CPU_FREQ_MHZ_40)
- 80 MHz (CONFIG_ESP_DEFAULT_CPU_FREQ_MHZ_80)
- 120 MHz (CONFIG_ESP_DEFAULT_CPU_FREQ_MHZ_120)
- 160 MHz (CONFIG_ESP_DEFAULT_CPU_FREQ_MHZ_160)

Cache config Contains:

- [CONFIG_CACHE_L2_CACHE_LINE_SIZE](#)
- [CONFIG_CACHE_L2_CACHE_SIZE](#)

CONFIG_CACHE_L2_CACHE_SIZE

L2 cache size

Found in: [Component config](#) > [ESP System Settings](#) > [Cache config](#)

L2 cache size to be set on application startup.

Available options:

- 128KB (CONFIG_CACHE_L2_CACHE_128KB)
- 256KB (CONFIG_CACHE_L2_CACHE_256KB)
- 512KB (CONFIG_CACHE_L2_CACHE_512KB)

CONFIG_CACHE_L2_CACHE_LINE_SIZE

L2 cache line size

Found in: [Component config](#) > [ESP System Settings](#) > [Cache config](#)

L2 cache line size to be set on application startup.

Available options:

- 64 Bytes (CONFIG_CACHE_L2_CACHE_LINE_64B)
- 128 Bytes (CONFIG_CACHE_L2_CACHE_LINE_128B)

CONFIG_ESP_SYSTEM_PANIC

Panic handler behaviour

Found in: [Component config](#) > [ESP System Settings](#)

If FreeRTOS detects unexpected behaviour or an unhandled exception, the panic handler is invoked. Configure the panic handler's action here.

Available options:

- Print registers and halt (CONFIG_ESP_SYSTEM_PANIC_PRINT_HALT)
Outputs the relevant registers over the serial port and halt the processor. Needs a manual reset to restart.
- Print registers and reboot (CONFIG_ESP_SYSTEM_PANIC_PRINT_REBOOT)
Outputs the relevant registers over the serial port and immediately reset the processor.
- Silent reboot (CONFIG_ESP_SYSTEM_PANIC_SILENT_REBOOT)
Just resets the processor without outputting anything
- GDBStub on panic (CONFIG_ESP_SYSTEM_PANIC_GDBSTUB)
Invoke gdbstub on the serial port, allowing for gdb to attach to it to do a postmortem of the crash.

CONFIG_ESP_SYSTEM_PANIC_REBOOT_DELAY_SECONDS

Panic reboot delay (Seconds)

Found in: [Component config](#) > [ESP System Settings](#)

After the panic handler executes, you can specify a number of seconds to wait before the device reboots.

Range:

- from 0 to 99

Default value:

- 0

CONFIG_ESP_SYSTEM_RTC_EXT_XTAL_BOOTSTRAP_CYCLES

Bootstrap cycles for external 32kHz crystal

Found in: [Component config](#) > [ESP System Settings](#)

To reduce the startup time of an external RTC crystal, we bootstrap it with a 32kHz square wave for a fixed number of cycles. Setting 0 will disable bootstrapping (if disabled, the crystal may take longer to start up or fail to oscillate under some conditions).

If this value is too high, a faulty crystal may initially start and then fail. If this value is too low, an otherwise good crystal may not start.

To accurately determine if the crystal has started, set a larger "Number of cycles for RTC_SLOW_CLK calibration" (about 3000).

CONFIG_ESP_SYSTEM_ALLOW_RTC_FAST_MEM_AS_HEAP

Enable RTC fast memory for dynamic allocations

Found in: [Component config](#) > [ESP System Settings](#)

This config option allows to add RTC fast memory region to system heap with capability similar to that of DRAM region but without DMA. This memory will be consumed first per heap initialization order by early startup services and scheduler related code. Speed wise RTC fast memory operates on APB clock and hence does not have much performance impact.

Default value:

- Yes (enabled)

CONFIG_ESP_SYSTEM_USE_EH_FRAME

Generate and use eh_frame for backtracing

Found in: [Component config](#) > [ESP System Settings](#)

Generate DWARF information for each function of the project. These information will be parsed and used to perform backtracing when panics occur. Activating this option will activate asynchronous frame unwinding and generation of both .eh_frame and .eh_frame_hdr sections, resulting in a bigger binary size (20% to 100% larger). The main purpose of this option is to be able to have a backtrace parsed and printed by the program itself, regardless of the serial monitor used. This option shall NOT be used for production.

Default value:

- No (disabled)

Memory protection Contains:

- [CONFIG_ESP_SYSTEM_PMP_IDRAM_SPLIT](#)
- [CONFIG_ESP_SYSTEM_MEMPROT_FEATURE](#)

CONFIG_ESP_SYSTEM_PMP_IDRAM_SPLIT

Enable IRAM/DRAM split protection

Found in: [Component config](#) > [ESP System Settings](#) > [Memory protection](#)

If enabled, the CPU watches all the memory access and raises an exception in case of any memory violation. This feature automatically splits the SRAM memory, using PMP, into data and instruction segments and sets Read/Execute permissions for the instruction part (below given splitting address) and Read/Write permissions for the data part (above the splitting address). The memory protection is effective on all access through the IRAM0 and DRAM0 buses.

Default value:

- Yes (enabled)

CONFIG_ESP_SYSTEM_MEMPROT_FEATURE

Enable memory protection

Found in: [Component config](#) > [ESP System Settings](#) > [Memory protection](#)

If enabled, the permission control module watches all the memory access and fires the panic handler if a permission violation is detected. This feature automatically splits the SRAM memory into data and instruction segments and sets Read/Execute permissions for the instruction part (below given splitting address) and Read/Write permissions for the data part (above the splitting address). The memory protection is effective on all access through the IRAM0 and DRAM0 buses.

Default value:

- Yes (enabled) if SOC_MEMPROT_SUPPORTED

CONFIG_ESP_SYSTEM_MEMPROT_FEATURE_LOCK

Lock memory protection settings

Found in: [Component config](#) > [ESP System Settings](#) > [Memory protection](#) > [CONFIG_ESP_SYSTEM_MEMPROT_FEATURE](#)

Once locked, memory protection settings cannot be changed anymore. The lock is reset only on the chip startup.

Default value:

- Yes (enabled) if `CONFIG_ESP_SYSTEM_MEMPROT_FEATURE`

CONFIG_ESP_SYSTEM_EVENT_QUEUE_SIZE

System event queue size

Found in: [Component config](#) > [ESP System Settings](#)

Config system event queue size in different application.

Default value:

- 32

CONFIG_ESP_SYSTEM_EVENT_TASK_STACK_SIZE

Event loop task stack size

Found in: [Component config](#) > [ESP System Settings](#)

Config system event task stack size in different application.

Default value:

- 2304

CONFIG_ESP_MAIN_TASK_STACK_SIZE

Main task stack size

Found in: [Component config](#) > [ESP System Settings](#)

Configure the "main task" stack size. This is the stack of the task which calls `app_main()`. If `app_main()` returns then this task is deleted and its stack memory is freed.

Default value:

- 3584

CONFIG_ESP_MAIN_TASK_AFFINITY

Main task core affinity

Found in: [Component config](#) > [ESP System Settings](#)

Configure the "main task" core affinity. This is the used core of the task which calls `app_main()`. If `app_main()` returns then this task is deleted.

Available options:

- CPU0 (`CONFIG_ESP_MAIN_TASK_AFFINITY_CPU0`)
- CPU1 (`CONFIG_ESP_MAIN_TASK_AFFINITY_CPU1`)
- No affinity (`CONFIG_ESP_MAIN_TASK_AFFINITY_NO_AFFINITY`)

CONFIG_ESP_MINIMAL_SHARED_STACK_SIZE

Minimal allowed size for shared stack

Found in: [Component config](#) > [ESP System Settings](#)

Minimal value of size, in bytes, accepted to execute a expression with shared stack.

Default value:

- 2048

CONFIG_ESP_CONSOLE_UART

Channel for console output

Found in: [Component config](#) > [ESP System Settings](#)

Select where to send console output (through stdout and stderr).

- Default is to use UART0 on pre-defined GPIOs.
- If "Custom" is selected, UART0 or UART1 can be chosen, and any pins can be selected.
- If "None" is selected, there will be no console output on any UART, except for initial output from ROM bootloader. This ROM output can be suppressed by GPIO strapping or EFUSE, refer to chip datasheet for details.
- On chips with USB OTG peripheral, "USB CDC" option redirects output to the CDC port. This option uses the CDC driver in the chip ROM. This option is incompatible with TinyUSB stack.
- On chips with an USB serial/JTAG debug controller, selecting the option for that redirects output to the CDC/ACM (serial port emulation) component of that device.

Available options:

- Default: UART0 (CONFIG_ESP_CONSOLE_UART_DEFAULT)
- USB CDC (CONFIG_ESP_CONSOLE_USB_CDC)
- USB Serial/JTAG Controller (CONFIG_ESP_CONSOLE_USB_SERIAL_JTAG)
- Custom UART (CONFIG_ESP_CONSOLE_UART_CUSTOM)
- None (CONFIG_ESP_CONSOLE_NONE)

CONFIG_ESP_CONSOLE_SECONDARY

Channel for console secondary output

Found in: [Component config](#) > [ESP System Settings](#)

This secondary option supports output through other specific port like USB_SERIAL_JTAG when UART0 port as a primary is selected but not connected. This secondary output currently only supports non-blocking mode without using REPL. If you want to output in blocking mode with REPL or input through this secondary port, please change the primary config to this port in *Channel for console output* menu.

Available options:

- No secondary console (CONFIG_ESP_CONSOLE_SECONDARY_NONE)
- USB_SERIAL_JTAG PORT (CONFIG_ESP_CONSOLE_SECONDARY_USB_SERIAL_JTAG)
This option supports output through USB_SERIAL_JTAG port when the UART0 port is not connected. The output currently only supports non-blocking mode without using the console. If you want to output in blocking mode with REPL or input through USB_SERIAL_JTAG port, please change the primary config to ESP_CONSOLE_USB_SERIAL_JTAG above.

CONFIG_ESP_CONSOLE_UART_NUM

UART peripheral to use for console output (0-1)

Found in: [Component config](#) > [ESP System Settings](#)

This UART peripheral is used for console output from the ESP-IDF Bootloader and the app.

If the configuration is different in the Bootloader binary compared to the app binary, UART is reconfigured after the bootloader exits and the app starts.

Due to an ESP32 ROM bug, UART2 is not supported for console output via `esp_rom_printf`.

Available options:

- UART0 (CONFIG_ESP_CONSOLE_UART_CUSTOM_NUM_0)
- UART1 (CONFIG_ESP_CONSOLE_UART_CUSTOM_NUM_1)

CONFIG_ESP_CONSOLE_UART_TX_GPIO

UART TX on GPIO#

Found in: [Component config](#) > [ESP System Settings](#)

This GPIO is used for console UART TX output in the ESP-IDF Bootloader and the app (including boot log output and default standard output and standard error of the app).

If the configuration is different in the Bootloader binary compared to the app binary, UART is reconfigured after the bootloader exits and the app starts.

Range:

- from 0 to 56 if [CONFIG_ESP_CONSOLE_UART_CUSTOM](#)

Default value:

- 37 if [CONFIG_ESP_CONSOLE_UART_CUSTOM](#)
- 43 if [CONFIG_ESP_CONSOLE_UART_CUSTOM](#)

CONFIG_ESP_CONSOLE_UART_RX_GPIO

UART RX on GPIO#

Found in: [Component config](#) > [ESP System Settings](#)

This GPIO is used for UART RX input in the ESP-IDF Bootloader and the app (including default standard input of the app).

Note: The default ESP-IDF Bootloader configures this pin but doesn't read anything from the UART.

If the configuration is different in the Bootloader binary compared to the app binary, UART is reconfigured after the bootloader exits and the app starts.

Range:

- from 0 to 56 if [CONFIG_ESP_CONSOLE_UART_CUSTOM](#)

Default value:

- 38 if [CONFIG_ESP_CONSOLE_UART_CUSTOM](#)
- 44 if [CONFIG_ESP_CONSOLE_UART_CUSTOM](#)

CONFIG_ESP_CONSOLE_UART_BAUDRATE

UART console baud rate

Found in: [Component config](#) > [ESP System Settings](#)

This baud rate is used by both the ESP-IDF Bootloader and the app (including boot log output and default standard input/output/error of the app).

The app's maximum baud rate depends on the UART clock source. If Power Management is disabled, the UART clock source is the APB clock and all baud rates in the available range will be sufficiently accurate. If Power Management is enabled, REF_TICK clock source is used so the baud rate is divided from 1MHz. Baud rates above 1Mbps are not possible and values between 500Kbps and 1Mbps may not be accurate.

If the configuration is different in the Bootloader binary compared to the app binary, UART is reconfigured after the bootloader exits and the app starts.

Range:

- from 1200 to 1000000 if [CONFIG_PM_ENABLE](#)

Default value:

- 115200

CONFIG_ESP_INT_WDT

Interrupt watchdog

Found in: [Component config](#) > [ESP System Settings](#)

This watchdog timer can detect if the FreeRTOS tick interrupt has not been called for a certain time, either because a task turned off interrupts and did not turn them on for a long time, or because an interrupt handler did not return. It will try to invoke the panic handler first and failing that reset the SoC.

Default value:

- Yes (enabled)

CONFIG_ESP_INT_WDT_TIMEOUT_MS

Interrupt watchdog timeout (ms)

Found in: [Component config](#) > [ESP System Settings](#) > [CONFIG_ESP_INT_WDT](#)

The timeout of the watchdog, in milliseconds. Make this higher than the FreeRTOS tick rate.

Range:

- from 10 to 10000

Default value:

- 300

CONFIG_ESP_INT_WDT_CHECK_CPU1

Also watch CPU1 tick interrupt

Found in: [Component config](#) > [ESP System Settings](#) > [CONFIG_ESP_INT_WDT](#)

Also detect if interrupts on CPU 1 are disabled for too long.

CONFIG_ESP_TASK_WDT_EN

Enable Task Watchdog Timer

Found in: [Component config](#) > [ESP System Settings](#)

The Task Watchdog Timer can be used to make sure individual tasks are still running. Enabling this option will enable the Task Watchdog Timer. It can be either initialized automatically at startup or initialized after startup (see Task Watchdog Timer API Reference)

Default value:

- Yes (enabled)

CONFIG_ESP_TASK_WDT_INIT

Initialize Task Watchdog Timer on startup

Found in: [Component config](#) > [ESP System Settings](#) > [CONFIG_ESP_TASK_WDT_EN](#)

Enabling this option will cause the Task Watchdog Timer to be initialized automatically at startup.

Default value:

- Yes (enabled)

CONFIG_ESP_TASK_WDT_PANIC

Invoke panic handler on Task Watchdog timeout

Found in: [Component config](#) > [ESP System Settings](#) > [CONFIG_ESP_TASK_WDT_EN](#) > [CONFIG_ESP_TASK_WDT_INIT](#)

If this option is enabled, the Task Watchdog Timer will be configured to trigger the panic handler when it times out. This can also be configured at run time (see Task Watchdog Timer API Reference)

Default value:

- No (disabled)

CONFIG_ESP_TASK_WDT_TIMEOUT_S

Task Watchdog timeout period (seconds)

Found in: [Component config](#) > [ESP System Settings](#) > [CONFIG_ESP_TASK_WDT_EN](#) > [CONFIG_ESP_TASK_WDT_INIT](#)

Timeout period configuration for the Task Watchdog Timer in seconds. This is also configurable at run time (see Task Watchdog Timer API Reference)

Range:

- from 1 to 60

Default value:

- 5

CONFIG_ESP_TASK_WDT_CHECK_IDLE_TASK_CPU0

Watch CPU0 Idle Task

Found in: [Component config](#) > [ESP System Settings](#) > [CONFIG_ESP_TASK_WDT_EN](#) > [CONFIG_ESP_TASK_WDT_INIT](#)

If this option is enabled, the Task Watchdog Timer will watch the CPU0 Idle Task. Having the Task Watchdog watch the Idle Task allows for detection of CPU starvation as the Idle Task not being called is usually a symptom of CPU starvation. Starvation of the Idle Task is detrimental as FreeRTOS household tasks depend on the Idle Task getting some runtime every now and then.

Default value:

- Yes (enabled)

CONFIG_ESP_TASK_WDT_CHECK_IDLE_TASK_CPU1

Watch CPU1 Idle Task

Found in: [Component config](#) > [ESP System Settings](#) > [CONFIG_ESP_TASK_WDT_EN](#) > [CONFIG_ESP_TASK_WDT_INIT](#)

If this option is enabled, the Task Watchdog Timer will watch the CPU1 Idle Task.

CONFIG_ESP_XT_WDT

Initialize XTAL32K watchdog timer on startup

Found in: [Component config](#) > [ESP System Settings](#)

This watchdog timer can detect oscillation failure of the XTAL32K_CLK. When such a failure is detected the hardware can be set up to automatically switch to BACKUP32K_CLK and generate an interrupt.

CONFIG_ESP_XT_WDT_TIMEOUT

XTAL32K watchdog timeout period

Found in: [Component config](#) > [ESP System Settings](#) > [CONFIG_ESP_XT_WDT](#)

Timeout period configuration for the XTAL32K watchdog timer based on RTC_CLK.

Range:

- from 1 to 255 if [CONFIG_ESP_XT_WDT](#)

Default value:

- 200 if [CONFIG_ESP_XT_WDT](#)

CONFIG_ESP_XT_WDT_BACKUP_CLK_ENABLE

Automatically switch to BACKUP32K_CLK when timer expires

Found in: [Component config](#) > [ESP System Settings](#) > [CONFIG_ESP_XT_WDT](#)

Enable this to automatically switch to BACKUP32K_CLK as the source of RTC_SLOW_CLK when the watchdog timer expires.

Default value:

- Yes (enabled) if [CONFIG_ESP_XT_WDT](#)

CONFIG_ESP_PANIC_HANDLER_IRAM

Place panic handler code in IRAM

Found in: [Component config](#) > [ESP System Settings](#)

If this option is disabled (default), the panic handler code is placed in flash not IRAM. This means that if ESP-IDF crashes while flash cache is disabled, the panic handler will automatically re-enable flash cache before running GDB Stub or Core Dump. This adds some minor risk, if the flash cache status is also corrupted during the crash.

If this option is enabled, the panic handler code (including required UART functions) is placed in IRAM. This may be necessary to debug some complex issues with crashes while flash cache is disabled (for example, when writing to SPI flash) or when flash cache is corrupted when an exception is triggered.

Default value:

- No (disabled)

CONFIG_ESP_DEBUG_STUBS_ENABLE

OpenOCD debug stubs

Found in: [Component config](#) > [ESP System Settings](#)

Debug stubs are used by OpenOCD to execute pre-compiled onboard code which does some useful debugging stuff, e.g. GCOV data dump.

CONFIG_ESP_DEBUG_OCDAWARE

Make exception and panic handlers JTAG/OCD aware

Found in: [Component config](#) > [ESP System Settings](#)

The FreeRTOS panic and unhandled exception handlers can detect a JTAG OCD debugger and instead of panicking, have the debugger stop on the offending instruction.

Default value:

- Yes (enabled)

CONFIG_ESP_SYSTEM_CHECK_INT_LEVEL

Interrupt level to use for Interrupt Watchdog and other system checks

Found in: [Component config](#) > [ESP System Settings](#)

Interrupt level to use for Interrupt Watchdog, IPC_ISR and other system checks.

Available options:

- Level 5 interrupt (CONFIG_ESP_SYSTEM_CHECK_INT_LEVEL_5)
Using level 5 interrupt for Interrupt Watchdog, IPC_ISR and other system checks.
- Level 4 interrupt (CONFIG_ESP_SYSTEM_CHECK_INT_LEVEL_4)
Using level 4 interrupt for Interrupt Watchdog, IPC_ISR and other system checks.

CONFIG_ESP_SYSTEM_HW_STACK_GUARD

Hardware stack guard

Found in: [Component config](#) > [ESP System Settings](#)

This config allows to trigger a panic interrupt when Stack Pointer register goes out of allocated stack memory bounds.

Default value:

- Yes (enabled) if SOC_ASSIST_DEBUG_SUPPORTED

IPC (Inter-Processor Call) Contains:

- [CONFIG_ESP_IPC_TASK_STACK_SIZE](#)
- [CONFIG_ESP_IPC_USES_CALLERS_PRIORITY](#)

CONFIG_ESP_IPC_TASK_STACK_SIZE

Inter-Processor Call (IPC) task stack size

Found in: [Component config](#) > [IPC \(Inter-Processor Call\)](#)

Configure the IPC tasks stack size. An IPC task runs on each core (in dual core mode), and allows for cross-core function calls. See IPC documentation for more details. The default IPC stack size should be enough for most common simple use cases. However, users can increase/decrease the stack size to their needs.

Range:

- from 512 to 65536

Default value:

- 1024

CONFIG_ESP_IPC_USES_CALLERS_PRIORITY

IPC runs at caller's priority

Found in: [Component config](#) > [IPC \(Inter-Processor Call\)](#)

If this option is not enabled then the IPC task will keep behavior same as prior to that of ESP-IDF v4.0, hence IPC task will run at (configMAX_PRIORITIES - 1) priority.

High resolution timer (esp_timer) Contains:

- [CONFIG_ESP_TIMER_PROFILING](#)
- [CONFIG_ESP_TIMER_TASK_AFFINITY](#)
- [CONFIG_ESP_TIMER_TASK_STACK_SIZE](#)
- [CONFIG_ESP_TIMER_INTERRUPT_LEVEL](#)
- [CONFIG_ESP_TIMER_SHOW_EXPERIMENTAL](#)
- [CONFIG_ESP_TIMER_SUPPORTS_ISR_DISPATCH_METHOD](#)
- [CONFIG_ESP_TIMER_ISR_AFFINITY](#)

CONFIG_ESP_TIMER_PROFILING

Enable esp_timer profiling features

Found in: [Component config](#) > [High resolution timer \(esp_timer\)](#)

If enabled, esp_timer_dump will dump information such as number of times the timer was started, number of times the timer has triggered, and the total time it took for the callback to run. This option has some effect on timer performance and the amount of memory used for timer storage, and should only be used for debugging/testing purposes.

Default value:

- No (disabled)

CONFIG_ESP_TIMER_TASK_STACK_SIZE

High-resolution timer task stack size

Found in: [Component config](#) > [High resolution timer \(esp_timer\)](#)

Configure the stack size of "timer_task" task. This task is used to dispatch callbacks of timers created using ets_timer and esp_timer APIs. If you are seeing stack overflow errors in timer task, increase this value.

Note that this is not the same as FreeRTOS timer task. To configure FreeRTOS timer task size, see "FreeRTOS timer task stack size" option in "FreeRTOS".

Range:

- from 2048 to 65536

Default value:

- 3584

CONFIG_ESP_TIMER_INTERRUPT_LEVEL

Interrupt level

Found in: [Component config](#) > [High resolution timer \(esp_timer\)](#)

It sets the interrupt level for esp_timer ISR in range 1..3. A higher level (3) helps to decrease the ISR esp_timer latency.

Range:

- from 1 to 1

Default value:

- 1

CONFIG_ESP_TIMER_SHOW_EXPERIMENTAL

show esp_timer's experimental features

Found in: [Component config](#) > [High resolution timer \(esp_timer\)](#)

This shows some hidden features of esp_timer. Note that they may break other features, use them with care.

CONFIG_ESP_TIMER_TASK_AFFINITY

esp_timer task core affinity

Found in: [Component config](#) > [High resolution timer \(esp_timer\)](#)

The default settings: timer TASK on CPU0 and timer ISR on CPU0. Other settings may help in certain cases, but note that they may break other features, use them with care. - "CPU0": (default) esp_timer task is processed by CPU0. - "CPU1": esp_timer task is processed by CPU1. - "No affinity": esp_timer task can be processed by any CPU.

Available options:

- CPU0 (CONFIG_ESP_TIMER_TASK_AFFINITY_CPU0)
- CPU1 (CONFIG_ESP_TIMER_TASK_AFFINITY_CPU1)
- No affinity (CONFIG_ESP_TIMER_TASK_AFFINITY_NO_AFFINITY)

CONFIG_ESP_TIMER_ISR_AFFINITY

timer interrupt core affinity

Found in: [Component config](#) > [High resolution timer \(esp_timer\)](#)

The default settings: timer TASK on CPU0 and timer ISR on CPU0. Other settings may help in certain cases, but note that they may break other features, use them with care. - "CPU0": (default) timer interrupt is processed by CPU0. - "CPU1": timer interrupt is processed by CPU1. - "No affinity": timer interrupt can be processed by any CPU. It helps to reduce latency but there is a disadvantage it leads to the timer ISR running on every core. It increases the CPU time usage for timer ISRs by N on an N-core system.

Available options:

- CPU0 (CONFIG_ESP_TIMER_ISR_AFFINITY_CPU0)
- CPU1 (CONFIG_ESP_TIMER_ISR_AFFINITY_CPU1)
- No affinity (CONFIG_ESP_TIMER_ISR_AFFINITY_NO_AFFINITY)

CONFIG_ESP_TIMER_SUPPORTS_ISR_DISPATCH_METHOD

Support ISR dispatch method

Found in: [Component config](#) > [High resolution timer \(esp_timer\)](#)

Allows using ESP_TIMER_ISR dispatch method (ESP_TIMER_TASK dispatch method is also available). - ESP_TIMER_TASK - Timer callbacks are dispatched from a high-priority esp_timer task. - ESP_TIMER_ISR - Timer callbacks are dispatched directly from the timer interrupt handler. The ISR dispatch can be used, in some cases, when a callback is very simple or need a lower-latency.

Default value:

- No (disabled)

Wi-Fi Contains:

- [CONFIG_ESP_WIFI_WPS_SOFTAP_REGISTRAR](#)
- [CONFIG_ESP_WIFI_11KV_SUPPORT](#)
- [CONFIG_ESP_WIFI_11R_SUPPORT](#)
- [CONFIG_ESP_WIFI_MODEM_RF_FLAG_UPDATE_DEBUG](#)
- [CONFIG_ESP_WIFI_DPP_SUPPORT](#)
- [CONFIG_ESP_WIFI_ENTERPRISE_SUPPORT](#)
- [CONFIG_ESP_WIFI_MBO_SUPPORT](#)
- [CONFIG_ESP_WIFI_SUITE_B_192](#)
- [CONFIG_ESP_WIFI_ENABLE_WPA3_OWE_STA](#)
- [CONFIG_ESP_WIFI_P256_ACCEL](#)
- [CONFIG_ESP_WIFI_WAPI_PSK](#)
- [CONFIG_ESP_WIFI_ENABLE_WIFI_RX_STATS](#)
- [CONFIG_ESP_WIFI_ENABLE_WIFI_TX_STATS](#)
- [CONFIG_ESP_WIFI_ENABLE_WPA3_SAE](#)
- [CONFIG_ESP_WIFI_SCAN_CACHE](#)
- [CONFIG_ESP_WIFI_SOFTAP_BEACON_MAX_LEN](#)
- [CONFIG_ESP_WIFI_CACHE_TX_BUFFER_NUM](#)

- `CONFIG_ESP_WIFI_DYNAMIC_RX_BUFFER_NUM`
- `CONFIG_ESP_WIFI_DYNAMIC_TX_BUFFER_NUM`
- `CONFIG_ESP_WIFI_RX_MGMT_BUF_NUM_DEF`
- `CONFIG_ESP_WIFI_STATIC_RX_BUFFER_NUM`
- `CONFIG_ESP_WIFI_STATIC_TX_BUFFER_NUM`
- `CONFIG_ESP_WIFI_ESPNOW_MAX_ENCRYPT_NUM`
- `CONFIG_ESP_WIFI_SLP_DEFAULT_MAX_ACTIVE_TIME`
- `CONFIG_ESP_WIFI_SLP_DEFAULT_MIN_ACTIVE_TIME`
- `CONFIG_ESP_WIFI_SLP_DEFAULT_WAIT_BROADCAST_DATA_TIME`
- `CONFIG_ESP_WIFI_STA_DISCONNECTED_PM_ENABLE`
- `CONFIG_ESP_WIFI_DEBUG_PRINT`
- `CONFIG_ESP_WIFI_MGMT_RX_BUFFER`
- `CONFIG_ESP_WIFI_TX_BUFFER`
- `CONFIG_ESP_WIFI_MBEDTLS_CRYPT`
- `CONFIG_ESP_WIFI_AMPDU_RX_ENABLED`
- `CONFIG_ESP_WIFI_AMPDU_TX_ENABLED`
- `CONFIG_ESP_WIFI_AMSDU_TX_ENABLED`
- `CONFIG_ESP_WIFI_NAN_ENABLE`
- `CONFIG_ESP_WIFI_CSI_ENABLED`
- `CONFIG_ESP_WIFI_EXTRA_IRAM_OPT`
- `CONFIG_ESP_WIFI_FTM_ENABLE`
- `CONFIG_ESP_WIFI_GCMP_SUPPORT`
- `CONFIG_ESP_WIFI_GMAC_SUPPORT`
- `CONFIG_ESP_WIFI_IRAM_OPT`
- `CONFIG_ESP_WIFI_MGMT_SBUF_NUM`
- `CONFIG_ESP_WIFI_ENHANCED_LIGHT_SLEEP`
- `CONFIG_ESP_WIFI_NVS_ENABLED`
- `CONFIG_ESP_WIFI_RX_IRAM_OPT`
- `CONFIG_ESP_WIFI_SLP_BEACON_LOST_OPT`
- `CONFIG_ESP_WIFI_SLP_IRAM_OPT`
- `CONFIG_ESP_WIFI_SOFTAP_SUPPORT`
- `CONFIG_ESP_WIFI_TASK_CORE_ID`
- *WPS Configuration Options*

CONFIG_ESP_WIFI_STATIC_RX_BUFFER_NUM

Max number of WiFi static RX buffers

Found in: [Component config](#) > [Wi-Fi](#)

Set the number of WiFi static RX buffers. Each buffer takes approximately 1.6KB of RAM. The static rx buffers are allocated when `esp_wifi_init` is called, they are not freed until `esp_wifi_deinit` is called.

WiFi hardware use these buffers to receive all 802.11 frames. A higher number may allow higher throughput but increases memory use. If `ESP_WIFI_AMPDU_RX_ENABLED` is enabled, this value is recommended to set equal or bigger than `ESP_WIFI_RX_BA_WIN` in order to achieve better throughput and compatibility with both stations and APs.

Range:

- from 2 to 128 if `SOC_WIFI_HE_SUPPORT`

Default value:

- 10

CONFIG_ESP_WIFI_DYNAMIC_RX_BUFFER_NUM

Max number of WiFi dynamic RX buffers

Found in: [Component config](#) > [Wi-Fi](#)

Set the number of WiFi dynamic RX buffers, 0 means unlimited RX buffers will be allocated (provided sufficient free RAM). The size of each dynamic RX buffer depends on the size of the received data

frame.

For each received data frame, the WiFi driver makes a copy to an RX buffer and then delivers it to the high layer TCP/IP stack. The dynamic RX buffer is freed after the higher layer has successfully received the data frame.

For some applications, WiFi data frames may be received faster than the application can process them. In these cases we may run out of memory if RX buffer number is unlimited (0).

If a dynamic RX buffer limit is set, it should be at least the number of static RX buffers.

Range:

- from 0 to 1024 if [CONFIG_LWIP_WND_SCALE](#)

Default value:

- 32

CONFIG_ESP_WIFI_TX_BUFFER

Type of WiFi TX buffers

Found in: [Component config](#) > [Wi-Fi](#)

Select type of WiFi TX buffers:

If "Static" is selected, WiFi TX buffers are allocated when WiFi is initialized and released when WiFi is de-initialized. The size of each static TX buffer is fixed to about 1.6KB.

If "Dynamic" is selected, each WiFi TX buffer is allocated as needed when a data frame is delivered to the Wifi driver from the TCP/IP stack. The buffer is freed after the data frame has been sent by the WiFi driver. The size of each dynamic TX buffer depends on the length of each data frame sent by the TCP/IP layer.

If PSRAM is enabled, "Static" should be selected to guarantee enough WiFi TX buffers. If PSRAM is disabled, "Dynamic" should be selected to improve the utilization of RAM.

Available options:

- Static (CONFIG_ESP_WIFI_STATIC_TX_BUFFER)
- Dynamic (CONFIG_ESP_WIFI_DYNAMIC_TX_BUFFER)

CONFIG_ESP_WIFI_STATIC_TX_BUFFER_NUM

Max number of WiFi static TX buffers

Found in: [Component config](#) > [Wi-Fi](#)

Set the number of WiFi static TX buffers. Each buffer takes approximately 1.6KB of RAM. The static RX buffers are allocated when `esp_wifi_init()` is called, they are not released until `esp_wifi_deinit()` is called.

For each transmitted data frame from the higher layer TCP/IP stack, the WiFi driver makes a copy of it in a TX buffer. For some applications especially UDP applications, the upper layer can deliver frames faster than WiFi layer can transmit. In these cases, we may run out of TX buffers.

Range:

- from 1 to 64 if [CONFIG_ESP_WIFI_STATIC_TX_BUFFER](#)

Default value:

- 16 if [CONFIG_ESP_WIFI_STATIC_TX_BUFFER](#)

CONFIG_ESP_WIFI_CACHE_TX_BUFFER_NUM

Max number of WiFi cache TX buffers

Found in: [Component config](#) > [Wi-Fi](#)

Set the number of WiFi cache TX buffer number.

For each TX packet from uplayer, such as LWIP etc, WiFi driver needs to allocate a static TX buffer and makes a copy of uplayer packet. If WiFi driver fails to allocate the static TX buffer, it caches the uplayer packets to a dedicated buffer queue, this option is used to configure the size of the cached TX queue.

CONFIG_ESP_WIFI_DYNAMIC_TX_BUFFER_NUM

Max number of WiFi dynamic TX buffers

Found in: [Component config](#) > [Wi-Fi](#)

Set the number of WiFi dynamic TX buffers. The size of each dynamic TX buffer is not fixed, it depends on the size of each transmitted data frame.

For each transmitted frame from the higher layer TCP/IP stack, the WiFi driver makes a copy of it in a TX buffer. For some applications, especially UDP applications, the upper layer can deliver frames faster than WiFi layer can transmit. In these cases, we may run out of TX buffers.

Range:

- from 1 to 128 if [CONFIG_ESP_WIFI_DYNAMIC_TX_BUFFER](#)

Default value:

- 32 if [CONFIG_ESP_WIFI_DYNAMIC_TX_BUFFER](#)

CONFIG_ESP_WIFI_MGMT_RX_BUFFER

Type of WiFi RX MGMT buffers

Found in: [Component config](#) > [Wi-Fi](#)

Select type of WiFi RX MGMT buffers:

If "Static" is selected, WiFi RX MGMT buffers are allocated when WiFi is initialized and released when WiFi is de-initialized. The size of each static RX MGMT buffer is fixed to about 500 Bytes.

If "Dynamic" is selected, each WiFi RX MGMT buffer is allocated as needed when a MGMT data frame is received. The MGMT buffer is freed after the MGMT data frame has been processed by the WiFi driver.

Available options:

- Static ([CONFIG_ESP_WIFI_STATIC_RX_MGMT_BUFFER](#))
- Dynamic ([CONFIG_ESP_WIFI_DYNAMIC_RX_MGMT_BUFFER](#))

CONFIG_ESP_WIFI_RX_MGMT_BUF_NUM_DEF

Max number of WiFi RX MGMT buffers

Found in: [Component config](#) > [Wi-Fi](#)

Set the number of WiFi RX_MGMT buffers.

For Management buffers, the number of dynamic and static management buffers is the same. In order to prevent memory fragmentation, the management buffer type should be set to static first.

Range:

- from 1 to 10

Default value:

- 5

CONFIG_ESP_WIFI_CSI_ENABLED

WiFi CSI(Channel State Information)

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable CSI(Channel State Information) feature. CSI takes about CONFIG_ESP_WIFI_STATIC_RX_BUFFER_NUM KB of RAM. If CSI is not used, it is better to disable this feature in order to save memory.

Default value:

- No (disabled) if SOC_WIFI_CSI_SUPPORT

CONFIG_ESP_WIFI_AMPDU_TX_ENABLED

WiFi AMPDU TX

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable AMPDU TX feature

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_TX_BA_WIN

WiFi AMPDU TX BA window size

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_AMPDU_TX_ENABLED](#)

Set the size of WiFi Block Ack TX window. Generally a bigger value means higher throughput but more memory. Most of time we should NOT change the default value unless special reason, e.g. test the maximum UDP TX throughput with iperf etc. For iperf test in shieldbox, the recommended value is 9~12.

Range:

- from 2 to 64 if SOC_WIFI_HE_SUPPORT && [CONFIG_ESP_WIFI_AMPDU_TX_ENABLED](#)

Default value:

- 6

CONFIG_ESP_WIFI_AMPDU_RX_ENABLED

WiFi AMPDU RX

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable AMPDU RX feature

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_RX_BA_WIN

WiFi AMPDU RX BA window size

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_AMPDU_RX_ENABLED](#)

Set the size of WiFi Block Ack RX window. Generally a bigger value means higher throughput and better compatibility but more memory. Most of time we should NOT change the default value unless special reason, e.g. test the maximum UDP RX throughput with iperf etc. For iperf test in shieldbox, the recommended value is 9~12. If PSRAM is used and WiFi memory is preferred to allocate in PSRAM first, the default and minimum value should be 16 to achieve better throughput and compatibility with both stations and APs.

Range:

- from 2 to 64 if SOC_WIFI_HE_SUPPORT && [CONFIG_ESP_WIFI_AMPDU_RX_ENABLED](#)

Default value:

- 6

CONFIG_ESP_WIFI_AMSDU_TX_ENABLED

WiFi AMSDU TX

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable AMSDU TX feature

Default value:

- No (disabled) if [CONFIG_ESP_WIFI_CACHE_TX_BUFFER_NUM](#) >= 2

CONFIG_ESP_WIFI_NVS_ENABLED

WiFi NVS flash

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable WiFi NVS flash

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_TASK_CORE_ID

WiFi Task Core ID

Found in: [Component config](#) > [Wi-Fi](#)

Pinned WiFi task to core 0 or core 1.

Available options:

- Core 0 ([CONFIG_ESP_WIFI_TASK_PINNED_TO_CORE_0](#))
- Core 1 ([CONFIG_ESP_WIFI_TASK_PINNED_TO_CORE_1](#))

CONFIG_ESP_WIFI_SOFTAP_BEACON_MAX_LEN

Max length of WiFi SoftAP Beacon

Found in: [Component config](#) > [Wi-Fi](#)

ESP-MESH utilizes beacon frames to detect and resolve root node conflicts (see documentation). However the default length of a beacon frame can simultaneously hold only five root node identifier structures, meaning that a root node conflict of up to five nodes can be detected at one time. In the occurrence of more root nodes conflict involving more than five root nodes, the conflict resolution process will detect five of the root nodes, resolve the conflict, and re-detect more root nodes. This process will repeat until all root node conflicts are resolved. However this process can generally take a very long time.

To counter this situation, the beacon frame length can be increased such that more root nodes can be detected simultaneously. Each additional root node will require 36 bytes and should be added on top of the default beacon frame length of 752 bytes. For example, if you want to detect 10 root nodes simultaneously, you need to set the beacon frame length as 932 (752+36*5).

Setting a longer beacon length also assists with debugging as the conflicting root nodes can be identified more quickly.

Range:

- from 752 to 1256

Default value:

- 752

CONFIG_ESP_WIFI_MGMT_SBUF_NUM

WiFi mgmt short buffer number

Found in: [Component config](#) > [Wi-Fi](#)

Set the maximum number of Wi-Fi management short buffers. These buffers are dynamically allocated, with their size determined by the length of the management packet to be sent. When a management packet is less than 64 bytes, the Wi-Fi driver classifies it as a short management packet and assigns it to one of these buffers.

Range:

- from 6 to 32

Default value:

- 32

CONFIG_ESP_WIFI_IRAM_OPT

WiFi IRAM speed optimization

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to place frequently called Wi-Fi library functions in IRAM. When this option is disabled, more than 10Kbytes of IRAM memory will be saved but Wi-Fi throughput will be reduced.

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_EXTRA_IRAM_OPT

WiFi EXTRA IRAM speed optimization

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to place additional frequently called Wi-Fi library functions in IRAM. When this option is disabled, more than 5Kbytes of IRAM memory will be saved but Wi-Fi throughput will be reduced.

Default value:

- No (disabled)

CONFIG_ESP_WIFI_RX_IRAM_OPT

WiFi RX IRAM speed optimization

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to place frequently called Wi-Fi library RX functions in IRAM. When this option is disabled, more than 17Kbytes of IRAM memory will be saved but Wi-Fi performance will be reduced.

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_ENABLE_WPA3_SAE

Enable WPA3-Personal

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to allow the device to establish a WPA3-Personal connection with eligible AP's. PMF (Protected Management Frames) is a prerequisite feature for a WPA3 connection, it needs to be explicitly configured before attempting connection. Please refer to the Wi-Fi Driver API Guide for details.

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_ENABLE_SAE_PK

Enable SAE-PK

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_ENABLE_WPA3_SAE](#)

Select this option to enable SAE-PK

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_SOFTAP_SAE_SUPPORT

Enable WPA3 Personal(SAE) SoftAP

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_ENABLE_WPA3_SAE](#)

Select this option to enable SAE support in softAP mode.

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_ENABLE_WPA3_OWE_STA

Enable OWE STA

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to allow the device to establish OWE connection with eligible AP's. PMF (Protected Management Frames) is a prerequisite feature for a WPA3 connection, it needs to be explicitly configured before attempting connection. Please refer to the Wi-Fi Driver API Guide for details.

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_SLP_IRAM_OPT

WiFi SLP IRAM speed optimization

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to place called Wi-Fi library TBTT process and receive beacon functions in IRAM. Some functions can be put in IRAM either by ESP_WIFI_IRAM_OPT and ESP_WIFI_RX_IRAM_OPT, or this one. If already enabled ESP_WIFI_IRAM_OPT, the other 7.3KB IRAM memory would be taken by this option. If already enabled ESP_WIFI_RX_IRAM_OPT, the other 1.3KB IRAM memory would be taken by this option. If neither of them are enabled, the other 7.4KB IRAM memory would be taken by this option. Wi-Fi power-save mode average current would be reduced if this option is enabled.

CONFIG_ESP_WIFI_SLP_DEFAULT_MIN_ACTIVE_TIME

Minimum active time

Found in: [Component config](#) > [Wi-Fi](#)

Only for station in WIFI_PS_MIN_MODEM or WIFI_PS_MAX_MODEM. When the station enters the active state, it will work for at least ESP_WIFI_SLP_DEFAULT_MIN_ACTIVE_TIME. If a data

packet is received or sent during this period, the time will be refreshed. If the time is up, but the station still has packets to receive or send, the time will also be refreshed. unit: milliseconds.

Range:

- from 8 to 60

Default value:

- 50

CONFIG_ESP_WIFI_SLP_DEFAULT_MAX_ACTIVE_TIME

Maximum keep alive time

Found in: [Component config](#) > [Wi-Fi](#)

Only for station in WIFI_PS_MIN_MODEM or WIFI_PS_MAX_MODEM. If no packet has been sent within ESP_WIFI_SLP_DEFAULT_MAX_ACTIVE_TIME, a null data packet will be sent to maintain the connection with the AP. unit: seconds.

Range:

- from 10 to 60

Default value:

- 10

CONFIG_ESP_WIFI_SLP_DEFAULT_WAIT_BROADCAST_DATA_TIME

Minimum wait broadcast data time

Found in: [Component config](#) > [Wi-Fi](#)

Only for station in WIFI_PS_MIN_MODEM or WIFI_PS_MAX_MODEM. When the station knows through the beacon that AP will send broadcast packet, it will wait for ESP_WIFI_SLP_DEFAULT_WAIT_BROADCAST_DATA_TIME before entering the sleep process. If a broadcast packet is received with more data bits, the time will refreshed. unit: milliseconds.

Range:

- from 10 to 30

Default value:

- 15

CONFIG_ESP_WIFI_FTM_ENABLE

WiFi FTM

Found in: [Component config](#) > [Wi-Fi](#)

Enable feature Fine Timing Measurement for calculating WiFi Round-Trip-Time (RTT).

Default value:

- No (disabled) if SOC_WIFI_FTM_SUPPORT

CONFIG_ESP_WIFI_FTM_INITIATOR_SUPPORT

FTM Initiator support

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_FTM_ENABLE](#)

Default value:

- Yes (enabled) if [CONFIG_ESP_WIFI_FTM_ENABLE](#)

CONFIG_ESP_WIFI_FTM_RESPONDER_SUPPORT

FTM Responder support

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_FTM_ENABLE](#)

Default value:

- Yes (enabled) if [CONFIG_ESP_WIFI_FTM_ENABLE](#)

CONFIG_ESP_WIFI_STA_DISCONNECTED_PM_ENABLE

Power Management for station at disconnected

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable power_management for station when disconnected. Chip will do modem-sleep when rf module is not in use any more.

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_GCMP_SUPPORT

WiFi GCMP Support(GCMP128 and GCMP256)

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable GCMP support. GCMP support is compulsory for WiFi Suite-B support.

Default value:

- No (disabled) if [SOC_WIFI_GCMP_SUPPORT](#)

CONFIG_ESP_WIFI_GMAC_SUPPORT

WiFi GMAC Support(GMAC128 and GMAC256)

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable GMAC support. GMAC support is compulsory for WiFi 192 bit certification.

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_SOFTAP_SUPPORT

WiFi SoftAP Support

Found in: [Component config](#) > [Wi-Fi](#)

WiFi module can be compiled without SoftAP to save code size.

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_ENHANCED_LIGHT_SLEEP

WiFi modem automatically receives the beacon

Found in: [Component config](#) > [Wi-Fi](#)

The wifi modem automatically receives the beacon frame during light sleep.

Default value:

- No (disabled) if [ESP_PHY_MAC_BB_PD](#) && [SOC_PM_SUPPORT_BEACON_WAKEUP](#)

CONFIG_ESP_WIFI_SLP_BEACON_LOST_OPT

Wifi sleep optimize when beacon lost

Found in: [Component config](#) > [Wi-Fi](#)

Enable wifi sleep optimization when beacon loss occurs and immediately enter sleep mode when the WiFi module detects beacon loss.

CONFIG_ESP_WIFI_SLP_BEACON_LOST_TIMEOUT

Beacon loss timeout

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_SLP_BEACON_LOST_OPT](#)

Timeout time for close rf phy when beacon loss occurs, Unit: 1024 microsecond.

Range:

- from 5 to 100 if [CONFIG_ESP_WIFI_SLP_BEACON_LOST_OPT](#)

Default value:

- 10 if [CONFIG_ESP_WIFI_SLP_BEACON_LOST_OPT](#)

CONFIG_ESP_WIFI_SLP_BEACON_LOST_THRESHOLD

Maximum number of consecutive lost beacons allowed

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_SLP_BEACON_LOST_OPT](#)

Maximum number of consecutive lost beacons allowed, WiFi keeps Rx state when the number of consecutive beacons lost is greater than the given threshold.

Range:

- from 0 to 8 if [CONFIG_ESP_WIFI_SLP_BEACON_LOST_OPT](#)

Default value:

- 3 if [CONFIG_ESP_WIFI_SLP_BEACON_LOST_OPT](#)

CONFIG_ESP_WIFI_SLP_PHY_ON_DELTA_EARLY_TIME

Delta early time for RF PHY on

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_SLP_BEACON_LOST_OPT](#)

Delta early time for rf phy on, When the beacon is lost, the next rf phy on will be earlier the time specified by the configuration item, Unit: 32 microsecond.

Range:

- from 0 to 100 if [CONFIG_ESP_WIFI_SLP_BEACON_LOST_OPT](#) &&
SOC_WIFI_SUPPORT_VARIABLE_BEACON_WINDOW

Default value:

- 2 if [CONFIG_ESP_WIFI_SLP_BEACON_LOST_OPT](#) &&
SOC_WIFI_SUPPORT_VARIABLE_BEACON_WINDOW

CONFIG_ESP_WIFI_SLP_PHY_OFF_DELTA_TIMEOUT_TIME

Delta timeout time for RF PHY off

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_SLP_BEACON_LOST_OPT](#)

Delta timeout time for rf phy off, When the beacon is lost, the next rf phy off will be delayed for the time specified by the configuration item. Unit: 1024 microsecond.

Range:

- from 0 to 8 if [CONFIG_ESP_WIFI_SLP_BEACON_LOST_OPT](#) &&
SOC_WIFI_SUPPORT_VARIABLE_BEACON_WINDOW

Default value:

- 2 if `CONFIG_ESP_WIFI_SLP_BEACON_LOST_OPT` &&
SOC_WIFI_SUPPORT_VARIABLE_BEACON_WINDOW

CONFIG_ESP_WIFI_ESPNOW_MAX_ENCRYPT_NUM

Maximum espnow encrypt peers number

Found in: [Component config](#) > [Wi-Fi](#)

Maximum number of encrypted peers supported by espnow. The number of hardware keys for encryption is fixed. And the espnow and SoftAP share the same hardware keys. So this configuration will affect the maximum connection number of SoftAP. Maximum espnow encrypted peers number + maximum number of connections of SoftAP = Max hardware keys number. When using ESP mesh, this value should be set to a maximum of 6.

Range:

- from 0 to 17

Default value:

- 7

CONFIG_ESP_WIFI_NAN_ENABLE

WiFi Aware

Found in: [Component config](#) > [Wi-Fi](#)

Enable WiFi Aware (NAN) feature.

Default value:

- No (disabled) if SOC_WIFI_NAN_SUPPORT

CONFIG_ESP_WIFI_ENABLE_WIFI_TX_STATS

Enable Wi-Fi transmission statistics

Found in: [Component config](#) > [Wi-Fi](#)

Enable Wi-Fi transmission statistics. Total support 4 access category. Each access category will use 346 bytes memory.

Default value:

- Yes (enabled) if SOC_WIFI_HE_SUPPORT

CONFIG_ESP_WIFI_MBEDTLS_CRYPT

Use MbedTLS crypto APIs

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable the use of MbedTLS crypto APIs. The internal crypto support within the supplicant is limited and may not suffice for all new security features, including WPA3.

It is recommended to always keep this option enabled. Additionally, note that MbedTLS can leverage hardware acceleration if available, resulting in significantly faster cryptographic operations.

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_MBEDTLS_TLS_CLIENT

Use MbedTLS TLS client for WiFi Enterprise connection

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_MBEDTLS_CRYPT](#)

Select this option to use MbedTLS TLS client for WPA2 enterprise connection. Please note that from MbedTLS-3.0 onwards, MbedTLS does not support SSL-3.0 TLS-v1.0, TLS-v1.1 versions. In case your server is using one of these versions, it is advisable to update your server. Please disable this option for compatibility with older TLS versions.

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_WAPI_PSK

Enable WAPI PSK support

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable WAPI-PSK which is a Chinese National Standard Encryption for Wireless LANs (GB 15629.11-2003).

Default value:

- No (disabled) if SOC_WIFI_WAPI_SUPPORT

CONFIG_ESP_WIFI_SUITE_B_192

Enable NSA suite B support with 192 bit key

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable 192 bit NSA suite-B. This is necessary to support WPA3 192 bit security.

Default value:

- No (disabled) if SOC_WIFI_GCMP_SUPPORT

CONFIG_ESP_WIFI_11KV_SUPPORT

Enable 802.11k, 802.11v APIs Support

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable 802.11k 802.11v APIs(RRM and BTM support). Only APIs which are helpful for network assisted roaming are supported for now. Enable this option with BTM and RRM enabled in sta config to make device ready for network assisted roaming. BTM: BSS transition management enables an AP to request a station to transition to a specific AP, or to indicate to a station a set of preferred APs. RRM: Radio measurements enable STAs to understand the radio environment, it enables STAs to observe and gather data on radio link performance and on the radio environment. Current implementation adds beacon report, link measurement, neighbor report.

Default value:

- No (disabled)

CONFIG_ESP_WIFI_RRM_SUPPORT

Enable 802.11k APIs Support

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_11KV_SUPPORT](#)

Select this option to enable 802.11k APIs(RRM support). Only APIs which are helpful for network assisted roaming are supported for now. Enable this option with RRM enabled in sta config to make device ready for network assisted roaming. RRM: Radio measurements enable STAs to understand the radio environment, it enables STAs to observe and gather data on radio link performance and on the radio environment. Current implementation adds beacon report, link measurement, neighbor report.

Default value:

- Yes (enabled) if [CONFIG_ESP_WIFI_11KV_SUPPORT](#)

CONFIG_ESP_WIFI_WNM_SUPPORT

Enable 802.11v APIs Support

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_11KV_SUPPORT](#)

Select this option to enable 802.11v APIs(BTM support). Only APIs which are helpful for network assisted roaming are supported for now. Enable this option with BTM enabled in sta config to make device ready for network assisted roaming. BTM: BSS transition management enables an AP to request a station to transition to a specific AP, or to indicate to a station a set of preferred APs.

Default value:

- Yes (enabled) if [CONFIG_ESP_WIFI_11KV_SUPPORT](#)

CONFIG_ESP_WIFI_SCAN_CACHE

Keep scan results in cache

Found in: [Component config](#) > [Wi-Fi](#)

Keep scan results in cache, if not enabled, those will be flushed immediately.

Default value:

- No (disabled) if [CONFIG_ESP_WIFI_RRM_SUPPORT](#)

CONFIG_ESP_WIFI_MBO_SUPPORT

Enable Multi Band Operation Certification Support

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable WiFi Multiband operation certification support.

Default value:

- No (disabled)

CONFIG_ESP_WIFI_DPP_SUPPORT

Enable DPP support

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable WiFi Easy Connect Support.

Default value:

- No (disabled)

CONFIG_ESP_WIFI_P256_ACCEL

Enable P-256 crypto acceleration

Found in: [Component config](#) > [Wi-Fi](#)

Enable Espressif-specific P-256 acceleration in the WPA supplicant crypto layer. This reduces SAE and DPP latency on supported targets at the cost of additional code size.

If disabled, the supplicant falls back to the generic Mbed TLS implementation.

Default value:

- No (disabled)

CONFIG_ESP_WIFI_11R_SUPPORT

Enable 802.11R (Fast Transition) Support

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable WiFi Fast Transition Support.

Default value:

- No (disabled)

CONFIG_ESP_WIFI_WPS_SOFTAP_REGISTRAR

Add WPS Registrar support in SoftAP mode

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to enable WPS registrar support in softAP mode.

Default value:

- No (disabled)

CONFIG_ESP_WIFI_ENABLE_WIFI_RX_STATS

Enable Wi-Fi reception statistics

Found in: [Component config](#) > [Wi-Fi](#)

Enable Wi-Fi reception statistics. Total support 2 access category. Each access category will use 190 bytes memory.

Default value:

- Yes (enabled) if SOC_WIFI_HE_SUPPORT

CONFIG_ESP_WIFI_ENABLE_WIFI_RX_MU_STATS

Enable Wi-Fi DL MU-MIMO and DL OFDMA reception statistics

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_ENABLE_WIFI_RX_STATS](#)

Enable Wi-Fi DL MU-MIMO and DL OFDMA reception statistics. Will use 10932 bytes memory.

Default value:

- Yes (enabled) if [CONFIG_ESP_WIFI_ENABLE_WIFI_RX_STATS](#)

WPS Configuration Options

 Contains:

- [CONFIG_ESP_WIFI_WPS_PASSPHRASE](#)
- [CONFIG_ESP_WIFI_WPS_STRICT](#)

CONFIG_ESP_WIFI_WPS_STRICT

Strictly validate all WPS attributes

Found in: [Component config](#) > [Wi-Fi](#) > [WPS Configuration Options](#)

Select this option to enable validate each WPS attribute rigorously. Disabling this add the workarounds with various APs. Enabling this may cause inter operability issues with some APs.

Default value:

- No (disabled)

CONFIG_ESP_WIFI_WPS_PASSPHRASE

Get WPA2 passphrase in WPS config

Found in: [Component config](#) > [Wi-Fi](#) > [WPS Configuration Options](#)

Select this option to get passphrase during WPS configuration. This option fakes the virtual display capabilities to get the configuration in passphrase mode. Not recommended to be used since WPS credentials should not be shared to other devices, making it in readable format increases that risk, also passphrase requires pbkdf2 to convert in psk.

Default value:

- No (disabled)

CONFIG_ESP_WIFI_DEBUG_PRINT

Print debug messages from WPA Supplicant

Found in: [Component config](#) > [Wi-Fi](#)

Select this option to print logging information from WPA supplicant, this includes handshake information and key hex dumps depending on the project logging level.

Enabling this could increase the build size ~60kb depending on the project logging level.

Default value:

- No (disabled)

CONFIG_ESP_WIFI_ENTERPRISE_SUPPORT

Enable enterprise option

Found in: [Component config](#) > [Wi-Fi](#)

Select this to enable/disable enterprise connection support.

disabling this will reduce binary size. disabling this will disable the use of any esp_wifi_sta_wpa2_ent_* (as APIs will be meaningless)

Default value:

- Yes (enabled)

CONFIG_ESP_WIFI_ENT_FREE_DYNAMIC_BUFFER

Free dynamic buffers during WiFi enterprise connection

Found in: [Component config](#) > [Wi-Fi](#) > [CONFIG_ESP_WIFI_ENTERPRISE_SUPPORT](#)

Select this configuration to free dynamic buffers during WiFi enterprise connection. This will enable chip to reduce heap consumption during WiFi enterprise connection.

Default value:

- No (disabled)

CONFIG_ESP_WIFI_MODEM_RF_FLAG_UPDATE_DEBUG

Enable debug assertions for modem RF flag update

Found in: [Component config](#) > [Wi-Fi](#)

Enable debug assertions to verify modem RF flag update operations. This option enables assert checks to verify that modem RF power state is correctly cleared before pmu sleep.

Default value:

- No (disabled) if [CONFIG_ESP_WIFI_ENHANCED_LIGHT_SLEEP](#)

Core dump Contains:

- [CONFIG_ESP_COREDUMP_CHECK_BOOT](#)
- [CONFIG_ESP_COREDUMP_DATA_FORMAT](#)
- [CONFIG_ESP_COREDUMP_CHECKSUM](#)
- [CONFIG_ESP_COREDUMP_TO_FLASH_OR_UART](#)
- [CONFIG_ESP_COREDUMP_UART_DELAY](#)
- [CONFIG_ESP_COREDUMP_LOGS](#)
- [CONFIG_ESP_COREDUMP_DECODE](#)
- [CONFIG_ESP_COREDUMP_MAX_TASKS_NUM](#)
- [CONFIG_ESP_COREDUMP_STACK_SIZE](#)
- [CONFIG_ESP_COREDUMP_SUMMARY_STACKDUMP_SIZE](#)

CONFIG_ESP_COREDUMP_TO_FLASH_OR_UART

Data destination

Found in: [Component config](#) > [Core dump](#)

Select place to store core dump: flash, uart or none (to disable core dumps generation).

Core dumps to Flash are not available if PSRAM is used for task stacks.

If core dump is configured to be stored in flash and custom partition table is used add corresponding entry to your CSV. For examples, please see predefined partition table CSV descriptions in the `components/partition_table` directory.

Available options:

- Flash ([CONFIG_ESP_COREDUMP_ENABLE_TO_FLASH](#))
- UART ([CONFIG_ESP_COREDUMP_ENABLE_TO_UART](#))
- None ([CONFIG_ESP_COREDUMP_ENABLE_TO_NONE](#))

CONFIG_ESP_COREDUMP_DATA_FORMAT

Core dump data format

Found in: [Component config](#) > [Core dump](#)

Select the data format for core dump.

Available options:

- Binary format ([CONFIG_ESP_COREDUMP_DATA_FORMAT_BIN](#))
- ELF format ([CONFIG_ESP_COREDUMP_DATA_FORMAT_ELF](#))

CONFIG_ESP_COREDUMP_CHECKSUM

Core dump data integrity check

Found in: [Component config](#) > [Core dump](#)

Select the integrity check for the core dump.

Available options:

- Use CRC32 for integrity verification ([CONFIG_ESP_COREDUMP_CHECKSUM_CRC32](#))
- Use SHA256 for integrity verification ([CONFIG_ESP_COREDUMP_CHECKSUM_SHA256](#))

CONFIG_ESP_COREDUMP_CHECK_BOOT

Check core dump data integrity on boot

Found in: [Component config](#) > [Core dump](#)

When enabled, if any data are found on the flash core dump partition, they will be checked by calculating their checksum.

Default value:

- Yes (enabled) if [CONFIG_ESP_COREDUMP_ENABLE_TO_FLASH](#)

CONFIG_ESP_COREDUMP_LOGS

Enable coredump logs for debugging

Found in: [Component config](#) > [Core dump](#)

Enable/disable coredump logs. Logs strings from espcoredump component are placed in DRAM. Disabling these helps to save ~5KB of internal memory.

CONFIG_ESP_COREDUMP_MAX_TASKS_NUM

Maximum number of tasks

Found in: [Component config](#) > [Core dump](#)

Maximum number of tasks that will be included in the core dump. Crashed task registers and stacks are always included. Other tasks are included in order of their priority. (Highest priority ready task first)

CONFIG_ESP_COREDUMP_UART_DELAY

Delay before print to UART

Found in: [Component config](#) > [Core dump](#)

Config delay (in ms) before printing core dump to UART. Delay can be interrupted by pressing Enter key.

Default value:

- 0 if [CONFIG_ESP_COREDUMP_ENABLE_TO_UART](#)

CONFIG_ESP_COREDUMP_STACK_SIZE

Reserved stack size

Found in: [Component config](#) > [Core dump](#)

Size of the memory to be reserved for core dump stack. If 0 core dump process will run on the stack of crashed task/ISR, otherwise special stack will be allocated. To ensure that core dump itself will not overflow task/ISR stack set this to the value above 800. NOTE: It eats DRAM.

CONFIG_ESP_COREDUMP_SUMMARY_STACKDUMP_SIZE

Size of the stack dump buffer

Found in: [Component config](#) > [Core dump](#)

Size of the buffer that would be reserved for extracting backtrace info summary. This buffer will contain the stack dump of the crashed task. This dump is useful in generating backtrace

Range:

- from 512 to 4096 if [CONFIG_ESP_COREDUMP_DATA_FORMAT_ELF](#) && [CONFIG_ESP_COREDUMP_ENABLE_TO_FLASH](#)

Default value:

- 1024 if `CONFIG_ESP_COREDUMP_DATA_FORMAT_ELF` && `CONFIG_ESP_COREDUMP_ENABLE_TO_FLASH`

CONFIG_ESP_COREDUMP_DECODE

Handling of UART core dumps in IDF Monitor

Found in: *Component config > Core dump*

Available options:

- Decode and show summary (info_corefile) (CONFIG_ESP_COREDUMP_DECODE_INFO)
- Don't decode (CONFIG_ESP_COREDUMP_DECODE_DISABLE)

FAT Filesystem support Contains:

- `CONFIG_FATFS_API_ENCODING`
- `CONFIG_FATFS_VFS_FSTAT_BLKSIZE`
- `CONFIG_FATFS_IMMEDIATE_FSYNC`
- `CONFIG_FATFS_USE_FASTSEEK`
- `CONFIG_FATFS_LONG_FILENAMES`
- `CONFIG_FATFS_MAX_LFN`
- `CONFIG_FATFS_FS_LOCK`
- `CONFIG_FATFS_VOLUME_COUNT`
- `CONFIG_FATFS_CHOOSE_CODEPAGE`
- `CONFIG_FATFS_ALLOC_PREFER_EXTRAM`
- `CONFIG_FATFS_SECTOR_SIZE`
- `CONFIG_FATFS_TIMEOUT_MS`
- `CONFIG_FATFS_PER_FILE_CACHE`

CONFIG_FATFS_VOLUME_COUNT

Number of volumes

Found in: *Component config > FAT Filesystem support*

Number of volumes (logical drives) to use.

Range:

- from 1 to 10

Default value:

- 2

CONFIG_FATFS_LONG_FILENAMES

Long filename support

Found in: *Component config > FAT Filesystem support*

Support long filenames in FAT. Long filename data increases memory usage. FATFS can be configured to store the buffer for long filename data in stack or heap.

Available options:

- No long filenames (CONFIG_FATFS_LFN_NONE)
- Long filename buffer in heap (CONFIG_FATFS_LFN_HEAP)
- Long filename buffer on stack (CONFIG_FATFS_LFN_STACK)

CONFIG_FATFS_SECTOR_SIZE

Sector size

Found in: [Component config > FAT Filesystem support](#)

Specify the size of the sector in bytes for FATFS partition generator.

Available options:

- 512 (CONFIG_FATFS_SECTOR_512)
- 4096 (CONFIG_FATFS_SECTOR_4096)

CONFIG_FATFS_CHOOSE_CODEPAGE

OEM Code Page

Found in: [Component config > FAT Filesystem support](#)

OEM code page used for file name encodings.

If "Dynamic" is selected, code page can be chosen at runtime using `f_setcp` function. Note that choosing this option will increase application size by ~480kB.

Available options:

- Dynamic (all code pages supported) (CONFIG_FATFS_CODEPAGE_DYNAMIC)
- US (CP437) (CONFIG_FATFS_CODEPAGE_437)
- Arabic (CP720) (CONFIG_FATFS_CODEPAGE_720)
- Greek (CP737) (CONFIG_FATFS_CODEPAGE_737)
- KBL (CP771) (CONFIG_FATFS_CODEPAGE_771)
- Baltic (CP775) (CONFIG_FATFS_CODEPAGE_775)
- Latin 1 (CP850) (CONFIG_FATFS_CODEPAGE_850)
- Latin 2 (CP852) (CONFIG_FATFS_CODEPAGE_852)
- Cyrillic (CP855) (CONFIG_FATFS_CODEPAGE_855)
- Turkish (CP857) (CONFIG_FATFS_CODEPAGE_857)
- Portugese (CP860) (CONFIG_FATFS_CODEPAGE_860)
- Icelandic (CP861) (CONFIG_FATFS_CODEPAGE_861)
- Hebrew (CP862) (CONFIG_FATFS_CODEPAGE_862)
- Canadian French (CP863) (CONFIG_FATFS_CODEPAGE_863)
- Arabic (CP864) (CONFIG_FATFS_CODEPAGE_864)
- Nordic (CP865) (CONFIG_FATFS_CODEPAGE_865)
- Russian (CP866) (CONFIG_FATFS_CODEPAGE_866)
- Greek 2 (CP869) (CONFIG_FATFS_CODEPAGE_869)
- Japanese (DBCS) (CP932) (CONFIG_FATFS_CODEPAGE_932)
- Simplified Chinese (DBCS) (CP936) (CONFIG_FATFS_CODEPAGE_936)
- Korean (DBCS) (CP949) (CONFIG_FATFS_CODEPAGE_949)
- Traditional Chinese (DBCS) (CP950) (CONFIG_FATFS_CODEPAGE_950)

CONFIG_FATFS_MAX_LFN

Max long filename length

Found in: [Component config > FAT Filesystem support](#)

Maximum long filename length. Can be reduced to save RAM.

CONFIG_FATFS_API_ENCODING

API character encoding

Found in: [Component config](#) > [FAT Filesystem support](#)

Choose encoding for character and string arguments/returns when using FATFS APIs. The encoding of arguments will usually depend on text editor settings.

Available options:

- API uses ANSI/OEM encoding (CONFIG_FATFS_API_ENCODING_ANSI_OEM)
- API uses UTF-8 encoding (CONFIG_FATFS_API_ENCODING_UTF_8)

CONFIG_FATFS_FS_LOCK

Number of simultaneously open files protected by lock function

Found in: [Component config](#) > [FAT Filesystem support](#)

This option sets the FATFS configuration value `_FS_LOCK`. The option `_FS_LOCK` switches file lock function to control duplicated file open and illegal operation to open objects.

* 0: Disable file lock function. To avoid volume corruption, application should avoid illegal open, remove and rename to the open objects.

* >0: Enable file lock function. The value defines how many files/sub-directories can be opened simultaneously under file lock control.

Note that the file lock control is independent of re-entrancy.

Range:

- from 0 to 65535

Default value:

- 0

CONFIG_FATFS_TIMEOUT_MS

Timeout for acquiring a file lock, ms

Found in: [Component config](#) > [FAT Filesystem support](#)

This option sets FATFS configuration value `_FS_TIMEOUT`, scaled to milliseconds. Sets the number of milliseconds FATFS will wait to acquire a mutex when operating on an open file. For example, if one task is performing a lengthy operation, another task will wait for the first task to release the lock, and time out after amount of time set by this option.

Default value:

- 10000

CONFIG_FATFS_PER_FILE_CACHE

Use separate cache for each file

Found in: [Component config](#) > [FAT Filesystem support](#)

This option affects FATFS configuration value `_FS_TINY`.

If this option is set, `_FS_TINY` is 0, and each open file has its own cache, size of the cache is equal to the `_MAX_SS` variable (512 or 4096 bytes). This option uses more RAM if more than 1 file is open, but needs less reads and writes to the storage for some operations.

If this option is not set, `_FS_TINY` is 1, and single cache is used for all open files, size is also equal to `_MAX_SS` variable. This reduces the amount of heap used when multiple files are open, but increases the number of read and write operations which FATFS needs to make.

Default value:

- Yes (enabled)

CONFIG_FATFS_ALLOC_PREFER_EXTRAM

Prefer external RAM when allocating FATFS buffers

Found in: [Component config](#) > [FAT Filesystem support](#)

When the option is enabled, internal buffers used by FATFS will be allocated from external RAM. If the allocation from external RAM fails, the buffer will be allocated from the internal RAM. Disable this option if optimizing for performance. Enable this option if optimizing for internal memory size.

Default value:

- Yes (enabled) if [CONFIG_SPIRAM_USE_CAPS_ALLOC](#) || [CONFIG_SPIRAM_USE_MALLOC](#)

CONFIG_FATFS_USE_FASTSEEK

Enable fast seek algorithm when using lseek function through VFS FAT

Found in: [Component config](#) > [FAT Filesystem support](#)

The fast seek feature enables fast backward/long seek operations without FAT access by using an in-memory CLMT (cluster link map table). Please note, fast-seek is only allowed for read-mode files, if a file is opened in write-mode, the seek mechanism will automatically fallback to the default implementation.

Default value:

- No (disabled)

CONFIG_FATFS_FAST_SEEK_BUFFER_SIZE

Fast seek CLMT buffer size

Found in: [Component config](#) > [FAT Filesystem support](#) > [CONFIG_FATFS_USE_FASTSEEK](#)

If fast seek algorithm is enabled, this defines the size of CLMT buffer used by this algorithm in 32-bit word units. This value should be chosen based on prior knowledge of maximum elements of each file entry would store.

Default value:

- 64 if [CONFIG_FATFS_USE_FASTSEEK](#)

CONFIG_FATFS_VFS_FSTAT_BLKSIZE

Default block size

Found in: [Component config](#) > [FAT Filesystem support](#)

If set to 0, the 'newlib' library's default size (BLKSIZ) is used (128 B). If set to a non-zero value, the value is used as the block size. Default file buffer size is set to this value and the buffer is allocated when first attempt of reading/writing to a file is made. Increasing this value improves fread() speed, however the heap usage is increased as well.

NOTE: The block size value is shared by all the filesystem functions accessing target media for given file descriptor! See 'Improving I/O performance' section of 'Maximizing Execution Speed' documentation page for more details.

Default value:

- 0

CONFIG_FATFS_IMMEDIATE_FSYNC

Enable automatic `f_sync`

Found in: *Component config > FAT Filesystem support*

Enables automatic calling of `f_sync()` to flush recent file changes after each call of `vfs_fat_write()`, `vfs_fat_pwrite()`, `vfs_fat_link()`, `vfs_fat_truncate()` and `vfs_fat_ftruncate()` functions. This feature improves file-consistency and size reporting accuracy for the FatFS, at a price on decreased performance due to frequent disk operations

Default value:

- No (disabled)

FreeRTOS Contains:

- *Kernel*
- *Port*

Kernel Contains:

- *CONFIG_FREERTOS_CHECK_STACKOVERFLOW*
- *CONFIG_FREERTOS_ENABLE_BACKWARD_COMPATIBILITY*
- *CONFIG_FREERTOS_GENERATE_RUN_TIME_STATS*
- *CONFIG_FREERTOS_MAX_TASK_NAME_LEN*
- *CONFIG_FREERTOS_IDLE_TASK_STACKSIZE*
- *CONFIG_FREERTOS_THREAD_LOCAL_STORAGE_POINTERS*
- *CONFIG_FREERTOS_QUEUE_REGISTRY_SIZE*
- *CONFIG_FREERTOS_TASK_NOTIFICATION_ARRAY_ENTRIES*
- *CONFIG_FREERTOS_HZ*
- *CONFIG_FREERTOS_USE_APPLICATION_TASK_TAG*
- *CONFIG_FREERTOS_USE_IDLE_HOOK*
- *CONFIG_FREERTOS_OPTIMIZED_SCHEDULER*
- *CONFIG_FREERTOS_USE_TICK_HOOK*
- *CONFIG_FREERTOS_USE_TICKLESS_IDLE*
- *CONFIG_FREERTOS_USE_TIMERS*
- *CONFIG_FREERTOS_USE_TRACE_FACILITY*
- *CONFIG_FREERTOS_UNICORE*

CONFIG_FREERTOS_UNICORE

Run FreeRTOS only on first core

Found in: *Component config > FreeRTOS > Kernel*

This version of FreeRTOS normally takes control of all cores of the CPU. Select this if you only want to start it on the first core. This is needed when e.g. another process needs complete control over the second core.

CONFIG_FREERTOS_HZ

`configTICK_RATE_HZ`

Found in: *Component config > FreeRTOS > Kernel*

Sets the FreeRTOS tick interrupt frequency in Hz (see `configTICK_RATE_HZ` documentation for more details).

Range:

- from 1 to 1000

Default value:

- 100

CONFIG_FREERTOS_OPTIMIZED_SCHEDULER

configUSE_PORT_OPTIMISED_TASK_SELECTION

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

Enables port specific task selection method. This option can speed up the search of ready tasks when scheduling (see configUSE_PORT_OPTIMISED_TASK_SELECTION documentation for more details).

Default value:

- Yes (enabled) if [CONFIG_FREERTOS_UNICORE](#)

CONFIG_FREERTOS_CHECK_STACKOVERFLOW

configCHECK_FOR_STACK_OVERFLOW

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

Enables FreeRTOS to check for stack overflows (see configCHECK_FOR_STACK_OVERFLOW documentation for more details).

Note: If users do not provide their own `vApplicationStackOverflowHook()` function, a default function will be provided by ESP-IDF.

Available options:

- No checking (CONFIG_FREERTOS_CHECK_STACKOVERFLOW_NONE)
Do not check for stack overflows (configCHECK_FOR_STACK_OVERFLOW = 0)
- Check by stack pointer value (Method 1) (CONFIG_FREERTOS_CHECK_STACKOVERFLOW_PTRVAL)
Check for stack overflows on each context switch by checking if the stack pointer is in a valid range. Quick but does not detect stack overflows that happened between context switches (configCHECK_FOR_STACK_OVERFLOW = 1)
- Check using canary bytes (Method 2) (CONFIG_FREERTOS_CHECK_STACKOVERFLOW_CANARY)
Places some magic bytes at the end of the stack area and on each context switch, check if these bytes are still intact. More thorough than just checking the pointer, but also slightly slower. (configCHECK_FOR_STACK_OVERFLOW = 2)

CONFIG_FREERTOS_THREAD_LOCAL_STORAGE_POINTERS

configNUM_THREAD_LOCAL_STORAGE_POINTERS

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

Set the number of thread local storage pointers in each task (see configNUM_THREAD_LOCAL_STORAGE_POINTERS documentation for more details).

Note: In ESP-IDF, this value must be at least 1. Index 0 is reserved for use by the pthreads API thread-local-storage. Other indexes can be used for any desired purpose.

Range:

- from 1 to 256

Default value:

- 1

CONFIG_FREERTOS_IDLE_TASK_STACKSIZE

configMINIMAL_STACK_SIZE (Idle task stack size)

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

Sets the idle task stack size in bytes (see `configMINIMAL_STACK_SIZE` documentation for more details).

Note:

- ESP-IDF specifies stack sizes in bytes instead of words.
- The default size is enough for most use cases.
- The stack size may need to be increased above the default if the app installs idle or thread local storage cleanup hooks that use a lot of stack memory.
- Conversely, the stack size can be reduced to the minimum if non of the idle features are used.

Range:

- from 768 to 32768

Default value:

- 1536

CONFIG_FREERTOS_USE_IDLE_HOOK

`configUSE_IDLE_HOOK`

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

Enables the idle task application hook (see `configUSE_IDLE_HOOK` documentation for more details).

Note:

- The application must provide the hook function `void vApplicationIdleHook( void ) ;`
- `vApplicationIdleHook()` is called from FreeRTOS idle task(s)
- The FreeRTOS idle hook is NOT the same as the ESP-IDF Idle Hook, but both can be enabled simultaneously.

Default value:

- No (disabled)

CONFIG_FREERTOS_USE_TICK_HOOK

`configUSE_TICK_HOOK`

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

Enables the tick hook (see `configUSE_TICK_HOOK` documentation for more details).

Note:

- The application must provide the hook function `void vApplicationTickHook( void ) ;`
- `vApplicationTickHook()` is called from FreeRTOS's tick handling function `xTaskIncrementTick()`
- The FreeRTOS tick hook is NOT the same as the ESP-IDF Tick Interrupt Hook, but both can be enabled simultaneously.

Default value:

- No (disabled)

CONFIG_FREERTOS_MAX_TASK_NAME_LEN

`configMAX_TASK_NAME_LEN`

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

Sets the maximum number of characters for task names (see `configMAX_TASK_NAME_LEN` documentation for more details).

Note: For most uses, the default of 16 characters is sufficient.

Range:

- from 1 to 256

Default value:

- 16

CONFIG_FREERTOS_ENABLE_BACKWARD_COMPATIBILITY

configENABLE_BACKWARD_COMPATIBILITY

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

Enable backward compatibility with APIs prior to FreeRTOS v8.0.0. (see configENABLE_BACKWARD_COMPATIBILITY documentation for more details).

Default value:

- No (disabled)

CONFIG_FREERTOS_USE_TIMERS

configUSE_TIMERS

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

Enable FreeRTOS Software Timers. Normally the timer task will only get pulled into the build and created if any software timer related functions are used. This is achieved through IDF defining a weak empty function for xTimerCreateTimerTask, which should take effect if timers.c is not pulled into the build.

In certain special cases (if you use configUSE_TRACE_FACILITY=y and event groups) the linker will still pull in the xTimerCreateTimerTask from timers.c even if the function that utilized it gets discarded due to not being used.

In these cases you can use this option to force the timer task to be disabled.

Default value:

- Yes (enabled)

CONFIG_FREERTOS_TIMER_SERVICE_TASK_NAME

configTIMER_SERVICE_TASK_NAME

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#) > [CONFIG_FREERTOS_USE_TIMERS](#)

Sets the timer task's name (see configTIMER_SERVICE_TASK_NAME documentation for more details).

Default value:

- "Tmr Svc"

CONFIG_FREERTOS_TIMER_TASK_PRIORITY

configTIMER_TASK_PRIORITY

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#) > [CONFIG_FREERTOS_USE_TIMERS](#)

Sets the timer task's priority (see configTIMER_TASK_PRIORITY documentation for more details).

Range:

- from 1 to 25

Default value:

- 1

CONFIG_FREERTOS_TIMER_TASK_STACK_DEPTH

configTIMER_TASK_STACK_DEPTH

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#) > [CONFIG_FREERTOS_USE_TIMERS](#)

Set the timer task's stack size (see configTIMER_TASK_STACK_DEPTH documentation for more details).

Range:

- from 1536 to 32768

Default value:

- 2048

CONFIG_FREERTOS_TIMER_QUEUE_LENGTH

configTIMER_QUEUE_LENGTH

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#) > [CONFIG_FREERTOS_USE_TIMERS](#)

Set the timer task's command queue length (see configTIMER_QUEUE_LENGTH documentation for more details).

Range:

- from 5 to 20

Default value:

- 10

CONFIG_FREERTOS_QUEUE_REGISTRY_SIZE

configQUEUE_REGISTRY_SIZE

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

Set the size of the queue registry (see configQUEUE_REGISTRY_SIZE documentation for more details).

Note: A value of 0 will disable queue registry functionality

Range:

- from 0 to 20

Default value:

- 0

CONFIG_FREERTOS_TASK_NOTIFICATION_ARRAY_ENTRIES

configTASK_NOTIFICATION_ARRAY_ENTRIES

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

Set the size of the task notification array of each task. When increasing this value, keep in mind that this means additional memory for each and every task on the system. However, task notifications in general are more light weight compared to alternatives such as semaphores.

Range:

- from 1 to 32

Default value:

- 1

CONFIG_FREERTOS_USE_TRACE_FACILITY

configUSE_TRACE_FACILITY

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

Enables additional structure members and functions to assist with execution visualization and tracing (see `configUSE_TRACE_FACILITY` documentation for more details).

Default value:

- No (disabled)

CONFIG_FREERTOS_USE_STATS_FORMATTING_FUNCTIONS

`configUSE_STATS_FORMATTING_FUNCTIONS`

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#) > [CONFIG_FREERTOS_USE_TRACE_FACILITY](#)

Set `configUSE_TRACE_FACILITY` and `configUSE_STATS_FORMATTING_FUNCTIONS` to 1 to include the `vTaskList()` and `vTaskGetRunTimeStats()` functions in the build (see `configUSE_STATS_FORMATTING_FUNCTIONS` documentation for more details).

Default value:

- No (disabled) if [CONFIG_FREERTOS_USE_TRACE_FACILITY](#)

CONFIG_FREERTOS_VTASKLIST_INCLUDE_COREID

Enable display of `xCoreID` in `vTaskList`

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#) > [CONFIG_FREERTOS_USE_TRACE_FACILITY](#) > [CONFIG_FREERTOS_USE_STATS_FORMATTING_FUNCTIONS](#)

If enabled, this will include an extra column when `vTaskList` is called to display the `CoreID` the task is pinned to (0,1) or -1 if not pinned.

Default value:

- No (disabled) if [CONFIG_FREERTOS_USE_STATS_FORMATTING_FUNCTIONS](#)

CONFIG_FREERTOS_GENERATE_RUN_TIME_STATS

`configGENERATE_RUN_TIME_STATS`

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

Enables collection of run time statistics for each task (see `configGENERATE_RUN_TIME_STATS` documentation for more details).

Note: The clock used for run time statistics can be configured in `FREERTOS_RUN_TIME_STATS_CLK`.

Default value:

- No (disabled)

CONFIG_FREERTOS_RUN_TIME_COUNTER_TYPE

`configRUN_TIME_COUNTER_TYPE`

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#) > [CONFIG_FREERTOS_GENERATE_RUN_TIME_STATS](#)

Sets the data type used for the FreeRTOS run time stats. A larger data type can be used to reduce the frequency of the counter overflowing.

Available options:

- `uint32_t` (`CONFIG_FREERTOS_RUN_TIME_COUNTER_TYPE_U32`)
`configRUN_TIME_COUNTER_TYPE` is set to `uint32_t`
- `uint64_t` (`CONFIG_FREERTOS_RUN_TIME_COUNTER_TYPE_U64`)
`configRUN_TIME_COUNTER_TYPE` is set to `uint64_t`

CONFIG_FREERTOS_USE_TICKLESS_IDLE

configUSE_TICKLESS_IDLE

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

If power management support is enabled, FreeRTOS will be able to put the system into light sleep mode when no tasks need to run for a number of ticks. This number can be set using FREERTOS_IDLE_TIME_BEFORE_SLEEP option. This feature is also known as "automatic light sleep".

Note that timers created using esp_timer APIs may prevent the system from entering sleep mode, even when no tasks need to run. To skip unnecessary wake-up initialize a timer with the "skip_unhandled_events" option as true.

If disabled, automatic light sleep support will be disabled.

Default value:

- No (disabled) if [CONFIG_PM_ENABLE](#)

CONFIG_FREERTOS_IDLE_TIME_BEFORE_SLEEP

configEXPECTED_IDLE_TIME_BEFORE_SLEEP

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#) > [CONFIG_FREERTOS_USE_TICKLESS_IDLE](#)

FreeRTOS will enter light sleep mode if no tasks need to run for this number of ticks. You can enable PM_PROFILING feature in esp_pm components and dump the sleep status with esp_pm_dump_locks, if the proportion of rejected sleeps is too high, please increase this value to improve scheduling efficiency

Range:

- from 2 to 4294967295 if [CONFIG_FREERTOS_USE_TICKLESS_IDLE](#)

Default value:

- 3 if [CONFIG_FREERTOS_USE_TICKLESS_IDLE](#)

CONFIG_FREERTOS_USE_APPLICATION_TASK_TAG

configUSE_APPLICATION_TASK_TAG

Found in: [Component config](#) > [FreeRTOS](#) > [Kernel](#)

Enables task tagging functionality and its associated API (see configUSE_APPLICATION_TASK_TAG documentation for more details).

Default value:

- No (disabled)

Port Contains:

- [CONFIG_FREERTOS_CHECK_MUTEX_GIVEN_BY_OWNER](#)
- [CONFIG_FREERTOS_RUN_TIME_STATS_CLK](#)
- [CONFIG_FREERTOS_INTERRUPT_BACKTRACE](#)
- [CONFIG_FREERTOS_WATCHPOINT_END_OF_STACK](#)
- [CONFIG_FREERTOS_ENABLE_STATIC_TASK_CLEAN_UP](#)
- [CONFIG_FREERTOS_TASK_PRE_DELETION_HOOK](#)
- [CONFIG_FREERTOS_TLSP_DELETION_CALLBACKS](#)
- [CONFIG_FREERTOS_ISR_STACKSIZE](#)
- [CONFIG_FREERTOS_PLACE_FUNCTIONS_INTO_FLASH](#)
- [CONFIG_FREERTOS_CHECK_PORT_CRITICAL_COMPLIANCE](#)
- [CONFIG_FREERTOS_CORETIMER](#)
- [CONFIG_FREERTOS_TASK_FUNCTION_WRAPPER](#)

CONFIG_FREERTOS_TASK_FUNCTION_WRAPPER

Wrap task functions

Found in: [Component config](#) > [FreeRTOS](#) > [Port](#)

If enabled, all FreeRTOS task functions will be enclosed in a wrapper function. If a task function mistakenly returns (i.e. does not delete), the call flow will return to the wrapper function. The wrapper function will then log an error and abort the application. This option is also required for GDB backtraces and C++ exceptions to work correctly inside top-level task functions.

Default value:

- Yes (enabled)

CONFIG_FREERTOS_WATCHPOINT_END_OF_STACK

Enable stack overflow debug watchpoint

Found in: [Component config](#) > [FreeRTOS](#) > [Port](#)

FreeRTOS can check if a stack has overflowed its bounds by checking either the value of the stack pointer or by checking the integrity of canary bytes. (See `CONFIG_FREERTOS_CHECK_STACKOVERFLOW` for more information.) These checks only happen on a context switch, and the situation that caused the stack overflow may already be long gone by then. This option will use the last debug memory watchpoint to allow breaking into the debugger (or panic'ing) as soon as any of the last 32 bytes on the stack of a task are overwritten. The side effect is that using gdb, you effectively have one hardware watchpoint less because the last one is overwritten as soon as a task switch happens.

Another consequence is that due to alignment requirements of the watchpoint, the usable stack size decreases by up to 60 bytes. This is because the watchpoint region has to be aligned to its size and the size for the stack watchpoint in IDF is 32 bytes.

This check only triggers if the stack overflow writes within 32 bytes near the end of the stack, rather than overshooting further, so it is worth combining this approach with one of the other stack overflow check methods.

When this watchpoint is hit, gdb will stop with a SIGTRAP message. When no JTAG OCD is attached, esp-idf will panic on an unhandled debug exception.

Default value:

- No (disabled)

CONFIG_FREERTOS_TLSP_DELETION_CALLBACKS

Enable thread local storage pointers deletion callbacks

Found in: [Component config](#) > [FreeRTOS](#) > [Port](#)

ESP-IDF provides users with the ability to free TLSP memory by registering TLSP deletion callbacks. These callbacks are automatically called by FreeRTOS when a task is deleted. When this option is turned on, the memory reserved for TLSPs in the TCB is doubled to make space for storing the deletion callbacks. If the user does not wish to use TLSP deletion callbacks then this option could be turned off to save space in the TCB memory.

Default value:

- Yes (enabled)

CONFIG_FREERTOS_TASK_PRE_DELETION_HOOK

Enable task pre-deletion hook

Found in: [Component config](#) > [FreeRTOS](#) > [Port](#)

Enable this option to make FreeRTOS call a user provided hook function right before it deletes a task (i.e., frees/releases a dynamically/statically allocated task's memory). This is useful if users want to know when a task is actually deleted (in case the task's deletion is delegated to the IDLE task).

If this config option is enabled, users must define a `void vTaskPreDeletionHook( void \* pxTCB )` hook function in their application.

CONFIG_FREERTOS_ENABLE_STATIC_TASK_CLEAN_UP

Enable static task clean up hook (DEPRECATED)

Found in: [Component config](#) > [FreeRTOS](#) > [Port](#)

THIS OPTION IS DEPRECATED. Use FREERTOS_TASK_PRE_DELETION_HOOK instead.

Enable this option to make FreeRTOS call the static task clean up hook when a task is deleted.

Note: Users will need to provide a `void vPortCleanUpTCB ( void \*pxTCB )` callback

Default value:

- No (disabled)

CONFIG_FREERTOS_CHECK_MUTEX_GIVEN_BY_OWNER

Check that mutex semaphore is given by owner task

Found in: [Component config](#) > [FreeRTOS](#) > [Port](#)

If enabled, assert that when a mutex semaphore is given, the task giving the semaphore is the task which is currently holding the mutex.

Default value:

- Yes (enabled)

CONFIG_FREERTOS_ISR_STACKSIZE

ISR stack size

Found in: [Component config](#) > [FreeRTOS](#) > [Port](#)

The interrupt handlers have their own stack. The size of the stack can be defined here. Each processor has its own stack, so the total size occupied will be twice this.

Range:

- from 2096 to 32768 if [CONFIG_ESP_COREDUMP_DATA_FORMAT_ELF](#)
- from 1536 to 32768

Default value:

- 2096 if [CONFIG_ESP_COREDUMP_DATA_FORMAT_ELF](#)
- 1536

CONFIG_FREERTOS_INTERRUPT_BACKTRACE

Enable backtrace from interrupt to task context

Found in: [Component config](#) > [FreeRTOS](#) > [Port](#)

If this option is enabled, interrupt stack frame will be modified to point to the code of the interrupted task as its return address. This helps the debugger (or the panic handler) show a backtrace from the interrupt to the task which was interrupted. This also works for nested interrupts: higher level interrupt stack can be traced back to the lower level interrupt. This option adds 4 instructions to the interrupt dispatching code.

Default value:

- Yes (enabled)

CONFIG_FREERTOS_CORETIMER

Tick timer source (Xtensa Only)

Found in: [Component config](#) > [FreeRTOS](#) > [Port](#)

FreeRTOS needs a timer with an associated interrupt to use as the main tick source to increase counters, run timers and do pre-emptive multitasking with. There are multiple timers available to do this, with different interrupt priorities.

Available options:

- Timer 0 (int 6, level 1) (CONFIG_FREERTOS_CORETIMER_0)
Select this to use timer 0
- Timer 1 (int 15, level 3) (CONFIG_FREERTOS_CORETIMER_1)
Select this to use timer 1
- SYSTIMER 0 (level 1) (CONFIG_FREERTOS_CORETIMER_SYSTIMER_LVL1)
Select this to use systimer with the 1 interrupt priority.
- SYSTIMER 0 (level 3) (CONFIG_FREERTOS_CORETIMER_SYSTIMER_LVL3)
Select this to use systimer with the 3 interrupt priority.

CONFIG_FREERTOS_RUN_TIME_STATS_CLK

Choose the clock source for run time stats

Found in: [Component config](#) > [FreeRTOS](#) > [Port](#)

Choose the clock source for FreeRTOS run time stats. Options are CPU0's CPU Clock or the ESP Timer. Both clock sources are 32 bits. The CPU Clock can run at a higher frequency hence provide a finer resolution but will overflow much quicker. Note that run time stats are only valid until the clock source overflows.

Available options:

- Use ESP TIMER for run time stats (CONFIG_FREERTOS_RUN_TIME_STATS_USING_ESP_TIMER)
ESP Timer will be used as the clock source for FreeRTOS run time stats. The ESP Timer runs at a frequency of 1MHz regardless of Dynamic Frequency Scaling. Therefore the ESP Timer will overflow in approximately 4290 seconds.
- Use CPU Clock for run time stats (CONFIG_FREERTOS_RUN_TIME_STATS_USING_CPU_CLK)
CPU Clock will be used as the clock source for the generation of run time stats. The CPU Clock has a frequency dependent on ESP_DEFAULT_CPU_FREQ_MHZ and Dynamic Frequency Scaling (DFS). Therefore the CPU Clock frequency can fluctuate between 80 to 240MHz. Run time stats generated using the CPU Clock represents the number of CPU cycles each task is allocated and DOES NOT reflect the amount of time each task runs for (as CPU clock frequency can change). If the CPU clock consistently runs at the maximum frequency of 240MHz, it will overflow in approximately 17 seconds.

CONFIG_FREERTOS_PLACE_FUNCTIONS_INTO_FLASH

Place FreeRTOS functions into Flash

Found in: [Component config](#) > [FreeRTOS](#) > [Port](#)

When enabled the selected Non-ISR FreeRTOS functions will be placed into Flash memory instead of IRAM. This saves up to 8KB of IRAM depending on which functions are used.

Default value:

- No (disabled)

CONFIG_FREERTOS_CHECK_PORT_CRITICAL_COMPLIANCE

Tests compliance with Vanilla FreeRTOS port*_CRITICAL calls

Found in: *Component config* > *FreeRTOS* > *Port*

If enabled, context of port*_CRITICAL calls (ISR or Non-ISR) would be checked to be in compliance with Vanilla FreeRTOS. e.g Calling port*_CRITICAL from ISR context would cause assert failure

Default value:

- No (disabled)

Hardware Abstraction Layer (HAL) and Low Level (LL) Contains:

- *CONFIG_HAL_DEFAULT_ASSERTION_LEVEL*
- *CONFIG_HAL_LOG_LEVEL*
- *CONFIG_HAL_SYSTIMER_USE_ROM_IMPL*
- *CONFIG_HAL_WDT_USE_ROM_IMPL*

CONFIG_HAL_DEFAULT_ASSERTION_LEVEL

Default HAL assertion level

Found in: *Component config* > *Hardware Abstraction Layer (HAL) and Low Level (LL)*

Set the assert behavior / level for HAL component. HAL component assert level can be set separately, but the level can't exceed the system assertion level. e.g. If the system assertion is disabled, then the HAL assertion can't be enabled either. If the system assertion is enable, then the HAL assertion can still be disabled by this Kconfig option.

Available options:

- Same as system assertion level (CONFIG_HAL_ASSERTION_EQUALS_SYSTEM)
- Disabled (CONFIG_HAL_ASSERTION_DISABLE)
- Silent (CONFIG_HAL_ASSERTION_SILENT)
- Enabled (CONFIG_HAL_ASSERTION_ENABLE)

CONFIG_HAL_LOG_LEVEL

HAL layer log verbosity

Found in: *Component config* > *Hardware Abstraction Layer (HAL) and Low Level (LL)*

Specify how much output to see in HAL logs.

Available options:

- No output (CONFIG_HAL_LOG_LEVEL_NONE)
- Error (CONFIG_HAL_LOG_LEVEL_ERROR)
- Warning (CONFIG_HAL_LOG_LEVEL_WARN)
- Info (CONFIG_HAL_LOG_LEVEL_INFO)
- Debug (CONFIG_HAL_LOG_LEVEL_DEBUG)
- Verbose (CONFIG_HAL_LOG_LEVEL_VERBOSE)

CONFIG_HAL_SYSTIMER_USE_ROM_IMPL

Use ROM implementation of SysTimer HAL driver

Found in: *Component config* > *Hardware Abstraction Layer (HAL) and Low Level (LL)*

Enable this flag to use HAL functions from ROM instead of ESP-IDF.

If keeping this as "n" in your project, you will have less free IRAM. If making this as "y" in your project, you will increase free IRAM, but you will lose the possibility to debug this module, and some new features will be added and bugs will be fixed in the IDF source but cannot be synced to ROM.

Default value:

- Yes (enabled)

CONFIG_HAL_WDT_USE_ROM_IMPL

Use ROM implementation of WDT HAL driver

Found in: [Component config](#) > [Hardware Abstraction Layer \(HAL\) and Low Level \(LL\)](#)

Enable this flag to use HAL functions from ROM instead of ESP-IDF.

If keeping this as "n" in your project, you will have less free IRAM. If making this as "y" in your project, you will increase free IRAM, but you will lose the possibility to debug this module, and some new features will be added and bugs will be fixed in the IDF source but cannot be synced to ROM.

Default value:

- Yes (enabled)

Heap memory debugging Contains:

- [CONFIG_HEAP_ABORT_WHEN_ALLOCATION_FAILS](#)
- [CONFIG_HEAP_TASK_TRACKING](#)
- [CONFIG_HEAP_PLACE_FUNCTION_INTO_FLASH](#)
- [CONFIG_HEAP_CORRUPTION_DETECTION](#)
- [CONFIG_HEAP_TRACING_DEST](#)
- [CONFIG_HEAP_TRACING_STACK_DEPTH](#)
- [CONFIG_HEAP_USE_HOOKS](#)
- [CONFIG_HEAP_TRACE_HASH_MAP](#)
- [CONFIG_HEAP_TLSF_USE_ROM_IMPL](#)

CONFIG_HEAP_CORRUPTION_DETECTION

Heap corruption detection

Found in: [Component config](#) > [Heap memory debugging](#)

Enable heap poisoning features to detect heap corruption caused by out-of-bounds access to heap memory.

See the "Heap Memory Debugging" page of the IDF documentation for a description of each level of heap corruption detection.

Available options:

- Basic (no poisoning) (CONFIG_HEAP_POISONING_DISABLED)
- Light impact (CONFIG_HEAP_POISONING_LIGHT)
- Comprehensive (CONFIG_HEAP_POISONING_COMPREHENSIVE)

CONFIG_HEAP_TRACING_DEST

Heap tracing

Found in: [Component config](#) > [Heap memory debugging](#)

Enables the heap tracing API defined in esp_heap_trace.h.

This function causes a moderate increase in IRAM code side and a minor increase in heap function (malloc/free/realloc) CPU overhead, even when the tracing feature is not used. So it's best to keep it disabled unless tracing is being used.

Available options:

- Disabled (CONFIG_HEAP_TRACING_OFF)
- Standalone (CONFIG_HEAP_TRACING_STANDALONE)
- Host-based (CONFIG_HEAP_TRACING_TOHOST)

CONFIG_HEAP_TRACING_STACK_DEPTH

Heap tracing stack depth

Found in: [Component config](#) > [Heap memory debugging](#)

Number of stack frames to save when tracing heap operation callers.

More stack frames uses more memory in the heap trace buffer (and slows down allocation), but can provide useful information.

CONFIG_HEAP_USE_HOOKS

Use allocation and free hooks

Found in: [Component config](#) > [Heap memory debugging](#)

Enable the user to implement function hooks triggered for each successful allocation and free.

CONFIG_HEAP_TASK_TRACKING

Enable heap task tracking

Found in: [Component config](#) > [Heap memory debugging](#)

Enables tracking the task responsible for each heap allocation.

This function depends on heap poisoning being enabled and adds four more bytes of overhead for each block allocated.

CONFIG_HEAP_TRACE_HASH_MAP

Use hash map mechanism to access heap trace records

Found in: [Component config](#) > [Heap memory debugging](#)

Enable this flag to use a hash map to increase performance in handling heap trace records.

Heap trace standalone supports storing records as a list, or a list + hash map.

Using only a list takes less memory, but calls to 'free' will get slower as the list grows. This is particularly affected when using HEAP_TRACE_ALL mode.

By using a list + hash map, calls to 'free' remain fast, at the cost of additional memory to store the hash map.

Default value:

- No (disabled) if [CONFIG_HEAP_TRACING_STANDALONE](#)

CONFIG_HEAP_TRACE_HASH_MAP_IN_EXT_RAM

Place hash map in external RAM

Found in: [Component config](#) > [Heap memory debugging](#) > [CONFIG_HEAP_TRACE_HASH_MAP](#)

When enabled this configuration forces the hash map to be placed in external RAM.

Default value:

- No (disabled) if [CONFIG_HEAP_TRACE_HASH_MAP](#)

CONFIG_HEAP_TRACE_HASH_MAP_SIZE

The number of entries in the hash map

Found in: [Component config](#) > [Heap memory debugging](#) > [CONFIG_HEAP_TRACE_HASH_MAP](#)

Defines the number of entries in the heap trace hashmap. Each entry takes 8 bytes. The bigger this number is, the better the performance. Recommended range: 200 - 2000.

Default value:

- 512 if [CONFIG_HEAP_TRACE_HASH_MAP](#)

CONFIG_HEAP_ABORT_WHEN_ALLOCATION_FAILS

Abort if memory allocation fails

Found in: [Component config](#) > [Heap memory debugging](#)

When enabled, if a memory allocation operation fails it will cause a system abort.

Default value:

- No (disabled)

CONFIG_HEAP_TLSF_USE_ROM_IMPL

Use ROM implementation of heap tlsf library

Found in: [Component config](#) > [Heap memory debugging](#)

Enable this flag to use heap functions from ROM instead of ESP-IDF.

If keeping this as "n" in your project, you will have less free IRAM. If making this as "y" in your project, you will increase free IRAM, but you will lose the possibility to debug this module, and some new features will be added and bugs will be fixed in the IDF source but cannot be synced to ROM.

Default value:

- Yes (enabled) if [ESP_ROM_HAS_HEAP_TLSF](#)

CONFIG_HEAP_PLACE_FUNCTION_INTO_FLASH

Force the entire heap component to be placed in flash memory

Found in: [Component config](#) > [Heap memory debugging](#)

Enable this flag to save up RAM space by placing the heap component in the flash memory

Note that it is only safe to enable this configuration if no functions from `esp_heap_caps.h` or `esp_heap_trace.h` are called from IRAM ISR which runs when cache is disabled.

IEEE 802.15.4 Contains:

- [CONFIG_IEEE802154_ENABLED](#)

CONFIG_IEEE802154_ENABLED

IEEE802154 Enable

Found in: *Component config* > *IEEE 802.15.4*

Default value:

- Yes (enabled) if SOC_IEEE802154_SUPPORTED

CONFIG_IEEE802154_RX_BUFFER_SIZE

The number of 802.15.4 receive buffers

Found in: *Component config* > *IEEE 802.15.4* > *CONFIG_IEEE802154_ENABLED*

Number of receive buffers allocated for IEEE 802.15.4 frames. Increasing this value allows more frames to be buffered before processing, but consumes more memory.

Range:

- from 2 to 100 if *CONFIG_IEEE802154_ENABLED*

Default value:

- 20 if *CONFIG_IEEE802154_ENABLED*

CONFIG_IEEE802154_CCA_MODE

Clear Channel Assessment (CCA) mode

Found in: *Component config* > *IEEE 802.15.4* > *CONFIG_IEEE802154_ENABLED*

Select the Clear Channel Assessment (CCA) mode used to determine if the channel is clear before transmission. CCA is required by the IEEE 802.15.4 standard to avoid collisions.

Available options:

- Carrier sense only (CONFIG_IEEE802154_CCA_CARRIER)
CCA reports channel busy only if a valid IEEE 802.15.4 signal is detected. This mode is less sensitive to non-802.15.4 interference.
- Energy above threshold (CONFIG_IEEE802154_CCA_ED)
CCA reports channel busy if the energy level exceeds the configured threshold. This is the default mode and provides good balance between sensitivity and reliability.
- Carrier sense OR energy above threshold (CONFIG_IEEE802154_CCA_CARRIER_OR_ED)
CCA reports channel busy if either a valid IEEE 802.15.4 signal is detected OR the energy level exceeds the threshold. This is the most sensitive mode.
- Carrier sense AND energy above threshold (CONFIG_IEEE802154_CCA_CARRIER_AND_ED)
CCA reports channel busy only if both a valid IEEE 802.15.4 signal is detected AND the energy level exceeds the threshold. This is the least sensitive mode.

CONFIG_IEEE802154_CCA_THRESHOLD

CCA detection threshold

Found in: *Component config* > *IEEE 802.15.4* > *CONFIG_IEEE802154_ENABLED*

Energy detection threshold for CCA, in dBm. The channel is considered busy if the detected energy exceeds this threshold. Lower values make the radio more sensitive to interference but may cause false positives.

Range:

- from -120 to 0 if *CONFIG_IEEE802154_ENABLED*

Default value:

- "-75" if *CONFIG_IEEE802154_ENABLED*

CONFIG_IEEE802154_PENDING_TABLE_SIZE

Pending table size

Found in: *Component config* > *IEEE 802.15.4* > *CONFIG_IEEE802154_ENABLED*

Size of the pending data table used to track frames waiting to be sent to sleeping devices. Each entry consumes memory, so adjust based on the number of sleeping child devices in your network.

Range:

- from 1 to 100 if *CONFIG_IEEE802154_ENABLED*

Default value:

- 20 if *CONFIG_IEEE802154_ENABLED*

CONFIG_IEEE802154_MULTI_PAN_ENABLE

Enable multi-pan feature for frame filter

Found in: *Component config* > *IEEE 802.15.4* > *CONFIG_IEEE802154_ENABLED*

Enable IEEE 802.15.4 multi-PAN (Personal Area Network) support. This allows the device to participate in multiple PANs simultaneously by filtering frames based on PAN ID.

Default value:

- No (disabled) if *CONFIG_IEEE802154_ENABLED*

CONFIG_IEEE802154_INTERFACE_NUM

Number of IEEE802154 interfaces

Found in: *Component config* > *IEEE 802.15.4* > *CONFIG_IEEE802154_ENABLED*

Number of IEEE802154 interfaces

Range:

- from 1 to 4 if *CONFIG_IEEE802154_MULTI_PAN_ENABLE* && *CONFIG_IEEE802154_ENABLED*

Default value:

- 2 if *CONFIG_IEEE802154_MULTI_PAN_ENABLE* && *CONFIG_IEEE802154_ENABLED*

CONFIG_IEEE802154_TIMING_OPTIMIZATION

Enable throughput optimization

Found in: *Component config* > *IEEE 802.15.4* > *CONFIG_IEEE802154_ENABLED*

Enabling this option increases throughput by ~5% at the expense of ~2.1k IRAM code size increase.

Default value:

- Yes (enabled) if *CONFIG_IEEE802154_ENABLED*

CONFIG_IEEE802154_SLEEP_ENABLE

Enable IEEE802154 light sleep

Found in: *Component config* > *IEEE 802.15.4* > *CONFIG_IEEE802154_ENABLED*

Enabling this option allows the IEEE802.15.4 module to be powered down during automatic light sleep, which reduces current consumption.

Default value:

- No (disabled) if *CONFIG_PM_ENABLE* && *CONFIG_IEEE802154_ENABLED*

CONFIG_IEEE802154_DEBUG

Enable IEEE802154 Debug

Found in: [Component config](#) > [IEEE 802.15.4](#) > [CONFIG_IEEE802154_ENABLED](#)

Enabling this option allows different kinds of IEEE802154 debug output. All IEEE802154 debug features increase the size of the final binary.

Default value:

- No (disabled) if [CONFIG_IEEE802154_ENABLED](#)

IEEE 802.15.4 Debug Parameters Contains:

- [CONFIG_IEEE802154_DEBUG_ASSERT_MONITOR](#)
- [CONFIG_IEEE802154_PRINT_PHY_REG](#)
- [CONFIG_IEEE802154_TXRX_STATISTIC](#)
- [CONFIG_IEEE802154_RECORD_TXRX_FRAME](#)
- [CONFIG_IEEE802154_ASSERT](#)
- [CONFIG_IEEE802154_RECORD](#)
- [CONFIG_IEEE802154_RX_BUFFER_STATISTIC](#)

CONFIG_IEEE802154_DEBUG_ASSERT_MONITOR

Enable IEEE802154 assert monitor

Found in: [Component config](#) > [IEEE 802.15.4](#) > [CONFIG_IEEE802154_ENABLED](#) > [CONFIG_IEEE802154_DEBUG](#) > [IEEE 802.15.4 Debug Parameters](#)

Enable monitoring and detection of abnormal or unexpected states during IEEE 802.15.4 operation. When enabled, performs additional runtime checks and assertions to catch potential driver issues early.

Default value:

- No (disabled) if [CONFIG_IEEE802154_DEBUG](#)

CONFIG_IEEE802154_RX_BUFFER_STATISTIC

Rx buffer statistic

Found in: [Component config](#) > [IEEE 802.15.4](#) > [CONFIG_IEEE802154_ENABLED](#) > [CONFIG_IEEE802154_DEBUG](#) > [IEEE 802.15.4 Debug Parameters](#)

Enable tracking of receive buffer allocation and deallocation. This provides statistics on buffer usage patterns, which can help optimize buffer sizing for your application.

Default value:

- No (disabled) if [CONFIG_IEEE802154_DEBUG](#)

CONFIG_IEEE802154_ASSERT

Enrich the assert information

Found in: [Component config](#) > [IEEE 802.15.4](#) > [CONFIG_IEEE802154_ENABLED](#) > [CONFIG_IEEE802154_DEBUG](#) > [IEEE 802.15.4 Debug Parameters](#)

Enable enhanced assert information. When an assertion fails, additional context including state, events, and commands will be printed to help diagnose the issue. Automatically enables IEEE802154_RECORD.

Default value:

- No (disabled) if [CONFIG_IEEE802154_DEBUG](#)

CONFIG_IEEE802154_RECORD

Record the information with IEEE802154 state and event

Found in: [Component config](#) > [IEEE 802.15.4](#) > [CONFIG_IEEE802154_ENABLED](#) > [CONFIG_IEEE802154_DEBUG](#) > [IEEE 802.15.4 Debug Parameters](#)

Enable recording of IEEE 802.15.4 driver state, events, and commands. Probe code is added throughout the driver to capture runtime information for debugging purposes. This is a prerequisite for other RECORD_* options.

Default value:

- No (disabled) if [CONFIG_IEEE802154_DEBUG](#)

CONFIG_IEEE802154_RECORD_EVENT

Enable record event information for debugging

Found in: [Component config](#) > [IEEE 802.15.4](#) > [CONFIG_IEEE802154_ENABLED](#) > [CONFIG_IEEE802154_DEBUG](#) > [IEEE 802.15.4 Debug Parameters](#) > [CONFIG_IEEE802154_RECORD](#)

Record IEEE 802.15.4 driver events in a circular buffer. When an assertion occurs, the recorded event history will be printed to help trace the sequence of events leading to the failure.

Default value:

- No (disabled) if [CONFIG_IEEE802154_RECORD](#) && [CONFIG_IEEE802154_DEBUG](#)

CONFIG_IEEE802154_RECORD_EVENT_SIZE

Record event table size

Found in: [Component config](#) > [IEEE 802.15.4](#) > [CONFIG_IEEE802154_ENABLED](#) > [CONFIG_IEEE802154_DEBUG](#) > [IEEE 802.15.4 Debug Parameters](#) > [CONFIG_IEEE802154_RECORD](#) > [CONFIG_IEEE802154_RECORD_EVENT](#)

Size of the circular buffer for recording events. Larger values provide more history but consume more memory.

Range:

- from 1 to 50 if [CONFIG_IEEE802154_RECORD_EVENT](#) && [CONFIG_IEEE802154_DEBUG](#)

Default value:

- 30 if [CONFIG_IEEE802154_RECORD_EVENT](#) && [CONFIG_IEEE802154_DEBUG](#)

CONFIG_IEEE802154_RECORD_STATE

Enable record state information for debugging

Found in: [Component config](#) > [IEEE 802.15.4](#) > [CONFIG_IEEE802154_ENABLED](#) > [CONFIG_IEEE802154_DEBUG](#) > [IEEE 802.15.4 Debug Parameters](#) > [CONFIG_IEEE802154_RECORD](#)

Record IEEE 802.15.4 driver state transitions in a circular buffer. When an assertion occurs, the recorded state history will be printed to show the sequence of state changes.

Default value:

- No (disabled) if [CONFIG_IEEE802154_RECORD](#) && [CONFIG_IEEE802154_DEBUG](#)

CONFIG_IEEE802154_RECORD_STATE_SIZE

Record state table size

Found in: [Component config](#) > [IEEE 802.15.4](#) > [CONFIG_IEEE802154_ENABLED](#) > [CONFIG_IEEE802154_DEBUG](#) > [IEEE 802.15.4 Debug Parameters](#) > [CONFIG_IEEE802154_RECORD](#) > [CONFIG_IEEE802154_RECORD_STATE](#)

Size of the circular buffer for recording state transitions. Larger values provide more history but consume more memory.

Range:

- from 1 to 50 if `CONFIG_IEEE802154_RECORD_STATE` && `CONFIG_IEEE802154_DEBUG`

Default value:

- 10 if `CONFIG_IEEE802154_RECORD_STATE` && `CONFIG_IEEE802154_DEBUG`

CONFIG_IEEE802154_RECORD_CMD

Enable record command information for debugging

Found in: `Component config > IEEE 802.15.4 > CONFIG_IEEE802154_ENABLED > CONFIG_IEEE802154_DEBUG > IEEE 802.15.4 Debug Parameters > CONFIG_IEEE802154_RECORD`

Record IEEE 802.15.4 driver commands in a circular buffer. When an assertion occurs, the recorded command history will be printed to show what commands were issued before the failure.

Default value:

- No (disabled) if `CONFIG_IEEE802154_RECORD` && `CONFIG_IEEE802154_DEBUG`

CONFIG_IEEE802154_RECORD_CMD_SIZE

Record command table size

Found in: `Component config > IEEE 802.15.4 > CONFIG_IEEE802154_ENABLED > CONFIG_IEEE802154_DEBUG > IEEE 802.15.4 Debug Parameters > CONFIG_IEEE802154_RECORD > CONFIG_IEEE802154_RECORD_CMD`

Size of the circular buffer for recording commands. Larger values provide more history but consume more memory.

Range:

- from 1 to 50 if `CONFIG_IEEE802154_RECORD_CMD` && `CONFIG_IEEE802154_DEBUG`

Default value:

- 10 if `CONFIG_IEEE802154_RECORD_CMD` && `CONFIG_IEEE802154_DEBUG`

CONFIG_IEEE802154_RECORD_ABORT

Enable record abort information for debugging

Found in: `Component config > IEEE 802.15.4 > CONFIG_IEEE802154_ENABLED > CONFIG_IEEE802154_DEBUG > IEEE 802.15.4 Debug Parameters > CONFIG_IEEE802154_RECORD`

Record IEEE 802.15.4 operation aborts in a circular buffer. When an assertion occurs, the recorded abort history will be printed to show what operations were aborted before the failure.

Default value:

- No (disabled) if `CONFIG_IEEE802154_RECORD` && `CONFIG_IEEE802154_DEBUG`

CONFIG_IEEE802154_RECORD_ABORT_SIZE

Record abort table size

Found in: `Component config > IEEE 802.15.4 > CONFIG_IEEE802154_ENABLED > CONFIG_IEEE802154_DEBUG > IEEE 802.15.4 Debug Parameters > CONFIG_IEEE802154_RECORD > CONFIG_IEEE802154_RECORD_ABORT`

Size of the circular buffer for recording aborts. Larger values provide more history but consume more memory.

Range:

- from 1 to 50 if `CONFIG_IEEE802154_RECORD_ABORT` && `CONFIG_IEEE802154_DEBUG`

Default value:

- 10 if `CONFIG_IEEE802154_RECORD_ABORT` && `CONFIG_IEEE802154_DEBUG`

CONFIG_IEEE802154_RECORD_TXRX_FRAME

Enable record txrx packets for debugging

Found in: `Component config > IEEE 802.15.4 > CONFIG_IEEE802154_ENABLED > CONFIG_IEEE802154_DEBUG > IEEE 802.15.4 Debug Parameters`

Record transmitted and received IEEE 802.15.4 frames in a circular buffer. When an assertion occurs, the recorded frame history will be printed to help diagnose frame-level issues. Note: This can consume significant memory depending on frame size and buffer size.

Default value:

- No (disabled) if `CONFIG_IEEE802154_DEBUG`

CONFIG_IEEE802154_RECORD_TXRX_FRAME_SIZE

Record frame table size

Found in: `Component config > IEEE 802.15.4 > CONFIG_IEEE802154_ENABLED > CONFIG_IEEE802154_DEBUG > IEEE 802.15.4 Debug Parameters > CONFIG_IEEE802154_RECORD_TXRX_FRAME`

Size of the circular buffer for recording frames. Larger values provide more frame history but consume significantly more memory (each frame can be up to 127 bytes plus metadata).

Range:

- from 1 to 50 if `CONFIG_IEEE802154_RECORD_TXRX_FRAME` && `CONFIG_IEEE802154_DEBUG`

Default value:

- 15 if `CONFIG_IEEE802154_RECORD_TXRX_FRAME` && `CONFIG_IEEE802154_DEBUG`

CONFIG_IEEE802154_TXRX_STATISTIC

Enable record tx/rx packets information for debugging

Found in: `Component config > IEEE 802.15.4 > CONFIG_IEEE802154_ENABLED > CONFIG_IEEE802154_DEBUG > IEEE 802.15.4 Debug Parameters`

Enable collection of transmit and receive packet statistics. This tracks counts, errors, and other metrics for transmitted and received frames. Statistics can be accessed via debug APIs for performance analysis.

Default value:

- No (disabled) if `CONFIG_IEEE802154_DEBUG`

CONFIG_IEEE802154_PRINT_PHY_REG

Enable print IEEE802154 PHY register command

Found in: `Component config > IEEE 802.15.4 > CONFIG_IEEE802154_ENABLED > CONFIG_IEEE802154_DEBUG > IEEE 802.15.4 Debug Parameters`

Enable the 'phyreg' console command to print IEEE 802.15.4 PHY register information for debugging purposes. This command displays register values and configuration settings of the IEEE 802.15.4 hardware.

Default value:

- No (disabled) if `CONFIG_IEEE802154_DEBUG`

Log output Contains:

- [CONFIG_LOG_DEFAULT_LEVEL](#)
- [CONFIG_LOG_MASTER_LEVEL](#)
- [CONFIG_LOG_TIMESTAMP_SOURCE](#)
- [CONFIG_LOG_MAXIMUM_LEVEL](#)
- [CONFIG_LOG_COLORS](#)

CONFIG_LOG_DEFAULT_LEVEL

Default log verbosity

Found in: [Component config](#) > [Log output](#)

Specify how much output to see in logs by default. You can set lower verbosity level at runtime using `esp_log_level_set` function.

By default, this setting limits which log statements are compiled into the program. For example, selecting "Warning" would mean that changing log level to "Debug" at runtime will not be possible. To allow increasing log level above the default at runtime, see the next option.

Available options:

- No output (`CONFIG_LOG_DEFAULT_LEVEL_NONE`)
- Error (`CONFIG_LOG_DEFAULT_LEVEL_ERROR`)
- Warning (`CONFIG_LOG_DEFAULT_LEVEL_WARN`)
- Info (`CONFIG_LOG_DEFAULT_LEVEL_INFO`)
- Debug (`CONFIG_LOG_DEFAULT_LEVEL_DEBUG`)
- Verbose (`CONFIG_LOG_DEFAULT_LEVEL_VERBOSE`)

CONFIG_LOG_MAXIMUM_LEVEL

Maximum log verbosity

Found in: [Component config](#) > [Log output](#)

This config option sets the highest log verbosity that it's possible to select at runtime by calling `esp_log_level_set()`. This level may be higher than the default verbosity level which is set when the app starts up.

This can be used enable debugging output only at a critical point, for a particular tag, or to minimize startup time but then enable more logs once the firmware has loaded.

Note that increasing the maximum available log level will increase the firmware binary size.

This option only applies to logging from the app, the bootloader log level is fixed at compile time to the separate "Bootloader log verbosity" setting.

Available options:

- Same as default (`CONFIG_LOG_MAXIMUM_EQUALS_DEFAULT`)
- Error (`CONFIG_LOG_MAXIMUM_LEVEL_ERROR`)
- Warning (`CONFIG_LOG_MAXIMUM_LEVEL_WARN`)
- Info (`CONFIG_LOG_MAXIMUM_LEVEL_INFO`)
- Debug (`CONFIG_LOG_MAXIMUM_LEVEL_DEBUG`)
- Verbose (`CONFIG_LOG_MAXIMUM_LEVEL_VERBOSE`)

CONFIG_LOG_MASTER_LEVEL

Enable global master log level

Found in: [Component config](#) > [Log output](#)

Enables an additional global "master" log level check that occurs before a log tag cache lookup. This is useful if you want to compile in a lot of logs that are selectable at runtime, but avoid the performance hit during periods where you don't want log output. Examples include remote log forwarding, or disabling logs during a time-critical or CPU-intensive section and re-enabling them later. Results in larger program size depending on number of logs compiled in.

If enabled, defaults to LOG_DEFAULT_LEVEL and can be set using esp_log_set_level_master(). This check takes precedence over ESP_LOG_LEVEL_LOCAL.

Default value:

- No (disabled)

CONFIG_LOG_COLORS

Use ANSI terminal colors in log output

Found in: [Component config](#) > [Log output](#)

Enable ANSI terminal color codes in bootloader output.

In order to view these, your terminal program must support ANSI color codes.

Default value:

- Yes (enabled)

CONFIG_LOG_TIMESTAMP_SOURCE

Log Timestamps

Found in: [Component config](#) > [Log output](#)

Choose what sort of timestamp is displayed in the log output:

- Milliseconds since boot is calculated from the RTOS tick count multiplied by the tick period. This time will reset after a software reboot. e.g. (90000)
- System time is taken from POSIX time functions which use the chip's RTC and high resolution timers to maintain an accurate time. The system time is initialized to 0 on startup, it can be set with an SNTP sync, or with POSIX time functions. This time will not reset after a software reboot. e.g. (00:01:30.000)
- NOTE: Currently this will not get used in logging from binary blobs (i.e WiFi & Bluetooth libraries), these will always print milliseconds since boot.

Available options:

- Milliseconds Since Boot (CONFIG_LOG_TIMESTAMP_SOURCE_RTOS)
- System Time (CONFIG_LOG_TIMESTAMP_SOURCE_SYSTEM)

LWIP Contains:

- [CONFIG_LWIP_CHECK_THREAD_SAFETY](#)
- [Checksums](#)
- [CONFIG_LWIP_DHCP_COARSE_TIMER_SECS](#)
- [DHCP server](#)
- [CONFIG_LWIP_DHCP_OPTIONS_LEN](#)
- [CONFIG_LWIP_DHCP_DISABLE_CLIENT_ID](#)
- [CONFIG_LWIP_DHCP_DISABLE_VENDOR_CLASS_ID](#)
- [CONFIG_LWIP_DHCP_DOES_ARP_CHECK](#)

- *CONFIG_LWIP_DHCP_RESTORE_LAST_IP*
- *DNS*
- *CONFIG_LWIP_PPP_CHAP_SUPPORT*
- *CONFIG_LWIP_L2_TO_L3_COPY*
- *CONFIG_LWIP_IPV6_DHCP6*
- *CONFIG_LWIP_IP4_FRAG*
- *CONFIG_LWIP_IP6_FRAG*
- *CONFIG_LWIP_IP_FORWARD*
- *CONFIG_LWIP_NETBUF_RECVINFO*
- *CONFIG_LWIP_IPV4*
- *CONFIG_LWIP_AUTOIP*
- *CONFIG_LWIP_IPV6*
- *CONFIG_LWIP_ENABLE_LCP_ECHO*
- *CONFIG_LWIP_ESP_LWIP_ASSERT*
- *CONFIG_LWIP_DEBUG*
- *CONFIG_LWIP_IRAM_OPTIMIZATION*
- *CONFIG_LWIP_EXTRA_IRAM_OPTIMIZATION*
- *CONFIG_LWIP_ENABLE*
- *CONFIG_LWIP_STATS*
- *CONFIG_LWIP_TIMERS_ONDEMAND*
- *CONFIG_LWIP_DNS_SUPPORT_MDNS_QUERIES*
- *CONFIG_LWIP_PPP_MPPE_SUPPORT*
- *CONFIG_LWIP_PPP_MSCHAP_SUPPORT*
- *CONFIG_LWIP_PPP_NOTIFY_PHASE_SUPPORT*
- *CONFIG_LWIP_PPP_PAP_SUPPORT*
- *CONFIG_LWIP_PPP_DEBUG_ON*
- *CONFIG_LWIP_PPP_SERVER_SUPPORT*
- *CONFIG_LWIP_PPP_SUPPORT*
- *CONFIG_LWIP_IP4_REASSEMBLY*
- *CONFIG_LWIP_IP6_REASSEMBLY*
- *CONFIG_LWIP_SLIP_SUPPORT*
- *CONFIG_LWIP_SO_LINGER*
- *CONFIG_LWIP_SO_RCVBUF*
- *CONFIG_LWIP_SO_REUSE*
- *CONFIG_LWIP_NETIF_STATUS_CALLBACK*
- *CONFIG_LWIP_TCPIP_CORE_LOCKING*
- *CONFIG_LWIP_NETIF_API*
- *CONFIG_LWIP_PPP_VJ_HEADER_COMPRESSION*
- *Hooks*
- *ICMP*
- *CONFIG_LWIP_LOCAL_HOSTNAME*
- *CONFIG_LWIP_ND6*
- *LWIP RAW API*
- *CONFIG_LWIP_TCPIP_TASK_PRIO*
- *CONFIG_LWIP_IPV6_ND6_NUM_ROUTERS*
- *CONFIG_LWIP_IPV6_ND6_NUM_DESTINATIONS*
- *CONFIG_LWIP_IPV6_ND6_NUM_NEIGHBORS*
- *CONFIG_LWIP_IPV6_ND6_NUM_PREFIXES*
- *CONFIG_LWIP_IPV6_MEMP_NUM_ND6_QUEUE*
- *CONFIG_LWIP_MAX_SOCKETS*
- *CONFIG_LWIP_BRIDGEIF_MAX_PORTS*
- *CONFIG_LWIP_NUM_NETIF_CLIENT_DATA*
- *CONFIG_LWIP_ESP_GRATUITOUS_ARP*
- *CONFIG_LWIP_ESP_MLDV6_REPORT*
- *SNTP*
- *CONFIG_LWIP_USE_ONLY_LWIP_SELECT*
- *CONFIG_LWIP_NETIF_LOOPBACK*
- *TCP*

- [CONFIG_LWIP_TCPIP_TASK_AFFINITY](#)
- [CONFIG_LWIP_TCPIP_TASK_STACK_SIZE](#)
- [CONFIG_LWIP_TCPIP_RECVMBOX_SIZE](#)
- [CONFIG_LWIP_IP_REASS_MAX_PBUFS](#)
- [CONFIG_LWIP_IP_DEFAULT_TTL](#)
- [UDP](#)
- [CONFIG_LWIP_IPV6_RDNSS_MAX_DNS_SERVERS](#)
- [CONFIG_LWIP_USE_EXTERNAL_MBEDTLS](#)

CONFIG_LWIP_ENABLE

Enable LwIP stack

Found in: [Component config](#) > [LWIP](#)

Builds normally if selected. Excludes LwIP from build if unselected, even if it is a dependency of a component or application. Some applications can switch their IP stacks, e.g., when switching between chip and Linux targets (LwIP stack vs. Linux IP stack). Since the LwIP dependency cannot easily be excluded based on a Kconfig option, it has to be a dependency in all cases. This switch allows the LwIP stack to be built selectively, even if it is a dependency.

Default value:

- Yes (enabled)

CONFIG_LWIP_LOCAL_HOSTNAME

Local netif hostname

Found in: [Component config](#) > [LWIP](#)

The default name this device will report to other devices on the network. Could be updated at runtime with `esp_netif_set_hostname()`

Default value:

- "espressif"

CONFIG_LWIP_NETIF_API

Enable usage of standard POSIX APIs in LWIP

Found in: [Component config](#) > [LWIP](#)

If this feature is enabled, standard POSIX APIs: `if_indextoname()`, `if_nametoindex()` could be used to convert network interface index to name instead of IDF specific esp-netif APIs (such as `esp_netif_get_netif_impl_name()`)

Default value:

- No (disabled)

CONFIG_LWIP_TCPIP_TASK_PRIO

LWIP TCP/IP Task Priority

Found in: [Component config](#) > [LWIP](#)

LWIP tcpip task priority. In case of high throughput, this parameter could be changed up to `(config-MAX_PRIORITIES-1)`.

Range:

- from 1 to 24

Default value:

- 18

CONFIG_LWIP_TCPIP_CORE_LOCKING

Enable tcpip core locking

Found in: [Component config](#) > [LWIP](#)

If Enable tcpip core locking, Creates a global mutex that is held during TCPIP thread operations. Can be locked by client code to perform lwIP operations without changing into TCPIP thread using callbacks. See LOCK_TCPIP_CORE() and UNLOCK_TCPIP_CORE().

If disable tcpip core locking, TCP IP will perform tasks through context switching

Default value:

- No (disabled)

CONFIG_LWIP_TCPIP_CORE_LOCKING_INPUT

Enable tcpip core locking input

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_TCPIP_CORE_LOCKING](#)

when LWIP_TCPIP_CORE_LOCKING is enabled, this lets tcpip_input() grab the mutex for input packets as well, instead of allocating a message and passing it to tcpip_thread.

Default value:

- No (disabled) if [CONFIG_LWIP_TCPIP_CORE_LOCKING](#)

CONFIG_LWIP_CHECK_THREAD_SAFETY

Checks that lwip API runs in expected context

Found in: [Component config](#) > [LWIP](#)

Enable to check that the project does not violate lwip thread safety. If enabled, all lwip functions that require thread awareness run an assertion to verify that the TCP/IP core functionality is either locked or accessed from the correct thread.

Default value:

- No (disabled)

CONFIG_LWIP_DNS_SUPPORT_MDNS_QUERIES

Enable mDNS queries in resolving host name

Found in: [Component config](#) > [LWIP](#)

If this feature is enabled, standard API such as gethostbyname support .local addresses by sending one shot multicast mDNS query

Default value:

- Yes (enabled)

CONFIG_LWIP_L2_TO_L3_COPY

Enable copy between Layer2 and Layer3 packets

Found in: [Component config](#) > [LWIP](#)

If this feature is enabled, all traffic from layer2(WIFI Driver) will be copied to a new buffer before sending it to layer3(LWIP stack), freeing the layer2 buffer. Please be notified that the total layer2 receiving buffer is fixed and ESP32 currently supports 25 layer2 receiving buffer, when layer2 buffer runs out of memory, then the incoming packets will be dropped in hardware. The layer3 buffer is allocated from the heap, so the total layer3 receiving buffer depends on the available heap size, when heap runs out of memory, no copy will be sent to layer3 and packet will be dropped in layer2. Please make sure you fully understand the impact of this feature before enabling it.

Default value:

- No (disabled)

CONFIG_LWIP_IRAM_OPTIMIZATION

Enable LWIP IRAM optimization

Found in: [Component config](#) > [LWIP](#)

If this feature is enabled, some functions relating to RX/TX in LWIP will be put into IRAM, it can improve UDP/TCP throughput by >10% for single core mode, it doesn't help too much for dual core mode. On the other hand, it needs about 10KB IRAM for these optimizations.

If this feature is disabled, all lwip functions will be put into FLASH.

Default value:

- No (disabled)

CONFIG_LWIP_EXTRA_IRAM_OPTIMIZATION

Enable LWIP IRAM optimization for TCP part

Found in: [Component config](#) > [LWIP](#)

If this feature is enabled, some tcp part functions relating to RX/TX in LWIP will be put into IRAM, it can improve TCP throughput. On the other hand, it needs about 17KB IRAM for these optimizations.

Default value:

- No (disabled)

CONFIG_LWIP_TIMERS_ONDEMAND

Enable LWIP Timers on demand

Found in: [Component config](#) > [LWIP](#)

If this feature is enabled, IGMP and MLD6 timers will be activated only when joining groups or receiving QUERY packets.

This feature will reduce the power consumption for applications which do not use IGMP and MLD6.

Default value:

- Yes (enabled)

CONFIG_LWIP_ND6

LWIP NDP6 Enable/Disable

Found in: [Component config](#) > [LWIP](#)

This option is used to disable the Network Discovery Protocol (NDP) if it is not required. Please use this option with caution, as the NDP is essential for IPv6 functionality within a local network.

Default value:

- Yes (enabled)

CONFIG_LWIP_FORCE_ROUTER_FORWARDING

LWIP Force Router Forwarding Enable/Disable

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_ND6](#)

This option is used to set the router flag for the NA packets. When enabled, the router flag in NA packet will always set to 1, otherwise, never set router flag for NA packets.

Default value:

- No (disabled)

CONFIG_LWIP_MAX_SOCKETS

Max number of open sockets

Found in: [Component config](#) > [LWIP](#)

The practical maximum limit is determined by available heap memory at runtime.

Sockets take up a certain amount of memory, and allowing fewer sockets to be open at the same time conserves memory. Specify the maximum amount of sockets here. The valid value is from 1 to 253. If using value above 61, update CMakeLists defining FD_SETSIZE to the number of sockets used plus the expected open files (minimum of +3 for stdout, stderr and stdin).

Range:

- from 1 to 253

Default value:

- 10

CONFIG_LWIP_USE_ONLY_LWIP_SELECT

Support LWIP socket select() only (DEPRECATED)

Found in: [Component config](#) > [LWIP](#)

This option is deprecated. Do not use this option, use VFS_SUPPORT_SELECT instead.

Default value:

- No (disabled)

CONFIG_LWIP_SO_LINGER

Enable SO_LINGER processing

Found in: [Component config](#) > [LWIP](#)

Enabling this option allows SO_LINGER processing. l_onoff = 1, l_linger can set the timeout.

If l_linger=0, When a connection is closed, TCP will terminate the connection. This means that TCP will discard any data packets stored in the socket send buffer and send an RST to the peer.

If l_linger!=0, Then closesocket() calls to block the process until the remaining data packets has been sent or timed out.

Default value:

- No (disabled)

CONFIG_LWIP_SO_REUSE

Enable SO_REUSEADDR option

Found in: [Component config](#) > [LWIP](#)

Enabling this option allows binding to a port which remains in TIME_WAIT.

Default value:

- Yes (enabled)

CONFIG_LWIP_SO_REUSE_RXTOALL

SO_REUSEADDR copies broadcast/multicast to all matches

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_SO_REUSE](#)

Enabling this option means that any incoming broadcast or multicast packet will be copied to all of the local sockets that it matches (may be more than one if SO_REUSEADDR is set on the socket.)

This increases memory overhead as the packets need to be copied, however they are only copied per matching socket. You can safely disable it if you don't plan to receive broadcast or multicast traffic on more than one socket at a time.

Default value:

- Yes (enabled)

CONFIG_LWIP_SO_RCVBUF

Enable SO_RCVBUF option

Found in: [Component config](#) > [LWIP](#)

Enabling this option allows checking for available data on a netconn.

Default value:

- No (disabled)

CONFIG_LWIP_NETBUF_RECVINFO

Enable IP_PKTINFO option

Found in: [Component config](#) > [LWIP](#)

Enabling this option allows checking for the destination address of a received IPv4 Packet.

Default value:

- No (disabled)

CONFIG_LWIP_IP_DEFAULT_TTL

The value for Time-To-Live used by transport layers

Found in: [Component config](#) > [LWIP](#)

Set value for Time-To-Live used by transport layers.

Range:

- from 1 to 255

Default value:

- 64

CONFIG_LWIP_IP4_FRAG

Enable fragment outgoing IP4 packets

Found in: [Component config](#) > [LWIP](#)

Enabling this option allows fragmenting outgoing IP4 packets if their size exceeds MTU.

Default value:

- Yes (enabled)

CONFIG_LWIP_IP6_FRAG

Enable fragment outgoing IP6 packets

Found in: [Component config](#) > [LWIP](#)

Enabling this option allows fragmenting outgoing IP6 packets if their size exceeds MTU.

Default value:

- Yes (enabled)

CONFIG_LWIP_IP4_REASSEMBLY

Enable reassembly incoming fragmented IP4 packets

Found in: [Component config](#) > [LWIP](#)

Enabling this option allows reassembling incoming fragmented IP4 packets.

Default value:

- No (disabled)

CONFIG_LWIP_IP6_REASSEMBLY

Enable reassembly incoming fragmented IP6 packets

Found in: [Component config](#) > [LWIP](#)

Enabling this option allows reassembling incoming fragmented IP6 packets.

Default value:

- No (disabled)

CONFIG_LWIP_IP_REASS_MAX_PBUFS

The maximum amount of pbufs waiting to be reassembled

Found in: [Component config](#) > [LWIP](#)

Set the maximum amount of pbufs waiting to be reassembled.

Range:

- from 10 to 100

Default value:

- 10

CONFIG_LWIP_IP_FORWARD

Enable IP forwarding

Found in: [Component config](#) > [LWIP](#)

Enabling this option allows packets forwarding across multiple interfaces.

Default value:

- No (disabled)

CONFIG_LWIP_IPV4_NAPT

Enable NAT (new/experimental)

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_IP_FORWARD](#)

Enabling this option allows Network Address and Port Translation.

Default value:

- No (disabled) if [CONFIG_LWIP_IP_FORWARD](#)

CONFIG_LWIP_IPV4_NAPT_PORTMAP

Enable NAT Port Mapping (new/experimental)

Found in: *Component config* > *LWIP* > *CONFIG_LWIP_IP_FORWARD* > *CONFIG_LWIP_IPV4_NAPT*

Enabling this option allows Port Forwarding or Port mapping.

Default value:

- Yes (enabled) if *CONFIG_LWIP_IPV4_NAPT*

CONFIG_LWIP_STATS

Enable LWIP statistics

Found in: *Component config* > *LWIP*

Enabling this option allows LWIP statistics

Default value:

- No (disabled)

CONFIG_LWIP_ESP_GRATUITOUS_ARP

Send gratuitous ARP periodically

Found in: *Component config* > *LWIP*

Enable this option allows to send gratuitous ARP periodically.

This option solve the compatibility issues.If the ARP table of the AP is old, and the AP doesn't send ARP request to update it's ARP table, this will lead to the STA sending IP packet fail. Thus we send gratuitous ARP periodically to let AP update it's ARP table.

Default value:

- Yes (enabled)

CONFIG_LWIP_GARP_TMR_INTERVAL

GARP timer interval(seconds)

Found in: *Component config* > *LWIP* > *CONFIG_LWIP_ESP_GRATUITOUS_ARP*

Set the timer interval for gratuitous ARP. The default value is 60s

Default value:

- 60

CONFIG_LWIP_ESP_MLDV6_REPORT

Send mldv6 report periodically

Found in: *Component config* > *LWIP*

Enable this option allows to send mldv6 report periodically.

This option solve the issue that failed to receive multicast data. Some routers fail to forward multicast packets. To solve this problem, send multicast mldv6 report to routers regularly.

Default value:

- Yes (enabled)

CONFIG_LWIP_MLDV6_TMR_INTERVAL

mldv6 report timer interval(seconds)

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_ESP_MLDV6_REPORT](#)

Set the timer interval for mldv6 report. The default value is 30s

Default value:

- 40

CONFIG_LWIP_TCPIP_RECVMBOX_SIZE

TCPIP task receive mail box size

Found in: [Component config](#) > [LWIP](#)

Set TCPIP task receive mail box size. Generally bigger value means higher throughput but more memory. The value should be bigger than UDP/TCP mail box size.

Range:

- from 6 to 1024 if [CONFIG_LWIP_WND_SCALE](#)

Default value:

- 32

CONFIG_LWIP_DHCP_DOES_ARP_CHECK

DHCP: Perform ARP check on any offered address

Found in: [Component config](#) > [LWIP](#)

Enabling this option performs a check (via ARP request) if the offered IP address is not already in use by another host on the network.

Default value:

- Yes (enabled)

CONFIG_LWIP_DHCP_DISABLE_CLIENT_ID

DHCP: Disable Use of HW address as client identification

Found in: [Component config](#) > [LWIP](#)

This option could be used to disable DHCP client identification with its MAC address. (Client id is used by DHCP servers to uniquely identify clients and are included in the DHCP packets as an option 61) Set this option to "y" in order to exclude option 61 from DHCP packets.

Default value:

- No (disabled)

CONFIG_LWIP_DHCP_DISABLE_VENDOR_CLASS_ID

DHCP: Disable Use of vendor class identification

Found in: [Component config](#) > [LWIP](#)

This option could be used to disable DHCP client vendor class identification. Set this option to "y" in order to exclude option 60 from DHCP packets.

Default value:

- Yes (enabled)

CONFIG_LWIP_DHCP_RESTORE_LAST_IP

DHCP: Restore last IP obtained from DHCP server

Found in: [Component config](#) > [LWIP](#)

When this option is enabled, DHCP client tries to re-obtain last valid IP address obtained from DHCP server. Last valid DHCP configuration is stored in nvs and restored after reset/power-up. If IP is still available, there is no need for sending discovery message to DHCP server and save some time.

Default value:

- No (disabled)

CONFIG_LWIP_DHCP_OPTIONS_LEN

DHCP total option length

Found in: [Component config](#) > [LWIP](#)

Set total length of outgoing DHCP option msg. Generally bigger value means it can carry more options and values. If your code meets LWIP_ASSERT due to option value is too long. Please increase the LWIP_DHCP_OPTIONS_LEN value.

Range:

- from 68 to 255

Default value:

- 68

CONFIG_LWIP_NUM_NETIF_CLIENT_DATA

Number of clients store data in netif

Found in: [Component config](#) > [LWIP](#)

Number of clients that may store data in client_data member array of struct netif.

Range:

- from 0 to 256

Default value:

- 0

CONFIG_LWIP_DHCP_COARSE_TIMER_SECS

DHCP coarse timer interval(s)

Found in: [Component config](#) > [LWIP](#)

Set DHCP coarse interval in seconds. A higher value will be less precise but cost less power consumption.

Range:

- from 1 to 10

Default value:

- 1

DHCP server Contains:

- [CONFIG_LWIP_DHCPS](#)

CONFIG_LWIP_DHCPS

DHCPS: Enable IPv4 Dynamic Host Configuration Protocol Server (DHCPS)

Found in: [Component config](#) > [LWIP](#) > [DHCP server](#)

Enabling this option allows the device to run the DHCP server (to dynamically assign IPv4 addresses to clients).

Default value:

- Yes (enabled)

CONFIG_LWIP_DHCPS_LEASE_UNIT

Multiplier for lease time, in seconds

Found in: [Component config](#) > [LWIP](#) > [DHCP server](#) > [CONFIG_LWIP_DHCPS](#)

The DHCP server is calculating lease time multiplying the sent and received times by this number of seconds per unit. The default is 60, that equals one minute.

Range:

- from 1 to 3600

Default value:

- 60

CONFIG_LWIP_DHCPS_MAX_STATION_NUM

Maximum number of stations

Found in: [Component config](#) > [LWIP](#) > [DHCP server](#) > [CONFIG_LWIP_DHCPS](#)

The maximum number of DHCP clients that are connected to the server. After this number is exceeded, DHCP server removes of the oldest device from it's address pool, without notification.

Range:

- from 1 to 64

Default value:

- 8

CONFIG_LWIP_DHCPS_STATIC_ENTRIES

Enable ARP static entries

Found in: [Component config](#) > [LWIP](#) > [DHCP server](#) > [CONFIG_LWIP_DHCPS](#)

Enabling this option allows DHCP server to support temporary static ARP entries for DHCP Client. This will help the DHCP server to send the DHCP OFFER and DHCP ACK using IP unicast.

Default value:

- Yes (enabled)

CONFIG_LWIP_DHCPS_ADD_DNS

Always add DNS option in DHCP responses

Found in: [Component config](#) > [LWIP](#) > [DHCP server](#) > [CONFIG_LWIP_DHCPS](#)

This allows the DNS option to be optional in the DHCP offers, depending on the server's runtime configuration. When enabled, the DHCP server will always add the DNS option to DHCP responses. If a DNS server is not explicitly configured, the server's IP address will be used as the fallback for the DNS option. When disabled, the DHCP server will only include the DNS option in responses if a DNS server has been explicitly configured. This option will be removed in IDF v6.x

Default value:

- Yes (enabled)

CONFIG_LWIP_AUTOIP

Enable IPV4 Link-Local Addressing (AUTOIP)

Found in: [Component config](#) > [LWIP](#)

Enabling this option allows the device to self-assign an address in the 169.256/16 range if none is assigned statically or via DHCP.

See RFC 3927.

Default value:

- No (disabled)

Contains:

- [CONFIG_LWIP_AUTOIP_TRIES](#)
- [CONFIG_LWIP_AUTOIP_MAX_CONFLICTS](#)
- [CONFIG_LWIP_AUTOIP_RATE_LIMIT_INTERVAL](#)

CONFIG_LWIP_AUTOIP_TRIES

DHCP Probes before self-assigning IPv4 LL address

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_AUTOIP](#)

DHCP client will send this many probes before self-assigning a link local address.

From LWIP help: "This can be set as low as 1 to get an AutoIP address very quickly, but you should be prepared to handle a changing IP address when DHCP overrides AutoIP." (In the case of ESP-IDF, this means multiple SYSTEM_EVENT_STA_GOT_IP events.)

Range:

- from 1 to 100 if [CONFIG_LWIP_AUTOIP](#)

Default value:

- 2 if [CONFIG_LWIP_AUTOIP](#)

CONFIG_LWIP_AUTOIP_MAX_CONFLICTS

Max IP conflicts before rate limiting

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_AUTOIP](#)

If the AUTOIP functionality detects this many IP conflicts while self-assigning an address, it will go into a rate limited mode.

Range:

- from 1 to 100 if [CONFIG_LWIP_AUTOIP](#)

Default value:

- 9 if [CONFIG_LWIP_AUTOIP](#)

CONFIG_LWIP_AUTOIP_RATE_LIMIT_INTERVAL

Rate limited interval (seconds)

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_AUTOIP](#)

If rate limiting self-assignment requests, wait this long between each request.

Range:

- from 5 to 120 if [CONFIG_LWIP_AUTOIP](#)

Default value:

- 20 if [CONFIG_LWIP_AUTOIP](#)

CONFIG_LWIP_IPV4

Enable IPv4

Found in: [Component config](#) > [LWIP](#)

Enable IPv4 stack. If you want to use IPv6 only TCP/IP stack, disable this.

Default value:

- Yes (enabled)

CONFIG_LWIP_IPV6

Enable IPv6

Found in: [Component config](#) > [LWIP](#)

Enable IPv6 function. If not use IPv6 function, set this option to n. If disabling LWIP_IPV6 then some other components (coap and asio) will no longer be available.

Default value:

- Yes (enabled)

CONFIG_LWIP_IPV6_AUTOCONFIG

Enable IPV6 stateless address autoconfiguration (SLAAC)

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_IPV6](#)

Enabling this option allows the devices to IPV6 stateless address autoconfiguration (SLAAC).

See RFC 4862.

Default value:

- No (disabled)

CONFIG_LWIP_IPV6_NUM_ADDRESSES

Number of IPv6 addresses on each network interface

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_IPV6](#)

The maximum number of IPv6 addresses on each interface. Any additional addresses will be discarded.

Default value:

- 3

CONFIG_LWIP_IPV6_FORWARD

Enable IPv6 forwarding between interfaces

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_IPV6](#)

Forwarding IPv6 packets between interfaces is only required when acting as a router.

Default value:

- No (disabled)

CONFIG_LWIP_IPV6_RDNSS_MAX_DNS_SERVERS

Use IPv6 Router Advertisement Recursive DNS Server Option

Found in: [Component config](#) > [LWIP](#)

Use IPv6 Router Advertisement Recursive DNS Server Option (as per RFC 6106) to copy a defined maximum number of DNS servers to the DNS module. Set this option to a number of desired DNS servers advertised in the RA protocol. This feature is disabled when set to 0.

Default value:

- 0 if [CONFIG_LWIP_IPV6_AUTOCONFIG](#)

CONFIG_LWIP_IPV6_DHCP6

Enable DHCPv6 stateless address autoconfiguration

Found in: [Component config](#) > [LWIP](#)

Enable DHCPv6 for IPv6 stateless address autoconfiguration. Note that the dhcpv6 client has to be started using `dhcp6_enable_stateless(netif)`; Note that the stateful address autoconfiguration is not supported.

Default value:

- No (disabled) if [CONFIG_LWIP_IPV6_AUTOCONFIG](#)

CONFIG_LWIP_NETIF_STATUS_CALLBACK

Enable status callback for network interfaces

Found in: [Component config](#) > [LWIP](#)

Enable callbacks when the network interface is up/down and addresses are changed.

Default value:

- No (disabled)

CONFIG_LWIP_NETIF_LOOPBACK

Support per-interface loopback

Found in: [Component config](#) > [LWIP](#)

Enabling this option means that if a packet is sent with a destination address equal to the interface's own IP address, it will "loop back" and be received by this interface. Disabling this option disables support of loopback interface in lwIP

Default value:

- Yes (enabled)

Contains:

- [CONFIG_LWIP_LOOPBACK_MAX_PBUFS](#)

CONFIG_LWIP_LOOPBACK_MAX_PBUFS

Max queued loopback packets per interface

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_NETIF_LOOPBACK](#)

Configure the maximum number of packets which can be queued for loopback on a given interface. Reducing this number may cause packets to be dropped, but will avoid filling memory with queued packet data.

Range:

- from 0 to 16

Default value:

- 8

TCP Contains:

- [*CONFIG_LWIP_TCP_WND_DEFAULT*](#)
- [*CONFIG_LWIP_TCP_SND_BUF_DEFAULT*](#)
- [*CONFIG_LWIP_TCP_RECVMBOX_SIZE*](#)
- [*CONFIG_LWIP_TCP_RTO_TIME*](#)
- [*CONFIG_LWIP_MAX_ACTIVE_TCP*](#)
- [*CONFIG_LWIP_TCP_FIN_WAIT_TIMEOUT*](#)
- [*CONFIG_LWIP_MAX_LISTENING_TCP*](#)
- [*CONFIG_LWIP_TCP_MAXRTX*](#)
- [*CONFIG_LWIP_TCP_SYNMAXRTX*](#)
- [*CONFIG_LWIP_TCP_MSL*](#)
- [*CONFIG_LWIP_TCP_MSS*](#)
- [*CONFIG_LWIP_TCP_OVERSIZE*](#)
- [*CONFIG_LWIP_TCP_QUEUE_OOSEQ*](#)
- [*CONFIG_LWIP_WND_SCALE*](#)
- [*CONFIG_LWIP_TCP_HIGH_SPEED_RETRANSMISSION*](#)
- [*CONFIG_LWIP_TCP_TMR_INTERVAL*](#)

CONFIG_LWIP_MAX_ACTIVE_TCP

Maximum active TCP Connections

Found in: [*Component config*](#) > [*LWIP*](#) > [*TCP*](#)

The maximum number of simultaneously active TCP connections. The practical maximum limit is determined by available heap memory at runtime.

Changing this value by itself does not substantially change the memory usage of LWIP, except for preventing new TCP connections after the limit is reached.

Range:

- from 1 to 1024

Default value:

- 16

CONFIG_LWIP_MAX_LISTENING_TCP

Maximum listening TCP Connections

Found in: [*Component config*](#) > [*LWIP*](#) > [*TCP*](#)

The maximum number of simultaneously listening TCP connections. The practical maximum limit is determined by available heap memory at runtime.

Changing this value by itself does not substantially change the memory usage of LWIP, except for preventing new listening TCP connections after the limit is reached.

Range:

- from 1 to 1024

Default value:

- 16

CONFIG_LWIP_TCP_HIGH_SPEED_RETRANSMISSION

TCP high speed retransmissions

Found in: [*Component config*](#) > [*LWIP*](#) > [*TCP*](#)

Speed up the TCP retransmission interval. If disabled, it is recommended to change the number of SYN retransmissions to 6, and TCP initial rto time to 3000.

Default value:

- Yes (enabled)

CONFIG_LWIP_TCP_MAXRTX

Maximum number of retransmissions of data segments

Found in: [Component config](#) > [LWIP](#) > [TCP](#)

Set maximum number of retransmissions of data segments.

Range:

- from 3 to 12

Default value:

- 12

CONFIG_LWIP_TCP_SYNMAXRTX

Maximum number of retransmissions of SYN segments

Found in: [Component config](#) > [LWIP](#) > [TCP](#)

Set maximum number of retransmissions of SYN segments.

Range:

- from 3 to 12

Default value:

- 12

CONFIG_LWIP_TCP_MSS

Maximum Segment Size (MSS)

Found in: [Component config](#) > [LWIP](#) > [TCP](#)

Set maximum segment size for TCP transmission.

Can be set lower to save RAM, the default value 1460(ipv4)/1440(ipv6) will give best throughput. IPv4 TCP_MSS Range: 576 <= TCP_MSS <= 1460 IPv6 TCP_MSS Range: 1220<= TCP_MSS <= 1440

Range:

- from 536 to 1460

Default value:

- 1440

CONFIG_LWIP_TCP_TMR_INTERVAL

TCP timer interval(ms)

Found in: [Component config](#) > [LWIP](#) > [TCP](#)

Set TCP timer interval in milliseconds.

Can be used to speed connections on bad networks. A lower value will redeliver unacked packets faster.

Default value:

- 250

CONFIG_LWIP_TCP_MSL

Maximum segment lifetime (MSL)

Found in: [Component config](#) > [LWIP](#) > [TCP](#)

Set maximum segment lifetime in milliseconds.

Default value:

- 60000

CONFIG_LWIP_TCP_FIN_WAIT_TIMEOUT

Maximum FIN segment lifetime

Found in: [Component config](#) > [LWIP](#) > [TCP](#)

Set maximum segment lifetime in milliseconds.

Default value:

- 20000

CONFIG_LWIP_TCP_SND_BUF_DEFAULT

Default send buffer size

Found in: [Component config](#) > [LWIP](#) > [TCP](#)

Set default send buffer size for new TCP sockets.

Per-socket send buffer size can be changed at runtime with `lwip_setsockopt(s, TCP_SNDBUF, ...)`.

This value must be at least 2x the MSS size, and the default is 4x the default MSS size.

Setting a smaller default SNDBUF size can save some RAM, but will decrease performance.

Range:

- from 2440 to 1024000 if [CONFIG_LWIP_WND_SCALE](#)

Default value:

- 5760

CONFIG_LWIP_TCP_WND_DEFAULT

Default receive window size

Found in: [Component config](#) > [LWIP](#) > [TCP](#)

Set default TCP receive window size for new TCP sockets.

Per-socket receive window size can be changed at runtime with `lwip_setsockopt(s, TCP_WINDOW, ...)`.

Setting a smaller default receive window size can save some RAM, but will significantly decrease performance.

Range:

- from 2440 to 1024000 if [CONFIG_LWIP_WND_SCALE](#)

Default value:

- 5760

CONFIG_LWIP_TCP_RECVMBOX_SIZE

Default TCP receive mail box size

Found in: [Component config](#) > [LWIP](#) > [TCP](#)

Set TCP receive mail box size. Generally bigger value means higher throughput but more memory. The recommended value is: $\text{LWIP\_TCP\_WND\_DEFAULT}/\text{TCP\_MSS} + 2$, e.g. if $\text{LWIP\_TCP\_WND\_DEFAULT}=14360$, $\text{TCP\_MSS}=1436$, then the recommended receive mail box size is $(14360/1436 + 2) = 12$.

TCP receive mail box is a per socket mail box, when the application receives packets from TCP socket, LWIP core firstly posts the packets to TCP receive mail box and the application then fetches the packets from mail box. It means LWIP can cache maximum `LWIP_TCP_RECVMBOX_SIZE` packets for each TCP socket, so the maximum possible cached TCP packets for all TCP sockets is

LWIP_TCP_RECCVMBOX_SIZE multiplies the maximum TCP socket number. In other words, the bigger LWIP_TCP_RECCVMBOX_SIZE means more memory. On the other hand, if the receive mail box is too small, the mail box may be full. If the mail box is full, the LWIP drops the packets. So generally we need to make sure the TCP receive mail box is big enough to avoid packet drop between LWIP core and application.

Range:

- from 6 to 1024 if [CONFIG_LWIP_WND_SCALE](#)

Default value:

- 6

CONFIG_LWIP_TCP_QUEUE_OOSEQ

Queue incoming out-of-order segments

Found in: [Component config](#) > [LWIP](#) > [TCP](#)

Queue incoming out-of-order segments for later use.

Disable this option to save some RAM during TCP sessions, at the expense of increased retransmissions if segments arrive out of order.

Default value:

- Yes (enabled)

CONFIG_LWIP_TCP_OOSEQ_TIMEOUT

Timeout for each pbuf queued in TCP OOSEQ, in RTOs.

Found in: [Component config](#) > [LWIP](#) > [TCP](#) > [CONFIG_LWIP_TCP_QUEUE_OOSEQ](#)

The timeout value is TCP_OOSEQ_TIMEOUT * RTO.

Range:

- from 1 to 30

Default value:

- 6

CONFIG_LWIP_TCP_OOSEQ_MAX_PBUFS

The maximum number of pbufs queued on OOSEQ per pcb

Found in: [Component config](#) > [LWIP](#) > [TCP](#) > [CONFIG_LWIP_TCP_QUEUE_OOSEQ](#)

If LWIP_TCP_OOSEQ_MAX_PBUFS = 0, TCP will not control the number of OOSEQ pbufs.

In a poor network environment, many out-of-order tcp pbufs will be received. These out-of-order pbufs will be cached in the TCP out-of-order queue which will cause Wi-Fi/Ethernet fail to release RX buffer in time. It is possible that all RX buffers for MAC layer are used by OOSEQ.

Control the number of out-of-order pbufs to ensure that the MAC layer has enough RX buffer to receive packets.

In the Wi-Fi scenario, recommended OOSEQ PBUFS Range: $0 \leq \text{TCP\_OOSEQ\_MAX\_PBUFS} \leq \text{CONFIG\_ESP\_WIFI\_DYNAMIC\_RX\_BUFFER\_NUM}/(\text{MAX\_TCP\_NUMBER} + 1)$

In the Ethernet scenario, recommended Ethernet OOSEQ PBUFS Range: $0 \leq \text{TCP\_OOSEQ\_MAX\_PBUFS} \leq \text{CONFIG\_ETH\_DMA\_RX\_BUFFER\_NUM}/(\text{MAX\_TCP\_NUMBER} + 1)$

Within the recommended value range, the larger the value, the better the performance.

MAX_TCP_NUMBER represent Maximum number of TCP connections in Wi-Fi(STA+SoftAP) and Ethernet scenario.

Range:

- from 0 to 12

Default value:

- 4

CONFIG_LWIP_TCP_SACK_OUT

Support sending selective acknowledgements

Found in: [Component config](#) > [LWIP](#) > [TCP](#) > [CONFIG_LWIP_TCP_QUEUE_OOSEQ](#)

TCP will support sending selective acknowledgements (SACKs).

Default value:

- No (disabled)

CONFIG_LWIP_TCP_OVERSIZE

Pre-allocate transmit PBUF size

Found in: [Component config](#) > [LWIP](#) > [TCP](#)

Allows enabling "oversize" allocation of TCP transmission pbufs ahead of time, which can reduce the length of pbuf chains used for transmission.

This will not make a difference to sockets where Nagle's algorithm is disabled.

Default value of MSS is fine for most applications, 25% MSS may save some RAM when only transmitting small amounts of data. Disabled will have worst performance and fragmentation characteristics, but uses least RAM overall.

Available options:

- MSS (CONFIG_LWIP_TCP_OVERSIZE_MSS)
- 25% MSS (CONFIG_LWIP_TCP_OVERSIZE_QUARTER_MSS)
- Disabled (CONFIG_LWIP_TCP_OVERSIZE_DISABLE)

CONFIG_LWIP_WND_SCALE

Support TCP window scale

Found in: [Component config](#) > [LWIP](#) > [TCP](#)

Enable this feature to support TCP window scaling.

CONFIG_LWIP_TCP_RCV_SCALE

Set TCP receiving window scaling factor

Found in: [Component config](#) > [LWIP](#) > [TCP](#) > [CONFIG_LWIP_WND_SCALE](#)

Enable this feature to support TCP window scaling.

Range:

- from 0 to 14 if [CONFIG_LWIP_WND_SCALE](#)

Default value:

- 0 if [CONFIG_LWIP_WND_SCALE](#)

CONFIG_LWIP_TCP_RTO_TIME

Default TCP rto time

Found in: [Component config](#) > [LWIP](#) > [TCP](#)

Set default TCP rto time for a reasonable initial rto. In bad network environment, recommend set value of rto time to 1500.

Default value:

- 1500

UDP Contains:

- [CONFIG_LWIP_UDP_RECVMBOX_SIZE](#)
- [CONFIG_LWIP_MAX_UDP_PCBS](#)

CONFIG_LWIP_MAX_UDP_PCBS

Maximum active UDP control blocks

Found in: [Component config](#) > [LWIP](#) > [UDP](#)

The maximum number of active UDP "connections" (ie UDP sockets sending/receiving data). The practical maximum limit is determined by available heap memory at runtime.

Range:

- from 1 to 1024

Default value:

- 16

CONFIG_LWIP_UDP_RECVMBOX_SIZE

Default UDP receive mail box size

Found in: [Component config](#) > [LWIP](#) > [UDP](#)

Set UDP receive mail box size. The recommended value is 6.

UDP receive mail box is a per socket mail box, when the application receives packets from UDP socket, LWIP core firstly posts the packets to UDP receive mail box and the application then fetches the packets from mail box. It means LWIP can caches maximum UDP_RECVMBOX_SIZE packets for each UDP socket, so the maximum possible cached UDP packets for all UDP sockets is UDP_RECVMBOX_SIZE multiplies the maximum UDP socket number. In other words, the bigger UDP_RECVMBOX_SIZE means more memory. On the other hand, if the receive mail box is too small, the mail box may be full. If the mail box is full, the LWIP drops the packets. So generally we need to make sure the UDP receive mail box is big enough to avoid packet drop between LWIP core and application.

Range:

- from 6 to 64

Default value:

- 6

Checksums Contains:

- [CONFIG_LWIP_CHECKSUM_CHECK_ICMP](#)
- [CONFIG_LWIP_CHECKSUM_CHECK_IP](#)
- [CONFIG_LWIP_CHECKSUM_CHECK_UDP](#)

CONFIG_LWIP_CHECKSUM_CHECK_IP

Enable LWIP IP checksums

Found in: [Component config](#) > [LWIP](#) > [Checksums](#)

Enable checksum checking for received IP messages

Default value:

- No (disabled)

CONFIG_LWIP_CHECKSUM_CHECK_UDP

Enable LWIP UDP checksums

Found in: [Component config](#) > [LWIP](#) > [Checksums](#)

Enable checksum checking for received UDP messages

Default value:

- No (disabled)

CONFIG_LWIP_CHECKSUM_CHECK_ICMP

Enable LWIP ICMP checksums

Found in: [Component config](#) > [LWIP](#) > [Checksums](#)

Enable checksum checking for received ICMP messages

Default value:

- Yes (enabled)

CONFIG_LWIP_TCPIP_TASK_STACK_SIZE

TCP/IP Task Stack Size

Found in: [Component config](#) > [LWIP](#)

Configure TCP/IP task stack size, used by LWIP to process multi-threaded TCP/IP operations. Setting this stack too small will result in stack overflow crashes.

Range:

- from 2048 to 65536

Default value:

- 3072

CONFIG_LWIP_TCPIP_TASK_AFFINITY

TCP/IP task affinity

Found in: [Component config](#) > [LWIP](#)

Allows setting LwIP tasks affinity, i.e. whether the task is pinned to CPU0, pinned to CPU1, or allowed to run on any CPU. Currently this applies to "TCP/IP" task and "Ping" task.

Available options:

- No affinity (CONFIG_LWIP_TCPIP_TASK_AFFINITY_NO_AFFINITY)
- CPU0 (CONFIG_LWIP_TCPIP_TASK_AFFINITY_CPU0)
- CPU1 (CONFIG_LWIP_TCPIP_TASK_AFFINITY_CPU1)

CONFIG_LWIP_IPV6_ND6_NUM_PREFIXES

Max number of entries in IPv6 on-link prefixes cache

Found in: [Component config](#) > [LWIP](#)

Maximum number of entries in IPv6 on-link prefixes cache

Default value:

- 5

CONFIG_LWIP_IPV6_ND6_NUM_ROUTERS

Max number of entries in IPv6 default routers cache

Found in: [Component config](#) > [LWIP](#)

Maximum number of entries in IPv6 default routers cache

Default value:

- 3

CONFIG_LWIP_IPV6_ND6_NUM_DESTINATIONS

Max number of entries in IPv6 destinations cache

Found in: [Component config](#) > [LWIP](#)

Maximum number of entries in IPv6 destinations cache

Default value:

- 10

CONFIG_LWIP_PPP_SUPPORT

Enable PPP support

Found in: [Component config](#) > [LWIP](#)

Enable PPP stack. Now only PPP over serial is possible.

Default value:

- No (disabled)

Contains:

- [CONFIG_LWIP_PPP_ENABLE_IPV6](#)

CONFIG_LWIP_PPP_ENABLE_IPV6

Enable IPV6 support for PPP connections (IPV6CP)

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_PPP_SUPPORT](#)

Enable IPV6 support in PPP for the local link between the DTE (processor) and DCE (modem). There are some modems which do not support the IPV6 addressing in the local link. If they are requested for IPV6CP negotiation, they may time out. This would in turn fail the configuration for the whole link. If your modem is not responding correctly to PPP Phase Network, try to disable IPV6 support.

Default value:

- Yes (enabled) if [CONFIG_LWIP_PPP_SUPPORT](#) && [CONFIG_LWIP_IPV6](#)

CONFIG_LWIP_IPV6_MEMP_NUM_ND6_QUEUE

Max number of IPv6 packets to queue during MAC resolution

Found in: [Component config](#) > [LWIP](#)

Config max number of IPv6 packets to queue during MAC resolution.

Range:

- from 3 to 20

Default value:

- 3

CONFIG_LWIP_IPV6_ND6_NUM_NEIGHBORS

Max number of entries in IPv6 neighbor cache

Found in: [Component config](#) > [LWIP](#)

Config max number of entries in IPv6 neighbor cache

Range:

- from 3 to 10

Default value:

- 5

CONFIG_LWIP_PPP_NOTIFY_PHASE_SUPPORT

Enable Notify Phase Callback

Found in: [Component config](#) > [LWIP](#)

Enable to set a callback which is called on change of the internal PPP state machine.

Default value:

- No (disabled) if [CONFIG_LWIP_PPP_SUPPORT](#)

CONFIG_LWIP_PPP_PAP_SUPPORT

Enable PAP support

Found in: [Component config](#) > [LWIP](#)

Enable Password Authentication Protocol (PAP) support

Default value:

- No (disabled) if [CONFIG_LWIP_PPP_SUPPORT](#)

CONFIG_LWIP_PPP_CHAP_SUPPORT

Enable CHAP support

Found in: [Component config](#) > [LWIP](#)

Enable Challenge Handshake Authentication Protocol (CHAP) support

Default value:

- No (disabled) if [CONFIG_LWIP_PPP_SUPPORT](#)

CONFIG_LWIP_PPP_MSCHAP_SUPPORT

Enable MSCHAP support

Found in: [Component config](#) > [LWIP](#)

Enable Microsoft version of the Challenge-Handshake Authentication Protocol (MSCHAP) support

Default value:

- No (disabled) if [CONFIG_LWIP_PPP_SUPPORT](#)

CONFIG_LWIP_PPP_MPPE_SUPPORT

Enable MPPE support

Found in: [Component config](#) > [LWIP](#)

Enable Microsoft Point-to-Point Encryption (MPPE) support

Default value:

- No (disabled) if [CONFIG_LWIP_PPP_SUPPORT](#)

CONFIG_LWIP_PPP_SERVER_SUPPORT

Enable PPP server support

Found in: [Component config](#) > [LWIP](#)

Enable to use PPP server

Default value:

- No (disabled) if [CONFIG_LWIP_PPP_SUPPORT](#)

CONFIG_LWIP_PPP_VJ_HEADER_COMPRESSION

Enable VJ IP Header compression

Found in: [Component config](#) > [LWIP](#)

Enable support for VJ header compression. Please disable this if you're using NAT on PPP interface, since the compressed IP header might not be correctly interpreted in NAT causing the compressed packet to be dropped.

Default value:

- Yes (enabled) if [CONFIG_LWIP_PPP_SUPPORT](#)

CONFIG_LWIP_ENABLE_LCP_ECHO

Enable LCP ECHO

Found in: [Component config](#) > [LWIP](#)

Enable LCP echo keepalive requests

Default value:

- No (disabled) if [CONFIG_LWIP_PPP_SUPPORT](#)

CONFIG_LWIP_LCP_ECHOINTERVAL

Echo interval (s)

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_ENABLE_LCP_ECHO](#)

Interval in seconds between keepalive LCP echo requests, 0 to disable.

Range:

- from 0 to 1000000 if [CONFIG_LWIP_ENABLE_LCP_ECHO](#)

Default value:

- 3 if [CONFIG_LWIP_ENABLE_LCP_ECHO](#)

CONFIG_LWIP_LCP_MAXECHOFAILS

Maximum echo failures

Found in: *Component config* > *LWIP* > *CONFIG_LWIP_ENABLE_LCP_ECHO*

Number of consecutive unanswered echo requests before failure is indicated.

Range:

- from 0 to 100000 if *CONFIG_LWIP_ENABLE_LCP_ECHO*

Default value:

- 3 if *CONFIG_LWIP_ENABLE_LCP_ECHO*

CONFIG_LWIP_PPP_DEBUG_ON

Enable PPP debug log output

Found in: *Component config* > *LWIP*

Enable PPP debug log output

Default value:

- No (disabled) if *CONFIG_LWIP_PPP_SUPPORT*

CONFIG_LWIP_USE_EXTERNAL_MBEDTLS

Use mbedTLS instead of internal polarSSL

Found in: *Component config* > *LWIP*

This option uses mbedTLS crypto functions (instead of internal PolarSSL implementation) for PPP authentication modes (PAP, CHAP, etc.). You can use this option to address symbol duplication issues, since the internal functions are not namespaced (e.g. md5_init()).

CONFIG_LWIP_SLIP_SUPPORT

Enable SLIP support (new/experimental)

Found in: *Component config* > *LWIP*

Enable SLIP stack. Now only SLIP over serial is possible.

SLIP over serial support is experimental and unsupported.

Default value:

- No (disabled)

Contains:

- *CONFIG_LWIP_SLIP_DEBUG_ON*

CONFIG_LWIP_SLIP_DEBUG_ON

Enable SLIP debug log output

Found in: *Component config* > *LWIP* > *CONFIG_LWIP_SLIP_SUPPORT*

Enable SLIP debug log output

Default value:

- No (disabled) if *CONFIG_LWIP_SLIP_SUPPORT*

ICMP Contains:

- [CONFIG_LWIP_ICMP](#)
- [CONFIG_LWIP_BROADCAST_PING](#)
- [CONFIG_LWIP_MULTICAST_PING](#)

CONFIG_LWIP_ICMP

ICMP: Enable ICMP

Found in: [Component config](#) > [LWIP](#) > [ICMP](#)

Enable ICMP module for check network stability

Default value:

- Yes (enabled)

CONFIG_LWIP_MULTICAST_PING

Respond to multicast pings

Found in: [Component config](#) > [LWIP](#) > [ICMP](#)

Default value:

- No (disabled)

CONFIG_LWIP_BROADCAST_PING

Respond to broadcast pings

Found in: [Component config](#) > [LWIP](#) > [ICMP](#)

Default value:

- No (disabled)

LWIP RAW API Contains:

- [CONFIG_LWIP_MAX_RAW_PCBS](#)

CONFIG_LWIP_MAX_RAW_PCBS

Maximum LWIP RAW PCBs

Found in: [Component config](#) > [LWIP](#) > [LWIP RAW API](#)

The maximum number of simultaneously active LWIP RAW protocol control blocks. The practical maximum limit is determined by available heap memory at runtime.

Range:

- from 1 to 1024

Default value:

- 16

SNTP Contains:

- [CONFIG_LWIP_SNTP_MAX_SERVERS](#)
- [CONFIG_LWIP_SNTP_UPDATE_DELAY](#)
- [CONFIG_LWIP_DHCP_GET_NTP_SRV](#)

CONFIG_LWIP_SNTP_MAX_SERVERS

Maximum number of NTP servers

Found in: *Component config > LWIP > SNTP*

Set maximum number of NTP servers used by LwIP SNTP module. First argument of `sntp_setserver/sntp_setservername` functions is limited to this value.

Range:

- from 1 to 16

Default value:

- 1

CONFIG_LWIP_DHCP_GET_NTP_SRV

Request NTP servers from DHCP

Found in: *Component config > LWIP > SNTP*

If enabled, LWIP will add 'NTP' to Parameter-Request Option sent via DHCP-request. DHCP server might reply with an NTP server address in option 42. SNTP callback for such replies should be set accordingly (see `sntp_servermode_dhcp()` func.)

Default value:

- No (disabled)

CONFIG_LWIP_DHCP_MAX_NTP_SERVERS

Maximum number of NTP servers acquired via DHCP

Found in: *Component config > LWIP > SNTP > CONFIG_LWIP_DHCP_GET_NTP_SRV*

Set maximum number of NTP servers acquired via DHCP-offer. Should be less or equal to "Maximum number of NTP servers", any extra servers would be just ignored.

Range:

- from 1 to 16 if *CONFIG_LWIP_DHCP_GET_NTP_SRV*

Default value:

- 1 if *CONFIG_LWIP_DHCP_GET_NTP_SRV*

CONFIG_LWIP_SNTP_UPDATE_DELAY

Request interval to update time (ms)

Found in: *Component config > LWIP > SNTP*

This option allows you to set the time update period via SNTP. Default is 1 hour. Must not be below 15 seconds by specification. (SNTPv4 RFC 4330 enforces a minimum update time of 15 seconds).

Range:

- from 15000 to 4294967295

Default value:

- 3600000

DNS Contains:

- *CONFIG_LWIP_FALLBACK_DNS_SERVER_SUPPORT*
- *CONFIG_LWIP_DNS_SETSERVER_WITH_NETIF*
- *CONFIG_LWIP_DNS_MAX_SERVERS*
- *CONFIG_LWIP_DNS_MAX_HOST_IP*

CONFIG_LWIP_DNS_MAX_HOST_IP

Maximum number of IP addresses per host

Found in: [Component config](#) > [LWIP](#) > [DNS](#)

Maximum number of IP addresses that can be returned by DNS queries for a single host.

Default value:

- 1

CONFIG_LWIP_DNS_MAX_SERVERS

Maximum number of DNS servers

Found in: [Component config](#) > [LWIP](#) > [DNS](#)

Set maximum number of DNS servers. If fallback DNS servers are supported, the number of DNS servers needs to be greater than or equal to 3.

Range:

- from 1 to 4

Default value:

- 3

CONFIG_LWIP_FALLBACK_DNS_SERVER_SUPPORT

Enable DNS fallback server support

Found in: [Component config](#) > [LWIP](#) > [DNS](#)

Enable this feature to support DNS fallback server.

Default value:

- No (disabled)

CONFIG_LWIP_FALLBACK_DNS_SERVER_ADDRESS

DNS fallback server address

Found in: [Component config](#) > [LWIP](#) > [DNS](#) > [CONFIG_LWIP_FALLBACK_DNS_SERVER_SUPPORT](#)

This option allows you to config dns fallback server address.

Default value:

- "114.114.114.114" if [CONFIG_LWIP_FALLBACK_DNS_SERVER_SUPPORT](#)

CONFIG_LWIP_DNS_SETSERVER_WITH_NETIF

Enable DNS server settings with netif

Found in: [Component config](#) > [LWIP](#) > [DNS](#)

This option allows collecting DNS server settings per netif using configurable callback function. It's typically used with [CONFIG_ESP_NETIF_SET_DNS_PER_DEFAULT_NETIF](#) which configures a callback to collect the DNS info on esp_netif layer.

Default value:

- No (disabled)

CONFIG_LWIP_BRIDGEIF_MAX_PORTS

Maximum number of bridge ports

Found in: [Component config](#) > [LWIP](#)

Set maximum number of ports a bridge can consists of.

Range:

- from 1 to 63

Default value:

- 7

CONFIG_LWIP_ESP_LWIP_ASSERT

Enable LWIP ASSERT checks

Found in: [Component config](#) > [LWIP](#)

Enable this option keeps LWIP assertion checks enabled. It is recommended to keep this option enabled.

If asserts are disabled for the entire project, they are also disabled for LWIP and this option is ignored.

Hooks Contains:

- [CONFIG_LWIP_HOOK_DNS_EXTERNAL_RESOLVE](#)
- [CONFIG_LWIP_HOOK_ND6_GET_GW](#)
- [CONFIG_LWIP_HOOK_IP6_INPUT](#)
- [CONFIG_LWIP_HOOK_IP6_ROUTE](#)
- [CONFIG_LWIP_HOOK_IP6_SELECT_SRC_ADDR](#)
- [CONFIG_LWIP_HOOK_NETCONN_EXTERNAL_RESOLVE](#)
- [CONFIG_LWIP_HOOK_TCP_ISN](#)

CONFIG_LWIP_HOOK_TCP_ISN

TCP ISN Hook

Found in: [Component config](#) > [LWIP](#) > [Hooks](#)

Enables to define a TCP ISN hook to randomize initial sequence number in TCP connection. The default TCP ISN algorithm used in IDF (standardized in RFC 6528) produces ISN by combining an MD5 of the new TCP id and a stable secret with the current time. This is because the lwIP implementation (*tcp_next_isn*) is not very strong, as it does not take into consideration any platform specific entropy source.

Set to `LWIP_HOOK_TCP_ISN_CUSTOM` to provide custom implementation. Set to `LWIP_HOOK_TCP_ISN_NONE` to use lwIP implementation.

Available options:

- No hook declared (`CONFIG_LWIP_HOOK_TCP_ISN_NONE`)
- Default implementation (`CONFIG_LWIP_HOOK_TCP_ISN_DEFAULT`)
- Custom implementation (`CONFIG_LWIP_HOOK_TCP_ISN_CUSTOM`)

CONFIG_LWIP_HOOK_IP6_ROUTE

IPv6 route Hook

Found in: [Component config](#) > [LWIP](#) > [Hooks](#)

Enables custom IPv6 route hook. Setting this to "default" provides weak implementation stub that could be overwritten in application code. Setting this to "custom" provides hook's declaration only and expects the application to implement it.

Available options:

- No hook declared (CONFIG_LWIP_HOOK_IP6_ROUTE_NONE)
- Default (weak) implementation (CONFIG_LWIP_HOOK_IP6_ROUTE_DEFAULT)
- Custom implementation (CONFIG_LWIP_HOOK_IP6_ROUTE_CUSTOM)

CONFIG_LWIP_HOOK_ND6_GET_GW

IPv6 get gateway Hook

Found in: [Component config](#) > [LWIP](#) > [Hooks](#)

Enables custom IPv6 route hook. Setting this to "default" provides weak implementation stub that could be overwritten in application code. Setting this to "custom" provides hook's declaration only and expects the application to implement it.

Available options:

- No hook declared (CONFIG_LWIP_HOOK_ND6_GET_GW_NONE)
- Default (weak) implementation (CONFIG_LWIP_HOOK_ND6_GET_GW_DEFAULT)
- Custom implementation (CONFIG_LWIP_HOOK_ND6_GET_GW_CUSTOM)

CONFIG_LWIP_HOOK_IP6_SELECT_SRC_ADDR

IPv6 source address selection Hook

Found in: [Component config](#) > [LWIP](#) > [Hooks](#)

Enables custom IPv6 source address selection. Setting this to "default" provides weak implementation stub that could be overwritten in application code. Setting this to "custom" provides hook's declaration only and expects the application to implement it.

Available options:

- No hook declared (CONFIG_LWIP_HOOK_IP6_SELECT_SRC_ADDR_NONE)
- Default (weak) implementation (CONFIG_LWIP_HOOK_IP6_SELECT_SRC_ADDR_DEFAULT)
- Custom implementation (CONFIG_LWIP_HOOK_IP6_SELECT_SRC_ADDR_CUSTOM)

CONFIG_LWIP_HOOK_NETCONN_EXTERNAL_RESOLVE

Netconn external resolve Hook

Found in: [Component config](#) > [LWIP](#) > [Hooks](#)

Enables custom DNS resolve hook (without callback). Setting this to "default" provides weak implementation stub that could be overwritten in application code. Setting this to "custom" provides hook's declaration only and expects the application to implement it.

Available options:

- No hook declared (CONFIG_LWIP_HOOK_NETCONN_EXT_RESOLVE_NONE)
- Default (weak) implementation (CONFIG_LWIP_HOOK_NETCONN_EXT_RESOLVE_DEFAULT)
- Custom implementation (CONFIG_LWIP_HOOK_NETCONN_EXT_RESOLVE_CUSTOM)

CONFIG_LWIP_HOOK_DNS_EXTERNAL_RESOLVE

DNS external resolve Hook

Found in: [Component config](#) > [LWIP](#) > [Hooks](#)

Enables custom DNS resolve hook (with callback). Setting this to "custom" provides hook's declaration only and expects the application to implement it.

Available options:

- No hook declared (CONFIG_LWIP_HOOK_DNS_EXT_RESOLVE_NONE)
- Custom implementation (CONFIG_LWIP_HOOK_DNS_EXT_RESOLVE_CUSTOM)

CONFIG_LWIP_HOOK_IP6_INPUT

IPv6 packet input

Found in: [Component config](#) > [LWIP](#) > [Hooks](#)

Enables custom IPv6 packet input. Setting this to "default" provides weak IDF implementation, which drops all incoming IPv6 traffic if the interface has no link local address. (this default implementation is "weak" and could be still overwritten in the application if some additional IPv6 input packet filtering is needed) Setting this to "none" removes this default filter and conforms to the lwIP implementation (which accepts multicasts even if the interface has no link local address) Setting this to "custom" provides hook's declaration only and expects the application to implement it.

Available options:

- No hook declared (CONFIG_LWIP_HOOK_IP6_INPUT_NONE)
- Default (weak) implementation (CONFIG_LWIP_HOOK_IP6_INPUT_DEFAULT)
- Custom implementation (CONFIG_LWIP_HOOK_IP6_INPUT_CUSTOM)

CONFIG_LWIP_DEBUG

Enable LWIP Debug

Found in: [Component config](#) > [LWIP](#)

Enabling this option allows different kinds of lwIP debug output.

All lwIP debug features increase the size of the final binary.

Default value:

- No (disabled)

Contains:

- [CONFIG_LWIP_API_LIB_DEBUG](#)
- [CONFIG_LWIP_BRIDGEIF_FDB_DEBUG](#)
- [CONFIG_LWIP_BRIDGEIF_FW_DEBUG](#)
- [CONFIG_LWIP_BRIDGEIF_DEBUG](#)
- [CONFIG_LWIP_DHCP_DEBUG](#)
- [CONFIG_LWIP_DHCP_STATE_DEBUG](#)
- [CONFIG_LWIP_DNS_DEBUG](#)
- [CONFIG_LWIP_ETHARP_DEBUG](#)
- [CONFIG_LWIP_ICMP_DEBUG](#)
- [CONFIG_LWIP_ICMP6_DEBUG](#)
- [CONFIG_LWIP_IP_DEBUG](#)
- [CONFIG_LWIP_IP6_DEBUG](#)
- [CONFIG_LWIP_NAPT_DEBUG](#)
- [CONFIG_LWIP_NETIF_DEBUG](#)

- [*CONFIG_LWIP_PBUF_DEBUG*](#)
- [*CONFIG_LWIP_SNTP_DEBUG*](#)
- [*CONFIG_LWIP_SOCKETS_DEBUG*](#)
- [*CONFIG_LWIP_TCP_DEBUG*](#)
- [*CONFIG_LWIP_UDP_DEBUG*](#)
- [*CONFIG_LWIP_DEBUG_ESP_LOG*](#)

CONFIG_LWIP_DEBUG_ESP_LOG

Route LWIP debugs through ESP_LOG interface

Found in: [*Component config*](#) > [*LWIP*](#) > [*CONFIG_LWIP_DEBUG*](#)

Enabling this option routes all enabled LWIP debugs through ESP_LOGD.

Default value:

- No (disabled) if [*CONFIG_LWIP_DEBUG*](#)

CONFIG_LWIP_NETIF_DEBUG

Enable netif debug messages

Found in: [*Component config*](#) > [*LWIP*](#) > [*CONFIG_LWIP_DEBUG*](#)

Default value:

- No (disabled) if [*CONFIG_LWIP_DEBUG*](#)

CONFIG_LWIP_PBUF_DEBUG

Enable pbuf debug messages

Found in: [*Component config*](#) > [*LWIP*](#) > [*CONFIG_LWIP_DEBUG*](#)

Default value:

- No (disabled) if [*CONFIG_LWIP_DEBUG*](#)

CONFIG_LWIP_ETHARP_DEBUG

Enable etharp debug messages

Found in: [*Component config*](#) > [*LWIP*](#) > [*CONFIG_LWIP_DEBUG*](#)

Default value:

- No (disabled) if [*CONFIG_LWIP_DEBUG*](#)

CONFIG_LWIP_API_LIB_DEBUG

Enable api lib debug messages

Found in: [*Component config*](#) > [*LWIP*](#) > [*CONFIG_LWIP_DEBUG*](#)

Default value:

- No (disabled) if [*CONFIG_LWIP_DEBUG*](#)

CONFIG_LWIP_SOCKETS_DEBUG

Enable socket debug messages

Found in: [*Component config*](#) > [*LWIP*](#) > [*CONFIG_LWIP_DEBUG*](#)

Default value:

- No (disabled) if [*CONFIG_LWIP_DEBUG*](#)

CONFIG_LWIP_IP_DEBUG

Enable IP debug messages

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_DEBUG](#)

Default value:

- No (disabled) if [CONFIG_LWIP_DEBUG](#)

CONFIG_LWIP_ICMP_DEBUG

Enable ICMP debug messages

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_DEBUG](#)

Default value:

- No (disabled) if [CONFIG_LWIP_DEBUG](#) && [CONFIG_LWIP_ICMP](#)

CONFIG_LWIP_DHCP_STATE_DEBUG

Enable DHCP state tracking

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_DEBUG](#)

Default value:

- No (disabled) if [CONFIG_LWIP_DEBUG](#)

CONFIG_LWIP_DHCP_DEBUG

Enable DHCP debug messages

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_DEBUG](#)

Default value:

- No (disabled) if [CONFIG_LWIP_DEBUG](#)

CONFIG_LWIP_IP6_DEBUG

Enable IP6 debug messages

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_DEBUG](#)

Default value:

- No (disabled) if [CONFIG_LWIP_DEBUG](#)

CONFIG_LWIP_ICMP6_DEBUG

Enable ICMP6 debug messages

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_DEBUG](#)

Default value:

- No (disabled) if [CONFIG_LWIP_DEBUG](#)

CONFIG_LWIP_TCP_DEBUG

Enable TCP debug messages

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_DEBUG](#)

Default value:

- No (disabled) if [CONFIG_LWIP_DEBUG](#)

CONFIG_LWIP_UDP_DEBUG

Enable UDP debug messages

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_DEBUG](#)

Default value:

- No (disabled) if [CONFIG_LWIP_DEBUG](#)

CONFIG_LWIP_SNTP_DEBUG

Enable SNTP debug messages

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_DEBUG](#)

Default value:

- No (disabled) if [CONFIG_LWIP_DEBUG](#)

CONFIG_LWIP_DNS_DEBUG

Enable DNS debug messages

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_DEBUG](#)

Default value:

- No (disabled) if [CONFIG_LWIP_DEBUG](#)

CONFIG_LWIP_NAPT_DEBUG

Enable NAPT debug messages

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_DEBUG](#)

Default value:

- No (disabled) if [CONFIG_LWIP_DEBUG](#) && [CONFIG_LWIP_IPV4_NAPT](#)

CONFIG_LWIP_BRIDGEIF_DEBUG

Enable bridge generic debug messages

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_DEBUG](#)

Default value:

- No (disabled) if [CONFIG_LWIP_DEBUG](#)

CONFIG_LWIP_BRIDGEIF_FDB_DEBUG

Enable bridge FDB debug messages

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_DEBUG](#)

Default value:

- No (disabled) if [CONFIG_LWIP_DEBUG](#)

CONFIG_LWIP_BRIDGEIF_FW_DEBUG

Enable bridge forwarding debug messages

Found in: [Component config](#) > [LWIP](#) > [CONFIG_LWIP_DEBUG](#)

Default value:

- No (disabled) if [CONFIG_LWIP_DEBUG](#)

MBEDTLS Contains:

- `CONFIG_MBEDTLS_ALLOW_WEAK_CERTIFICATE_VERIFICATION`
- `CONFIG_MBEDTLS_ASYMMETRIC_CONTENT_LEN`
- *Certificate Bundle*
- *Certificates*
- `CONFIG_MBEDTLS_CHACHA20_C`
- `CONFIG_MBEDTLS_DHM_C`
- `CONFIG_MBEDTLS_ECP_C`
- `CONFIG_MBEDTLS_ECDH_C`
- `CONFIG_MBEDTLS_ECJPAKE_C`
- `CONFIG_MBEDTLS_AES_USE_PSEUDO_ROUND_FUNC`
- `CONFIG_MBEDTLS_ECP_DP_BP256R1_ENABLED`
- `CONFIG_MBEDTLS_ECP_DP_BP384R1_ENABLED`
- `CONFIG_MBEDTLS_ECP_DP_BP512R1_ENABLED`
- `CONFIG_MBEDTLS_CMAC_C`
- `CONFIG_MBEDTLS_ECP_DP_CURVE25519_ENABLED`
- `CONFIG_MBEDTLS_ECDSA_DETERMINISTIC`
- `CONFIG_MBEDTLS_HARDWARE_ECDSA_VERIFY`
- `CONFIG_MBEDTLS_HARDWARE_ECDSA_SIGN`
- `CONFIG_MBEDTLS_ERROR_STRINGS`
- `CONFIG_MBEDTLS_ECP_FIXED_POINT_OPTIM`
- `CONFIG_MBEDTLS_FS_IO`
- `CONFIG_MBEDTLS_HARDWARE_AES`
- `CONFIG_MBEDTLS_HARDWARE_ECC`
- `CONFIG_MBEDTLS_ATCA_HW_ECDSA_SIGN`
- `CONFIG_MBEDTLS_ATCA_HW_ECDSA_VERIFY`
- `CONFIG_MBEDTLS_HARDWARE_MPI`
- `CONFIG_MBEDTLS_HARDWARE_SHA`
- `CONFIG_MBEDTLS_DEBUG`
- `CONFIG_MBEDTLS_ECP_RESTARTABLE`
- `CONFIG_MBEDTLS_HAVE_TIME`
- `CONFIG_MBEDTLS_HARDWARE_GCM`
- `CONFIG_MBEDTLS_RIPEMD160_C`
- `CONFIG_MBEDTLS_ECP_DP_SECP192K1_ENABLED`
- `CONFIG_MBEDTLS_ECP_DP_SECP192R1_ENABLED`
- `CONFIG_MBEDTLS_ECP_DP_SECP224K1_ENABLED`
- `CONFIG_MBEDTLS_ECP_DP_SECP224R1_ENABLED`
- `CONFIG_MBEDTLS_ECP_DP_SECP256K1_ENABLED`
- `CONFIG_MBEDTLS_ECP_DP_SECP256R1_ENABLED`
- `CONFIG_MBEDTLS_ECP_DP_SECP384R1_ENABLED`
- `CONFIG_MBEDTLS_ECP_DP_SECP521R1_ENABLED`
- `CONFIG_MBEDTLS_GCM_SUPPORT_NON_AES_CIPHER`
- `CONFIG_MBEDTLS_SHA512_C`
- `CONFIG_MBEDTLS_SHA3_C`
- `CONFIG_MBEDTLS_THREADING_C`
- `CONFIG_MBEDTLS_HKDF_C`
- *MBEDTLS v3.x related*
- `CONFIG_MBEDTLS_MEM_ALLOC_MODE`
- `CONFIG_MBEDTLS_ECP_NIST_OPTIM`
- `CONFIG_MBEDTLS_POLY1305_C`
- `CONFIG_MBEDTLS_SSL_ALPN`
- `CONFIG_MBEDTLS_SSL_PROTO_DTLS`
- `CONFIG_MBEDTLS_SSL_PROTO_GMTSSL1_1`
- `CONFIG_MBEDTLS_SSL_PROTO_TLS1_2`
- `CONFIG_MBEDTLS_SSL_RENEGOTIATION`
- *Symmetric Ciphers*
- *TLS Key Exchange Methods*
- `CONFIG_MBEDTLS_SSL_MAX_CONTENT_LEN`

- [CONFIG_MBEDTLS_TLS_MODE](#)
- [CONFIG_MBEDTLS_CLIENT_SSL_SESSION_TICKETS](#)
- [CONFIG_MBEDTLS_SERVER_SSL_SESSION_TICKETS](#)
- [CONFIG_MBEDTLS_ROM_MD5](#)
- [CONFIG_MBEDTLS_USE_CRYPTO_ROM_IMPL](#)
- [CONFIG_MBEDTLS_DYNAMIC_BUFFER](#)

CONFIG_MBEDTLS_MEM_ALLOC_MODE

Memory allocation strategy

Found in: [Component config > mbedTLS](#)

Allocation strategy for mbedTLS, essentially provides ability to allocate all required dynamic allocations from,

- Internal DRAM memory only
- External SPIRAM memory only
- Either internal or external memory based on default malloc() behavior in ESP-IDF
- Custom allocation mode, by overwriting calloc()/free() using mbedtls_platform_set_calloc_free() function
- Internal IRAM memory wherever applicable else internal DRAM

Recommended mode here is always internal (*), since that is most preferred from security perspective. But if application requirement does not allow sufficient free internal memory then alternate mode can be selected.

(*) In case of ESP32-S2/ESP32-S3, hardware allows encryption of external SPIRAM contents provided hardware flash encryption feature is enabled. In that case, using external SPIRAM allocation strategy is also safe choice from security perspective.

Available options:

- Internal memory (CONFIG_MBEDTLS_INTERNAL_MEM_ALLOC)
- External SPIRAM (CONFIG_MBEDTLS_EXTERNAL_MEM_ALLOC)
- Default alloc mode (CONFIG_MBEDTLS_DEFAULT_MEM_ALLOC)
- Custom alloc mode (CONFIG_MBEDTLS_CUSTOM_MEM_ALLOC)
- Internal IRAM (CONFIG_MBEDTLS_IRAM_8BIT_MEM_ALLOC)

Allows to use IRAM memory region as 8bit accessible region.

TLS input and output buffers will be allocated in IRAM section which is 32bit aligned memory. Every unaligned (8bit or 16bit) access will result in an exception and incur penalty of certain clock cycles per unaligned read/write.

CONFIG_MBEDTLS_SSL_MAX_CONTENT_LEN

TLS maximum message content length

Found in: [Component config > mbedTLS](#)

Maximum TLS message length (in bytes) supported by mbedTLS.

16384 is the default and this value is required to comply fully with TLS standards.

However you can set a lower value in order to save RAM. This is safe if the other end of the connection supports Maximum Fragment Length Negotiation Extension (max_fragment_length, see RFC6066) or you know for certain that it will never send a message longer than a certain number of bytes.

If the value is set too low, symptoms are a failed TLS handshake or a return value of MBEDTLS_ERR_SSL_INVALID_RECORD (-0x7200).

CONFIG_MBEDTLS_ASYMMETRIC_CONTENT_LEN

Asymmetric in/out fragment length

Found in: *Component config* > *mbedTLS*

If enabled, this option allows customizing TLS in/out fragment length in asymmetric way. Please note that enabling this with default values saves 12KB of dynamic memory per TLS connection.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_SSL_IN_CONTENT_LEN

TLS maximum incoming fragment length

Found in: *Component config* > *mbedTLS* > *CONFIG_MBEDTLS_ASYMMETRIC_CONTENT_LEN*

This defines maximum incoming fragment length, overriding default maximum content length (MBEDTLS_SSL_MAX_CONTENT_LEN).

Range:

- from 512 to 16384

Default value:

- 16384

CONFIG_MBEDTLS_SSL_OUT_CONTENT_LEN

TLS maximum outgoing fragment length

Found in: *Component config* > *mbedTLS* > *CONFIG_MBEDTLS_ASYMMETRIC_CONTENT_LEN*

This defines maximum outgoing fragment length, overriding default maximum content length (MBEDTLS_SSL_MAX_CONTENT_LEN).

Range:

- from 512 to 16384

Default value:

- 4096

CONFIG_MBEDTLS_DYNAMIC_BUFFER

Using dynamic TX/RX buffer

Found in: *Component config* > *mbedTLS*

Using dynamic TX/RX buffer. After enabling this option, mbedTLS will allocate TX buffer when need to send data and then free it if all data is sent, allocate RX buffer when need to receive data and then free it when all data is used or read by upper layer.

By default, when SSL is initialized, mbedTLS also allocate TX and RX buffer with the default value of "MBEDTLS_SSL_OUT_CONTENT_LEN" or "MBEDTLS_SSL_IN_CONTENT_LEN", so to save more heap, users can set the options to be an appropriate value.

CONFIG_MBEDTLS_DYNAMIC_FREE_CONFIG_DATA

Free private key and DHM data after its usage

Found in: *Component config* > *mbedTLS* > *CONFIG_MBEDTLS_DYNAMIC_BUFFER*

Free private key and DHM data after its usage in handshake process.

The option will decrease heap cost when handshake, but also lead to problem:

Because all certificate, private key and DHM data are freed so users should register certificate and private key to ssl config object again.

Default value:

- No (disabled) if [CONFIG_MBEDTLS_DYNAMIC_BUFFER](#)

CONFIG_MBEDTLS_DYNAMIC_FREE_CA_CERT

Free SSL CA certificate after its usage

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_DYNAMIC_BUFFER](#) > [CONFIG_MBEDTLS_DYNAMIC_FREE_CONFIG_DATA](#)

Free CA certificate after its usage in the handshake process. This option will decrease the heap footprint for the TLS handshake, but may lead to a problem: If the respective ssl object needs to perform the TLS handshake again, the CA certificate should once again be registered to the ssl object.

Default value:

- Yes (enabled) if [CONFIG_MBEDTLS_DYNAMIC_FREE_CONFIG_DATA](#)

CONFIG_MBEDTLS_DEBUG

Enable mbedtls debugging

Found in: [Component config](#) > [mbedtls](#)

Enable mbedtls debugging functions at compile time.

If this option is enabled, you can include "mbedtls/esp_debug.h" and call `mbedtls_esp_enable_debug_log()` at runtime in order to enable mbedtls debug output via the ESP log mechanism.

Default value:

- No (disabled)

CONFIG_MBEDTLS_DEBUG_LEVEL

Set mbedtls debugging level

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_DEBUG](#)

Set mbedtls debugging level

Available options:

- Warning ([CONFIG_MBEDTLS_DEBUG_LEVEL_WARN](#))
- Info ([CONFIG_MBEDTLS_DEBUG_LEVEL_INFO](#))
- Debug ([CONFIG_MBEDTLS_DEBUG_LEVEL_DEBUG](#))
- Verbose ([CONFIG_MBEDTLS_DEBUG_LEVEL_VERBOSE](#))

mbedtls v3.x related Contains:

- [DTLS-based configurations](#)
- [CONFIG_MBEDTLS_SSL_KEYING_MATERIAL_EXPORT](#)
- [CONFIG_MBEDTLS_PKCS7_C](#)
- [CONFIG_MBEDTLS_SSL_CONTEXT_SERIALIZATION](#)
- [CONFIG_MBEDTLS_X509_TRUSTED_CERT_CALLBACK](#)
- [CONFIG_MBEDTLS_SSL_KEEP_PEER_CERTIFICATE](#)
- [CONFIG_MBEDTLS_SSL_CID_PADDING_GRANULARITY](#)
- [CONFIG_MBEDTLS_SSL_PROTO_TLS1_3](#)
- [CONFIG_MBEDTLS_ECDH_LEGACY_CONTEXT](#)
- [CONFIG_MBEDTLS_SSL_VARIABLE_BUFFER_LENGTH](#)

CONFIG_MBEDTLS_SSL_PROTO_TLS1_3

Support TLS 1.3 protocol

Found in: *Component config > mbedTLS > mbedTLS v3.x related*

TLS 1.3 related configurations Contains:

- [CONFIG_MBEDTLS_SSL_TLS1_3_KEXM_EPHEMERAL](#)
- [CONFIG_MBEDTLS_SSL_TLS1_3_COMPATIBILITY_MODE](#)
- [CONFIG_MBEDTLS_SSL_TLS1_3_KEXM_PSK_EPHEMERAL](#)
- [CONFIG_MBEDTLS_SSL_TLS1_3_KEXM_PSK](#)

CONFIG_MBEDTLS_SSL_TLS1_3_COMPATIBILITY_MODE

TLS 1.3 middlebox compatibility mode

Found in: *Component config > mbedTLS > mbedTLS v3.x related > CONFIG_MBEDTLS_SSL_PROTO_TLS1_3 > TLS 1.3 related configurations*

Default value:

- Yes (enabled) if [CONFIG_MBEDTLS_SSL_PROTO_TLS1_3](#)

CONFIG_MBEDTLS_SSL_TLS1_3_KEXM_PSK

TLS 1.3 PSK key exchange mode

Found in: *Component config > mbedTLS > mbedTLS v3.x related > CONFIG_MBEDTLS_SSL_PROTO_TLS1_3 > TLS 1.3 related configurations*

Default value:

- Yes (enabled) if [CONFIG_MBEDTLS_SSL_PROTO_TLS1_3](#)

CONFIG_MBEDTLS_SSL_TLS1_3_KEXM_EPHEMERAL

TLS 1.3 ephemeral key exchange mode

Found in: *Component config > mbedTLS > mbedTLS v3.x related > CONFIG_MBEDTLS_SSL_PROTO_TLS1_3 > TLS 1.3 related configurations*

Default value:

- Yes (enabled) if [CONFIG_MBEDTLS_SSL_PROTO_TLS1_3](#)

CONFIG_MBEDTLS_SSL_TLS1_3_KEXM_PSK_EPHEMERAL

TLS 1.3 PSK ephemeral key exchange mode

Found in: *Component config > mbedTLS > mbedTLS v3.x related > CONFIG_MBEDTLS_SSL_PROTO_TLS1_3 > TLS 1.3 related configurations*

Default value:

- Yes (enabled) if [CONFIG_MBEDTLS_SSL_PROTO_TLS1_3](#)

CONFIG_MBEDTLS_SSL_VARIABLE_BUFFER_LENGTH

Variable SSL buffer length

Found in: *Component config > mbedTLS > mbedTLS v3.x related*

This enables the SSL buffer to be resized automatically based on the negotiated maximum fragment length in each direction.

Default value:

- No (disabled)

CONFIG_MBEDTLS_ECDH_LEGACY_CONTEXT

Use a backward compatible ECDH context (Experimental)

Found in: [Component config](#) > [mbedtls](#) > [mbedtls v3.x related](#)

Use the legacy ECDH context format. Define this option only if you enable MBEDTLS_ECP_RESTARTABLE or if you want to access ECDH context fields directly.

Default value:

- No (disabled) if [CONFIG_MBEDTLS_ECDH_C](#) && [CONFIG_MBEDTLS_ECP_RESTARTABLE](#)

CONFIG_MBEDTLS_X509_TRUSTED_CERT_CALLBACK

Enable trusted certificate callbacks

Found in: [Component config](#) > [mbedtls](#) > [mbedtls v3.x related](#)

Enables users to configure the set of trusted certificates through a callback instead of a linked list.

See mbedtls documentation for required API and more details.

Default value:

- No (disabled)

CONFIG_MBEDTLS_SSL_CONTEXT_SERIALIZATION

Enable serialization of the TLS context structures

Found in: [Component config](#) > [mbedtls](#) > [mbedtls v3.x related](#)

Enable serialization of the TLS context structures This is a local optimization in handling a single, potentially long-lived connection.

See mbedtls documentation for required API and more details. Disabling this option will save some code size.

Default value:

- No (disabled)

CONFIG_MBEDTLS_SSL_KEEP_PEER_CERTIFICATE

Keep peer certificate after handshake completion

Found in: [Component config](#) > [mbedtls](#) > [mbedtls v3.x related](#)

Keep the peer's certificate after completion of the handshake. Disabling this option will save about 4kB of heap and some code size.

See mbedtls documentation for required API and more details.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_SSL_KEYING_MATERIAL_EXPORT

Enable keying material export

Found in: [Component config](#) > [mbedtls](#) > [mbedtls v3.x related](#)

Enable shared symmetric keys export for TLS sessions using `mbedtls_ssl_export_keying_material()` after SSL handshake. The process for deriving the keys is specified in RFC 5705 for TLS 1.2 and in RFC 8446, Section 7.5, for TLS 1.3.

Default value:

- No (disabled)

CONFIG_MBEDTLS_PKCS7_C

Enable PKCS #7

Found in: [Component config](#) > [mbedtls](#) > [mbedtls v3.x related](#)

Enable PKCS #7 core for using PKCS #7-formatted signatures.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_SSL_CID_PADDING_GRANULARITY

Record plaintext padding

Found in: [Component config](#) > [mbedtls](#) > [mbedtls v3.x related](#)

Controls the use of record plaintext padding in TLS 1.3 and when using the Connection ID extension in DTLS 1.2.

The padding will always be chosen so that the length of the padded plaintext is a multiple of the value of this option.

Notes: A value of 1 means that no padding will be used for outgoing records. On systems lacking division instructions, a power of two should be preferred.

Range:

- from 0 to 32 if [CONFIG_MBEDTLS_SSL_PROTO_TLS1_3](#) || [CONFIG_MBEDTLS_SSL_DTLS_CONNECTION_ID](#)

Default value:

- 16 if [CONFIG_MBEDTLS_SSL_PROTO_TLS1_3](#) || [CONFIG_MBEDTLS_SSL_DTLS_CONNECTION_ID](#)

DTLS-based configurations

 Contains:

- [CONFIG_MBEDTLS_SSL_DTLS_SRTP](#)
- [CONFIG_MBEDTLS_SSL_DTLS_CONNECTION_ID](#)

CONFIG_MBEDTLS_SSL_DTLS_CONNECTION_ID

Support for the DTLS Connection ID extension

Found in: [Component config](#) > [mbedtls](#) > [mbedtls v3.x related](#) > [DTLS-based configurations](#)

Enable support for the DTLS Connection ID extension which allows to identify DTLS connections across changes in the underlying transport.

Default value:

- No (disabled) if [CONFIG_MBEDTLS_SSL_PROTO_DTLS](#)

CONFIG_MBEDTLS_SSL_CID_IN_LEN_MAX

Maximum length of CIDs used for incoming DTLS messages

Found in: [Component config](#) > [mbedtls](#) > [mbedtls v3.x related](#) > [DTLS-based configurations](#) > [CONFIG_MBEDTLS_SSL_DTLS_CONNECTION_ID](#)

Maximum length of CIDs used for incoming DTLS messages

Range:

- from 0 to 32 if [CONFIG_MBEDTLS_SSL_DTLS_CONNECTION_ID](#) && [CONFIG_MBEDTLS_SSL_PROTO_DTLS](#)

Default value:

- 32 if [CONFIG_MBEDTLS_SSL_DTLS_CONNECTION_ID](#) && [CONFIG_MBEDTLS_SSL_PROTO_DTLS](#)

CONFIG_MBEDTLS_SSL_CID_OUT_LEN_MAX

Maximum length of CIDs used for outgoing DTLS messages

Found in: [Component config](#) > [mbedtls](#) > [mbedtls v3.x related](#) > [DTLS-based configurations](#) > [CONFIG_MBEDTLS_SSL_DTLS_CONNECTION_ID](#)

Maximum length of CIDs used for outgoing DTLS messages

Range:

- from 0 to 32 if [CONFIG_MBEDTLS_SSL_DTLS_CONNECTION_ID](#) && [CONFIG_MBEDTLS_SSL_PROTO_DTLS](#)

Default value:

- 32 if [CONFIG_MBEDTLS_SSL_DTLS_CONNECTION_ID](#) && [CONFIG_MBEDTLS_SSL_PROTO_DTLS](#)

CONFIG_MBEDTLS_SSL_DTLS_SRTP

Enable support for negotiation of DTLS-SRTP (RFC 5764)

Found in: [Component config](#) > [mbedtls](#) > [mbedtls v3.x related](#) > [DTLS-based configurations](#)

Enable support for negotiation of DTLS-SRTP (RFC 5764) through the use_srtp extension.

See mbedtls documentation for required API and more details. Disabling this option will save some code size.

Default value:

- No (disabled) if [CONFIG_MBEDTLS_SSL_PROTO_DTLS](#)

Certificate Bundle Contains:

- [CONFIG_MBEDTLS_CERTIFICATE_BUNDLE](#)

CONFIG_MBEDTLS_CERTIFICATE_BUNDLE

Enable trusted root certificate bundle

Found in: [Component config](#) > [mbedtls](#) > [Certificate Bundle](#)

Enable support for large number of default root certificates

When enabled this option allows user to store default as well as customer specific root certificates in compressed format rather than storing full certificate. For the root certificates the public key and the subject name will be stored.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_DEFAULT_CERTIFICATE_BUNDLE

Default certificate bundle options

Found in: [Component config](#) > [mbedtls](#) > [Certificate Bundle](#) > [CONFIG_MBEDTLS_CERTIFICATE_BUNDLE](#)

Available options:

- Use the full default certificate bundle ([CONFIG_MBEDTLS_CERTIFICATE_BUNDLE_DEFAULT_FULL](#))
- Use only the most common certificates from the default bundles ([CONFIG_MBEDTLS_CERTIFICATE_BUNDLE_DEFAULT_CMN](#))
Use only the most common certificates from the default bundles, reducing the size with 50%, while still having around 99% coverage.

- Do not use the default certificate bundle (CONFIG_MBEDTLS_CERTIFICATE_BUNDLE_DEFAULT_NONE)

CONFIG_MBEDTLS_CUSTOM_CERTIFICATE_BUNDLE

Add custom certificates to the default bundle

Found in: [Component config](#) > [mbedtls](#) > [Certificate Bundle](#) > [CONFIG_MBEDTLS_CERTIFICATE_BUNDLE](#)

Default value:

- No (disabled)

CONFIG_MBEDTLS_CUSTOM_CERTIFICATE_BUNDLE_PATH

Custom certificate bundle path

Found in: [Component config](#) > [mbedtls](#) > [Certificate Bundle](#) > [CONFIG_MBEDTLS_CERTIFICATE_BUNDLE](#) > [CONFIG_MBEDTLS_CUSTOM_CERTIFICATE_BUNDLE](#)

Name of the custom certificate directory or file. This path is evaluated relative to the project root directory.

CONFIG_MBEDTLS_CERTIFICATE_BUNDLE_MAX_CERTS

Maximum no of certificates allowed in certificate bundle

Found in: [Component config](#) > [mbedtls](#) > [Certificate Bundle](#) > [CONFIG_MBEDTLS_CERTIFICATE_BUNDLE](#)

Default value:

- 200

CONFIG_MBEDTLS_ECP_RESTARTABLE

Enable mbedtls ecp restartable

Found in: [Component config](#) > [mbedtls](#)

Enable "non-blocking" ECC operations that can return early and be resumed.

Default value:

- No (disabled)

CONFIG_MBEDTLS_CMAC_C

Enable CMAC mode for block ciphers

Found in: [Component config](#) > [mbedtls](#)

Enable the CMAC (Cipher-based Message Authentication Code) mode for block ciphers.

Default value:

- No (disabled)

CONFIG_MBEDTLS_HARDWARE_AES

Enable hardware AES acceleration

Found in: [Component config](#) > [mbedtls](#)

Enable hardware accelerated AES encryption & decryption.

Note that if the ESP32 CPU is running at 240MHz, hardware AES does not offer any speed boost over software AES.

CONFIG_MBEDTLS_AES_USE_INTERRUPT

Use interrupt for long AES operations

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_HARDWARE_AES](#)

Use an interrupt to coordinate long AES operations.

This allows other code to run on the CPU while an AES operation is pending. Otherwise the CPU busy-waits.

Default value:

- Yes (enabled) if [CONFIG_MBEDTLS_HARDWARE_AES](#)

CONFIG_MBEDTLS_AES_INTERRUPT_LEVEL

AES hardware interrupt level

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_HARDWARE_AES](#) > [CONFIG_MBEDTLS_AES_USE_INTERRUPT](#)

This config helps to set the interrupt priority level for the AES peripheral. Value 0 (default) means that there is no preference regarding the interrupt priority level and any level from 1 to 3 can be selected (based on the availability). Note: Higher value indicates high interrupt priority.

Range:

- from 0 to 3 if [CONFIG_MBEDTLS_AES_USE_INTERRUPT](#)

Default value:

- 0 if [CONFIG_MBEDTLS_AES_USE_INTERRUPT](#)

CONFIG_MBEDTLS_AES_USE_PSEUDO_ROUND_FUNC

Enable AES hardware's pseudo round function

Found in: [Component config](#) > [mbedtls](#)

Enables the pseudo round function of the AES peripheral. Enabling this would impact the performance of the AES operations. For more info regarding the performance impact, please checkout the pseudo round function section of the security guide.

Default value:

- No (disabled) if [SOC_AES_SUPPORT_PSEUDO_ROUND_FUNCTION](#)

CONFIG_MBEDTLS_AES_USE_PSEUDO_ROUND_FUNC_STRENGTH

Strength of the pseudo rounds function

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_AES_USE_PSEUDO_ROUND_FUNC](#)

The strength of the pseudo rounds functions can be configured to low, medium and high. You can configure the strength of the pseudo rounds functions according to your use cases, for example, increasing the strength would provide higher security but would slow down the hardware AES encryption/decryption operations.

Available options:

- Low ([CONFIG_MBEDTLS_AES_USE_PSEUDO_ROUND_FUNC_STRENGTH_LOW](#))
- Medium ([CONFIG_MBEDTLS_AES_USE_PSEUDO_ROUND_FUNC_STRENGTH_MEDIUM](#))
- High ([CONFIG_MBEDTLS_AES_USE_PSEUDO_ROUND_FUNC_STRENGTH_HIGH](#))

CONFIG_MBEDTLS_HARDWARE_GCM

Enable partially hardware accelerated GCM

Found in: [Component config](#) > [mbedtls](#)

Enable partially hardware accelerated GCM. GHASH calculation is still done in software.

If MBEDTLS_HARDWARE_GCM is disabled and MBEDTLS_HARDWARE_AES is enabled then mbedtls will still use the hardware accelerated AES block operation, but on a single block at a time.

Default value:

- Yes (enabled) if SOC_AES_SUPPORT_GCM && [CONFIG_MBEDTLS_HARDWARE_AES](#)

CONFIG_MBEDTLS_GCM_SUPPORT_NON_AES_CIPHER

Enable support for non-AES ciphers in GCM operation

Found in: [Component config](#) > [mbedtls](#)

Enable this config to support fallback to software definitions for a non-AES cipher GCM operation as we support hardware acceleration only for AES cipher. Some of the non-AES ciphers used in a GCM operation are DES, ARIA, CAMELLIA, CHACHA20, BLOWFISH.

If this config is disabled, performing a non-AES cipher GCM operation with the config MBEDTLS_HARDWARE_AES enabled will result in calculation of an AES-GCM operation instead for the given input values and thus could lead to failure in certificate validation which would ultimately lead to a SSL handshake failure.

This config being by-default enabled leads to an increase in binary size footprint of ~2.5KB. In case you are sure that your use case (for example, client and server configurations in case of a TLS handshake) would not involve any GCM operations using a non-AES cipher, you can safely disable this config, leading to reduction in binary size footprint.

Default value:

- No (disabled) if [CONFIG_MBEDTLS_HARDWARE_AES](#)

CONFIG_MBEDTLS_HARDWARE_MPI

Enable hardware MPI (bignum) acceleration

Found in: [Component config](#) > [mbedtls](#)

Enable hardware accelerated multiple precision integer operations.

Hardware accelerated multiplication, modulo multiplication, and modular exponentiation for up to SOC_RSA_MAX_BIT_LEN bit results.

These operations are used by RSA.

CONFIG_MBEDTLS_LARGE_KEY_SOFTWARE_MPI

Fallback to software implementation for larger MPI values

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_HARDWARE_MPI](#)

Fallback to software implementation for RSA key lengths larger than SOC_RSA_MAX_BIT_LEN. If this is not active then the ESP will be unable to process keys greater than SOC_RSA_MAX_BIT_LEN.

Default value:

- No (disabled)

CONFIG_MBEDTLS_MPI_USE_INTERRUPT

Use interrupt for MPI exp-mod operations

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_HARDWARE_MPI](#)

Use an interrupt to coordinate long MPI operations.

This allows other code to run on the CPU while an MPI operation is pending. Otherwise the CPU busy-waits.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_MPI_INTERRUPT_LEVEL

MPI hardware interrupt level

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_HARDWARE_MPI](#) > [CONFIG_MBEDTLS_MPI_USE_INTERRUPT](#)

This config helps to set the interrupt priority level for the MPI peripheral. Value 0 (default) means that there is no preference regarding the interrupt priority level and any level from 1 to 3 can be selected (based on the availability). Note: Higher value indicates high interrupt priority.

Range:

- from 0 to 3

Default value:

- 0

CONFIG_MBEDTLS_HARDWARE_SHA

Enable hardware SHA acceleration

Found in: [Component config](#) > [mbedtls](#)

Enable hardware accelerated SHA1, SHA256, SHA384 & SHA512 in mbedtls.

Due to a hardware limitation, on the ESP32 hardware acceleration is only guaranteed if SHA digests are calculated one at a time. If more than one SHA digest is calculated at the same time, one will be calculated fully in hardware and the rest will be calculated (at least partially calculated) in software. This happens automatically.

SHA hardware acceleration is faster than software in some situations but slower in others. You should benchmark to find the best setting for you.

CONFIG_MBEDTLS_HARDWARE_ECC

Enable hardware ECC acceleration

Found in: [Component config](#) > [mbedtls](#)

Enable hardware accelerated ECC point multiplication and point verification for points on curve SECP192R1 and SECP256R1 in mbedtls

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_ECC_OTHER_CURVES_SOFT_FALLBACK

Fallback to software implementation for curves not supported in hardware

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_HARDWARE_ECC](#)

Fallback to software implementation of ECC point multiplication and point verification for curves not supported in hardware.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_ROM_MD5

Use MD5 implementation in ROM

Found in: [Component config](#) > [mbedtls](#)

Use ROM MD5 in mbedtls.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_HARDWARE_ECDSA_SIGN

Enable ECDSA signing using on-chip ECDSA peripheral

Found in: [Component config](#) > [mbedtls](#)

Enable hardware accelerated ECDSA peripheral to sign data on curve SECP192R1 and SECP256R1 in mbedtls.

Note that for signing, the private key has to be burnt in an efuse key block with key purpose set to ECDSA_KEY. If no key is burnt, it will report an error

The key should be burnt in little endian format. espfuse.py utility handles it internally but care needs to be taken while burning using esp_efuse APIs

Default value:

- No (disabled)

Enable Software Countermeasure for ECDSA signing using on-chip ECDSA peripheral**CONFIG_MBEDTLS_HARDWARE_ECDSA_VERIFY**

Enable ECDSA signature verification using on-chip ECDSA peripheral

Found in: [Component config](#) > [mbedtls](#)

Enable hardware accelerated ECDSA peripheral to verify signature on curve SECP192R1 and SECP256R1 in mbedtls.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_ATCA_HW_ECDSA_SIGN

Enable hardware ECDSA sign acceleration when using ATECC608A

Found in: [Component config](#) > [mbedtls](#)

This option enables hardware acceleration for ECDSA sign function, only when using ATECC608A cryptoauth chip (integrated with ESP32-WROOM-32SE)

Default value:

- No (disabled)

CONFIG_MBEDTLS_ATCA_HW_ECDSA_VERIFY

Enable hardware ECDSA verify acceleration when using ATECC608A

Found in: [Component config](#) > [mbedtls](#)

This option enables hardware acceleration for ECDSA sign function, only when using ATECC608A cryptoauth chip (integrated with ESP32-WROOM-32SE)

Default value:

- No (disabled)

CONFIG_MBEDTLS_HAVE_TIME

Enable mbedtls time support

Found in: [Component config](#) > [mbedtls](#)

Enable use of time.h functions (time() and gmtime()) by mbedtls.

This option doesn't require the system time to be correct, but enables functionality that requires relative timekeeping - for example periodic expiry of TLS session tickets or session cache entries.

Disabling this option will save some firmware size, particularly if the rest of the firmware doesn't call any standard timekeeping functions.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_PLATFORM_TIME_ALT

Enable mbedtls time support: platform-specific

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_HAVE_TIME](#)

Enabling this config will provide users with a function "mbedtls_platform_set_time()" that allows to set an alternative time function pointer.

Default value:

- No (disabled)

CONFIG_MBEDTLS_HAVE_TIME_DATE

Enable mbedtls certificate expiry check

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_HAVE_TIME](#)

Enables X.509 certificate expiry checks in mbedtls.

If this option is disabled (default) then X.509 certificate "valid from" and "valid to" timestamp fields are ignored.

If this option is enabled, these fields are compared with the current system date and time. The time is retrieved using the standard time() and gmtime() functions. If the certificate is not valid for the current system time then verification will fail with code MBEDTLS_X509_BADCERT_FUTURE or MBEDTLS_X509_BADCERT_EXPIRED.

Enabling this option requires adding functionality in the firmware to set the system clock to a valid timestamp before using TLS. The recommended way to do this is via ESP-IDF's SNTP functionality, but any method can be used.

In the case where only a small number of certificates are trusted by the device, please carefully consider the tradeoffs of enabling this option. There may be undesired consequences, for example if all trusted certificates expire while the device is offline and a TLS connection is required to update. Or if an issue with the SNTP server means that the system time is invalid for an extended period after a reset.

Default value:

- No (disabled)

CONFIG_MBEDTLS_ECDSA_DETERMINISTIC

Enable deterministic ECDSA

Found in: [Component config](#) > [mbedtls](#)

Standard ECDSA is "fragile" in the sense that lack of entropy when signing may result in a compromise of the long-term signing key.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_SHA512_C

Enable the SHA-384 and SHA-512 cryptographic hash algorithms

Found in: [Component config](#) > [mbedtls](#)

Enable MBEDTLS_SHA512_C adds support for SHA-384 and SHA-512.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_SHA3_C

Enable the SHA3 cryptographic hash algorithm

Found in: [Component config](#) > [mbedtls](#)

Enabling MBEDTLS_SHA3_C adds support for SHA3. Enabling this configuration option increases the flash footprint by almost 4KB.

Default value:

- No (disabled)

CONFIG_MBEDTLS_TLS_MODE

TLS Protocol Role

Found in: [Component config](#) > [mbedtls](#)

mbedtls can be compiled with protocol support for the TLS server, TLS client, or both server and client.

Reducing the number of TLS roles supported saves code size.

Available options:

- Server & Client (CONFIG_MBEDTLS_TLS_SERVER_AND_CLIENT)
- Server (CONFIG_MBEDTLS_TLS_SERVER_ONLY)
- Client (CONFIG_MBEDTLS_TLS_CLIENT_ONLY)
- None (CONFIG_MBEDTLS_TLS_DISABLED)

TLS Key Exchange Methods

 Contains:

- [CONFIG_MBEDTLS_KEY_EXCHANGE_DHE_RSA](#)
- [CONFIG_MBEDTLS_KEY_EXCHANGE_ECJPAKE](#)
- [CONFIG_MBEDTLS_PSK_MODES](#)
- [CONFIG_MBEDTLS_KEY_EXCHANGE_RSA](#)
- [CONFIG_MBEDTLS_KEY_EXCHANGE_ELLIPTIC_CURVE](#)

CONFIG_MBEDTLS_PSK_MODES

Enable pre-shared-key ciphersuites

Found in: [Component config](#) > [mbedtls](#) > [TLS Key Exchange Methods](#)

Enable to show configuration for different types of pre-shared-key TLS authentication methods.

Leaving this options disabled will save code size if they are not used.

Default value:

- No (disabled)

CONFIG_MBEDTLS_KEY_EXCHANGE_PSK

Enable PSK based ciphersuite modes

Found in: [Component config](#) > [mbedtls](#) > [TLS Key Exchange Methods](#) > [CONFIG_MBEDTLS_PSK_MODES](#)

Enable to support symmetric key PSK (pre-shared-key) TLS key exchange modes.

Default value:

- No (disabled) if [CONFIG_MBEDTLS_PSK_MODES](#)

CONFIG_MBEDTLS_KEY_EXCHANGE_DHE_PSK

Enable DHE-PSK based ciphersuite modes

Found in: [Component config](#) > [mbedtls](#) > [TLS Key Exchange Methods](#) > [CONFIG_MBEDTLS_PSK_MODES](#)

Enable to support Diffie-Hellman PSK (pre-shared-key) TLS authentication modes.

Default value:

- Yes (enabled) if [CONFIG_MBEDTLS_PSK_MODES](#) && [CONFIG_MBEDTLS_DHM_C](#)

CONFIG_MBEDTLS_KEY_EXCHANGE_ECDHE_PSK

Enable ECDHE-PSK based ciphersuite modes

Found in: [Component config](#) > [mbedtls](#) > [TLS Key Exchange Methods](#) > [CONFIG_MBEDTLS_PSK_MODES](#)

Enable to support Elliptic-Curve-Diffie-Hellman PSK (pre-shared-key) TLS authentication modes.

Default value:

- Yes (enabled) if [CONFIG_MBEDTLS_PSK_MODES](#) && [CONFIG_MBEDTLS_ECDH_C](#)

CONFIG_MBEDTLS_KEY_EXCHANGE_RSA_PSK

Enable RSA-PSK based ciphersuite modes

Found in: [Component config](#) > [mbedtls](#) > [TLS Key Exchange Methods](#) > [CONFIG_MBEDTLS_PSK_MODES](#)

Enable to support RSA PSK (pre-shared-key) TLS authentication modes.

Default value:

- Yes (enabled) if [CONFIG_MBEDTLS_PSK_MODES](#)

CONFIG_MBEDTLS_KEY_EXCHANGE_RSA

Enable RSA-only based ciphersuite modes

Found in: [Component config](#) > [mbedtls](#) > [TLS Key Exchange Methods](#)

Enable to support ciphersuites with prefix TLS-RSA-WITH-

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_KEY_EXCHANGE_DHE_RSA

Enable DHE-RSA based ciphersuite modes

Found in: [Component config](#) > [mbedtls](#) > [TLS Key Exchange Methods](#)

Enable to support ciphersuites with prefix TLS-DHE-RSA-WITH-

Default value:

- Yes (enabled) if [CONFIG_MBEDTLS_DHM_C](#)

CONFIG_MBEDTLS_KEY_EXCHANGE_ELLIPTIC_CURVE

Support Elliptic Curve based ciphersuites

Found in: [Component config](#) > [mbedtls](#) > [TLS Key Exchange Methods](#)

Enable to show Elliptic Curve based ciphersuite mode options.

Disabling all Elliptic Curve ciphersuites saves code size and can give slightly faster TLS handshakes, provided the server supports RSA-only ciphersuite modes.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_KEY_EXCHANGE_ECDHE_RSA

Enable ECDHE-RSA based ciphersuite modes

Found in: [Component config](#) > [mbedtls](#) > [TLS Key Exchange Methods](#) > [CONFIG_MBEDTLS_KEY_EXCHANGE_ELLIPTIC_CURVE](#)

Enable to support ciphersuites with prefix TLS-ECDHE-RSA-WITH-

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_KEY_EXCHANGE_ECDHE_ECDSA

Enable ECDHE-ECDSA based ciphersuite modes

Found in: [Component config](#) > [mbedtls](#) > [TLS Key Exchange Methods](#) > [CONFIG_MBEDTLS_KEY_EXCHANGE_ELLIPTIC_CURVE](#)

Enable to support ciphersuites with prefix TLS-ECDHE-RSA-WITH-

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_KEY_EXCHANGE_ECDH_ECDSA

Enable ECDH-ECDSA based ciphersuite modes

Found in: [Component config > mbedTLS > TLS Key Exchange Methods > CONFIG_MBEDTLS_KEY_EXCHANGE_ELLIPTIC_CURVE](#)

Enable to support ciphersuites with prefix TLS-ECDHE-RSA-WITH-

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_KEY_EXCHANGE_ECDH_RSA

Enable ECDH-RSA based ciphersuite modes

Found in: [Component config > mbedTLS > TLS Key Exchange Methods > CONFIG_MBEDTLS_KEY_EXCHANGE_ELLIPTIC_CURVE](#)

Enable to support ciphersuites with prefix TLS-ECDHE-RSA-WITH-

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_KEY_EXCHANGE_ECJPAKE

Enable ECJPAKE based ciphersuite modes

Found in: [Component config > mbedTLS > TLS Key Exchange Methods](#)

Enable to support ciphersuites with prefix TLS-ECJPAKE-WITH-

Default value:

- No (disabled) if [CONFIG_MBEDTLS_ECJPAKE_C](#) && [CONFIG_MBEDTLS_ECP_DP_SECP256R1_ENABLED](#)

CONFIG_MBEDTLS_SSL_RENEGOTIATION

Support TLS renegotiation

Found in: [Component config > mbedTLS](#)

The two main uses of renegotiation are (1) refresh keys on long-lived connections and (2) client authentication after the initial handshake. If you don't need renegotiation, disabling it will save code size and reduce the possibility of abuse/vulnerability.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_SSL_PROTO_TLS1_2

Support TLS 1.2 protocol

Found in: [Component config > mbedTLS](#)

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_SSL_PROTO_GMTSSL1_1

Support GM/T SSL 1.1 protocol

Found in: [Component config > mbedTLS](#)

Provisions for GM/T SSL 1.1 support

Default value:

- No (disabled)

CONFIG_MBEDTLS_SSL_PROTO_DTLS

Support DTLS protocol (all versions)

Found in: [Component config](#) > [mbedtls](#)

Requires TLS 1.2 to be enabled for DTLS 1.2

Default value:

- No (disabled)

CONFIG_MBEDTLS_SSL_ALPN

Support ALPN (Application Layer Protocol Negotiation)

Found in: [Component config](#) > [mbedtls](#)

Disabling this option will save some code size if it is not needed.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_CLIENT_SSL_SESSION_TICKETS

TLS: Client Support for RFC 5077 SSL session tickets

Found in: [Component config](#) > [mbedtls](#)

Client support for RFC 5077 session tickets. See mbedtls documentation for more details. Disabling this option will save some code size.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_SERVER_SSL_SESSION_TICKETS

TLS: Server Support for RFC 5077 SSL session tickets

Found in: [Component config](#) > [mbedtls](#)

Server support for RFC 5077 session tickets. See mbedtls documentation for more details. Disabling this option will save some code size.

Default value:

- Yes (enabled)

Symmetric Ciphers Contains:

- [CONFIG_MBEDTLS_AES_C](#)
- [CONFIG_MBEDTLS_BLOWFISH_C](#)
- [CONFIG_MBEDTLS_CAMELLIA_C](#)
- [CONFIG_MBEDTLS_CCM_C](#)
- [CONFIG_MBEDTLS_DES_C](#)
- [CONFIG_MBEDTLS_GCM_C](#)
- [CONFIG_MBEDTLS_NIST_KW_C](#)
- [CONFIG_MBEDTLS_XTEA_C](#)

CONFIG_MBEDTLS_AES_C

AES block cipher

Found in: [Component config](#) > [mbedtls](#) > [Symmetric Ciphers](#)

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_CAMELLIA_C

Camellia block cipher

Found in: [Component config](#) > [mbedtls](#) > [Symmetric Ciphers](#)

Default value:

- No (disabled)

CONFIG_MBEDTLS_DES_C

DES block cipher (legacy, insecure)

Found in: [Component config](#) > [mbedtls](#) > [Symmetric Ciphers](#)

Enables the DES block cipher to support 3DES-based TLS ciphersuites.

3DES is vulnerable to the Sweet32 attack and should only be enabled if absolutely necessary.

Default value:

- No (disabled)

CONFIG_MBEDTLS_BLOWFISH_C

Blowfish block cipher (read help)

Found in: [Component config](#) > [mbedtls](#) > [Symmetric Ciphers](#)

Enables the Blowfish block cipher (not used for TLS sessions.)

The Blowfish cipher is not used for mbedtls TLS sessions but can be used for other purposes. Read up on the limitations of Blowfish (including Sweet32) before enabling.

Default value:

- No (disabled)

CONFIG_MBEDTLS_XTEA_C

XTEA block cipher

Found in: [Component config](#) > [mbedtls](#) > [Symmetric Ciphers](#)

Enables the XTEA block cipher.

Default value:

- No (disabled)

CONFIG_MBEDTLS_CCM_C

CCM (Counter with CBC-MAC) block cipher modes

Found in: [Component config](#) > [mbedtls](#) > [Symmetric Ciphers](#)

Enable Counter with CBC-MAC (CCM) modes for AES and/or Camellia ciphers.

Disabling this option saves some code size.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_GCM_C

GCM (Galois/Counter) block cipher modes

Found in: [Component config](#) > [mbedtls](#) > [Symmetric Ciphers](#)

Enable Galois/Counter Mode for AES and/or Camellia ciphers.

This option is generally faster than CCM.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_NIST_KW_C

NIST key wrapping (KW) and KW padding (KWP)

Found in: [Component config](#) > [mbedtls](#) > [Symmetric Ciphers](#)

Enable NIST key wrapping and key wrapping padding.

Default value:

- No (disabled)

CONFIG_MBEDTLS_RIPEMD160_C

Enable RIPEMD-160 hash algorithm

Found in: [Component config](#) > [mbedtls](#)

Enable the RIPEMD-160 hash algorithm.

Default value:

- No (disabled)

Certificates Contains:

- [CONFIG_MBEDTLS_PEM_PARSE_C](#)
- [CONFIG_MBEDTLS_PEM_WRITE_C](#)
- [CONFIG_MBEDTLS_X509_CRL_PARSE_C](#)
- [CONFIG_MBEDTLS_X509_CSR_PARSE_C](#)

CONFIG_MBEDTLS_PEM_PARSE_C

Read & Parse PEM formatted certificates

Found in: [Component config](#) > [mbedtls](#) > [Certificates](#)

Enable decoding/parsing of PEM formatted certificates.

If your certificates are all in the simpler DER format, disabling this option will save some code size.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_PEM_WRITE_C

Write PEM formatted certificates

Found in: [Component config](#) > [mbedtls](#) > [Certificates](#)

Enable writing of PEM formatted certificates.

If writing certificate data only in DER format, disabling this option will save some code size.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_X509_CRL_PARSE_C

X.509 CRL parsing

Found in: [Component config](#) > [mbedtls](#) > [Certificates](#)

Support for parsing X.509 Certificate Revocation Lists.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_X509_CSR_PARSE_C

X.509 CSR parsing

Found in: [Component config](#) > [mbedtls](#) > [Certificates](#)

Support for parsing X.509 Certificate Signing Requests

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_ECP_C

Elliptic Curve Ciphers

Found in: [Component config](#) > [mbedtls](#)

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_DHM_C

Diffie-Hellman-Merkle key exchange (DHM)

Found in: [Component config](#) > [mbedtls](#)

Enable DHM. Needed to use DHE-xxx TLS ciphersuites.

Note that the security of Diffie-Hellman key exchanges depends on a suitable prime being used for the exchange. Please see detailed warning text about this in file *mbedtls/dhm.h* file.

Default value:

- No (disabled)

CONFIG_MBEDTLS_ECDH_C

Elliptic Curve Diffie-Hellman (ECDH)

Found in: [Component config](#) > [mbedtls](#)

Enable ECDH. Needed to use ECDHE-xxx TLS ciphersuites.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_ECDSA_C

Elliptic Curve DSA

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_ECDH_C](#)

Enable ECDSA. Needed to use ECDSA-xxx TLS ciphersuites.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_ECJPAKE_C

Elliptic curve J-PAKE

Found in: [Component config](#) > [mbedtls](#)

Enable ECJPAKE. Needed to use ECJPAKE-xxx TLS ciphersuites.

Default value:

- No (disabled)

CONFIG_MBEDTLS_ECP_DP_SECP192R1_ENABLED

Enable SECP192R1 curve

Found in: [Component config](#) > [mbedtls](#)

Enable support for SECP192R1 Elliptic Curve.

CONFIG_MBEDTLS_ECP_DP_SECP224R1_ENABLED

Enable SECP224R1 curve

Found in: [Component config](#) > [mbedtls](#)

Enable support for SECP224R1 Elliptic Curve.

CONFIG_MBEDTLS_ECP_DP_SECP256R1_ENABLED

Enable SECP256R1 curve

Found in: [Component config](#) > [mbedtls](#)

Enable support for SECP256R1 Elliptic Curve.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_ECP_DP_SECP384R1_ENABLED

Enable SECP384R1 curve

Found in: [Component config](#) > [mbedtls](#)

Enable support for SECP384R1 Elliptic Curve.

CONFIG_MBEDTLS_ECP_DP_SECP521R1_ENABLED

Enable SECP521R1 curve

Found in: [Component config](#) > [mbedtls](#)

Enable support for SECP521R1 Elliptic Curve.

CONFIG_MBEDTLS_ECP_DP_SECP192K1_ENABLED

Enable SECP192K1 curve

Found in: [Component config](#) > [mbedtls](#)

Enable support for SECP192K1 Elliptic Curve.

CONFIG_MBEDTLS_ECP_DP_SECP224K1_ENABLED

Enable SECP224K1 curve

Found in: [Component config](#) > [mbedtls](#)

Enable support for SECP224K1 Elliptic Curve.

CONFIG_MBEDTLS_ECP_DP_SECP256K1_ENABLED

Enable SECP256K1 curve

Found in: [Component config](#) > [mbedtls](#)

Enable support for SECP256K1 Elliptic Curve.

CONFIG_MBEDTLS_ECP_DP_BP256R1_ENABLED

Enable BP256R1 curve

Found in: [Component config](#) > [mbedtls](#)

support for DP Elliptic Curve.

CONFIG_MBEDTLS_ECP_DP_BP384R1_ENABLED

Enable BP384R1 curve

Found in: [Component config](#) > [mbedtls](#)

support for DP Elliptic Curve.

CONFIG_MBEDTLS_ECP_DP_BP512R1_ENABLED

Enable BP512R1 curve

Found in: [Component config](#) > [mbedtls](#)

support for DP Elliptic Curve.

CONFIG_MBEDTLS_ECP_DP_CURVE25519_ENABLED

Enable CURVE25519 curve

Found in: [Component config](#) > [mbedtls](#)

Enable support for CURVE25519 Elliptic Curve.

CONFIG_MBEDTLS_ECP_NIST_OPTIM

NIST 'modulo p' optimisations

Found in: [Component config](#) > [mbedtls](#)

NIST 'modulo p' optimisations increase Elliptic Curve operation performance.

Disabling this option saves some code size.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_ECP_FIXED_POINT_OPTIM

Enable fixed-point multiplication optimisations

Found in: [Component config](#) > [mbedtls](#)

This configuration option enables optimizations to speedup (about 3 ~ 4 times) the ECP fixed point multiplication using pre-computed tables in the flash memory. Enabling this configuration option increases the flash footprint (about 29KB if all Elliptic Curve selected) in the application binary.

end of Elliptic Curve options

Default value:

- No (disabled)

CONFIG_MBEDTLS_POLY1305_C

Poly1305 MAC algorithm

Found in: [Component config](#) > [mbedtls](#)

Enable support for Poly1305 MAC algorithm.

Default value:

- No (disabled)

CONFIG_MBEDTLS_CHACHA20_C

Chacha20 stream cipher

Found in: [Component config](#) > [mbedtls](#)

Enable support for Chacha20 stream cipher.

Default value:

- No (disabled)

CONFIG_MBEDTLS_CHACHAPOLY_C

ChaCha20-Poly1305 AEAD algorithm

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_CHACHA20_C](#)

Enable support for ChaCha20-Poly1305 AEAD algorithm.

Default value:

- No (disabled) if [CONFIG_MBEDTLS_CHACHA20_C](#) && [CONFIG_MBEDTLS_POLY1305_C](#)

CONFIG_MBEDTLS_HKDF_C

HKDF algorithm (RFC 5869)

Found in: [Component config](#) > [mbedtls](#)

Enable support for the Hashed Message Authentication Code (HMAC)-based key derivation function (HKDF).

Default value:

- No (disabled)

CONFIG_MBEDTLS_THREADING_C

Enable the threading abstraction layer

Found in: [Component config](#) > [mbedtls](#)

If you do intend to use contexts between threads, you will need to enable this layer to prevent race conditions.

Default value:

- No (disabled)

CONFIG_MBEDTLS_THREADING_ALT

Enable threading alternate implementation

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_THREADING_C](#)

Enable threading alt to allow your own alternate threading implementation.

Default value:

- Yes (enabled) if [CONFIG_MBEDTLS_THREADING_C](#)

CONFIG_MBEDTLS_THREADING_PTHREAD

Enable threading pthread implementation

Found in: [Component config](#) > [mbedtls](#) > [CONFIG_MBEDTLS_THREADING_C](#)

Enable the pthread wrapper layer for the threading layer.

Default value:

- No (disabled) if [CONFIG_MBEDTLS_THREADING_C](#)

CONFIG_MBEDTLS_ERROR_STRINGS

Enable error code to error string conversion

Found in: [Component config](#) > [mbedtls](#)

Enables mbedtls_strerror() for converting error codes to error strings. Disabling this config can save some code/rodata size as the error string conversion implementation is replaced with an empty stub.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_USE_CRYPTOROM_IMPL

Use ROM implementation of the crypto algorithm

Found in: [Component config](#) > [mbedtls](#)

Enable this flag to use mbedtls crypto algorithm from ROM instead of ESP-IDF.

This configuration option saves flash footprint in the application binary. Note that the version of mbedtls crypto algorithm library in ROM(ECO1~ECO3) is v2.16.12, and the version of mbedtls crypto algorithm library in ROM(ECO4) is v3.6.0. We have done the security analysis of the mbedtls revision in ROM (ECO1~ECO4) and ensured that affected symbols have been patched (removed). If in the future mbedtls revisions there are security issues that also affects the version in ROM (ECO1~ECO4) then we shall patch the relevant symbols. This would increase the flash footprint and hence care must be taken to keep some reserved space for the application binary in flash layout.

Default value:

- No (disabled) if ESP_ROM_HAS_MBEDTLS_CRYPTOLIB

CONFIG_MBEDTLS_FS_IO

Enable functions that use the filesystem

Found in: *Component config > mbedTLS*

This option enables functions in mbedTLS that use the filesystem. It uses the default filesystem support for the target, which is added through vfs component for ESP32 based targets or by the host system when the target is Linux.

Default value:

- Yes (enabled)

CONFIG_MBEDTLS_ALLOW_WEAK_CERTIFICATE_VERIFICATION

Allow weak certificate verification

Found in: *Component config > mbedTLS*

This options allows weak certificate verification by skipping the hostname verification. It is not recommended to use this option.

Default value:

- No (disabled)

ESP-MQTT Configurations Contains:

- *CONFIG_MQTT_CUSTOM_OUTBOX*
- *CONFIG_MQTT_TRANSPORT_SSL*
- *CONFIG_MQTT_TRANSPORT_WEBSOCKET*
- *CONFIG_MQTT_PROTOCOL_311*
- *CONFIG_MQTT_PROTOCOL_5*
- *CONFIG_MQTT_TASK_CORE_SELECTION_ENABLED*
- *CONFIG_MQTT_USE_CUSTOM_CONFIG*
- *CONFIG_MQTT_OUTBOX_EXPIRED_TIMEOUT_MS*
- *CONFIG_MQTT_REPORT_DELETED_MESSAGES*
- *CONFIG_MQTT_SKIP_PUBLISH_IF_DISCONNECTED*
- *CONFIG_MQTT_OUTBOX_DATA_ON_EXTERNAL_MEMORY*
- *CONFIG_MQTT_MSG_ID_INCREMENTAL*

CONFIG_MQTT_PROTOCOL_311

Enable MQTT protocol 3.1.1

Found in: *Component config > ESP-MQTT Configurations*

If not, this library will use MQTT protocol 3.1

Default value:

- Yes (enabled)

CONFIG_MQTT_PROTOCOL_5

Enable MQTT protocol 5.0

Found in: *Component config > ESP-MQTT Configurations*

If not, this library will not support MQTT 5.0

Default value:

- No (disabled)

CONFIG_MQTT_TRANSPORT_SSL

Enable MQTT over SSL

Found in: [Component config](#) > [ESP-MQTT Configurations](#)

Enable MQTT transport over SSL with mbedtls

Default value:

- Yes (enabled)

CONFIG_MQTT_TRANSPORT_WEBSOCKET

Enable MQTT over Websocket

Found in: [Component config](#) > [ESP-MQTT Configurations](#)

Enable MQTT transport over Websocket.

Default value:

- Yes (enabled)

CONFIG_MQTT_TRANSPORT_WEBSOCKET_SECURE

Enable MQTT over Websocket Secure

Found in: [Component config](#) > [ESP-MQTT Configurations](#) > [CONFIG_MQTT_TRANSPORT_WEBSOCKET_SECURE](#)

Enable MQTT transport over Websocket Secure.

Default value:

- Yes (enabled)

CONFIG_MQTT_MSG_ID_INCREMENTAL

Use Incremental Message Id

Found in: [Component config](#) > [ESP-MQTT Configurations](#)

Set this to true for the message id (2.3.1 Packet Identifier) to be generated as an incremental number rather than a random value (used by default)

Default value:

- No (disabled)

CONFIG_MQTT_SKIP_PUBLISH_IF_DISCONNECTED

Skip publish if disconnected

Found in: [Component config](#) > [ESP-MQTT Configurations](#)

Set this to true to avoid publishing (enqueueing messages) if the client is disconnected. The MQTT client tries to publish all messages by default, even in the disconnected state (where the qos1 and qos2 packets are stored in the internal outbox to be published later) The MQTT_SKIP_PUBLISH_IF_DISCONNECTED option allows applications to override this behaviour and not enqueue publish packets in the disconnected state.

Default value:

- No (disabled)

CONFIG_MQTT_REPORT_DELETED_MESSAGES

Report deleted messages

Found in: *Component config* > *ESP-MQTT Configurations*

Set this to true to post events for all messages which were deleted from the outbox before being correctly sent and confirmed.

Default value:

- No (disabled)

CONFIG_MQTT_USE_CUSTOM_CONFIG

MQTT Using custom configurations

Found in: *Component config* > *ESP-MQTT Configurations*

Custom MQTT configurations.

Default value:

- No (disabled)

CONFIG_MQTT_TCP_DEFAULT_PORT

Default MQTT over TCP port

Found in: *Component config* > *ESP-MQTT Configurations* > *CONFIG_MQTT_USE_CUSTOM_CONFIG*

Default MQTT over TCP port

Default value:

- 1883 if *CONFIG_MQTT_USE_CUSTOM_CONFIG*

CONFIG_MQTT_SSL_DEFAULT_PORT

Default MQTT over SSL port

Found in: *Component config* > *ESP-MQTT Configurations* > *CONFIG_MQTT_USE_CUSTOM_CONFIG*

Default MQTT over SSL port

Default value:

- 8883 if *CONFIG_MQTT_USE_CUSTOM_CONFIG* && *CONFIG_MQTT_TRANSPORT_SSL*

CONFIG_MQTT_WS_DEFAULT_PORT

Default MQTT over Websocket port

Found in: *Component config* > *ESP-MQTT Configurations* > *CONFIG_MQTT_USE_CUSTOM_CONFIG*

Default MQTT over Websocket port

Default value:

- 80 if *CONFIG_MQTT_USE_CUSTOM_CONFIG* && *CONFIG_MQTT_TRANSPORT_WEBSOCKET*

CONFIG_MQTT_WSS_DEFAULT_PORT

Default MQTT over Websocket Secure port

Found in: *Component config* > *ESP-MQTT Configurations* > *CONFIG_MQTT_USE_CUSTOM_CONFIG*

Default MQTT over Websocket Secure port

Default value:

- 443 if `CONFIG_MQTT_USE_CUSTOM_CONFIG` && `CONFIG_MQTT_TRANSPORT_WEBSOCKET` && `CONFIG_MQTT_TRANSPORT_WEBSOCKET_SECURE`

CONFIG_MQTT_BUFFER_SIZE

Default MQTT Buffer Size

Found in: *Component config* > *ESP-MQTT Configurations* > `CONFIG_MQTT_USE_CUSTOM_CONFIG`

This buffer size using for both transmit and receive

Default value:

- 1024 if `CONFIG_MQTT_USE_CUSTOM_CONFIG`

CONFIG_MQTT_TASK_STACK_SIZE

MQTT task stack size

Found in: *Component config* > *ESP-MQTT Configurations* > `CONFIG_MQTT_USE_CUSTOM_CONFIG`

MQTT task stack size

Default value:

- 6144 if `CONFIG_MQTT_USE_CUSTOM_CONFIG`

CONFIG_MQTT_DISABLE_API_LOCKS

Disable API locks

Found in: *Component config* > *ESP-MQTT Configurations* > `CONFIG_MQTT_USE_CUSTOM_CONFIG`

Default config employs API locks to protect internal structures. It is possible to disable these locks if the user code doesn't access MQTT API from multiple concurrent tasks

Default value:

- No (disabled) if `CONFIG_MQTT_USE_CUSTOM_CONFIG`

CONFIG_MQTT_TASK_PRIORITY

MQTT task priority

Found in: *Component config* > *ESP-MQTT Configurations* > `CONFIG_MQTT_USE_CUSTOM_CONFIG`

MQTT task priority. Higher number denotes higher priority.

Default value:

- 5 if `CONFIG_MQTT_USE_CUSTOM_CONFIG`

CONFIG_MQTT_POLL_READ_TIMEOUT_MS

MQTT transport poll read timeout

Found in: *Component config* > *ESP-MQTT Configurations* > `CONFIG_MQTT_USE_CUSTOM_CONFIG`

Timeout when polling underlying transport for read.

Default value:

- 1000 if `CONFIG_MQTT_USE_CUSTOM_CONFIG`

CONFIG_MQTT_EVENT_QUEUE_SIZE

Number of queued events.

Found in: *Component config > ESP-MQTT Configurations > CONFIG_MQTT_USE_CUSTOM_CONFIG*

A value higher than 1 enables multiple queued events.

Default value:

- 1 if *CONFIG_MQTT_USE_CUSTOM_CONFIG*

CONFIG_MQTT_TASK_CORE_SELECTION_ENABLED

Enable MQTT task core selection

Found in: *Component config > ESP-MQTT Configurations*

This will enable core selection

CONFIG_MQTT_TASK_CORE_SELECTION

Core to use ?

Found in: *Component config > ESP-MQTT Configurations > CONFIG_MQTT_TASK_CORE_SELECTION_ENABLED*

Available options:

- Core 0 (CONFIG_MQTT_USE_CORE_0)
- Core 1 (CONFIG_MQTT_USE_CORE_1)

CONFIG_MQTT_OUTBOX_DATA_ON_EXTERNAL_MEMORY

Use external memory for outbox data

Found in: *Component config > ESP-MQTT Configurations*

Set to true to use external memory for outbox data.

Default value:

- No (disabled) if *CONFIG_MQTT_USE_CUSTOM_CONFIG*

CONFIG_MQTT_CUSTOM_OUTBOX

Enable custom outbox implementation

Found in: *Component config > ESP-MQTT Configurations*

Set to true if a specific implementation of message outbox is needed (e.g. persistent outbox in NVM or similar). Note: Implementation of the custom outbox must be added to the mqtt component. These CMake commands could be used to append the custom implementation to lib-mqtt sources: `idf_component_get_property(mqtt mqtt COMPONENT_LIB) set_property(TARGET ${mqtt} PROPERTY SOURCES ${PROJECT_DIR}/custom_outbox.c APPEND)`

Default value:

- No (disabled)

CONFIG_MQTT_OUTBOX_EXPIRED_TIMEOUT_MS

Outbox message expired timeout[ms]

Found in: *Component config > ESP-MQTT Configurations*

Messages which stays in the outbox longer than this value before being published will be discarded.

Default value:

- 30000 if [CONFIG_MQTT_USE_CUSTOM_CONFIG](#)

Newlib Contains:

- [CONFIG_NEWLIB_NANO_FORMAT](#)
- [CONFIG_NEWLIB_STDIN_LINE_ENDING](#)
- [CONFIG_NEWLIB_STDOUT_LINE_ENDING](#)
- [CONFIG_NEWLIB_TIME_SYSCALL](#)

CONFIG_NEWLIB_STDOUT_LINE_ENDING

Line ending for UART output

Found in: [Component config](#) > [Newlib](#)

This option allows configuring the desired line endings sent to UART when a newline ('n', LF) appears on stdout. Three options are possible:

CRLF: whenever LF is encountered, prepend it with CR

LF: no modification is applied, stdout is sent as is

CR: each occurrence of LF is replaced with CR

This option doesn't affect behavior of the UART driver (drivers/uart.h).

Available options:

- CRLF (CONFIG_NEWLIB_STDOUT_LINE_ENDING_CRLF)
- LF (CONFIG_NEWLIB_STDOUT_LINE_ENDING_LF)
- CR (CONFIG_NEWLIB_STDOUT_LINE_ENDING_CR)

CONFIG_NEWLIB_STDIN_LINE_ENDING

Line ending for UART input

Found in: [Component config](#) > [Newlib](#)

This option allows configuring which input sequence on UART produces a newline ('n', LF) on stdin. Three options are possible:

CRLF: CRLF is converted to LF

LF: no modification is applied, input is sent to stdin as is

CR: each occurrence of CR is replaced with LF

This option doesn't affect behavior of the UART driver (drivers/uart.h).

Available options:

- CRLF (CONFIG_NEWLIB_STDIN_LINE_ENDING_CRLF)
- LF (CONFIG_NEWLIB_STDIN_LINE_ENDING_LF)
- CR (CONFIG_NEWLIB_STDIN_LINE_ENDING_CR)

CONFIG_NEWLIB_NANO_FORMAT

Enable 'nano' formatting options for printf/scanf family

Found in: [Component config](#) > [Newlib](#)

In most chips the ROM contains parts of newlib C library, including printf/scanf family of functions. These functions have been compiled with so-called "nano" formatting option. This option doesn't support 64-bit integer formats and C99 features, such as positional arguments.

For more details about "nano" formatting option, please see newlib readme file, search for '--enable-newlib-nano-formatted-io': <https://sourceware.org/newlib/README>

If this option is enabled and the ROM contains functions from newlib-nano, the build system will use functions available in ROM, reducing the application binary size. Functions available in ROM run faster than functions which run from flash. Functions available in ROM can also run when flash instruction cache is disabled.

Some chips (e.g. ESP32-C6) has the full formatting versions of printf/scanf in ROM instead of the nano versions and in this building with newlib nano might actually increase the size of the binary. Which functions are present in ROM can be seen from ROM caps: ESP_ROM_HAS_NEWLIB_NANO_FORMAT and ESP_ROM_HAS_NEWLIB_NORMAL_FORMAT.

If you need 64-bit integer formatting support or C99 features, keep this option disabled.

CONFIG_NEWLIB_TIME_SYSCALL

Timers used for gettimeofday function

Found in: [Component config > Newlib](#)

This setting defines which hardware timers are used to implement 'gettimeofday' and 'time' functions in C library.

- **If both high-resolution (systimer for all targets except ESP32) and RTC timers are used**, timekeeping will continue in deep sleep. Time will be reported at 1 microsecond resolution. This is the default, and the recommended option.
- **If only high-resolution timer (systimer) is used, gettimeofday will** provide time at microsecond resolution. Time will not be preserved when going into deep sleep mode.
- **If only RTC timer is used, timekeeping will continue in** deep sleep, but time will be measured at 6.(6) microsecond resolution. Also the gettimeofday function itself may take longer to run.
- **If no timers are used, gettimeofday and time functions** return -1 and set errno to ENOSYS.
- **When RTC is used for timekeeping, two RTC_STORE registers are** used to keep time in deep sleep mode.

Available options:

- RTC and high-resolution timer (CONFIG_NEWLIB_TIME_SYSCALL_USE_RTC_HRT)
- RTC (CONFIG_NEWLIB_TIME_SYSCALL_USE_RTC)
- High-resolution timer (CONFIG_NEWLIB_TIME_SYSCALL_USE_HRT)
- None (CONFIG_NEWLIB_TIME_SYSCALL_USE_NONE)

NVS Contains:

- [CONFIG_NVS_LEGACY_DUP_KEYS_COMPATIBILITY](#)
- [CONFIG_NVS_ENCRYPTION](#)
- [CONFIG_NVS_COMPATIBLE_PRE_V4_3_ENCRYPTION_FLAG](#)
- [CONFIG_NVS_ASSERT_ERROR_CHECK](#)

CONFIG_NVS_ENCRYPTION

Enable NVS encryption

Found in: [Component config > NVS](#)

This option enables encryption for NVS. When enabled, XTS-AES is used to encrypt the complete NVS data, except the page headers. It requires XTS encryption keys to be stored in an encrypted partition (enabling flash encryption is mandatory here) or to be derived from an HMAC key burnt in eFuse.

Default value:

- Yes (enabled) if [CONFIG_SECURE_FLASH_ENC_ENABLED](#)

CONFIG_NVS_COMPATIBLE_PRE_V4_3_ENCRYPTION_FLAG

NVS partition encrypted flag compatible with ESP-IDF before v4.3

Found in: [Component config](#) > [NVS](#)

Enabling this will ignore "encrypted" flag for NVS partitions. NVS encryption scheme is different than hardware flash encryption and hence it is not recommended to have "encrypted" flag for NVS partitions. This was not being checked in pre v4.3 IDF. Hence, if you have any devices where this flag is kept enabled in partition table then enabling this config will allow to have same behavior as pre v4.3 IDF.

CONFIG_NVS_ASSERT_ERROR_CHECK

Use assertions for error checking

Found in: [Component config](#) > [NVS](#)

This option switches error checking type between assertions (y) or return codes (n).

Default value:

- No (disabled)

CONFIG_NVS_LEGACY_DUP_KEYS_COMPATIBILITY

Enable legacy nvs_set function behavior when same key is reused with different data types

Found in: [Component config](#) > [NVS](#)

Enabling this option will switch the nvs_set() family of functions to the legacy mode: when called repeatedly with the same key but different data type, the existing value in the NVS remains active and the new value is just stored, actually not accessible through corresponding nvs_get() call for the key given. Use this option only when your application relies on such NVS API behaviour.

Default value:

- No (disabled)

NVS Security Provider Contains:

- [CONFIG_NVS_SEC_HMAC_EFUSE_KEY_ID](#)
- [CONFIG_NVS_SEC_KEY_PROTECTION_SCHEME](#)

CONFIG_NVS_SEC_KEY_PROTECTION_SCHEME

NVS Encryption: Key Protection Scheme

Found in: [Component config](#) > [NVS Security Provider](#)

This choice defines the default NVS encryption keys protection scheme; which will be used for the default NVS partition. Users can use the corresponding scheme registration APIs to register other schemes for the default as well as other NVS partitions.

Available options:

- Using Flash Encryption (`CONFIG_NVS_SEC_KEY_PROTECT_USING_FLASH_ENC`)
Protect the NVS Encryption Keys using Flash Encryption Requires a separate 'nvs_keys' partition (which will be encrypted by flash encryption) for storing the NVS encryption keys
- Using HMAC peripheral (`CONFIG_NVS_SEC_KEY_PROTECT_USING_HMAC`)
Derive and protect the NVS Encryption Keys using the HMAC peripheral Requires the specified eFuse block (`NVS_SEC_HMAC_EFUSE_KEY_ID` or the v2 API argument) to be empty or pre-written with a key with the purpose `ESP_EFUSE_KEY_PURPOSE_HMAC_UP`

CONFIG_NVS_SEC_HMAC_EFUSE_KEY_ID

eFuse key ID storing the HMAC key

Found in: *Component config > NVS Security Provider*

eFuse block key ID storing the HMAC key for deriving the NVS encryption keys

Note: The eFuse block key ID required by the HMAC scheme (`CONFIG_NVS_SEC_KEY_PROTECT_USING_HMAC`) is set using this config when the default NVS partition is initialized with `nvs_flash_init()`. The eFuse block key ID can also be set at runtime by passing the appropriate value to the NVS security scheme registration APIs.

Range:

- from 0 to 6 if *CONFIG_NVS_SEC_KEY_PROTECT_USING_HMAC*

Default value:

- 6 if *CONFIG_NVS_SEC_KEY_PROTECT_USING_HMAC*

OpenThread Contains:

- *CONFIG_OPENTHREAD_DEBUG*
- *CONFIG_OPENTHREAD_ENABLED*
- *OpenThread Spinel*

CONFIG_OPENTHREAD_ENABLED

OpenThread

Found in: *Component config > OpenThread*

Select this option to enable OpenThread and show the submenu with OpenThread configuration choices.

Default value:

- No (disabled)

Thread Task Parameters Contains:

- *CONFIG_OPENTHREAD_TASK_NAME*
- *CONFIG_OPENTHREAD_TASK_PRIORITY*
- *CONFIG_OPENTHREAD_TASK_SIZE*

CONFIG_OPENTHREAD_TASK_NAME

OpenThread task name

Found in: *Component config > OpenThread > CONFIG_OPENTHREAD_ENABLED > Thread Task Parameters*

The OpenThread task name.

Default value:

- "ot_main" if *CONFIG_OPENTHREAD_ENABLED*

CONFIG_OPENTHREAD_TASK_SIZE

Size of OpenThread task

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Task Parameters](#)

The size in bytes of OpenThread task.

Default value:

- 8192 if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_TASK_PRIORITY

Priority of OpenThread task

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Task Parameters](#)

The priority of OpenThread task.

Default value:

- 5 if [CONFIG_OPENTHREAD_ENABLED](#)

Thread Version Message Contains:

- [CONFIG_OPENTHREAD_PACKAGE_NAME](#)
- [CONFIG_OPENTHREAD_PLATFORM_INFO](#)

CONFIG_OPENTHREAD_PACKAGE_NAME

OpenThread package name

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Version Message](#)

The OpenThread package name.

Default value:

- "openthread-esp32" if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_PLATFORM_INFO

platform information

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Version Message](#)

The OpenThread platform information.

Default value:

- "esp32p4" if [CONFIG_OPENTHREAD_ENABLED](#)

Thread Console Contains:

- [CONFIG_OPENTHREAD_CLI](#)
- [CONFIG_OPENTHREAD_CONSOLE_ENABLE](#)
- [CONFIG_OPENTHREAD_CONSOLE_COMMAND_PREFIX](#)

CONFIG_OPENTHREAD_CONSOLE_ENABLE

Enable OpenThread console

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Console](#)

Enable the OpenThread-specific console provided by the SDK. This only controls whether the SDK sets up a dedicated console for OpenThread. Even if disabled, the default ESP-IDF console (if initialized elsewhere) can still be used independently.

Default value:

- Yes (enabled) if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_CONSOLE_TYPE

OpenThread console type

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Console](#) > [CONFIG_OPENTHREAD_CONSOLE_ENABLE](#)

Select OpenThread console type

Available options:

- OpenThread console type UART ([CONFIG_OPENTHREAD_CONSOLE_TYPE_UART](#))
- OpenThread console type USB Serial/JTAG Controller ([CONFIG_OPENTHREAD_CONSOLE_TYPE_USB_SERIAL_JTAG](#))

CONFIG_OPENTHREAD_CLI

Enable OpenThread Command-Line Interface

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Console](#)

Select this option to enable Command-Line Interface in OpenThread.

Default value:

- Yes (enabled) if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_CONSOLE_COMMAND_PREFIX

The prefix of the openthread CLI command registered on the esp console

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Console](#)

A prefix string used before a Thread CLI command, allowing the ESP console to identify it and delegate the remaining command to the OpenThread callback for processing.

Default value:

- "ot" if [CONFIG_OPENTHREAD_ENABLED](#)

Thread Core Features

 Contains:

- [CSL Configurations](#)
- [CONFIG_OPENTHREAD_BORDER_AGENT_ENABLE](#)
- [CONFIG_OPENTHREAD_BORDER_ROUTER](#)
- [CONFIG_OPENTHREAD_COMMISSIONER](#)
- [CONFIG_OPENTHREAD_CSL_ENABLE](#)
- [CONFIG_OPENTHREAD_DIAG](#)
- [CONFIG_OPENTHREAD_DNS_CLIENT](#)
- [CONFIG_OPENTHREAD_DNS64_CLIENT](#)
- [CONFIG_OPENTHREAD_JOINER](#)
- [CONFIG_OPENTHREAD_LINK_METRICS](#)

- [CONFIG_OPENTHREAD_MACFILTER_ENABLE](#)
- [CONFIG_OPENTHREAD_RX_ON_WHEN_IDLE](#)
- [CONFIG_OPENTHREAD_PARENT_SEARCH_MTD](#)
- [CONFIG_OPENTHREAD_RADIO_STATS_ENABLE](#)
- [CONFIG_OPENTHREAD_SRP_CLIENT](#)
- [CONFIG_OPENTHREAD_TIME_SYNC](#)
- [CONFIG_OPENTHREAD_TIMING_OPTIMIZATION](#)
- [OpenThread Stack Parameters](#)
- [Parent Search Configurations](#)
- [Thread 15.4 Radio Link](#)
- [CONFIG_OPENTHREAD_DEVICE_TYPE](#)
- [Thread Memory Allocation](#)
- [Thread Operational Dataset](#)
- [Thread Radio Co-Processor Feature](#)
- [Thread Trel Radio Link](#)

Thread Operational Dataset Contains:

- [CONFIG_OPENTHREAD_NETWORK_EXTPANID](#)
- [CONFIG_OPENTHREAD_MESH_LOCAL_PREFIX](#)
- [CONFIG_OPENTHREAD_NETWORK_CHANNEL](#)
- [CONFIG_OPENTHREAD_NETWORK_MASTERKEY](#)
- [CONFIG_OPENTHREAD_NETWORK_NAME](#)
- [CONFIG_OPENTHREAD_NETWORK_PANID](#)
- [CONFIG_OPENTHREAD_NETWORK_PSKC](#)

CONFIG_OPENTHREAD_NETWORK_NAME

OpenThread network name

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Thread Operational Dataset](#)

Default value:

- "OpenThread-ESP" if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_MESH_LOCAL_PREFIX

OpenThread mesh local prefix, format <address>/<plen>

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Thread Operational Dataset](#)

A string in the format "<address>/<plen>", where <address> is an IPv6 address and <plen> is a prefix length. For example "fd00:db8:a0:0::/64"

Default value:

- "fd00:db8:a0:0::/64" if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_NETWORK_CHANNEL

OpenThread network channel

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Thread Operational Dataset](#)

Range:

- from 11 to 26 if [CONFIG_OPENTHREAD_ENABLED](#)

Default value:

- 15 if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_NETWORK_PANID

OpenThread network pan id

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Thread Operational Dataset](#)

Range:

- from 0 to 0xFFFFE if [CONFIG_OPENTHREAD_ENABLED](#)

Default value:

- "0x1234" if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_NETWORK_EXTPANID

OpenThread extended pan id

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Thread Operational Dataset](#)

The OpenThread network extended pan id in hex string format

Default value:

- dead00beef00cafe if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_NETWORK_MASTERKEY

OpenThread network key

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Thread Operational Dataset](#)

The OpenThread network network key in hex string format

Default value:

- 00112233445566778899aabbccddeeff if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_NETWORK_PSKC

OpenThread pre-shared commissioner key

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Thread Operational Dataset](#)

The OpenThread pre-shared commissioner key in hex string format

Default value:

- 104810e2315100afd6bc9215a6bfac53 if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_DEVICE_TYPE

Thread device type

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

OpenThread can be configured to different device types (FTD, MTD, Radio)

Available options:

- Full Thread Device (CONFIG_OPENTHREAD_FTD)
Select this to enable Full Thread Device which can act as router and leader in a Thread network.

- **Minimal Thread Device (CONFIG_OPENTHREAD_MTD)**
Select this to enable Minimal Thread Device which can only act as end device in a Thread network. This will reduce the code size of the OpenThread stack.
- **Radio Only Device (CONFIG_OPENTHREAD_RADIO)**
Select this to enable Radio Only Device which can only forward 15.4 packets to the host. The OpenThread stack will be run on the host and OpenThread will have minimal footprint on the radio only device.

Thread Trel Radio Link Contains:

- [CONFIG_OPENTHREAD_RADIO_TREL](#)

CONFIG_OPENTHREAD_RADIO_TREL

Enable Thread Radio Encapsulation Link (TREL)

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Thread Trel Radio Link](#)

Select this option to enable Thread Radio Encapsulation Link.

Default value:

- No (disabled) if (EXAMPLE_CONNECT_WIFI || EXAMPLE_CONNECT_ETHERNET) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_TREL_PORT

The port of openthread trel service

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Thread Trel Radio Link](#) > [CONFIG_OPENTHREAD_RADIO_TREL](#)

Configure the port number of TREL service.

Default value:

- 12390 if [CONFIG_OPENTHREAD_RADIO_TREL](#) && (EXAMPLE_CONNECT_WIFI || EXAMPLE_CONNECT_ETHERNET) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_TREL_BUFFER_SIZE

The receive buffer size of openthread trel

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Thread Trel Radio Link](#) > [CONFIG_OPENTHREAD_RADIO_TREL](#)

Configure the receive buffer size of TREL service.

Range:

- from 10 to 255 if [CONFIG_OPENTHREAD_RADIO_TREL](#) && (EXAMPLE_CONNECT_WIFI || EXAMPLE_CONNECT_ETHERNET) && [CONFIG_OPENTHREAD_ENABLED](#)

Default value:

- 50 if [CONFIG_OPENTHREAD_RADIO_TREL](#) && (EXAMPLE_CONNECT_WIFI || EXAMPLE_CONNECT_ETHERNET) && [CONFIG_OPENTHREAD_ENABLED](#)

Thread 15.4 Radio Link Contains:

- [CONFIG_OPENTHREAD_RADIO_TYPE](#)

CONFIG_OPENTHREAD_RADIO_TYPE

Config the Thread radio type with 15.4 link

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Thread 15.4 Radio Link](#)

Configure how OpenThread connects to the 15.4 radio

Available options:

- Native 15.4 radio (CONFIG_OPENTHREAD_RADIO_NATIVE)
Select this to use the native 15.4 radio.
- Connect via UART (CONFIG_OPENTHREAD_RADIO_SPINEL_UART)
Select this to connect to a Radio Co-Processor via UART.
- Connect via SPI (CONFIG_OPENTHREAD_RADIO_SPINEL_SPI)
Select this to connect to a Radio Co-Processor via SPI.
- Disable the Thread radio based on 15.4 link (CONFIG_OPENTHREAD_RADIO_154_NONE)
Select this to disable the Thread radio based on 15.4 link.

Thread Radio Co-Processor Feature Contains:

- [CONFIG_OPENTHREAD_RCP_SPINEL_CONSOLE](#)
- [CONFIG_OPENTHREAD_NCP_VENDOR_HOOK](#)
- [CONFIG_OPENTHREAD_RCP_TRANSPORT](#)

CONFIG_OPENTHREAD_RCP_TRANSPORT

The RCP transport type

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Thread Radio Co-Processor Feature](#)

Available options:

- UART RCP (CONFIG_OPENTHREAD_RCP_UART)
Select this to enable UART connection to host.
- SPI RCP (CONFIG_OPENTHREAD_RCP_SPI)
Select this to enable SPI connection to host.

CONFIG_OPENTHREAD_NCP_VENDOR_HOOK

Enable vendor command for RCP

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Thread Radio Co-Processor Feature](#)

Select this to enable OpenThread NCP vendor commands.

Default value:

- Yes (enabled) if [CONFIG_OPENTHREAD_RADIO](#) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_RCP_SPINEL_CONSOLE

Enable RCP console via Spinel

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Thread Radio Co-Processor Feature](#)

Select this to enable sending console commands to OpenThread RCP via Spinel.

Default value:

- No (disabled) if `CONFIG_OPENTHREAD_RADIO` && `CONFIG_OPENTHREAD_ENABLED`

CONFIG_OPENTHREAD_BORDER_ROUTER

Enable Border Router

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable border router features in OpenThread.

Default value:

- No (disabled) if `CONFIG_OPENTHREAD_FTD` && `CONFIG_OPENTHREAD_ENABLED`

CONFIG_OPENTHREAD_COMMISSIONER

Enable Commissioner

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable commissioner in OpenThread. This will enable the device to act as a commissioner in the Thread network. A commissioner checks the pre-shared key from a joining device with the Thread commissioning protocol and shares the network parameter with the joining device upon success.

Default value:

- No (disabled) if `CONFIG_OPENTHREAD_ENABLED`

Commissioner Configurations Contains:

- `CONFIG_OPENTHREAD_COMM_MAX_JOINER_ENTRIES`

CONFIG_OPENTHREAD_COMM_MAX_JOINER_ENTRIES

The size of max commissioning joiner entries

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [CONFIG_OPENTHREAD_COMMISSIONER](#) > [Commissioner Configurations](#)

Default value:

- 2 if `CONFIG_OPENTHREAD_COMMISSIONER` && `CONFIG_OPENTHREAD_ENABLED`

CONFIG_OPENTHREAD_JOINER

Enable Joiner

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable Joiner in OpenThread. This allows a device to join the Thread network with a pre-shared key using the Thread commissioning protocol.

Default value:

- No (disabled) if `CONFIG_OPENTHREAD_ENABLED`

CONFIG_OPENTHREAD_SRP_CLIENT

Enable SRP Client

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable SRP Client in OpenThread. This allows a device to register SRP services to SRP Server.

Default value:

- Yes (enabled) if [CONFIG_OPENTHREAD_ENABLED](#)

SRP Client Configurations

 Contains:

- [CONFIG_OPENTHREAD_SRP_CLIENT_MAX_SERVICES](#)

CONFIG_OPENTHREAD_SRP_CLIENT_MAX_SERVICES

Specifies number of service entries in the SRP client service pool

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [CONFIG_OPENTHREAD_SRP_CLIENT](#) > [SRP Client Configurations](#)

Set the max buffer size of service entries in the SRP client service pool.

Default value:

- 5 if [CONFIG_OPENTHREAD_SRP_CLIENT](#) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_DNS_CLIENT

Enable DNS Client

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable DNS Client in OpenThread.

Default value:

- Yes (enabled) if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_DNS64_CLIENT

Enable DNS64 Client

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to acquire NAT64 address from dns servers.

Default value:

- No (disabled) if [CONFIG_LWIP_IPV4](#) && [CONFIG_OPENTHREAD_ENABLED](#)

DNS64 Client Configurations

 Contains:

- [CONFIG_OPENTHREAD_DNS_SERVER_ADDR](#)

CONFIG_OPENTHREAD_DNS_SERVER_ADDR

DNS server address (IPv4)

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [CONFIG_OPENTHREAD_DNS64_CLIENT](#) > [DNS64 Client Configurations](#)

Set the DNS server IPv4 address.

Default value:

- "8.8.8.8" if [CONFIG_OPENTHREAD_DNS64_CLIENT](#) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_TIMING_OPTIMIZATION

Enable timing optimization

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable timing optimization for link metrics / CSL features.

Default value:

- No (disabled) if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_LINK_METRICS

Enable link metrics feature

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable link metrics feature

Default value:

- No (disabled) if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_BORDER_AGENT_ENABLE

Enable border agent feature

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable border agent feature

Default value:

- Yes (enabled) if [CONFIG_OPENTHREAD_BORDER_ROUTER](#) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_MACFILTER_ENABLE

Enable mac filter feature

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable mac filter feature

Default value:

- No (disabled) if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_CSL_ENABLE

Enable CSL feature

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable CSL feature

Default value:

- No (disabled) if [CONFIG_OPENTHREAD_ENABLED](#)

CSL Configurations Contains:

- [CONFIG_OPENTHREAD_CSL_DEBUG_ENABLE](#)
- [CONFIG_OPENTHREAD_CSL_UNCERTAIN](#)
- [CONFIG_OPENTHREAD_CSL_ACCURACY](#)

CONFIG_OPENTHREAD_CSL_ACCURACY

The current CSL rx/tx scheduling drift, in units of $\pm$ ppm

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [CSL Configurations](#)

The current accuracy of the clock used for scheduling CSL operations

Default value:

- 50 if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_CSL_UNCERTAIN

The CSL Uncertainty in units of 10 us.

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [CSL Configurations](#)

The fixed uncertainty of the Device for scheduling CSL Transmissions in units of 10 microseconds.

Default value:

- 50 if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_CSL_DEBUG_ENABLE

Enable CSL debug

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [CSL Configurations](#)

Select this option to set rx on when sleep in CSL feature, only for debug

Default value:

- No (disabled) if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_TIME_SYNC

Enable the time synchronization service feature

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable time synchronization feature, the devices in the same Thread network could sync to the same network time.

Default value:

- No (disabled) if `CONFIG_OPENTHREAD_ENABLED`

CONFIG_OPENTHREAD_RADIO_STATS_ENABLE

Enable Radio Statistics feature

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable the radio statistics feature, you can use radio command to print some radio Statistics information.

Default value:

- No (disabled) if `(CONFIG_OPENTHREAD_FTD || CONFIG_OPENTHREAD_MTD) && CONFIG_OPENTHREAD_ENABLED`

CONFIG_OPENTHREAD_RX_ON_WHEN_IDLE

Enable OpenThread radio capability rx on when idle

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable OpenThread radio capability rx on when idle. Do not support this feature when SW coexistence is enabled.

Default value:

- No (disabled) if `(CONFIG_ESP_COEX_SW_COEXIST_ENABLE || CONFIG_ESP_COEX_EXTERNAL_COEXIST_ENABLE) && CONFIG_OPENTHREAD_ENABLED`

CONFIG_OPENTHREAD_DIAG

Enable diag

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable Diag in OpenThread. This will enable diag mode and a series of diag commands in the OpenThread command line. These commands allow users to manipulate low-level features of the storage and 15.4 radio.

Default value:

- Yes (enabled) if `CONFIG_OPENTHREAD_ENABLED`

CONFIG_OPENTHREAD_PARENT_SEARCH_MTD

Enable Parent Search

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#)

Select this option to enable "Periodic Parent Search" function for MTD. This checks the average RSS to its current parent every periodically and starts a parent search process if the average RSS is below `OPENTHREAD_PARENT_SEARCH_RSS_THRESHOLD`. This feature is always enabled for FTDs.

Default value:

- Yes (enabled) if `CONFIG_OPENTHREAD_MTD && CONFIG_OPENTHREAD_ENABLED`

Parent Search Configurations Contains:

- [CONFIG_OPENTHREAD_PARENT_SEARCH_BACKOFF_INTERVAL_MINS](#)
- [CONFIG_OPENTHREAD_PARENT_SEARCH_CHECK_INTERVAL_MINS](#)
- [CONFIG_OPENTHREAD_PARENT_SEARCH_RESELECT_TIMEOUT_MINS](#)
- [CONFIG_OPENTHREAD_PARENT_SEARCH_RSS_MARGIN](#)
- [CONFIG_OPENTHREAD_PARENT_SEARCH_RSS_THRESHOLD](#)

CONFIG_OPENTHREAD_PARENT_SEARCH_CHECK_INTERVAL_MINS

The interval in minutes for a child to check the trigger condition to perform a parent search

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Parent Search Configurations](#)

Default value:

- 9 if ([CONFIG_OPENTHREAD_PARENT_SEARCH_MTD](#) || [CONFIG_OPENTHREAD_FTD](#)) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_PARENT_SEARCH_BACKOFF_INTERVAL_MINS

The backoff interval in minutes for a child to not perform a parent search after triggering one

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Parent Search Configurations](#)

Default value:

- 600 if ([CONFIG_OPENTHREAD_PARENT_SEARCH_MTD](#) || [CONFIG_OPENTHREAD_FTD](#)) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_PARENT_SEARCH_RSS_THRESHOLD

The RSS threshold used to trigger a parent search

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Parent Search Configurations](#)

Default value:

- "-65" if ([CONFIG_OPENTHREAD_PARENT_SEARCH_MTD](#) || [CONFIG_OPENTHREAD_FTD](#)) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_PARENT_SEARCH_RESELECT_TIMEOUT_MINS

The parent reselect timeout duration in minutes used on FTD child devices

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Parent Search Configurations](#)

Default value:

- 90 if ([CONFIG_OPENTHREAD_PARENT_SEARCH_MTD](#) || [CONFIG_OPENTHREAD_FTD](#)) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_PARENT_SEARCH_RSS_MARGIN

The RSS margin over the current parent RSS used on FTD child devices

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [Parent Search Configurations](#)

Default value:

- 7 if ([CONFIG_OPENTHREAD_PARENT_SEARCH_MTD](#) || [CONFIG_OPENTHREAD_FTD](#)) && [CONFIG_OPENTHREAD_ENABLED](#)

Thread Memory Allocation Contains:

- [CONFIG_OPENTHREAD_PLATFORM_MALLOC_CAP_SPIRAM](#)

CONFIG_OPENTHREAD_PLATFORM_MALLOC_CAP_SPIRAM

Allocate memory from PSRAM

Found in: Component config > OpenThread > CONFIG_OPENTHREAD_ENABLED > Thread Core Features > Thread Memory Allocation

Select this option to allocate buffer from PSRAM for Thread

Default value:

- Yes (enabled) if ([CONFIG_SPIRAM_USE_CAPS_ALLOC](#) || [CONFIG_SPIRAM_USE_MALLOC](#)) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_PLATFORM_MSGPOOL_MANAGEMENT

Allocate message pool buffer from PSRAM

Found in: Component config > OpenThread > CONFIG_OPENTHREAD_ENABLED > Thread Core Features > Thread Memory Allocation > CONFIG_OPENTHREAD_PLATFORM_MALLOC_CAP_SPIRAM

If enabled, the message pool is managed by platform defined logic.

Default value:

- Yes (enabled) if [CONFIG_OPENTHREAD_PLATFORM_MALLOC_CAP_SPIRAM](#) && ([CONFIG_SPIRAM_USE_CAPS_ALLOC](#) || [CONFIG_SPIRAM_USE_MALLOC](#)) && [CONFIG_OPENTHREAD_ENABLED](#)

OpenThread Stack Parameters Contains:

- [CONFIG_OPENTHREAD_MAC_MAX_CSMA_BACKOFFS_DIRECT](#)
- [CONFIG_OPENTHREAD_PREFERRED_CHANNEL_MASK](#)
- [CONFIG_OPENTHREAD_SUPPORTED_CHANNEL_MASK](#)
- [CONFIG_OPENTHREAD_XTAL_ACCURACY](#)
- [CONFIG_OPENTHREAD_BUS_LATENCY](#)
- [CONFIG_OPENTHREAD_NUM_MESSAGE_BUFFERS](#)
- [CONFIG_OPENTHREAD_MLE_MAX_CHILDREN](#)
- [CONFIG_OPENTHREAD_TMF_ADDR_CACHE_ENTRIES](#)
- [CONFIG_OPENTHREAD_UART_BUFFER_SIZE](#)
- [Thread Address Query Config](#)

Thread Address Query Config Contains:

- [CONFIG_OPENTHREAD_ADDRESS_QUERY_RETRY_DELAY](#)
- [CONFIG_OPENTHREAD_ADDRESS_QUERY_MAX_RETRY_DELAY](#)
- [CONFIG_OPENTHREAD_ADDRESS_QUERY_TIMEOUT](#)

CONFIG_OPENTHREAD_ADDRESS_QUERY_TIMEOUT

Timeout (in seconds) for a address notification response after sending an address query.

Found in: Component config > OpenThread > CONFIG_OPENTHREAD_ENABLED > Thread Core Features > OpenThread Stack Parameters > Thread Address Query Config

Default value:

- 3 if ([CONFIG_OPENTHREAD_FTD](#) || [CONFIG_OPENTHREAD_MTD](#)) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_ADDRESS_QUERY_RETRY_DELAY

Initial retry delay for address query (in seconds).

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [OpenThread Stack Parameters](#) > [Thread Address Query Config](#)

Default value:

- 15 if ([CONFIG_OPENTHREAD_FTD](#) || [CONFIG_OPENTHREAD_MTD](#)) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_ADDRESS_QUERY_MAX_RETRY_DELAY

Maximum retry delay for address query (in seconds).

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [OpenThread Stack Parameters](#) > [Thread Address Query Config](#)

Default value:

- 120 if ([CONFIG_OPENTHREAD_FTD](#) || [CONFIG_OPENTHREAD_MTD](#)) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_PREFERRED_CHANNEL_MASK

Preferred channel mask

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [OpenThread Stack Parameters](#)

Default value:

- "0x7fff800" if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_SUPPORTED_CHANNEL_MASK

Supported channel mask

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [OpenThread Stack Parameters](#)

Default value:

- "0x7fff800" if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_NUM_MESSAGE_BUFFERS

The number of openthread message buffers

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [OpenThread Stack Parameters](#)

Default value:

- 1024 if [CONFIG_OPENTHREAD_PLATFORM_MSGPOOL_MANAGEMENT](#) && [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_XTAL_ACCURACY

The accuracy of the XTAL

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [OpenThread Stack Parameters](#)

The device's XTAL accuracy, in ppm.

Default value:

- 10 if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_BUS_LATENCY

The bus latency between host and radio chip

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [OpenThread Stack Parameters](#)

The device's bus latency, in us.

Default value:

- 4000 if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_MLE_MAX_CHILDREN

The size of max MLE children entries

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [OpenThread Stack Parameters](#)

Default value:

- 10 if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_TMF_ADDR_CACHE_ENTRIES

The size of max TMF address cache entries

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [OpenThread Stack Parameters](#)

Default value:

- 20 if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_UART_BUFFER_SIZE

The uart received buffer size of openthread

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [OpenThread Stack Parameters](#)

Set the OpenThread UART buffer size.

Default value:

- 2048 if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_MAC_MAX_CSMA_BACKOFFS_DIRECT

Maximum backoffs times before declaring a channel access failure.

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Core Features](#) > [OpenThread Stack Parameters](#)

The maximum number of backoffs the CSMA-CA algorithm will attempt before declaring a channel access failure.

Default value:

- 4 if [CONFIG_OPENTHREAD_ENABLED](#)

Thread Log Contains:

- [CONFIG_OPENTHREAD_LOG_LEVEL_DYNAMIC](#)
- [CONFIG_OPENTHREAD_LOG_LEVEL](#)

CONFIG_OPENTHREAD_LOG_LEVEL_DYNAMIC

Enable dynamic log level control

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Log](#)

Select this option to enable dynamic log level control for OpenThread

Default value:

- Yes (enabled) if [CONFIG_OPENTHREAD_ENABLED](#)

CONFIG_OPENTHREAD_LOG_LEVEL

OpenThread log verbosity

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Log](#)

Select OpenThread log level.

Available options:

- No logs (CONFIG_OPENTHREAD_LOG_LEVEL_NONE)
- Error logs (CONFIG_OPENTHREAD_LOG_LEVEL_CRIT)
- Warning logs (CONFIG_OPENTHREAD_LOG_LEVEL_WARN)
- Notice logs (CONFIG_OPENTHREAD_LOG_LEVEL_NOTE)
- Info logs (CONFIG_OPENTHREAD_LOG_LEVEL_INFO)
- Debug logs (CONFIG_OPENTHREAD_LOG_LEVEL_DEBUG)

Thread Extensioned Features

 Contains:

- [CONFIG_OPENTHREAD_HEADER_CUSTOM](#)

CONFIG_OPENTHREAD_HEADER_CUSTOM

Use a header file defined by customer

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Extensioned Features](#)

This option allows users to tailor the values of openthread macros according to their requirements. The openthread submodule contains numerous macros, each with a default value set. In the Kconfig of ESP openthread, users can set specific openthread parameters, which will be applied for certain openthread macros in the openthread-core-esp32x-xxx-config.h file. During compilation, the values specified in openthread-core-esp32x-xxx-config.h will replace the default settings in the openthread submodule. However, Kconfig does not cover all openthread macros, particularly those typically using default values. For such macros, users can enable the OPENTHREAD_HEADER_CUSTOM option in the Kconfig and provide a custom header file. Macros defined in the custom header file will have the highest priority.

Default value:

- No (disabled) if [CONFIG_OPENTHREAD_ENABLED](#)

OpenThread Custom Header Config

 Contains:

- [CONFIG_OPENTHREAD_CUSTOM_HEADER_FILE_NAME](#)
- [CONFIG_OPENTHREAD_CUSTOM_HEADER_PATH](#)

CONFIG_OPENTHREAD_CUSTOM_HEADER_PATH

Path of custom header file

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Extended Features](#) > [CONFIG_OPENTHREAD_HEADER_CUSTOM](#) > [OpenThread Custom Header Config](#)

Please use relative paths with respect to the project folder.

Default value:

- if `CONFIG_OPENTHREAD_HEADER_CUSTOM` && `CONFIG_OPENTHREAD_HEADER_CUSTOM` && `CONFIG_OPENTHREAD_ENABLED`

CONFIG_OPENTHREAD_CUSTOM_HEADER_FILE_NAME

Name of custom header file

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_ENABLED](#) > [Thread Extended Features](#) > [CONFIG_OPENTHREAD_HEADER_CUSTOM](#) > [OpenThread Custom Header Config](#)

Name of custom header file.

Default value:

- "esp_ot_custom_config.h" if `CONFIG_OPENTHREAD_HEADER_CUSTOM` && `CONFIG_OPENTHREAD_HEADER_CUSTOM` && `CONFIG_OPENTHREAD_ENABLED`

OpenThread Spinel Contains:

- [CONFIG_OPENTHREAD_SPINEL_ONLY](#)
- [CONFIG_OPENTHREAD_SPINEL_MAC_MAX_CSMA_BACKOFFS_DIRECT](#)
- [CONFIG_OPENTHREAD_SPINEL_RX_FRAME_BUFFER_SIZE](#)
- [CONFIG_OPENTHREAD_SPINEL_UART_DRIVER_BUFFER_SIZE](#)

CONFIG_OPENTHREAD_SPINEL_ONLY

Enable OpenThread External Radio Spinel feature

Found in: [Component config](#) > [OpenThread](#) > [OpenThread Spinel](#)

Select this option to enable the OpenThread Radio Spinel for external protocol stack, such as Zigbee.

Default value:

- No (disabled)

CONFIG_OPENTHREAD_SPINEL_RX_FRAME_BUFFER_SIZE

The size of openthread spinel rx frame buffer

Found in: [Component config](#) > [OpenThread](#) > [OpenThread Spinel](#)

Default value:

- 1024 if (`CONFIG_OPENTHREAD_MTD` || `CONFIG_OPENTHREAD_RADIO`) && (`CONFIG_OPENTHREAD_ENABLED` || `CONFIG_OPENTHREAD_SPINEL_ONLY`)
- 2048 if (`CONFIG_OPENTHREAD_FTD` || `CONFIG_OPENTHREAD_SPINEL_ONLY`) && (`CONFIG_OPENTHREAD_ENABLED` || `CONFIG_OPENTHREAD_SPINEL_ONLY`)

CONFIG_OPENTHREAD_SPINEL_MAC_MAX_CSMA_BACKOFFS_DIRECT

Maximum backoffs times before declaring a channel access failure.

Found in: [Component config](#) > [OpenThread](#) > [OpenThread Spinel](#)

The maximum number of backoffs the CSMA-CA algorithm will attempt before declaring a channel access failure.

Default value:

- 4 if `CONFIG_OPENTHREAD_SPINEL_ONLY`

CONFIG_OPENTHREAD_SPINEL_UART_DRIVER_BUFFER_SIZE

UART driver buffer size for radio spinel interface

Found in: [Component config](#) > [OpenThread](#) > [OpenThread Spinel](#)

RX buffer size passed to `uart_driver_install()` for the radio spinel UART interface.

Default value:

- 1024 if `CONFIG_OPENTHREAD_SPINEL_ONLY` || `CONFIG_OPENTHREAD_RADIO_SPINEL_UART`

CONFIG_OPENTHREAD_DEBUG

Enable ESP OpenThread Debug

Found in: [Component config](#) > [OpenThread](#)

Enable additional debugging support for ESP OpenThread integration. This includes various diagnostic tools and logs to help track down issues in OpenThread networking or system integration

Default value:

- No (disabled)

Contains:

- [CONFIG_OPENTHREAD_DUMP_MAC_ON_ASSERT](#)

CONFIG_OPENTHREAD_DUMP_MAC_ON_ASSERT

Dump 802.15.4 MAC debug info on OpenThread assert

Found in: [Component config](#) > [OpenThread](#) > [CONFIG_OPENTHREAD_DEBUG](#)

When enabled, this option triggers the printing of 802.15.4 MAC layer debug information whenever an OpenThread assert occurs. This can help developers analyze unexpected failures by providing additional MAC layer context.

Default value:

- No (disabled) if `CONFIG_OPENTHREAD_DEBUG` && `CONFIG_OPENTHREAD_RADIO_NATIVE` && `CONFIG_IEEE802154_DEBUG`

Protocomm Contains:

- [CONFIG_ESP_PROTOCOMM_SUPPORT_SECURITY_VERSION_0](#)
- [CONFIG_ESP_PROTOCOMM_SUPPORT_SECURITY_VERSION_1](#)
- [CONFIG_ESP_PROTOCOMM_SUPPORT_SECURITY_VERSION_2](#)

CONFIG_ESP_PROTOCOMM_SUPPORT_SECURITY_VERSION_0

Support protocomm security version 0 (no security)

Found in: [Component config](#) > [Protocomm](#)

Enable support of security version 0. Disabling this option saves some code size. Consult the Enabling protocomm security version section of the Protocomm documentation in ESP-IDF Programming guide for more details.

Default value:

- Yes (enabled)

CONFIG_ESP_PROTOCOMM_SUPPORT_SECURITY_VERSION_1

Support protocomm security version 1 (Curve25519 key exchange + AES-CTR encryption/decryption)

Found in: [Component config](#) > [Protocomm](#)

Enable support of security version 1. Disabling this option saves some code size. Consult the Enabling protocomm security version section of the Protocomm documentation in ESP-IDF Programming guide for more details.

Default value:

- Yes (enabled)

CONFIG_ESP_PROTOCOMM_SUPPORT_SECURITY_VERSION_2

Support protocomm security version 2 (SRP6a-based key exchange + AES-GCM encryption/decryption)

Found in: [Component config](#) > [Protocomm](#)

Enable support of security version 2. Disabling this option saves some code size. Consult the Enabling protocomm security version section of the Protocomm documentation in ESP-IDF Programming guide for more details.

Default value:

- Yes (enabled)

PThreads Contains:

- [CONFIG_PTHREAD_TASK_NAME_DEFAULT](#)
- [CONFIG_PTHREAD_TASK_CORE_DEFAULT](#)
- [CONFIG_PTHREAD_TASK_PRIO_DEFAULT](#)
- [CONFIG_PTHREAD_TASK_STACK_SIZE_DEFAULT](#)
- [CONFIG_PTHREAD_STACK_MIN](#)

CONFIG_PTHREAD_TASK_PRIO_DEFAULT

Default task priority

Found in: [Component config](#) > [PThreads](#)

Priority used to create new tasks with default pthread parameters.

Range:

- from 0 to 255

Default value:

- 5

CONFIG_PTHREAD_TASK_STACK_SIZE_DEFAULT

Default task stack size

Found in: [Component config](#) > [PThreads](#)

Stack size used to create new tasks with default pthread parameters.

Default value:

- 3072

CONFIG_PTHREAD_STACK_MIN

Minimum allowed pthread stack size

Found in: [Component config](#) > [PThreads](#)

Minimum allowed pthread stack size set in attributes passed to pthread_create

Default value:

- 768

CONFIG_PTHREAD_TASK_CORE_DEFAULT

Default pthread core affinity

Found in: [Component config](#) > [PThreads](#)

The default core to which pthreads are pinned.

Available options:

- No affinity (CONFIG_PTHREAD_DEFAULT_CORE_NO_AFFINITY)
- Core 0 (CONFIG_PTHREAD_DEFAULT_CORE_0)
- Core 1 (CONFIG_PTHREAD_DEFAULT_CORE_1)

CONFIG_PTHREAD_TASK_NAME_DEFAULT

Default name of pthreads

Found in: [Component config](#) > [PThreads](#)

The default name of pthreads.

Default value:

- "pthread"

SoC Settings Contains:

- [MMU Config](#)

MMU Config

Main Flash configuration Contains:

- [Optional and Experimental Features \(READ DOCS FIRST\)](#)
- [SPI Flash behavior when brownout](#)

SPI Flash behavior when brownout Contains:

- [CONFIG_SPI_FLASH_BROWNOUT_RESET_XMC](#)

CONFIG_SPI_FLASH_BROWNOUT_RESET_XMC

Enable sending reset when brownout for XMC flash chips

Found in: [Component config](#) > [Main Flash configuration](#) > [SPI Flash behavior when brownout](#)

When this option is selected, the patch will be enabled for XMC. Follow the recommended flow by XMC for better stability.

DO NOT DISABLE UNLESS YOU KNOW WHAT YOU ARE DOING.

Optional and Experimental Features (READ DOCS FIRST) Contains:

- [CONFIG_SPI_FLASH_AUTO_SUSPEND](#)
- [CONFIG_SPI_FLASH_FORCE_ENABLE_C6_H2_SUSPEND](#)
- [CONFIG_SPI_FLASH_FORCE_ENABLE_XMC_C_SUSPEND](#)
- [CONFIG_SPI_FLASH_HPM_DC](#)

CONFIG_SPI_FLASH_HPM_DC

Support HPM using DC (READ DOCS FIRST)

Found in: [Component config](#) > [Main Flash configuration](#) > [Optional and Experimental Features \(READ DOCS FIRST\)](#)

This feature needs your bootloader to be compiled DC-aware (BOOT-LOADER_FLASH_DC_AWARE=y). Otherwise the chip will not be able to boot after a reset.

Available options:

- Auto (Enable when bootloader support enabled (BOOT-LOADER_FLASH_DC_AWARE)) (CONFIG_SPI_FLASH_HPM_DC_AUTO)
- Disable (READ DOCS FIRST) (CONFIG_SPI_FLASH_HPM_DC_DISABLE)

CONFIG_SPI_FLASH_AUTO_SUSPEND

Auto suspend long erase/write operations (READ DOCS FIRST)

Found in: [Component config](#) > [Main Flash configuration](#) > [Optional and Experimental Features \(READ DOCS FIRST\)](#)

This option is disabled by default because it is supported only for specific flash chips and for specific Espressif chips. To evaluate if you can use this feature refer to *Optional Features for Flash > Auto Suspend & Resume* of the *ESP-IDF Programming Guide*.

CAUTION: If you want to OTA to an app with this feature turned on, please make sure the bootloader has the support for it. (later than IDF v4.3)

If you are using an official Espressif module, please contact Espressif Business support to check if the module has the flash that support this feature installed. Also refer to *Concurrency Constraints for Flash on SPI1 > Flash Auto Suspend Feature* before enabling this option.

CONFIG_SPI_FLASH_FORCE_ENABLE_XMC_C_SUSPEND

Enable XMC-C series flash chip suspend feature anyway

Found in: [Component config](#) > [Main Flash configuration](#) > [Optional and Experimental Features \(READ DOCS FIRST\)](#)

XMC-C series is regarded as not qualified for the Suspend feature, since its specification has a tRS >= 1ms restriction. We strongly do not suggest using it for the Suspend feature. However, if your product in field has enabled this feature, you may still enable this config option to keep the legacy behavior.

For new users, DO NOT enable this config.

CONFIG_SPI_FLASH_FORCE_ENABLE_C6_H2_SUSPEND

Enable chip suspend feature on c6 or h2 anyway (DO NOT ENABLE FOR NEW USERS OR APPLICATIONS)

Found in: [Component config](#) > [Main Flash configuration](#) > [Optional and Experimental Features \(READ DOCS FIRST\)](#)

Flash suspend has defect on ESP32C6 until v0.2 and ESP32H2 until v1.2. If you already use suspend feature for massive protection, you can enable this for bypassing check after knowing the risk. But if you are new users, or developing new applications, or producing a new batch, please DO NOT enable this config option.

For more information, please refer to errata or connect to Espressif business support team.

SPI Flash driver Contains:

- *Auto-detect flash chips*
- *CONFIG_SPI_FLASH_BYPASS_BLOCK_ERASE*
- *CONFIG_SPI_FLASH_ENABLE_ENCRYPTED_READ_WRITE*
- *CONFIG_SPI_FLASH_ENABLE_COUNTERS*
- *CONFIG_SPI_FLASH_ROM_DRIVER_PATCH*
- *CONFIG_SPI_FLASH_YIELD_DURING_ERASE*
- *CONFIG_SPI_FLASH_CHECK_ERASE_TIMEOUT_DISABLED*
- *CONFIG_SPI_FLASH_WRITE_CHUNK_SIZE*
- *CONFIG_SPI_FLASH_OVERRIDE_CHIP_DRIVER_LIST*
- *CONFIG_SPI_FLASH_SIZE_OVERRIDE*
- *CONFIG_SPI_FLASH_ROM_IMPL*
- *CONFIG_SPI_FLASH_VERIFY_WRITE*
- *CONFIG_SPI_FLASH_DANGEROUS_WRITE*

CONFIG_SPI_FLASH_VERIFY_WRITE

Verify SPI flash writes

Found in: [Component config](#) > [SPI Flash driver](#)

If this option is enabled, any time SPI flash is written then the data will be read back and verified. This can catch hardware problems with SPI flash, or flash which was not erased before verification.

This will slightly influence the write performance.

CONFIG_SPI_FLASH_LOG_FAILED_WRITE

Log errors if verification fails

Found in: [Component config](#) > [SPI Flash driver](#) > [CONFIG_SPI_FLASH_VERIFY_WRITE](#)

If this option is enabled, if SPI flash write verification fails then a log error line will be written with the address, expected & actual values. This can be useful when debugging hardware SPI flash problems.

CONFIG_SPI_FLASH_WARN_SETTING_ZERO_TO_ONE

Log warning if writing zero bits to ones

Found in: [Component config](#) > [SPI Flash driver](#) > [CONFIG_SPI_FLASH_VERIFY_WRITE](#)

If this option is enabled, any SPI flash write which tries to set zero bits in the flash to ones will log a warning. Such writes will not result in the requested data appearing identically in flash once written, as SPI NOR flash can only set bits to one when an entire sector is erased. After erasing, individual bits can only be written from one to zero.

Note that some software (such as SPIFFS) which is aware of SPI NOR flash may write one bits as an optimisation, relying on the data in flash becoming a bitwise AND of the new data and any existing data. Such software will log spurious warnings if this option is enabled.

CONFIG_SPI_FLASH_ENABLE_COUNTERS

Enable operation counters

Found in: [Component config](#) > [SPI Flash driver](#)

This option enables the following APIs:

- `esp_flash_reset_counters`
- `esp_flash_dump_counters`
- `esp_flash_get_counters`

These APIs may be used to collect performance data for spi_flash APIs and to help understand behaviour of libraries which use SPI flash.

CONFIG_SPI_FLASH_ROM_DRIVER_PATCH

Enable SPI flash ROM driver patched functions

Found in: [Component config](#) > [SPI Flash driver](#)

Enable this flag to use patched versions of SPI flash ROM driver functions. This option should be enabled, if any one of the following is true: (1) need to write to flash on ESP32-D2WD; (2) main SPI flash is connected to non-default pins; (3) main SPI flash chip is manufactured by ISSI.

CONFIG_SPI_FLASH_ROM_IMPL

Use esp_flash implementation in ROM

Found in: [Component config](#) > [SPI Flash driver](#)

Enable this flag to use new SPI flash driver functions from ROM instead of ESP-IDF.

If keeping this as "n" in your project, you will have less free IRAM. But you can use all of our flash features.

If making this as "y" in your project, you will increase free IRAM. But you may miss out on some flash features and support for new flash chips.

Currently the ROM cannot support the following features:

- `SPI_FLASH_AUTO_SUSPEND` (C3, S3)

CONFIG_SPI_FLASH_DANGEROUS_WRITE

Writing to dangerous flash regions

Found in: [Component config](#) > [SPI Flash driver](#)

SPI flash APIs can optionally abort or return a failure code if erasing or writing addresses that fall at the beginning of flash (covering the bootloader and partition table) or that overlap the app partition that contains the running app.

It is not recommended to ever write to these regions from an IDF app, and this check prevents logic errors or corrupted firmware memory from damaging these regions.

Note that this feature **does not** check calls to the `esp_rom_XXX` SPI flash ROM functions. These functions should not be called directly from IDF applications.

Available options:

- Aborts (`CONFIG_SPI_FLASH_DANGEROUS_WRITE_ABORTS`)
- Fails (`CONFIG_SPI_FLASH_DANGEROUS_WRITE_FAILS`)
- Allowed (`CONFIG_SPI_FLASH_DANGEROUS_WRITE_ALLOWED`)

CONFIG_SPI_FLASH_BYPASS_BLOCK_ERASE

Bypass a block erase and always do sector erase

Found in: [Component config](#) > [SPI Flash driver](#)

Some flash chips can have very high "max" erase times, especially for block erase (32KB or 64KB). This option allows to bypass "block erase" and always do sector erase commands. This will be much slower overall in most cases, but improves latency for other code to run.

CONFIG_SPI_FLASH_YIELD_DURING_ERASE

Enables yield operation during flash erase

Found in: [Component config](#) > [SPI Flash driver](#)

This allows to yield the CPUs between erase commands. Prevents starvation of other tasks. Please use this configuration together with `SPI\_FLASH\_ERASE\_YIELD\_DURATION\_MS` and `SPI\_FLASH\_ERASE\_YIELD\_TICKS` after carefully checking flash datasheet to avoid a watchdog timeout. For more information, please check *SPI Flash API* reference documentation under section *OS Function*.

CONFIG_SPI_FLASH_ERASE_YIELD_DURATION_MS

Duration of erasing to yield CPUs (ms)

Found in: [Component config](#) > [SPI Flash driver](#) > [CONFIG_SPI_FLASH_YIELD_DURING_ERASE](#)

If a duration of one erase command is large then it will yield CPUs after finishing a current command.

CONFIG_SPI_FLASH_ERASE_YIELD_TICKS

CPU release time (tick) for an erase operation

Found in: [Component config](#) > [SPI Flash driver](#) > [CONFIG_SPI_FLASH_YIELD_DURING_ERASE](#)

Defines how many ticks will be before returning to continue a erasing.

CONFIG_SPI_FLASH_WRITE_CHUNK_SIZE

Flash write chunk size

Found in: [Component config](#) > [SPI Flash driver](#)

Flash write is broken down in terms of multiple (smaller) write operations. This configuration options helps to set individual write chunk size, smaller value here ensures that cache (and non-IRAM resident interrupts) remains disabled for shorter duration.

CONFIG_SPI_FLASH_SIZE_OVERRIDE

Override flash size in bootloader header by `ESPTOOLPY_FLASHSIZE`

Found in: [Component config](#) > [SPI Flash driver](#)

SPI Flash driver uses the flash size configured in bootloader header by default. Enable this option to override flash size with latest `ESPTOOLPY_FLASHSIZE` value from the app header if the size in the bootloader header is incorrect.

CONFIG_SPI_FLASH_CHECK_ERASE_TIMEOUT_DISABLED

Flash timeout checkout disabled

Found in: [Component config](#) > [SPI Flash driver](#)

This option is helpful if you are using a flash chip whose timeout is quite large or unpredictable.

CONFIG_SPI_FLASH_OVERRIDE_CHIP_DRIVER_LIST

Override default chip driver list

Found in: [Component config](#) > [SPI Flash driver](#)

This option allows the chip driver list to be customized, instead of using the default list provided by ESP-IDF.

When this option is enabled, the default list is no longer compiled or linked. Instead, the *default_registered_chips* structure must be provided by the user.

See example: `custom_chip_driver` under `examples/storage` for more details.

Auto-detect flash chips Contains:

- [CONFIG_SPI_FLASH_SUPPORT_BOYA_CHIP](#)
- [CONFIG_SPI_FLASH_SUPPORT_GD_CHIP](#)
- [CONFIG_SPI_FLASH_SUPPORT_ISSI_CHIP](#)
- [CONFIG_SPI_FLASH_SUPPORT_MXIC_CHIP](#)
- [CONFIG_SPI_FLASH_SUPPORT_TH_CHIP](#)
- [CONFIG_SPI_FLASH_SUPPORT_WINBOND_CHIP](#)

CONFIG_SPI_FLASH_SUPPORT_ISSI_CHIP

ISSI

Found in: [Component config](#) > [SPI Flash driver](#) > [Auto-detect flash chips](#)

Enable this to support auto detection of ISSI chips if chip vendor not directly given by `chip\_drv` member of the chip struct. This adds support for variant chips, however will extend detecting time.

CONFIG_SPI_FLASH_SUPPORT_MXIC_CHIP

MXIC

Found in: [Component config](#) > [SPI Flash driver](#) > [Auto-detect flash chips](#)

Enable this to support auto detection of MXIC chips if chip vendor not directly given by `chip\_drv` member of the chip struct. This adds support for variant chips, however will extend detecting time.

CONFIG_SPI_FLASH_SUPPORT_GD_CHIP

GigaDevice

Found in: [Component config](#) > [SPI Flash driver](#) > [Auto-detect flash chips](#)

Enable this to support auto detection of GD (GigaDevice) chips if chip vendor not directly given by `chip\_drv` member of the chip struct. If you are using Wrover modules, please don't disable this, otherwise your flash may not work in 4-bit mode.

This adds support for variant chips, however will extend detecting time and image size. Note that the default chip driver supports the GD chips with product ID 60H.

CONFIG_SPI_FLASH_SUPPORT_WINBOND_CHIP

Winbond

Found in: [Component config](#) > [SPI Flash driver](#) > [Auto-detect flash chips](#)

Enable this to support auto detection of Winbond chips if chip vendor not directly given by `chip\_drv` member of the chip struct. This adds support for variant chips, however will extend detecting time.

CONFIG_SPI_FLASH_SUPPORT_BOYA_CHIP

BOYA

Found in: [Component config](#) > [SPI Flash driver](#) > [Auto-detect flash chips](#)

Enable this to support auto detection of BOYA chips if chip vendor not directly given by `chip\_drv` member of the chip struct. This adds support for variant chips, however will extend detecting time.

CONFIG_SPI_FLASH_SUPPORT_TH_CHIP

TH

Found in: [Component config](#) > [SPI Flash driver](#) > [Auto-detect flash chips](#)

Enable this to support auto detection of TH chips if chip vendor not directly given by `chip\_drv` member of the chip struct. This adds support for variant chips, however will extend detecting time.

CONFIG_SPI_FLASH_ENABLE_ENCRYPTED_READ_WRITE

Enable encrypted partition read/write operations

Found in: [Component config](#) > [SPI Flash driver](#)

This option enables flash read/write operations to encrypted partition/s. This option is kept enabled irrespective of state of flash encryption feature. However, in case application is not using flash encryption feature and is in need of some additional memory from IRAM region (~1KB) then this config can be disabled.

SPIFFS Configuration Contains:

- [Debug Configuration](#)
- [CONFIG_SPIFFS_USE_MAGIC](#)
- [CONFIG_SPIFFS_GC_STATS](#)
- [CONFIG_SPIFFS_PAGE_CHECK](#)
- [CONFIG_SPIFFS_FOLLOW_SYMLINKS](#)
- [CONFIG_SPIFFS_MAX_PARTITIONS](#)
- [CONFIG_SPIFFS_USE_MTIME](#)
- [CONFIG_SPIFFS_GC_MAX_RUNS](#)
- [CONFIG_SPIFFS_OBJ_NAME_LEN](#)
- [CONFIG_SPIFFS_META_LENGTH](#)
- [SPIFFS Cache Configuration](#)
- [CONFIG_SPIFFS_PAGE_SIZE](#)
- [CONFIG_SPIFFS_MTIME_WIDE_64_BITS](#)

CONFIG_SPIFFS_MAX_PARTITIONS

Maximum Number of Partitions

Found in: [Component config](#) > [SPIFFS Configuration](#)

Define maximum number of partitions that can be mounted.

Range:

- from 1 to 10

Default value:

- 3

SPIFFS Cache Configuration Contains:

- [CONFIG_SPIFFS_CACHE](#)

CONFIG_SPIFFS_CACHE

Enable SPIFFS Cache

Found in: [Component config](#) > [SPIFFS Configuration](#) > [SPIFFS Cache Configuration](#)

Enables/disable memory read caching of nucleus file system operations.

Default value:

- Yes (enabled)

CONFIG_SPIFFS_CACHE_WR

Enable SPIFFS Write Caching

Found in: [Component config](#) > [SPIFFS Configuration](#) > [SPIFFS Cache Configuration](#) > [CONFIG_SPIFFS_CACHE](#)

Enables memory write caching for file descriptors in hydrogen.

Default value:

- Yes (enabled)

CONFIG_SPIFFS_CACHE_STATS

Enable SPIFFS Cache Statistics

Found in: [Component config](#) > [SPIFFS Configuration](#) > [SPIFFS Cache Configuration](#) > [CONFIG_SPIFFS_CACHE](#)

Enable/disable statistics on caching. Debug/test purpose only.

Default value:

- No (disabled)

CONFIG_SPIFFS_PAGE_CHECK

Enable SPIFFS Page Check

Found in: [Component config](#) > [SPIFFS Configuration](#)

Always check header of each accessed page to ensure consistent state. If enabled it will increase number of reads from flash, especially if cache is disabled.

Default value:

- Yes (enabled)

CONFIG_SPIFFS_GC_MAX_RUNS

Set Maximum GC Runs

Found in: [Component config](#) > [SPIFFS Configuration](#)

Define maximum number of GC runs to perform to reach desired free pages.

Range:

- from 1 to 10000

Default value:

- 10

CONFIG_SPIFFS_GC_STATS

Enable SPIFFS GC Statistics

Found in: [Component config](#) > [SPIFFS Configuration](#)

Enable/disable statistics on gc. Debug/test purpose only.

Default value:

- No (disabled)

CONFIG_SPIFFS_PAGE_SIZE

SPIFFS logical page size

Found in: [Component config](#) > [SPIFFS Configuration](#)

Logical page size of SPIFFS partition, in bytes. Must be multiple of flash page size (which is usually 256 bytes). Larger page sizes reduce overhead when storing large files, and improve filesystem performance when reading large files. Smaller page sizes reduce overhead when storing small (< page size) files.

Range:

- from 256 to 1024

Default value:

- 256

CONFIG_SPIFFS_OBJ_NAME_LEN

Set SPIFFS Maximum Name Length

Found in: [Component config](#) > [SPIFFS Configuration](#)

Object name maximum length. Note that this length include the zero-termination character, meaning maximum string of characters can at most be SPIFFS_OBJ_NAME_LEN - 1.

SPIFFS_OBJ_NAME_LEN + SPIFFS_META_LENGTH should not exceed SPIFFS_PAGE_SIZE - 64.

Range:

- from 1 to 256

Default value:

- 32

CONFIG_SPIFFS_FOLLOW_SYMLINKS

Enable symbolic links for image creation

Found in: [Component config](#) > [SPIFFS Configuration](#)

If this option is enabled, symbolic links are taken into account during partition image creation.

Default value:

- No (disabled)

CONFIG_SPIFFS_USE_MAGIC

Enable SPIFFS Filesystem Magic

Found in: [Component config](#) > [SPIFFS Configuration](#)

Enable this to have an identifiable spiffs filesystem. This will look for a magic in all sectors to determine if this is a valid spiffs system or not at mount time.

Default value:

- Yes (enabled)

CONFIG_SPIFFS_USE_MAGIC_LENGTH

Enable SPIFFS Filesystem Length Magic

Found in: [Component config](#) > [SPIFFS Configuration](#) > [CONFIG_SPIFFS_USE_MAGIC](#)

If this option is enabled, the magic will also be dependent on the length of the filesystem. For example, a filesystem configured and formatted for 4 megabytes will not be accepted for mounting with a configuration defining the filesystem as 2 megabytes.

Default value:

- Yes (enabled)

CONFIG_SPIFFS_META_LENGTH

Size of per-file metadata field

Found in: [Component config](#) > [SPIFFS Configuration](#)

This option sets the number of extra bytes stored in the file header. These bytes can be used in an application-specific manner. Set this to at least 4 bytes to enable support for saving file modification time.

SPIFFS_OBJ_NAME_LEN + SPIFFS_META_LENGTH should not exceed SPIFFS_PAGE_SIZE - 64.

Default value:

- 4

CONFIG_SPIFFS_USE_MTIME

Save file modification time

Found in: [Component config](#) > [SPIFFS Configuration](#)

If enabled, then the first 4 bytes of per-file metadata will be used to store file modification time (mtime), accessible through stat/fstat functions. Modification time is updated when the file is opened.

Default value:

- Yes (enabled)

CONFIG_SPIFFS_MTIME_WIDE_64_BITS

The time field occupies 64 bits in the image instead of 32 bits

Found in: [Component config](#) > [SPIFFS Configuration](#)

If this option is not set, the time field is 32 bits (up to 2106 year), otherwise it is 64 bits and make sure it matches SPIFFS_META_LENGTH. If the chip already has the spiffs image with the time field = 32 bits then this option cannot be applied in this case. Erase it first before using this option. To resolve the Y2K38 problem for the spiffs, use a toolchain with 64-bit time_t support.

Default value:

- No (disabled) if [CONFIG_SPIFFS_META_LENGTH](#) >= 8

Debug Configuration Contains:

- [*CONFIG_SPIFFS_DBG*](#)
- [*CONFIG_SPIFFS_API_DBG*](#)
- [*CONFIG_SPIFFS_CACHE_DBG*](#)
- [*CONFIG_SPIFFS_CHECK_DBG*](#)
- [*CONFIG_SPIFFS_TEST_VISUALISATION*](#)
- [*CONFIG_SPIFFS_GC_DBG*](#)

CONFIG_SPIFFS_DBG

Enable general SPIFFS debug

Found in: [Component config](#) > [SPIFFS Configuration](#) > [Debug Configuration](#)

Enabling this option will print general debug messages to the console.

Default value:

- No (disabled)

CONFIG_SPIFFS_API_DBG

Enable SPIFFS API debug

Found in: [Component config](#) > [SPIFFS Configuration](#) > [Debug Configuration](#)

Enabling this option will print API debug messages to the console.

Default value:

- No (disabled)

CONFIG_SPIFFS_GC_DBG

Enable SPIFFS Garbage Cleaner debug

Found in: [Component config](#) > [SPIFFS Configuration](#) > [Debug Configuration](#)

Enabling this option will print GC debug messages to the console.

Default value:

- No (disabled)

CONFIG_SPIFFS_CACHE_DBG

Enable SPIFFS Cache debug

Found in: [Component config](#) > [SPIFFS Configuration](#) > [Debug Configuration](#)

Enabling this option will print cache debug messages to the console.

Default value:

- No (disabled)

CONFIG_SPIFFS_CHECK_DBG

Enable SPIFFS Filesystem Check debug

Found in: [Component config](#) > [SPIFFS Configuration](#) > [Debug Configuration](#)

Enabling this option will print Filesystem Check debug messages to the console.

Default value:

- No (disabled)

CONFIG_SPIFFS_TEST_VISUALISATION

Enable SPIFFS Filesystem Visualization

Found in: [Component config](#) > [SPIFFS Configuration](#) > [Debug Configuration](#)

Enable this option to enable SPIFFS_vis function in the API.

Default value:

- No (disabled)

TCP Transport Contains:

- [Websocket](#)

Websocket Contains:

- [CONFIG_WS_TRANSPORT](#)

CONFIG_WS_TRANSPORT

Enable Websocket Transport

Found in: [Component config](#) > [TCP Transport](#) > [Websocket](#)

Enable support for creating websocket transport.

Default value:

- Yes (enabled)

CONFIG_WS_BUFFER_SIZE

Websocket transport buffer size

Found in: [Component config](#) > [TCP Transport](#) > [Websocket](#) > [CONFIG_WS_TRANSPORT](#)

Size of the buffer used for constructing the HTTP Upgrade request during connect

Default value:

- 1024

CONFIG_WS_DYNAMIC_BUFFER

Using dynamic websocket transport buffer

Found in: [Component config](#) > [TCP Transport](#) > [Websocket](#) > [CONFIG_WS_TRANSPORT](#)

If enable this option, websocket transport buffer will be freed after connection succeed to save more heap.

Default value:

- No (disabled)

Ultra Low Power (ULP) Co-processor Contains:

- [CONFIG_ULP_COPROC_ENABLED](#)
- [ULP RISC-V Settings](#)

CONFIG_ULP_COPROC_ENABLED

Enable Ultra Low Power (ULP) Co-processor

Found in: [Component config](#) > [Ultra Low Power \(ULP\) Co-processor](#)

Enable this feature if you plan to use the ULP Co-processor. Once this option is enabled, further ULP co-processor configuration will appear in the menu.

Default value:

- No (disabled) if `SOC_ULP_SUPPORTED || SOC_RISCV_COPROC_SUPPORTED || SOC_LP_CORE_SUPPORTED`

CONFIG_ULP_COPROC_TYPE

ULP Co-processor type

Found in: [Component config](#) > [Ultra Low Power \(ULP\) Co-processor](#) > [CONFIG_ULP_COPROC_ENABLED](#)

Choose the ULP Coprocessor type: ULP FSM (Finite State Machine) or ULP RISC-V.

Available options:

- ULP FSM (Finite State Machine) (`CONFIG_ULP_COPROC_TYPE_FSM`)
- ULP RISC-V (`CONFIG_ULP_COPROC_TYPE_RISCV`)
- LP core RISC-V (`CONFIG_ULP_COPROC_TYPE_LP_CORE`)

CONFIG_ULP_COPROC_RESERVE_MEM

RTC slow memory reserved for coprocessor

Found in: [Component config](#) > [Ultra Low Power \(ULP\) Co-processor](#) > [CONFIG_ULP_COPROC_ENABLED](#)

Bytes of memory to reserve for ULP Co-processor firmware & data. Data is reserved at the beginning of RTC slow memory.

Range:

- from 32 to 8176 if `CONFIG_ULP_COPROC_ENABLED` && (`SOC_ULP_SUPPORTED || SOC_RISCV_COPROC_SUPPORTED || SOC_LP_CORE_SUPPORTED`)

Default value:

- 4096 if `CONFIG_ULP_COPROC_ENABLED` && (`SOC_ULP_SUPPORTED || SOC_RISCV_COPROC_SUPPORTED || SOC_LP_CORE_SUPPORTED`)

ULP RISC-V Settings

 Contains:

- [CONFIG_ULP_RISCV_UART_BAUDRATE](#)
- [CONFIG_ULP_RISCV_I2C_RW_TIMEOUT](#)

CONFIG_ULP_RISCV_UART_BAUDRATE

Baudrate used by the bitbanged ULP RISC-V UART driver

Found in: [Component config](#) > [Ultra Low Power \(ULP\) Co-processor](#) > [ULP RISC-V Settings](#)

The accuracy of the bitbanged UART driver is limited, it is not recommend to increase the value above 19200.

Default value:

- 9600 if `CONFIG_ULP_COPROC_TYPE_RISCV` && (`SOC_ULP_SUPPORTED || SOC_RISCV_COPROC_SUPPORTED || SOC_LP_CORE_SUPPORTED`)

CONFIG_ULP_RISCV_I2C_RW_TIMEOUT

Set timeout for ULP RISC-V I2C transaction timeout in ticks.

Found in: [Component config](#) > [Ultra Low Power \(ULP\) Co-processor](#) > [ULP RISC-V Settings](#)

Set the ULP RISC-V I2C read/write timeout. Set this value to -1 if the ULP RISC-V I2C read and write APIs should wait forever. Please note that the tick rate of the ULP co-processor would be different than the OS tick rate of the main core and therefore can have different timeout value depending on which core the API is invoked on.

Range:

- from -1 to 4294967295 if [CONFIG_ULP_COPROC_TYPE_RISCV](#) && (SOC_ULP_SUPPORTED || SOC_RISCV_COPROC_SUPPORTED || SOC_LP_CORE_SUPPORTED)

Default value:

- 500 if [CONFIG_ULP_COPROC_TYPE_RISCV](#) && (SOC_ULP_SUPPORTED || SOC_RISCV_COPROC_SUPPORTED || SOC_LP_CORE_SUPPORTED)

Unity unit testing library Contains:

- [CONFIG_UNITY_ENABLE_COLOR](#)
- [CONFIG_UNITY_ENABLE_IDF_TEST_RUNNER](#)
- [CONFIG_UNITY_ENABLE_FIXTURE](#)
- [CONFIG_UNITY_ENABLE_BACKTRACE_ON_FAIL](#)
- [CONFIG_UNITY_ENABLE_64BIT](#)
- [CONFIG_UNITY_ENABLE_DOUBLE](#)
- [CONFIG_UNITY_ENABLE_FLOAT](#)

CONFIG_UNITY_ENABLE_FLOAT

Support for float type

Found in: [Component config](#) > [Unity unit testing library](#)

If not set, assertions on float arguments will not be available.

Default value:

- Yes (enabled)

CONFIG_UNITY_ENABLE_DOUBLE

Support for double type

Found in: [Component config](#) > [Unity unit testing library](#)

If not set, assertions on double arguments will not be available.

Default value:

- Yes (enabled)

CONFIG_UNITY_ENABLE_64BIT

Support for 64-bit integer types

Found in: [Component config](#) > [Unity unit testing library](#)

If not set, assertions on 64-bit integer types will always fail. If this feature is enabled, take care not to pass pointers (which are 32 bit) to UNITY_ASSERT_EQUAL, as that will cause pointer-to-int-cast warnings.

Default value:

- No (disabled)

CONFIG_UNITY_ENABLE_COLOR

Colorize test output

Found in: [Component config](#) > [Unity unit testing library](#)

If set, Unity will colorize test results using console escape sequences.

Default value:

- No (disabled)

CONFIG_UNITY_ENABLE_IDF_TEST_RUNNER

Include ESP-IDF test registration/running helpers

Found in: [Component config](#) > [Unity unit testing library](#)

If set, then the following features will be available:

- TEST_CASE macro which performs automatic registration of test functions
- Functions to run registered test functions: `unity_run_all_tests`, `unity_run_tests_with_filter`, `unity_run_single_test_by_name`.
- Interactive menu which lists test cases and allows choosing the tests to be run, available via `unity_run_menu` function.

Disable if a different test registration mechanism is used.

Default value:

- Yes (enabled)

CONFIG_UNITY_ENABLE_FIXTURE

Include Unity test fixture

Found in: [Component config](#) > [Unity unit testing library](#)

If set, `unity_fixture.h` header file and associated source files are part of the build. These provide an optional set of macros and functions to implement test groups.

Default value:

- No (disabled)

CONFIG_UNITY_ENABLE_BACKTRACE_ON_FAIL

Print a backtrace when a unit test fails

Found in: [Component config](#) > [Unity unit testing library](#)

If set, the unity framework will print the backtrace information before jumping back to the test menu. The jumping is usually occurs in assert functions such as `TEST_ASSERT`, `TEST_FAIL` etc.

Default value:

- No (disabled)

USB-OTG Contains:

- [CONFIG_USB_HOST_ENABLE_ENUM_FILTER_CALLBACK](#)
- [CONFIG_USB_HOST_HW_BUFFER_BIAS](#)
- [Hub Driver Configuration](#)
- [CONFIG_USB_HOST_CONTROL_TRANSFER_MAX_SIZE](#)

CONFIG_USB_HOST_CONTROL_TRANSFER_MAX_SIZE

Largest size (in bytes) of transfers to/from default endpoints

Found in: *Component config > USB-OTG*

Each USB device attached is allocated a dedicated buffer for its OUT/IN transfers to/from the device's control endpoint. The maximum size of that buffer is determined by this option. The limited size of the transfer buffer have the following implications: - The maximum length of control transfers is limited - Device's with configuration descriptors larger than this limit cannot be supported

Default value:

- 256 if SOC_USB_OTG_SUPPORTED

CONFIG_USB_HOST_HW_BUFFER_BIAS

Hardware FIFO size biasing

Found in: *Component config > USB-OTG*

The underlying hardware has size adjustable FIFOs to cache USB packets on reception (IN) or for transmission (OUT). The size of these FIFOs will affect the largest MPS (maximum packet size) and the maximum number of packets that can be cached at any one time. The hardware contains the following FIFOs: RX (for all IN packets), Non-periodic TX (for Bulk and Control OUT packets), and Periodic TX (for Interrupt and Isochronous OUT packets). This configuration option allows biasing the FIFO sizes towards a particular use case, which may be necessary for devices that have endpoints with large MPS. The MPS limits for each biasing are listed below:

Balanced: - IN (all transfer types), 408 bytes - OUT non-periodic (Bulk/Control), 192 bytes (i.e., 3 x 64 byte packets) - OUT periodic (Interrupt/Isochronous), 192 bytes

Bias IN: - IN (all transfer types), 600 bytes - OUT non-periodic (Bulk/Control), 64 bytes (i.e., 1 x 64 byte packets) - OUT periodic (Interrupt/Isochronous), 128 bytes

Bias Periodic OUT: - IN (all transfer types), 128 bytes - OUT non-periodic (Bulk/Control), 64 bytes (i.e., 1 x 64 byte packets) - OUT periodic (Interrupt/Isochronous), 600 bytes

Available options:

- Balanced (CONFIG_USB_HOST_HW_BUFFER_BIAS_BALANCED)
- Bias IN (CONFIG_USB_HOST_HW_BUFFER_BIAS_IN)
- Periodic OUT (CONFIG_USB_HOST_HW_BUFFER_BIAS_PERIODIC_OUT)

Hub Driver Configuration Contains:

- *Root Port configuration*
- *CONFIG_USB_HOST_HUBS_SUPPORTED*

Root Port configuration Contains:

- *CONFIG_USB_HOST_DEBOUNCE_DELAY_MS*
- *CONFIG_USB_HOST_RESET_HOLD_MS*
- *CONFIG_USB_HOST_RESET_RECOVERY_MS*
- *CONFIG_USB_HOST_SET_ADDR_RECOVERY_MS*

CONFIG_USB_HOST_DEBOUNCE_DELAY_MS

Debounce delay in ms

Found in: *Component config > USB-OTG > Hub Driver Configuration > Root Port configuration*

On connection of a USB device, the USB 2.0 specification requires a "debounce interval with a minimum duration of 100ms" to allow the connection to stabilize (see USB 2.0 chapter 7.1.7.3 for more details). During the debounce interval, no new connection/disconnection events are registered.

The default value is set to 250 ms to be safe.

Default value:

- 250 if SOC_USB_OTG_SUPPORTED

CONFIG_USB_HOST_RESET_HOLD_MS

Reset hold in ms

Found in: [Component config](#) > [USB-OTG](#) > [Hub Driver Configuration](#) > [Root Port configuration](#)

The reset signaling can be generated on any Hub or Host Controller port by request from the USB System Software. The USB 2.0 specification requires that "the reset signaling must be driven for a minimum of 10ms" (see USB 2.0 chapter 7.1.7.5 for more details). After the reset, the hub port will transition to the Enabled state (refer to Section 11.5).

The default value is set to 30 ms to be safe.

Default value:

- 30 if SOC_USB_OTG_SUPPORTED

CONFIG_USB_HOST_RESET_RECOVERY_MS

Reset recovery delay in ms

Found in: [Component config](#) > [USB-OTG](#) > [Hub Driver Configuration](#) > [Root Port configuration](#)

After a port stops driving the reset signal, the USB 2.0 specification requires that the "USB System Software guarantees a minimum of 10 ms for reset recovery" before the attached device is expected to respond to data transfers (see USB 2.0 chapter 7.1.7.3 for more details). The device may ignore any data transfers during the recovery interval.

The default value is set to 30 ms to be safe.

Default value:

- 30 if SOC_USB_OTG_SUPPORTED

CONFIG_USB_HOST_SET_ADDR_RECOVERY_MS

SetAddress() recovery time in ms

Found in: [Component config](#) > [USB-OTG](#) > [Hub Driver Configuration](#) > [Root Port configuration](#)

"After successful completion of the Status stage, the device is allowed a SetAddress() recovery interval of 2 ms. At the end of this interval, the device must be able to accept Setup packets addressed to the new address. Also, at the end of the recovery interval, the device must not respond to tokens sent to the old address (unless, of course, the old and new address is the same)." See USB 2.0 chapter 9.2.6.3 for more details.

The default value is set to 10 ms to be safe.

Default value:

- 10 if SOC_USB_OTG_SUPPORTED

CONFIG_USB_HOST_HUBS_SUPPORTED

Support Hubs

Found in: [Component config](#) > [USB-OTG](#) > [Hub Driver Configuration](#)

Enables support of external Hubs.

Default value:

- No (disabled) if SOC_USB_OTG_SUPPORTED

CONFIG_USB_HOST_HUB_MULTI_LEVEL

Support multiple Hubs

Found in: [Component config](#) > [USB-OTG](#) > [Hub Driver Configuration](#) > [CONFIG_USB_HOST_HUBS_SUPPORTED](#)

Enables support for connecting multiple Hubs simultaneously.

Default value:

- Yes (enabled) if [CONFIG_USB_HOST_HUBS_SUPPORTED](#) && SOC_USB_OTG_SUPPORTED

Downstream Port configuration Contains:

- [CONFIG_USB_HOST_EXT_PORT_CUSTOM_POWER_ON_DELAY_ENABLE](#)
- [CONFIG_USB_HOST_EXT_PORT_RESET_RECOVERY_DELAY_MS](#)

CONFIG_USB_HOST_EXT_PORT_RESET_RECOVERY_DELAY_MS

Reset recovery delay in ms

Found in: [Component config](#) > [USB-OTG](#) > [Hub Driver Configuration](#) > [CONFIG_USB_HOST_HUBS_SUPPORTED](#) > [Downstream Port configuration](#)

After a port stops driving the reset signal, the USB 2.0 specification requires that the "USB System Software guarantees a minimum of 10 ms for reset recovery" before the attached device is expected to respond to data transfers (see USB 2.0 chapter 7.1.7.3 for more details). The device may ignore any data transfers during the recovery interval.

The default value is set to 30 ms to be safe.

Default value:

- 30 if [CONFIG_USB_HOST_HUBS_SUPPORTED](#) && SOC_USB_OTG_SUPPORTED

CONFIG_USB_HOST_EXT_PORT_CUSTOM_POWER_ON_DELAY_ENABLE

Custom bPwrOn2PwrGood value

Found in: [Component config](#) > [USB-OTG](#) > [Hub Driver Configuration](#) > [CONFIG_USB_HOST_HUBS_SUPPORTED](#) > [Downstream Port configuration](#)

Enables the possibility to configure custom time for the power-on sequence on a port until power is good on that port.

When enabled, applies the custom PwrOn2PwrGood delay. When disabled, applies the PwrOn2PwrGood value from the Hub Descriptor.

Default value:

- No (disabled) if [CONFIG_USB_HOST_HUBS_SUPPORTED](#) && SOC_USB_OTG_SUPPORTED

CONFIG_USB_HOST_EXT_PORT_CUSTOM_POWER_ON_DELAY_MS

PwrOn2PwrGood delay in ms

Found in: [Component config](#) > [USB-OTG](#) > [Hub Driver Configuration](#) > [CONFIG_USB_HOST_HUBS_SUPPORTED](#) > [Downstream Port configuration](#) > [CONFIG_USB_HOST_EXT_PORT_CUSTOM_POWER_ON_DELAY_ENABLE](#)

Custom value of delay from the time the power-on sequence begins on a port until power is good on that port. Value 0 is used for a hub with no power switches.

The default value is 100 ms.

Range:

- from 0 to 5000 if `CONFIG_USB_HOST_EXT_PORT_CUSTOM_POWER_ON_DELAY_ENABLE` && `CONFIG_USB_HOST_HUBS_SUPPORTED` && `SOC_USB_OTG_SUPPORTED`

Default value:

- 100 if `CONFIG_USB_HOST_EXT_PORT_CUSTOM_POWER_ON_DELAY_ENABLE` && `CONFIG_USB_HOST_HUBS_SUPPORTED` && `SOC_USB_OTG_SUPPORTED`

CONFIG_USB_HOST_ENABLE_ENUM_FILTER_CALLBACK

Enable enumeration filter callback

Found in: [Component config](#) > [USB-OTG](#)

The enumeration filter callback is called before enumeration of each newly attached device. This callback allows users to control whether a device should be enumerated, and what configuration number to use when enumerating a device.

If enabled, the enumeration filter callback can be set via 'usb_host_config_t' when calling 'usb_host_install()'.

Default value:

- No (disabled) if `SOC_USB_OTG_SUPPORTED`

Virtual file system Contains:

- [CONFIG_VFS_SUPPORT_IO](#)

CONFIG_VFS_SUPPORT_IO

Provide basic I/O functions

Found in: [Component config](#) > [Virtual file system](#)

If enabled, the following functions are provided by the VFS component.

open, close, read, write, pread, pwrite, lseek, fstat, fsync, ioctl, fcntl

Filesystem drivers can then be registered to handle these functions for specific paths.

Disabling this option can save memory when the support for these functions is not required.

Note that the following functions can still be used with socket file descriptors when this option is disabled:

close, read, write, ioctl, fcntl.

Default value:

- Yes (enabled)

CONFIG_VFS_SUPPORT_DIR

Provide directory related functions

Found in: [Component config](#) > [Virtual file system](#) > [CONFIG_VFS_SUPPORT_IO](#)

If enabled, the following functions are provided by the VFS component.

stat, link, unlink, rename, utime, access, truncate, rmdir, mkdir, opendir, closedir, readdir, readdir_r, seekdir, telldir, rewinddir

Filesystem drivers can then be registered to handle these functions for specific paths.

Disabling this option can save memory when the support for these functions is not required.

Default value:

- Yes (enabled)

CONFIG_VFS_SUPPORT_SELECT

Provide select function

Found in: [Component config](#) > [Virtual file system](#) > [CONFIG_VFS_SUPPORT_IO](#)

If enabled, select function is provided by the VFS component, and can be used on peripheral file descriptors (such as UART) and sockets at the same time.

If disabled, the default select implementation will be provided by LWIP for sockets only.

Disabling this option can reduce code size if support for "select" on UART file descriptors is not required.

CONFIG_VFS_SUPPRESS_SELECT_DEBUG_OUTPUT

Suppress select() related debug outputs

Found in: [Component config](#) > [Virtual file system](#) > [CONFIG_VFS_SUPPORT_IO](#) > [CONFIG_VFS_SUPPORT_SELECT](#)

Select() related functions might produce an unconveniently lot of debug outputs when one sets the default log level to DEBUG or higher. It is possible to suppress these debug outputs by enabling this option.

Default value:

- Yes (enabled)

CONFIG_VFS_SELECT_IN_RAM

Make VFS driver select() callbacks IRAM-safe

Found in: [Component config](#) > [Virtual file system](#) > [CONFIG_VFS_SUPPORT_IO](#) > [CONFIG_VFS_SUPPORT_SELECT](#)

If enabled, VFS driver select() callback function will be placed in IRAM.

Default value:

- No (disabled)

CONFIG_VFS_SUPPORT_TERMIOS

Provide termios.h functions

Found in: [Component config](#) > [Virtual file system](#) > [CONFIG_VFS_SUPPORT_IO](#)

Disabling this option can save memory when the support for termios.h is not required.

Default value:

- Yes (enabled)

CONFIG_VFS_MAX_COUNT

Maximum Number of Virtual Filesystems

Found in: [Component config](#) > [Virtual file system](#) > [CONFIG_VFS_SUPPORT_IO](#)

Define maximum number of virtual filesystems that can be registered.

Range:

- from 1 to 20

Default value:

- 8

Host File System I/O (Semihosting) Contains:

- [CONFIG_VFS_SEMIHOSTFS_MAX_MOUNT_POINTS](#)

CONFIG_VFS_SEMIHOSTFS_MAX_MOUNT_POINTS

Host FS: Maximum number of the host filesystem mount points

Found in: [Component config](#) > [Virtual file system](#) > [CONFIG_VFS_SUPPORT_IO](#) > [Host File System I/O \(Semihosting\)](#)

Define maximum number of host filesystem mount points.

Default value:

- 1

Wear Levelling Contains:

- [CONFIG_WL_SECTOR_MODE](#)
- [CONFIG_WL_SECTOR_SIZE](#)

CONFIG_WL_SECTOR_SIZE

Wear Levelling library sector size

Found in: [Component config](#) > [Wear Levelling](#)

Sector size used by wear levelling library. You can set default sector size or size that will fit to the flash device sector size.

With sector size set to 4096 bytes, wear levelling library is more efficient. However if FAT filesystem is used on top of wear levelling library, it will need more temporary storage: 4096 bytes for each mounted filesystem and 4096 bytes for each opened file.

With sector size set to 512 bytes, wear levelling library will perform more operations with flash memory, but less RAM will be used by FAT filesystem library (512 bytes for the filesystem and 512 bytes for each file opened).

Available options:

- 512 (CONFIG_WL_SECTOR_SIZE_512)
- 4096 (CONFIG_WL_SECTOR_SIZE_4096)

CONFIG_WL_SECTOR_MODE

Sector store mode

Found in: [Component config](#) > [Wear Levelling](#)

Specify the mode to store data into flash:

- In Performance mode a data will be stored to the RAM and then stored back to the flash. Compared to the Safety mode, this operation is faster, but if power will be lost when erase sector operation is in progress, then the data from complete flash device sector will be lost.
- In Safety mode data from complete flash device sector will be read from flash, modified, and then stored back to flash. Compared to the Performance mode, this operation is slower, but if power is lost during erase sector operation, then the data from full flash device sector will not be lost.

Available options:

- Performance (CONFIG_WL_SECTOR_MODE_PERF)
- Safety (CONFIG_WL_SECTOR_MODE_SAFE)

Wi-Fi Provisioning Manager Contains:

- [*CONFIG_WIFI_PROV_BLE_NOTIFY*](#)
- [*CONFIG_WIFI_PROV_BLE_BONDING*](#)
- [*CONFIG_WIFI_PROV_BLE_SEC_CONN*](#)
- [*CONFIG_WIFI_PROV_BLE_FORCE_ENCRYPTION*](#)
- [*CONFIG_WIFI_PROV_KEEP_BLE_ON_AFTER_PROV*](#)
- [*CONFIG_WIFI_PROV_SCAN_MAX_ENTRIES*](#)
- [*CONFIG_WIFI_PROV_AUTOSTOP_TIMEOUT*](#)
- [*CONFIG_WIFI_PROV_STA_SCAN_METHOD*](#)

CONFIG_WIFI_PROV_SCAN_MAX_ENTRIES

Max Wi-Fi Scan Result Entries

Found in: [Component config](#) > [Wi-Fi Provisioning Manager](#)

This sets the maximum number of entries of Wi-Fi scan results that will be kept by the provisioning manager

Range:

- from 1 to 255

Default value:

- 16

CONFIG_WIFI_PROV_AUTOSTOP_TIMEOUT

Provisioning auto-stop timeout

Found in: [Component config](#) > [Wi-Fi Provisioning Manager](#)

Time (in seconds) after which the Wi-Fi provisioning manager will auto-stop after connecting to a Wi-Fi network successfully.

Range:

- from 5 to 600

Default value:

- 30

CONFIG_WIFI_PROV_BLE_BONDING

Enable BLE bonding

Found in: [Component config](#) > [Wi-Fi Provisioning Manager](#)

This option is applicable only when provisioning transport is BLE.

CONFIG_WIFI_PROV_BLE_SEC_CONN

Enable BLE Secure connection flag

Found in: [Component config](#) > [Wi-Fi Provisioning Manager](#)

Used to enable Secure connection support when provisioning transport is BLE.

Default value:

- Yes (enabled) if [*CONFIG_BT_NIMBLE_ENABLED*](#)

CONFIG_WIFI_PROV_BLE_FORCE_ENCRYPTION

Force Link Encryption during characteristic Read / Write

Found in: [Component config](#) > [Wi-Fi Provisioning Manager](#)

Used to enforce link encryption when attempting to read / write characteristic

CONFIG_WIFI_PROV_BLE_NOTIFY

Add support for Notification for provisioning BLE descriptors

Found in: [Component config](#) > [Wi-Fi Provisioning Manager](#)

Used to enable support Notification in BLE descriptors of prov* characteristics

CONFIG_WIFI_PROV_KEEP_BLE_ON_AFTER_PROV

Keep BT on after provisioning is done

Found in: [Component config](#) > [Wi-Fi Provisioning Manager](#)

CONFIG_WIFI_PROV_DISCONNECT_AFTER_PROV

Terminate connection after provisioning is done

Found in: [Component config](#) > [Wi-Fi Provisioning Manager](#) > [CONFIG_WIFI_PROV_KEEP_BLE_ON_AFTER_PROV](#)

Default value:

- Yes (enabled) if [CONFIG_WIFI_PROV_KEEP_BLE_ON_AFTER_PROV](#)

CONFIG_WIFI_PROV_STA_SCAN_METHOD

Wifi Provisioning Scan Method

Found in: [Component config](#) > [Wi-Fi Provisioning Manager](#)

Available options:

- All Channel Scan (CONFIG_WIFI_PROV_STA_ALL_CHANNEL_SCAN)
Scan will end after scanning the entire channel. This option is useful in Mesh WiFi Systems.
- Fast Scan (CONFIG_WIFI_PROV_STA_FAST_SCAN)
Scan will end after an AP matching with the SSID has been detected.

CONFIG_IDF_EXPERIMENTAL_FEATURES

Make experimental features visible

Found in:

By enabling this option, ESP-IDF experimental feature options will be visible.

Note you should still enable a certain experimental feature option to use it, and you should read the corresponding risk warning and known issue list carefully.

Current experimental feature list:

- CONFIG_ESPTOOLPY_FLASHFREQ_120M && CONFIG_ESPTOOLPY_FLASH_SAMPLE_MODE_DTR
- CONFIG_SPIRAM_SPEED_120M && CONFIG_SPIRAM_MODE_OCT
- CONFIG_BOOTLOADER_CACHE_32BIT_ADDR_QUAD_FLASH
- CONFIG_USB_HOST_EXT_PORT_RESET_ATTEMPTS

Default value:

- No (disabled)

Deprecated options and their replacements

- CONFIG_A2DP_ENABLE ([CONFIG_BT_A2DP_ENABLE](#))
- CONFIG_A2D_INITIAL_TRACE_LEVEL ([CONFIG_BT_LOG_A2D_TRACE_LEVEL](#))
 - CONFIG_A2D_TRACE_LEVEL_NONE
 - CONFIG_A2D_TRACE_LEVEL_ERROR
 - CONFIG_A2D_TRACE_LEVEL_WARNING
 - CONFIG_A2D_TRACE_LEVEL_API
 - CONFIG_A2D_TRACE_LEVEL_EVENT
 - CONFIG_A2D_TRACE_LEVEL_DEBUG
 - CONFIG_A2D_TRACE_LEVEL_VERBOSE
- CONFIG_ADC2_DISABLE_DAC ([CONFIG_ADC_DISABLE_DAC](#))
- CONFIG_APPL_INITIAL_TRACE_LEVEL ([CONFIG_BT_LOG_APPL_TRACE_LEVEL](#))
 - CONFIG_APPL_TRACE_LEVEL_NONE
 - CONFIG_APPL_TRACE_LEVEL_ERROR
 - CONFIG_APPL_TRACE_LEVEL_WARNING
 - CONFIG_APPL_TRACE_LEVEL_API
 - CONFIG_APPL_TRACE_LEVEL_EVENT
 - CONFIG_APPL_TRACE_LEVEL_DEBUG
 - CONFIG_APPL_TRACE_LEVEL_VERBOSE
- CONFIG_APP_ANTI_ROLLBACK ([CONFIG_BOOTLOADER_APP_ANTI_ROLLBACK](#))
- CONFIG_APP_ROLLBACK_ENABLE ([CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE](#))
- CONFIG_APP_SECURE_VERSION ([CONFIG_BOOTLOADER_APP_SECURE_VERSION](#))
- CONFIG_APP_SECURE_VERSION_SIZE_EFUSE_FIELD ([CONFIG_BOOTLOADER_APP_SEC_VER_SIZE_EFUSE_FIELD](#))
- CONFIG_AVCT_INITIAL_TRACE_LEVEL ([CONFIG_BT_LOG_AVCT_TRACE_LEVEL](#))
 - CONFIG_AVCT_TRACE_LEVEL_NONE
 - CONFIG_AVCT_TRACE_LEVEL_ERROR
 - CONFIG_AVCT_TRACE_LEVEL_WARNING
 - CONFIG_AVCT_TRACE_LEVEL_API
 - CONFIG_AVCT_TRACE_LEVEL_EVENT
 - CONFIG_AVCT_TRACE_LEVEL_DEBUG
 - CONFIG_AVCT_TRACE_LEVEL_VERBOSE
- CONFIG_AVDT_INITIAL_TRACE_LEVEL ([CONFIG_BT_LOG_AVDT_TRACE_LEVEL](#))
 - CONFIG_AVDT_TRACE_LEVEL_NONE
 - CONFIG_AVDT_TRACE_LEVEL_ERROR
 - CONFIG_AVDT_TRACE_LEVEL_WARNING
 - CONFIG_AVDT_TRACE_LEVEL_API
 - CONFIG_AVDT_TRACE_LEVEL_EVENT
 - CONFIG_AVDT_TRACE_LEVEL_DEBUG
 - CONFIG_AVDT_TRACE_LEVEL_VERBOSE
- CONFIG_AVRC_INITIAL_TRACE_LEVEL ([CONFIG_BT_LOG_AVRC_TRACE_LEVEL](#))
 - CONFIG_AVRC_TRACE_LEVEL_NONE
 - CONFIG_AVRC_TRACE_LEVEL_ERROR
 - CONFIG_AVRC_TRACE_LEVEL_WARNING
 - CONFIG_AVRC_TRACE_LEVEL_API
 - CONFIG_AVRC_TRACE_LEVEL_EVENT
 - CONFIG_AVRC_TRACE_LEVEL_DEBUG
 - CONFIG_AVRC_TRACE_LEVEL_VERBOSE
- CONFIG_BLE_ACTIVE_SCAN_REPORT_ADV_SCAN_RSP_INDIVIDUALLY ([CONFIG_BT_BLE_ACT_SCAN_REP_ADV_SCAN](#))
- CONFIG_BLE_ESTABLISH_LINK_CONNECTION_TIMEOUT ([CONFIG_BT_BLE_ESTAB_LINK_CONN_TOUT](#))
- CONFIG_BLE_MESH_GATT_PROXY ([CONFIG_BLE_MESH_GATT_PROXY_SERVER](#))
- CONFIG_BLE_SMP_ENABLE ([CONFIG_BT_BLE_SMP_ENABLE](#))

- **CONFIG_BLUEDROID_MEM_DEBUG** ([*CONFIG_BT_BLUEDROID_MEM_DEBUG*](#))
- **CONFIG_BLUEDROID_PINNED_TO_CORE_CHOICE** ([*CONFIG_BT_BLUEDROID_PINNED_TO_CORE_CHOICE*](#))
 - CONFIG_BLUEDROID_PINNED_TO_CORE_0
 - CONFIG_BLUEDROID_PINNED_TO_CORE_1
- **CONFIG_BLUFI_INITIAL_TRACE_LEVEL** ([*CONFIG_BT_LOG_BLUFI_TRACE_LEVEL*](#))
 - CONFIG_BLUFI_TRACE_LEVEL_NONE
 - CONFIG_BLUFI_TRACE_LEVEL_ERROR
 - CONFIG_BLUFI_TRACE_LEVEL_WARNING
 - CONFIG_BLUFI_TRACE_LEVEL_API
 - CONFIG_BLUFI_TRACE_LEVEL_EVENT
 - CONFIG_BLUFI_TRACE_LEVEL_DEBUG
 - CONFIG_BLUFI_TRACE_LEVEL_VERBOSE
- **CONFIG_BNEP_INITIAL_TRACE_LEVEL** ([*CONFIG_BT_LOG_BNEP_TRACE_LEVEL*](#))
- **CONFIG_BTC_INITIAL_TRACE_LEVEL** ([*CONFIG_BT_LOG_BTC_TRACE_LEVEL*](#))
 - CONFIG_BTC_TRACE_LEVEL_NONE
 - CONFIG_BTC_TRACE_LEVEL_ERROR
 - CONFIG_BTC_TRACE_LEVEL_WARNING
 - CONFIG_BTC_TRACE_LEVEL_API
 - CONFIG_BTC_TRACE_LEVEL_EVENT
 - CONFIG_BTC_TRACE_LEVEL_DEBUG
 - CONFIG_BTC_TRACE_LEVEL_VERBOSE
- **CONFIG_BTC_TASK_STACK_SIZE** ([*CONFIG_BT_BTC_TASK_STACK_SIZE*](#))
- **CONFIG_BTH_LOG_SDP_INITIAL_TRACE_LEVEL** ([*CONFIG_BT_LOG_SDP_TRACE_LEVEL*](#))
 - CONFIG_SDP_TRACE_LEVEL_NONE
 - CONFIG_SDP_TRACE_LEVEL_ERROR
 - CONFIG_SDP_TRACE_LEVEL_WARNING
 - CONFIG_SDP_TRACE_LEVEL_API
 - CONFIG_SDP_TRACE_LEVEL_EVENT
 - CONFIG_SDP_TRACE_LEVEL_DEBUG
 - CONFIG_SDP_TRACE_LEVEL_VERBOSE
- **CONFIG_BTIF_INITIAL_TRACE_LEVEL** ([*CONFIG_BT_LOG_BTIF_TRACE_LEVEL*](#))
 - CONFIG_BTIF_TRACE_LEVEL_NONE
 - CONFIG_BTIF_TRACE_LEVEL_ERROR
 - CONFIG_BTIF_TRACE_LEVEL_WARNING
 - CONFIG_BTIF_TRACE_LEVEL_API
 - CONFIG_BTIF_TRACE_LEVEL_EVENT
 - CONFIG_BTIF_TRACE_LEVEL_DEBUG
 - CONFIG_BTIF_TRACE_LEVEL_VERBOSE
- **CONFIG_BTM_INITIAL_TRACE_LEVEL** ([*CONFIG_BT_LOG_BTM_TRACE_LEVEL*](#))
 - CONFIG_BTM_TRACE_LEVEL_NONE
 - CONFIG_BTM_TRACE_LEVEL_ERROR
 - CONFIG_BTM_TRACE_LEVEL_WARNING
 - CONFIG_BTM_TRACE_LEVEL_API
 - CONFIG_BTM_TRACE_LEVEL_EVENT
 - CONFIG_BTM_TRACE_LEVEL_DEBUG
 - CONFIG_BTM_TRACE_LEVEL_VERBOSE
- **CONFIG_BTU_TASK_STACK_SIZE** ([*CONFIG_BT_BTU_TASK_STACK_SIZE*](#))
- **CONFIG_BT_NIMBLE_ACL_BUF_COUNT** ([*CONFIG_BT_NIMBLE_TRANSPORT_ACL_FROM_LL_COUNT*](#))
- **CONFIG_BT_NIMBLE_ACL_BUF_SIZE** ([*CONFIG_BT_NIMBLE_TRANSPORT_ACL_SIZE*](#))
- **CONFIG_BT_NIMBLE_HCI_EVT_BUF_SIZE** ([*CONFIG_BT_NIMBLE_TRANSPORT_EVT_SIZE*](#))
- **CONFIG_BT_NIMBLE_HCI_EVT_HI_BUF_COUNT** ([*CONFIG_BT_NIMBLE_TRANSPORT_EVT_COUNT*](#))
- **CONFIG_BT_NIMBLE_HCI_EVT_LO_BUF_COUNT** ([*CONFIG_BT_NIMBLE_TRANSPORT_EVT_DISCARD_COUNT*](#))
- **CONFIG_BT_NIMBLE_MSYS1_BLOCK_COUNT** ([*CONFIG_BT_NIMBLE_MSYS_1_BLOCK_COUNT*](#))
- **CONFIG_BT_NIMBLE_TASK_STACK_SIZE** ([*CONFIG_BT_NIMBLE_HOST_TASK_STACK_SIZE*](#))
- **CONFIG_CLASSIC_BT_ENABLED** ([*CONFIG_BT_CLASSIC_ENABLED*](#))
- **CONFIG_CONSOLE_UART** ([*CONFIG_ESP_CONSOLE_UART*](#))
 - CONFIG_CONSOLE_UART_DEFAULT

- CONFIG_CONSOLE_UART_CUSTOM
- CONFIG_CONSOLE_UART_NONE, CONFIG_ESP_CONSOLE_UART_NONE
- CONFIG_CONSOLE_UART_BAUDRATE ([CONFIG_ESP_CONSOLE_UART_BAUDRATE](#))
- **CONFIG_CONSOLE_UART_NUM** ([CONFIG_ESP_CONSOLE_UART_NUM](#))
 - CONFIG_CONSOLE_UART_CUSTOM_NUM_0
 - CONFIG_CONSOLE_UART_CUSTOM_NUM_1
- CONFIG_CONSOLE_UART_RX_GPIO ([CONFIG_ESP_CONSOLE_UART_RX_GPIO](#))
- CONFIG_CONSOLE_UART_TX_GPIO ([CONFIG_ESP_CONSOLE_UART_TX_GPIO](#))
- CONFIG_CXX_EXCEPTIONS ([CONFIG_COMPILER_CXX_EXCEPTIONS](#))
- CONFIG_CXX_EXCEPTIONS_EMG_POOL_SIZE ([CONFIG_COMPILER_CXX_EXCEPTIONS_EMG_POOL_SIZE](#))
- CONFIG_EFUSE_SECURE_VERSION_EMULATE ([CONFIG_BOOTLOADER_EFUSE_SECURE_VERSION_EMULATE](#))
- CONFIG_ENABLE_STATIC_TASK_CLEAN_UP_HOOK ([CONFIG_FREERTOS_ENABLE_STATIC_TASK_CLEAN_UP](#))
- CONFIG_ESP32_APPTRACE_ONPANIC_HOST_FLUSH_TMO (CON-
[FIG_APPTRACE_ONPANIC_HOST_FLUSH_TMO](#))
- CONFIG_ESP32_APPTRACE_PENDING_DATA_SIZE_MAX ([CONFIG_APPTRACE_PENDING_DATA_SIZE_MAX](#))
- CONFIG_ESP32_APPTRACE_POSTMORTEM_FLUSH_TRAX_THRESH (CON-
[FIG_APPTRACE_POSTMORTEM_FLUSH_THRESH](#))
- **CONFIG_ESP32_CORE_DUMP_DECODE** ([CONFIG_ESP_COREDUMP_DECODE](#))
 - CONFIG_ESP32_CORE_DUMP_DECODE_INFO
 - CONFIG_ESP32_CORE_DUMP_DECODE_DISABLE
- CONFIG_ESP32_CORE_DUMP_MAX_TASKS_NUM ([CONFIG_ESP_COREDUMP_MAX_TASKS_NUM](#))
- CONFIG_ESP32_CORE_DUMP_STACK_SIZE ([CONFIG_ESP_COREDUMP_STACK_SIZE](#))
- CONFIG_ESP32_CORE_DUMP_UART_DELAY ([CONFIG_ESP_COREDUMP_UART_DELAY](#))
- CONFIG_ESP32_DEBUG_STUBS_ENABLE ([CONFIG_ESP_DEBUG_STUBS_ENABLE](#))
- CONFIG_ESP32_GCOV_ENABLE ([CONFIG_APPTRACE_GCOV_ENABLE](#))
- CONFIG_ESP32_PTHREAD_STACK_MIN ([CONFIG_PTHREAD_STACK_MIN](#))
- **CONFIG_ESP32_PTHREAD_TASK_CORE_DEFAULT** ([CONFIG_PTHREAD_TASK_CORE_DEFAULT](#))
 - CONFIG_ESP32_DEFAULT_PTHREAD_CORE_NO_AFFINITY
 - CONFIG_ESP32_DEFAULT_PTHREAD_CORE_0
 - CONFIG_ESP32_DEFAULT_PTHREAD_CORE_1
- CONFIG_ESP32_PTHREAD_TASK_NAME_DEFAULT ([CONFIG_PTHREAD_TASK_NAME_DEFAULT](#))
- CONFIG_ESP32_PTHREAD_TASK_PRIO_DEFAULT ([CONFIG_PTHREAD_TASK_PRIO_DEFAULT](#))
- CONFIG_ESP32_PTHREAD_TASK_STACK_SIZE_DEFAULT ([CONFIG_PTHREAD_TASK_STACK_SIZE_DEFAULT](#))
- CONFIG_ESP32_RTC_XTAL_BOOTSTRAP_CYCLES ([CONFIG_ESP_SYSTEM_RTC_EXT_XTAL_BOOTSTRAP_CYCLES](#))
- CONFIG_ESP32_WIFI_AMPDU_RX_ENABLED ([CONFIG_ESP_WIFI_AMPDU_RX_ENABLED](#))
- CONFIG_ESP32_WIFI_AMPDU_TX_ENABLED ([CONFIG_ESP_WIFI_AMPDU_TX_ENABLED](#))
- CONFIG_ESP32_WIFI_AMSDU_TX_ENABLED ([CONFIG_ESP_WIFI_AMSDU_TX_ENABLED](#))
- CONFIG_ESP32_WIFI_CACHE_TX_BUFFER_NUM ([CONFIG_ESP_WIFI_CACHE_TX_BUFFER_NUM](#))
- CONFIG_ESP32_WIFI_CSI_ENABLED ([CONFIG_ESP_WIFI_CSI_ENABLED](#))
- CONFIG_ESP32_WIFI_DYNAMIC_RX_BUFFER_NUM ([CONFIG_ESP_WIFI_DYNAMIC_RX_BUFFER_NUM](#))
- CONFIG_ESP32_WIFI_DYNAMIC_TX_BUFFER_NUM ([CONFIG_ESP_WIFI_DYNAMIC_TX_BUFFER_NUM](#))
- CONFIG_ESP32_WIFI_ENABLE_WPA3_OWE_STA ([CONFIG_ESP_WIFI_ENABLE_WPA3_OWE_STA](#))
- CONFIG_ESP32_WIFI_ENABLE_WPA3_SAE ([CONFIG_ESP_WIFI_ENABLE_WPA3_SAE](#))
- CONFIG_ESP32_WIFI_IRAM_OPT ([CONFIG_ESP_WIFI_IRAM_OPT](#))
- CONFIG_ESP32_WIFI_MGMT_SBUF_NUM ([CONFIG_ESP_WIFI_MGMT_SBUF_NUM](#))
- CONFIG_ESP32_WIFI_NVS_ENABLED ([CONFIG_ESP_WIFI_NVS_ENABLED](#))
- CONFIG_ESP32_WIFI_RX_BA_WIN ([CONFIG_ESP_WIFI_RX_BA_WIN](#))
- CONFIG_ESP32_WIFI_RX_IRAM_OPT ([CONFIG_ESP_WIFI_RX_IRAM_OPT](#))
- CONFIG_ESP32_WIFI_SOFTAP_BEACON_MAX_LEN ([CONFIG_ESP_WIFI_SOFTAP_BEACON_MAX_LEN](#))
- CONFIG_ESP32_WIFI_STATIC_RX_BUFFER_NUM ([CONFIG_ESP_WIFI_STATIC_RX_BUFFER_NUM](#))
- CONFIG_ESP32_WIFI_STATIC_TX_BUFFER_NUM ([CONFIG_ESP_WIFI_STATIC_TX_BUFFER_NUM](#))
- CONFIG_ESP32_WIFI_SW_COEXIST_ENABLE ([CONFIG_ESP_COEX_SW_COEXIST_ENABLE](#))
- **CONFIG_ESP32_WIFI_TASK_CORE_ID** ([CONFIG_ESP_WIFI_TASK_CORE_ID](#))
 - CONFIG_ESP32_WIFI_TASK_PINNED_TO_CORE_0
 - CONFIG_ESP32_WIFI_TASK_PINNED_TO_CORE_1
- CONFIG_ESP32_WIFI_TX_BA_WIN ([CONFIG_ESP_WIFI_TX_BA_WIN](#))
- **CONFIG_ESP32_WIFI_TX_BUFFER** ([CONFIG_ESP_WIFI_TX_BUFFER](#))

- CONFIG_ESP32_WIFI_STATIC_TX_BUFFER
 - CONFIG_ESP32_WIFI_DYNAMIC_TX_BUFFER
- CONFIG_ESP_GRATUITOUS_ARP ([CONFIG_LWIP_ESP_GRATUITOUS_ARP](#))
- CONFIG_ESP_SYSTEM_PD_FLASH ([CONFIG_ESP_SLEEP_POWER_DOWN_FLASH](#))
- CONFIG_ESP_SYSTEM_PM_POWER_DOWN_CPU ([CONFIG_PM_POWER_DOWN_CPU_IN_LIGHT_SLEEP](#))
- CONFIG_ESP_TASK_WDT ([CONFIG_ESP_TASK_WDT_INIT](#))
- CONFIG_ESP_WIFI_EXTERNAL_COEXIST_ENABLE ([CONFIG_ESP_COEX_EXTERNAL_COEXIST_ENABLE](#))
- CONFIG_ESP_WIFI_SW_COEXIST_ENABLE ([CONFIG_ESP_COEX_SW_COEXIST_ENABLE](#))
- CONFIG_EVENT_LOOP_PROFILING ([CONFIG_ESP_EVENT_LOOP_PROFILING](#))
- CONFIG_EXTERNAL_COEX_ENABLE ([CONFIG_ESP_COEX_EXTERNAL_COEXIST_ENABLE](#))
- CONFIG_FLASH_ENCRYPTION_ENABLED ([CONFIG_SECURE_FLASH_ENC_ENABLED](#))
- CONFIG_FLASH_ENCRYPTION_UART_BOOTLOADER_ALLOW_CACHE (CON-
[FIG_SECURE_FLASH_UART_BOOTLOADER_ALLOW_CACHE](#))
- CONFIG_FLASH_ENCRYPTION_UART_BOOTLOADER_ALLOW_ENCRYPT (CON-
[FIG_SECURE_FLASH_UART_BOOTLOADER_ALLOW_ENC](#))
- **CONFIG_GAP_INITIAL_TRACE_LEVEL** ([CONFIG_BT_LOG_GAP_TRACE_LEVEL](#))
 - CONFIG_GAP_TRACE_LEVEL_NONE
 - CONFIG_GAP_TRACE_LEVEL_ERROR
 - CONFIG_GAP_TRACE_LEVEL_WARNING
 - CONFIG_GAP_TRACE_LEVEL_API
 - CONFIG_GAP_TRACE_LEVEL_EVENT
 - CONFIG_GAP_TRACE_LEVEL_DEBUG
 - CONFIG_GAP_TRACE_LEVEL_VERBOSE
- CONFIG_GARP_TMR_INTERVAL ([CONFIG_LWIP_GARP_TMR_INTERVAL](#))
- CONFIG_GATTC_CACHE_NVS_FLASH ([CONFIG_BT_GATTC_CACHE_NVS_FLASH](#))
- CONFIG_GATTC_ENABLE ([CONFIG_BT_GATTC_ENABLE](#))
- CONFIG_GATTS_ENABLE ([CONFIG_BT_GATTS_ENABLE](#))
- **CONFIG_GATTS_SEND_SERVICE_CHANGE_MODE** ([CONFIG_BT_GATTS_SEND_SERVICE_CHANGE_MODE](#))
 - CONFIG_GATTS_SEND_SERVICE_CHANGE_MANUAL
 - CONFIG_GATTS_SEND_SERVICE_CHANGE_AUTO
- **CONFIG_GATT_INITIAL_TRACE_LEVEL** ([CONFIG_BT_LOG_GATT_TRACE_LEVEL](#))
 - CONFIG_GATT_TRACE_LEVEL_NONE
 - CONFIG_GATT_TRACE_LEVEL_ERROR
 - CONFIG_GATT_TRACE_LEVEL_WARNING
 - CONFIG_GATT_TRACE_LEVEL_API
 - CONFIG_GATT_TRACE_LEVEL_EVENT
 - CONFIG_GATT_TRACE_LEVEL_DEBUG
 - CONFIG_GATT_TRACE_LEVEL_VERBOSE
- CONFIG_GDBSTUB_MAX_TASKS ([CONFIG_ESP_GDBSTUB_MAX_TASKS](#))
- CONFIG_GDBSTUB_SUPPORT_TASKS ([CONFIG_ESP_GDBSTUB_SUPPORT_TASKS](#))
- **CONFIG_HCI_INITIAL_TRACE_LEVEL** ([CONFIG_BT_LOG_HCI_TRACE_LEVEL](#))
 - CONFIG_HCI_TRACE_LEVEL_NONE
 - CONFIG_HCI_TRACE_LEVEL_ERROR
 - CONFIG_HCI_TRACE_LEVEL_WARNING
 - CONFIG_HCI_TRACE_LEVEL_API
 - CONFIG_HCI_TRACE_LEVEL_EVENT
 - CONFIG_HCI_TRACE_LEVEL_DEBUG
 - CONFIG_HCI_TRACE_LEVEL_VERBOSE
- CONFIG_HFP_AG_ENABLE ([CONFIG_BT_HFP_AG_ENABLE](#))
- **CONFIG_HFP_AUDIO_DATA_PATH** ([CONFIG_BT_HFP_AUDIO_DATA_PATH](#))
 - CONFIG_HFP_AUDIO_DATA_PATH_PCM
 - CONFIG_HFP_AUDIO_DATA_PATH_HCI
- CONFIG_HFP_CLIENT_ENABLE ([CONFIG_BT_HFP_CLIENT_ENABLE](#))
- CONFIG_HFP_ENABLE ([CONFIG_BT_HFP_ENABLE](#))
- **CONFIG_HID_INITIAL_TRACE_LEVEL** ([CONFIG_BT_LOG_HID_TRACE_LEVEL](#))
 - CONFIG_HID_TRACE_LEVEL_NONE
 - CONFIG_HID_TRACE_LEVEL_ERROR

- CONFIG_HID_TRACE_LEVEL_WARNING
 - CONFIG_HID_TRACE_LEVEL_API
 - CONFIG_HID_TRACE_LEVEL_EVENT
 - CONFIG_HID_TRACE_LEVEL_DEBUG
 - CONFIG_HID_TRACE_LEVEL_VERBOSE
- CONFIG_INT_WDT ([CONFIG_ESP_INT_WDT](#))
- CONFIG_INT_WDT_CHECK_CPU1 ([CONFIG_ESP_INT_WDT_CHECK_CPU1](#))
- CONFIG_INT_WDT_TIMEOUT_MS ([CONFIG_ESP_INT_WDT_TIMEOUT_MS](#))
- CONFIG_IPC_TASK_STACK_SIZE ([CONFIG_ESP_IPC_TASK_STACK_SIZE](#))
- **CONFIG_L2CAP_INITIAL_TRACE_LEVEL** ([CONFIG_BT_LOG_L2CAP_TRACE_LEVEL](#))
 - CONFIG_L2CAP_TRACE_LEVEL_NONE
 - CONFIG_L2CAP_TRACE_LEVEL_ERROR
 - CONFIG_L2CAP_TRACE_LEVEL_WARNING
 - CONFIG_L2CAP_TRACE_LEVEL_API
 - CONFIG_L2CAP_TRACE_LEVEL_EVENT
 - CONFIG_L2CAP_TRACE_LEVEL_DEBUG
 - CONFIG_L2CAP_TRACE_LEVEL_VERBOSE
- CONFIG_L2_TO_L3_COPY ([CONFIG_LWIP_L2_TO_L3_COPY](#))
- **CONFIG_LOG_BOOTLOADER_LEVEL** ([CONFIG_BOOTLOADER_LOG_LEVEL](#))
 - CONFIG_LOG_BOOTLOADER_LEVEL_NONE
 - CONFIG_LOG_BOOTLOADER_LEVEL_ERROR
 - CONFIG_LOG_BOOTLOADER_LEVEL_WARN
 - CONFIG_LOG_BOOTLOADER_LEVEL_INFO
 - CONFIG_LOG_BOOTLOADER_LEVEL_DEBUG
 - CONFIG_LOG_BOOTLOADER_LEVEL_VERBOSE
- CONFIG_MAIN_TASK_STACK_SIZE ([CONFIG_ESP_MAIN_TASK_STACK_SIZE](#))
- **CONFIG_MCA_INITIAL_TRACE_LEVEL** ([CONFIG_BT_LOG_MCA_TRACE_LEVEL](#))
 - CONFIG_MCA_TRACE_LEVEL_NONE
 - CONFIG_MCA_TRACE_LEVEL_ERROR
 - CONFIG_MCA_TRACE_LEVEL_WARNING
 - CONFIG_MCA_TRACE_LEVEL_API
 - CONFIG_MCA_TRACE_LEVEL_EVENT
 - CONFIG_MCA_TRACE_LEVEL_DEBUG
 - CONFIG_MCA_TRACE_LEVEL_VERBOSE
- CONFIG_MCPWM_ISR_IN_IRAM ([CONFIG_MCPWM_ISR_IRAM_SAFE](#))
- CONFIG_NIMBLE_ATT_PREFERRED_MTU ([CONFIG_BT_NIMBLE_ATT_PREFERRED_MTU](#))
- CONFIG_NIMBLE_CRYPTOSTACK_MBEDTLS ([CONFIG_BT_NIMBLE_CRYPTOSTACK_MBEDTLS](#))
- CONFIG_NIMBLE_DEBUG ([CONFIG_BT_NIMBLE_DEBUG](#))
- CONFIG_NIMBLE_GAP_DEVICE_NAME_MAX_LEN ([CONFIG_BT_NIMBLE_GAP_DEVICE_NAME_MAX_LEN](#))
- CONFIG_NIMBLE_HS_FLOW_CTRL ([CONFIG_BT_NIMBLE_HS_FLOW_CTRL](#))
- CONFIG_NIMBLE_HS_FLOW_CTRL_ITVL ([CONFIG_BT_NIMBLE_HS_FLOW_CTRL_ITVL](#))
- CONFIG_NIMBLE_HS_FLOW_CTRL_THRESH ([CONFIG_BT_NIMBLE_HS_FLOW_CTRL_THRESH](#))
- CONFIG_NIMBLE_HS_FLOW_CTRL_TX_ON_DISCONNECT ([CONFIG_BT_NIMBLE_HS_FLOW_CTRL_TX_ON_DISCONNECT](#))
- CONFIG_NIMBLE_L2CAP_COC_MAX_NUM ([CONFIG_BT_NIMBLE_L2CAP_COC_MAX_NUM](#))
- CONFIG_NIMBLE_MAX_BONDS ([CONFIG_BT_NIMBLE_MAX_BONDS](#))
- CONFIG_NIMBLE_MAX_CCCDS ([CONFIG_BT_NIMBLE_MAX_CCCDS](#))
- CONFIG_NIMBLE_MAX_CONNECTIONS ([CONFIG_BT_NIMBLE_MAX_CONNECTIONS](#))
- **CONFIG_NIMBLE_MEM_ALLOC_MODE** ([CONFIG_BT_NIMBLE_MEM_ALLOC_MODE](#))
 - CONFIG_NIMBLE_MEM_ALLOC_MODE_INTERNAL
 - CONFIG_NIMBLE_MEM_ALLOC_MODE_EXTERNAL
 - CONFIG_NIMBLE_MEM_ALLOC_MODE_DEFAULT
- CONFIG_NIMBLE_MESH ([CONFIG_BT_NIMBLE_MESH](#))
- CONFIG_NIMBLE_MESH_DEVICE_NAME ([CONFIG_BT_NIMBLE_MESH_DEVICE_NAME](#))
- CONFIG_NIMBLE_MESH_FRIEND ([CONFIG_BT_NIMBLE_MESH_FRIEND](#))
- CONFIG_NIMBLE_MESH_GATT_PROXY ([CONFIG_BT_NIMBLE_MESH_GATT_PROXY](#))
- CONFIG_NIMBLE_MESH_LOW_POWER ([CONFIG_BT_NIMBLE_MESH_LOW_POWER](#))
- CONFIG_NIMBLE_MESH_PB_ADV ([CONFIG_BT_NIMBLE_MESH_PB_ADV](#))

- **CONFIG_NIMBLE_MESH_PB_GATT** ([CONFIG_BT_NIMBLE_MESH_PB_GATT](#))
- **CONFIG_NIMBLE_MESH_PROV** ([CONFIG_BT_NIMBLE_MESH_PROV](#))
- **CONFIG_NIMBLE_MESH_PROXY** ([CONFIG_BT_NIMBLE_MESH_PROXY](#))
- **CONFIG_NIMBLE_MESH_RELAY** ([CONFIG_BT_NIMBLE_MESH_RELAY](#))
- **CONFIG_NIMBLE_NVS_PERSIST** ([CONFIG_BT_NIMBLE_NVS_PERSIST](#))
- **CONFIG_NIMBLE_PINNED_TO_CORE_CHOICE** ([CONFIG_BT_NIMBLE_PINNED_TO_CORE_CHOICE](#))
 - **CONFIG_NIMBLE_PINNED_TO_CORE_0**
 - **CONFIG_NIMBLE_PINNED_TO_CORE_1**
- **CONFIG_NIMBLE_ROLE_BROADCASTER** ([CONFIG_BT_NIMBLE_ROLE_BROADCASTER](#))
- **CONFIG_NIMBLE_ROLE_CENTRAL** ([CONFIG_BT_NIMBLE_ROLE_CENTRAL](#))
- **CONFIG_NIMBLE_ROLE_OBSERVER** ([CONFIG_BT_NIMBLE_ROLE_OBSERVER](#))
- **CONFIG_NIMBLE_ROLE_PERIPHERAL** ([CONFIG_BT_NIMBLE_ROLE_PERIPHERAL](#))
- **CONFIG_NIMBLE_RPA_TIMEOUT** ([CONFIG_BT_NIMBLE_RPA_TIMEOUT](#))
- **CONFIG_NIMBLE_SM_LEGACY** ([CONFIG_BT_NIMBLE_SM_LEGACY](#))
- **CONFIG_NIMBLE_SM_SC** ([CONFIG_BT_NIMBLE_SM_SC](#))
- **CONFIG_NIMBLE_SM_SC_DEBUG_KEYS** ([CONFIG_BT_NIMBLE_SM_SC_DEBUG_KEYS](#))
- **CONFIG_NIMBLE_SVC_GAP_APPEARANCE** ([CONFIG_BT_NIMBLE_SVC_GAP_APPEARANCE](#))
- **CONFIG_NIMBLE_SVC_GAP_DEVICE_NAME** ([CONFIG_BT_NIMBLE_SVC_GAP_DEVICE_NAME](#))
- **CONFIG_NIMBLE_TASK_STACK_SIZE** ([CONFIG_BT_NIMBLE_HOST_TASK_STACK_SIZE](#))
- **CONFIG_NO_BLOBS** ([CONFIG_APP_NO_BLOBS](#))
- **CONFIG_OPTIMIZATION_ASSERTION_LEVEL** ([CONFIG_COMPILER_OPTIMIZATION_ASSERTION_LEVEL](#))
 - **CONFIG_OPTIMIZATION_ASSERTIONS_ENABLED**
 - **CONFIG_OPTIMIZATION_ASSERTIONS_SILENT**
 - **CONFIG_OPTIMIZATION_ASSERTIONS_DISABLED**
- **CONFIG_OPTIMIZATION_COMPILER** ([CONFIG_COMPILER_OPTIMIZATION](#))
 - **CONFIG_OPTIMIZATION_LEVEL_DEBUG**, **CONFIG_COMPILER_OPTIMIZATION_LEVEL_DEBUG**, **CONFIG_COMPILER_OPTIMIZATION_DEFAULT**
 - **CONFIG_OPTIMIZATION_LEVEL_RELEASE**, **CONFIG_COMPILER_OPTIMIZATION_LEVEL_RELEASE**
- **CONFIG_OSI_INITIAL_TRACE_LEVEL** ([CONFIG_BT_LOG_OSI_TRACE_LEVEL](#))
 - **CONFIG_OSI_TRACE_LEVEL_NONE**
 - **CONFIG_OSI_TRACE_LEVEL_ERROR**
 - **CONFIG_OSI_TRACE_LEVEL_WARNING**
 - **CONFIG_OSI_TRACE_LEVEL_API**
 - **CONFIG_OSI_TRACE_LEVEL_EVENT**
 - **CONFIG_OSI_TRACE_LEVEL_DEBUG**
 - **CONFIG_OSI_TRACE_LEVEL_VERBOSE**
- **CONFIG_OTA_ALLOW_HTTP** ([CONFIG_ESP_HTTPS_OTA_ALLOW_HTTP](#))
- **CONFIG_PAN_INITIAL_TRACE_LEVEL** ([CONFIG_BT_LOG_PAN_TRACE_LEVEL](#))
 - **CONFIG_PAN_TRACE_LEVEL_NONE**
 - **CONFIG_PAN_TRACE_LEVEL_ERROR**
 - **CONFIG_PAN_TRACE_LEVEL_WARNING**
 - **CONFIG_PAN_TRACE_LEVEL_API**
 - **CONFIG_PAN_TRACE_LEVEL_EVENT**
 - **CONFIG_PAN_TRACE_LEVEL_DEBUG**
 - **CONFIG_PAN_TRACE_LEVEL_VERBOSE**
- **CONFIG_POST_EVENTS_FROM_IRAM_ISR** ([CONFIG_ESP_EVENT_POST_FROM_IRAM_ISR](#))
- **CONFIG_POST_EVENTS_FROM_ISR** ([CONFIG_ESP_EVENT_POST_FROM_ISR](#))
- **CONFIG_PPP_CHAP_SUPPORT** ([CONFIG_LWIP_PPP_CHAP_SUPPORT](#))
- **CONFIG_PPP_DEBUG_ON** ([CONFIG_LWIP_PPP_DEBUG_ON](#))
- **CONFIG_PPP_MPPE_SUPPORT** ([CONFIG_LWIP_PPP_MPPE_SUPPORT](#))
- **CONFIG_PPP_MSCHAP_SUPPORT** ([CONFIG_LWIP_PPP_MSCHAP_SUPPORT](#))
- **CONFIG_PPP_NOTIFY_PHASE_SUPPORT** ([CONFIG_LWIP_PPP_NOTIFY_PHASE_SUPPORT](#))
- **CONFIG_PPP_PAP_SUPPORT** ([CONFIG_LWIP_PPP_PAP_SUPPORT](#))
- **CONFIG_PPP_SUPPORT** ([CONFIG_LWIP_PPP_SUPPORT](#))
- **CONFIG_RFCOMM_INITIAL_TRACE_LEVEL** ([CONFIG_BT_LOG_RFCOMM_TRACE_LEVEL](#))
 - **CONFIG_RFCOMM_TRACE_LEVEL_NONE**

- CONFIG_RFCOMM_TRACE_LEVEL_ERROR
- CONFIG_RFCOMM_TRACE_LEVEL_WARNING
- CONFIG_RFCOMM_TRACE_LEVEL_API
- CONFIG_RFCOMM_TRACE_LEVEL_EVENT
- CONFIG_RFCOMM_TRACE_LEVEL_DEBUG
- CONFIG_RFCOMM_TRACE_LEVEL_VERBOSE
- CONFIG_SEMIHOSTFS_MAX_MOUNT_POINTS ([CONFIG_VFS_SEMIHOSTFS_MAX_MOUNT_POINTS](#))
- **CONFIG_SMP_INITIAL_TRACE_LEVEL** ([CONFIG_BT_LOG_SMP_TRACE_LEVEL](#))
 - CONFIG_SMP_TRACE_LEVEL_NONE
 - CONFIG_SMP_TRACE_LEVEL_ERROR
 - CONFIG_SMP_TRACE_LEVEL_WARNING
 - CONFIG_SMP_TRACE_LEVEL_API
 - CONFIG_SMP_TRACE_LEVEL_EVENT
 - CONFIG_SMP_TRACE_LEVEL_DEBUG
 - CONFIG_SMP_TRACE_LEVEL_VERBOSE
- CONFIG_SMP_SLAVE_CON_PARAMS_UPD_ENABLE ([CONFIG_BT_SMP_SLAVE_CON_PARAMS_UPD_ENABLE](#))
- **CONFIG_SPI_FLASH_WRITING_DANGEROUS_REGIONS** ([CONFIG_SPI_FLASH_DANGEROUS_WRITE](#))
 - CONFIG_SPI_FLASH_WRITING_DANGEROUS_REGIONS_ABORTS
 - CONFIG_SPI_FLASH_WRITING_DANGEROUS_REGIONS_FAILS
 - CONFIG_SPI_FLASH_WRITING_DANGEROUS_REGIONS_ALLOWED
- **CONFIG_STACK_CHECK_MODE** ([CONFIG_COMPILER_STACK_CHECK_MODE](#))
 - CONFIG_STACK_CHECK_NONE
 - CONFIG_STACK_CHECK_NORM
 - CONFIG_STACK_CHECK_STRONG
 - CONFIG_STACK_CHECK_ALL
- CONFIG_SUPPORT_TERMIOS ([CONFIG_VFS_SUPPORT_TERMIOS](#))
- CONFIG_SUPPRESS_SELECT_DEBUG_OUTPUT ([CONFIG_VFS_SUPPRESS_SELECT_DEBUG_OUTPUT](#))
- CONFIG_SW_COEXIST_ENABLE ([CONFIG_ESP_COEX_SW_COEXIST_ENABLE](#))
- CONFIG_SYSTEM_EVENT_QUEUE_SIZE ([CONFIG_ESP_SYSTEM_EVENT_QUEUE_SIZE](#))
- CONFIG_SYSTEM_EVENT_TASK_STACK_SIZE ([CONFIG_ESP_SYSTEM_EVENT_TASK_STACK_SIZE](#))
- CONFIG_SYSVIEW_BUF_WAIT_TMO ([CONFIG_APPTRACE_SV_BUF_WAIT_TMO](#))
- CONFIG_SYSVIEW_ENABLE ([CONFIG_APPTRACE_SV_ENABLE](#))
- CONFIG_SYSVIEW_EVT_IDLE_ENABLE ([CONFIG_APPTRACE_SV_EVT_IDLE_ENABLE](#))
- CONFIG_SYSVIEW_EVT_ISR_ENTER_ENABLE ([CONFIG_APPTRACE_SV_EVT_ISR_ENTER_ENABLE](#))
- CONFIG_SYSVIEW_EVT_ISR_EXIT_ENABLE ([CONFIG_APPTRACE_SV_EVT_ISR_EXIT_ENABLE](#))
- CONFIG_SYSVIEW_EVT_ISR_TO_SCHEDULER_ENABLE ([CONFIG_APPTRACE_SV_EVT_ISR_TO_SCHED_ENABLE](#))
- CONFIG_SYSVIEW_EVT_OVERFLOW_ENABLE ([CONFIG_APPTRACE_SV_EVT_OVERFLOW_ENABLE](#))
- CONFIG_SYSVIEW_EVT_TASK_CREATE_ENABLE ([CONFIG_APPTRACE_SV_EVT_TASK_CREATE_ENABLE](#))
- CONFIG_SYSVIEW_EVT_TASK_START_EXEC_ENABLE ([CONFIG_APPTRACE_SV_EVT_TASK_START_EXEC_ENABLE](#))
- CONFIG_SYSVIEW_EVT_TASK_START_READY_ENABLE ([CONFIG_APPTRACE_SV_EVT_TASK_START_READY_ENABLE](#))
- CONFIG_SYSVIEW_EVT_TASK_STOP_EXEC_ENABLE ([CONFIG_APPTRACE_SV_EVT_TASK_STOP_EXEC_ENABLE](#))
- CONFIG_SYSVIEW_EVT_TASK_STOP_READY_ENABLE ([CONFIG_APPTRACE_SV_EVT_TASK_STOP_READY_ENABLE](#))
- CONFIG_SYSVIEW_EVT_TASK_TERMINATE_ENABLE ([CONFIG_APPTRACE_SV_EVT_TASK_TERMINATE_ENABLE](#))
- CONFIG_SYSVIEW_EVT_TIMER_ENTER_ENABLE ([CONFIG_APPTRACE_SV_EVT_TIMER_ENTER_ENABLE](#))
- CONFIG_SYSVIEW_EVT_TIMER_EXIT_ENABLE ([CONFIG_APPTRACE_SV_EVT_TIMER_EXIT_ENABLE](#))
- CONFIG_SYSVIEW_MAX_TASKS ([CONFIG_APPTRACE_SV_MAX_TASKS](#))
- **CONFIG_SYSVIEW_TS_SOURCE** ([CONFIG_APPTRACE_SV_TS_SOURCE](#))
 - CONFIG_SYSVIEW_TS_SOURCE_CCOUNT
 - CONFIG_SYSVIEW_TS_SOURCE_ESP_TIMER
- CONFIG_TASK_WDT ([CONFIG_ESP_TASK_WDT_INIT](#))
- CONFIG_TASK_WDT_CHECK_IDLE_TASK_CPU0 ([CONFIG_ESP_TASK_WDT_CHECK_IDLE_TASK_CPU0](#))
- CONFIG_TASK_WDT_CHECK_IDLE_TASK_CPU1 ([CONFIG_ESP_TASK_WDT_CHECK_IDLE_TASK_CPU1](#))
- CONFIG_TASK_WDT_PANIC ([CONFIG_ESP_TASK_WDT_PANIC](#))
- CONFIG_TASK_WDT_TIMEOUT_S ([CONFIG_ESP_TASK_WDT_TIMEOUT_S](#))
- CONFIG_TCPIP_RECVMBOX_SIZE ([CONFIG_LWIP_TCPIP_RECVMBOX_SIZE](#))
- **CONFIG_TCPIP_TASK_AFFINITY** ([CONFIG_LWIP_TCPIP_TASK_AFFINITY](#))
 - CONFIG_TCPIP_TASK_AFFINITY_NO_AFFINITY

- CONFIG_TCPIP_TASK_AFFINITY_CPU0
- CONFIG_TCPIP_TASK_AFFINITY_CPU1
- CONFIG_TCPIP_TASK_STACK_SIZE ([CONFIG_LWIP_TCPIP_TASK_STACK_SIZE](#))
- CONFIG_TCP_MAXRTX ([CONFIG_LWIP_TCP_MAXRTX](#))
- CONFIG_TCP_MSL ([CONFIG_LWIP_TCP_MSL](#))
- CONFIG_TCP_MSS ([CONFIG_LWIP_TCP_MSS](#))
- **CONFIG_TCP_OVERSIZE** ([CONFIG_LWIP_TCP_OVERSIZE](#))
 - CONFIG_TCP_OVERSIZE_MSS
 - CONFIG_TCP_OVERSIZE_QUARTER_MSS
 - CONFIG_TCP_OVERSIZE_DISABLE
- CONFIG_TCP_QUEUE_OOSEQ ([CONFIG_LWIP_TCP_QUEUE_OOSEQ](#))
- CONFIG_TCP_RECVMBOX_SIZE ([CONFIG_LWIP_TCP_RECVMBOX_SIZE](#))
- CONFIG_TCP_SND_BUF_DEFAULT ([CONFIG_LWIP_TCP_SND_BUF_DEFAULT](#))
- CONFIG_TCP_SYNMAXRTX ([CONFIG_LWIP_TCP_SYNMAXRTX](#))
- CONFIG_TCP_WND_DEFAULT ([CONFIG_LWIP_TCP_WND_DEFAULT](#))
- CONFIG_TIMER_QUEUE_LENGTH ([CONFIG_FREERTOS_TIMER_QUEUE_LENGTH](#))
- CONFIG_TIMER_TASK_PRIORITY ([CONFIG_FREERTOS_TIMER_TASK_PRIORITY](#))
- CONFIG_TIMER_TASK_STACK_DEPTH ([CONFIG_FREERTOS_TIMER_TASK_STACK_DEPTH](#))
- CONFIG_TIMER_TASK_STACK_SIZE ([CONFIG_ESP_TIMER_TASK_STACK_SIZE](#))
- CONFIG_UDP_RECVMBOX_SIZE ([CONFIG_LWIP_UDP_RECVMBOX_SIZE](#))
- CONFIG_WARN_WRITE_STRINGS ([CONFIG_COMPILER_WARN_WRITE_STRINGS](#))
- CONFIG_WPA_11KV_SUPPORT ([CONFIG_ESP_WIFI_11KV_SUPPORT](#))
- CONFIG_WPA_11R_SUPPORT ([CONFIG_ESP_WIFI_11R_SUPPORT](#))
- CONFIG_WPA_DEBUG_PRINT ([CONFIG_ESP_WIFI_DEBUG_PRINT](#))
- CONFIG_WPA_DPP_SUPPORT ([CONFIG_ESP_WIFI_DPP_SUPPORT](#))
- CONFIG_WPA_MBEDTLS_CRYPT ([CONFIG_ESP_WIFI_MBEDTLS_CRYPT](#))
- CONFIG_WPA_MBEDTLS_TLS_CLIENT ([CONFIG_ESP_WIFI_MBEDTLS_TLS_CLIENT](#))
- CONFIG_WPA_MBO_SUPPORT ([CONFIG_ESP_WIFI_MBO_SUPPORT](#))
- CONFIG_WPA_SCAN_CACHE ([CONFIG_ESP_WIFI_SCAN_CACHE](#))
- CONFIG_WPA_SUITE_B_192 ([CONFIG_ESP_WIFI_SUITE_B_192](#))
- CONFIG_WPA_WAPI_PSK ([CONFIG_ESP_WIFI_WAPI_PSK](#))
- CONFIG_WPA_WPS_SOFTAP_REGISTRAR ([CONFIG_ESP_WIFI_WPS_SOFTAP_REGISTRAR](#))
- CONFIG_WPA_WPS_STRICT ([CONFIG_ESP_WIFI_WPS_STRICT](#))

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.7 Provisioning API

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.7.1 Protocol Communication

Overview

The Protocol Communication (protocomm) component manages secure sessions and provides the framework for multiple transports. The application can also use the protocomm layer directly to have application-specific extensions for the provisioning or non-provisioning use cases.

Following features are available for provisioning:

- Communication security at the application level
 - `protocomm_security0` (no security)
 - `protocomm_security1` (Curve25519 key exchange + AES-CTR encryption/decryption)
 - `protocomm_security2` (SRP6a-based key exchange + AES-GCM encryption/decryption)
- Proof-of-possession (support with `protocomm_security1` only)
- Salt and Verifier (support with `protocomm_security2` only)

Protocomm internally uses protobuf (protocol buffers) for secure session establishment. Users can choose to implement their own security (even without using protobuf). Protocomm can also be used without any security layer.

Protocomm provides the framework for various transports:

- Console, in which case the handler invocation is automatically taken care of on the device side. See Transport Examples below for code snippets.

Note that for `protocomm_security1` and `protocomm_security2`, the client still needs to establish sessions by performing the two-way handshake. See provisioning for more details about the secure handshake logic.

Enabling Protocomm Security Version

The protocomm component provides a project configuration menu to enable/disable support of respective security versions. The respective configuration options are as follows:

- Support `protocomm_security0`, with no security: [`CONFIG\_ESP\_PROTOCOMM\_SUPPORT\_SECURITY\_VERSION\_0`](#), this option is enabled by default.
- Support `protocomm_security1` with Curve25519 key exchange + AES-CTR encryption/decryption: [`CONFIG\_ESP\_PROTOCOMM\_SUPPORT\_SECURITY\_VERSION\_1`](#), this option is enabled by default.
- Support `protocomm_security2` with SRP6a-based key exchange + AES-GCM encryption/decryption: [`CONFIG\_ESP\_PROTOCOMM\_SUPPORT\_SECURITY\_VERSION\_2`](#).

Note: Enabling multiple security versions at once offers the ability to control them dynamically but also increases the firmware size.

API Reference

Header File

- [components/protocomm/include/common/protocomm.h](#)
- This header file can be included with:

```
#include "protocomm.h"
```

- This header file is a part of the API provided by the `protocomm` component. To declare that your component depends on `protocomm`, add the following to your `CMakeLists.txt`:

```
REQUIRES protocomm
```

or

```
PRIV_REQUIRES protocomm
```


Functions

protocomm_t ***protocomm_new** (void)

Create a new protocomm instance.

This API will return a new dynamically allocated protocomm instance with all elements of the *protocomm_t* structure initialized to NULL.

Returns

- *protocomm_t** : On success
- NULL : No memory for allocating new instance

void **protocomm_delete** (*protocomm_t* *pc)

Delete a protocomm instance.

This API will deallocate a protocomm instance that was created using *protocomm_new* ().

Parameters *pc* -- [in] Pointer to the protocomm instance to be deleted

esp_err_t **protocomm_add_endpoint** (*protocomm_t* *pc, const char *ep_name, *protocomm_req_handler_t* h, void *priv_data)

Add endpoint request handler for a protocomm instance.

This API will bind an endpoint handler function to the specified endpoint name, along with any private data that needs to be pass to the handler at the time of call.

Note:

- An endpoint must be bound to a valid protocomm instance, created using *protocomm_new* ().
 - This function internally calls the registered *add_endpoint* () function of the selected transport which is a member of the *protocomm_t* instance structure.
-

Parameters

- *pc* -- [in] Pointer to the protocomm instance
- *ep_name* -- [in] Endpoint identifier(name) string
- *h* -- [in] Endpoint handler function
- *priv_data* -- [in] Pointer to private data to be passed as a parameter to the handler function on call. Pass NULL if not needed.

Returns

- ESP_OK : Success
- ESP_FAIL : Error adding endpoint / Endpoint with this name already exists
- ESP_ERR_NO_MEM : Error allocating endpoint resource
- ESP_ERR_INVALID_ARG : Null instance/name/handler arguments

esp_err_t **protocomm_remove_endpoint** (*protocomm_t* *pc, const char *ep_name)

Remove endpoint request handler for a protocomm instance.

This API will remove a registered endpoint handler identified by an endpoint name.

Note:

- This function internally calls the registered *remove_endpoint* () function which is a member of the *protocomm_t* instance structure.
-

Parameters

- *pc* -- [in] Pointer to the protocomm instance
- *ep_name* -- [in] Endpoint identifier(name) string

Returns

- ESP_OK : Success
- ESP_ERR_NOT_FOUND : Endpoint with specified name doesn't exist
- ESP_ERR_INVALID_ARG : Null instance/name arguments

esp_err_t **protocomm_open_session** (*protocomm_t* *pc, uint32_t session_id)

Allocates internal resources for new transport session.

Note:

- An endpoint must be bound to a valid protocomm instance, created using `protocomm_new()`.
-

Parameters

- **pc** -- [in] Pointer to the protocomm instance
- **session_id** -- [in] Unique ID for a communication session

Returns

- ESP_OK : Request handled successfully
- ESP_ERR_NO_MEM : Error allocating internal resource
- ESP_ERR_INVALID_ARG : Null instance/name arguments

esp_err_t **protocomm_close_session** (*protocomm_t* *pc, uint32_t session_id)

Frees internal resources used by a transport session.

Note:

- An endpoint must be bound to a valid protocomm instance, created using `protocomm_new()`.
-

Parameters

- **pc** -- [in] Pointer to the protocomm instance
- **session_id** -- [in] Unique ID for a communication session

Returns

- ESP_OK : Request handled successfully
- ESP_ERR_INVALID_ARG : Null instance/name arguments

esp_err_t **protocomm_req_handle** (*protocomm_t* *pc, const char *ep_name, uint32_t session_id, const uint8_t *inbuf, ssize_t inlen, uint8_t **outbuf, ssize_t *outlen)

Calls the registered handler of an endpoint session for processing incoming data and generating the response.

Note:

- An endpoint must be bound to a valid protocomm instance, created using `protocomm_new()`.
 - Resulting output buffer must be deallocated by the caller.
-

Parameters

- **pc** -- [in] Pointer to the protocomm instance
- **ep_name** -- [in] Endpoint identifier(name) string
- **session_id** -- [in] Unique ID for a communication session
- **inbuf** -- [in] Input buffer contains input request data which is to be processed by the registered handler
- **inlen** -- [in] Length of the input buffer
- **outbuf** -- [out] Pointer to internally allocated output buffer, where the resulting response data output from the registered handler is to be stored
- **outlen** -- [out] Buffer length of the allocated output buffer

Returns

- ESP_OK : Request handled successfully
- ESP_FAIL : Internal error in execution of registered handler
- ESP_ERR_NO_MEM : Error allocating internal resource
- ESP_ERR_NOT_FOUND : Endpoint with specified name doesn't exist
- ESP_ERR_INVALID_ARG : Null instance/name arguments

esp_err_t **protocomm_set_security** (*protocomm_t* *pc, const char *ep_name, const *protocomm_security_t* *sec, const void *sec_params)

Add endpoint security for a protocomm instance.

This API will bind a security session establisher to the specified endpoint name, along with any proof of possession that may be required for authenticating a session client.

Note:

- An endpoint must be bound to a valid protocomm instance, created using `protocomm_new()`.
 - The choice of security can be any `protocomm_security_t` instance. Choices `protocomm_security0` and `protocomm_security1` and `protocomm_security2` are readily available.
-

Parameters

- **pc** -- **[in]** Pointer to the protocomm instance
- **ep_name** -- **[in]** Endpoint identifier(name) string
- **sec** -- **[in]** Pointer to endpoint security instance
- **sec_params** -- **[in]** Pointer to security params (NULL if not needed) The pointer should contain the security params struct of appropriate security version. For protocomm security version 1 and 2 `sec_params` should contain pointer to struct of type `protocomm_security1_params_t` and `protocomm_security2_params_t` respectively. The contents of this pointer must be valid till the security session has been running and is not closed.

Returns

- `ESP_OK` : Success
- `ESP_FAIL` : Error adding endpoint / Endpoint with this name already exists
- `ESP_ERR_INVALID_STATE` : Security endpoint already set
- `ESP_ERR_NO_MEM` : Error allocating endpoint resource
- `ESP_ERR_INVALID_ARG` : Null instance/name/handler arguments

esp_err_t **protocomm_unset_security** (*protocomm_t* *pc, const char *ep_name)

Remove endpoint security for a protocomm instance.

This API will remove a registered security endpoint identified by an endpoint name.

Parameters

- **pc** -- **[in]** Pointer to the protocomm instance
- **ep_name** -- **[in]** Endpoint identifier(name) string

Returns

- `ESP_OK` : Success
- `ESP_ERR_NOT_FOUND` : Endpoint with specified name doesn't exist
- `ESP_ERR_INVALID_ARG` : Null instance/name arguments

esp_err_t **protocomm_set_version** (*protocomm_t* *pc, const char *ep_name, const char *version)

Set endpoint for version verification.

This API can be used for setting an application specific protocol version which can be verified by clients through the endpoint.

Note:

- An endpoint must be bound to a valid protocomm instance, created using `protocomm_new()`.
-

Parameters

- **pc** -- **[in]** Pointer to the protocomm instance
- **ep_name** -- **[in]** Endpoint identifier(name) string
- **version** -- **[in]** Version identifier(name) string

Returns

- ESP_OK : Success
- ESP_FAIL : Error adding endpoint / Endpoint with this name already exists
- ESP_ERR_INVALID_STATE : Version endpoint already set
- ESP_ERR_NO_MEM : Error allocating endpoint resource
- ESP_ERR_INVALID_ARG : Null instance/name/handler arguments

esp_err_t **protocomm_unset_version** (*protocomm_t* *pc, const char *ep_name)

Remove version verification endpoint from a protocomm instance.

This API will remove a registered version endpoint identified by an endpoint name.

Parameters

- **pc** -- **[in]** Pointer to the protocomm instance
- **ep_name** -- **[in]** Endpoint identifier(name) string

Returns

- ESP_OK : Success
- ESP_ERR_NOT_FOUND : Endpoint with specified name doesn't exist
- ESP_ERR_INVALID_ARG : Null instance/name arguments

esp_err_t **protocomm_get_sec_version** (*protocomm_t* *pc, int *sec_ver, uint8_t *sec_patch_ver)

Get the security version of the protocomm instance.

This API will return the security version of the protocomm instance.

Parameters

- **pc** -- **[in]** Pointer to the protocomm instance
- **sec_ver** -- **[out]** Pointer to the security version
- **sec_patch_ver** -- **[out]** Pointer to the security patch version

Returns

- ESP_OK : Success
- ESP_ERR_INVALID_ARG : Null instance/name arguments

Type Definitions

typedef *esp_err_t* (***protocomm_req_handler_t**)(uint32_t session_id, const uint8_t *inbuf, ssize_t inlen, uint8_t **outbuf, ssize_t *outlen, void *priv_data)

Function prototype for protocomm endpoint handler.

typedef struct protocomm **protocomm_t**

This structure corresponds to a unique instance of protocomm returned when the API `protocomm_new()` is called. The remaining Protocomm APIs require this object as the first parameter.

Note: Structure of the protocomm object is kept private

Header File

- `components/protocomm/include/security/protocomm_security.h`
- This header file can be included with:

```
#include "protocomm_security.h"
```

- This header file is a part of the API provided by the `protocomm` component. To declare that your component depends on `protocomm`, add the following to your `CMakeLists.txt`:

```
REQUIRES protocomm
```

or

PRIV_REQUIRES protocomm

Structures

struct **protocomm_security1_params**

Protocomm Security 1 parameters: Proof Of Possession.

Public Members

const uint8_t ***data**

Pointer to buffer containing the proof of possession data

uint16_t **len**

Length (in bytes) of the proof of possession data

struct **protocomm_security2_params**

Protocomm Security 2 parameters: Salt and Verifier.

Public Members

const char ***salt**

Pointer to the buffer containing the salt

uint16_t **salt_len**

Length (in bytes) of the salt

const char ***verifier**

Pointer to the buffer containing the verifier

uint16_t **verifier_len**

Length (in bytes) of the verifier

struct **protocomm_security**

Protocomm security object structure.

The member functions are used for implementing secure protocomm sessions.

Note: This structure should not have any dynamic members to allow re-entrancy

Public Members

int **ver**

Unique version number of security implementation

uint8_t **patch_ver**

Patch version number of security implementation

esp_err_t (***init**)(*protocomm_security_handle_t* *handle)

Function for initializing/allocating security infrastructure

esp_err_t (***cleanup**)(*protocomm_security_handle_t* handle)

Function for deallocating security infrastructure

esp_err_t (***new_transport_session**)(*protocomm_security_handle_t* handle, uint32_t session_id)

Starts new secure transport session with specified ID

esp_err_t (***close_transport_session**)(*protocomm_security_handle_t* handle, uint32_t session_id)

Closes a secure transport session with specified ID

esp_err_t (***security_req_handler**)(*protocomm_security_handle_t* handle, const void *sec_params, uint32_t session_id, const uint8_t *inbuf, ssize_t inlen, uint8_t **outbuf, ssize_t *outlen, void *priv_data)

Handler function for authenticating connection request and establishing secure session

esp_err_t (***encrypt**)(*protocomm_security_handle_t* handle, uint32_t session_id, const uint8_t *inbuf, ssize_t inlen, uint8_t **outbuf, ssize_t *outlen)

Function which implements the encryption algorithm

esp_err_t (***decrypt**)(*protocomm_security_handle_t* handle, uint32_t session_id, const uint8_t *inbuf, ssize_t inlen, uint8_t **outbuf, ssize_t *outlen)

Function which implements the decryption algorithm

Type Definitions

typedef struct *protocomm_security1_params* **protocomm_security1_params_t**

Protocomm Security 1 parameters: Proof Of Possession.

typedef *protocomm_security1_params_t* **protocomm_security_pop_t**

typedef struct *protocomm_security2_params* **protocomm_security2_params_t**

Protocomm Security 2 parameters: Salt and Verifier.

typedef void ***protocomm_security_handle_t**

typedef struct *protocomm_security* **protocomm_security_t**

Protocomm security object structure.

The member functions are used for implementing secure protocomm sessions.

Note: This structure should not have any dynamic members to allow re-entrancy

Enumerations

enum **protocomm_security_session_event_t**

Events generated by the protocomm security layer.

These events are generated while establishing secured session.

Values:

enumerator **PROTOCOLM_SECURITY_SESSION_SETUP_OK**

Secured session established successfully

enumerator **PROTOCOLM_SECURITY_SESSION_INVALID_SECURITY_PARAMS**

Received invalid (NULL) security parameters (username / client public-key)

enumerator **PROTOCOLM_SECURITY_SESSION_CREDENTIALS_MISMATCH**

Received incorrect credentials (username / PoP)

Header File

- [components/protocomm/include/security/protocomm_security0.h](#)
- This header file can be included with:

```
#include "protocomm_security0.h"
```

- This header file is a part of the API provided by the `protocomm` component. To declare that your component depends on `protocomm`, add the following to your `CMakeLists.txt`:

```
REQUIRES protocomm
```

or

```
PRIV_REQUIRES protocomm
```

Header File

- [components/protocomm/include/security/protocomm_security1.h](#)
- This header file can be included with:

```
#include "protocomm_security1.h"
```

- This header file is a part of the API provided by the `protocomm` component. To declare that your component depends on `protocomm`, add the following to your `CMakeLists.txt`:

```
REQUIRES protocomm
```

or

```
PRIV_REQUIRES protocomm
```

Header File

- [components/protocomm/include/security/protocomm_security2.h](#)
- This header file can be included with:

```
#include "protocomm_security2.h"
```

- This header file is a part of the API provided by the `protocomm` component. To declare that your component depends on `protocomm`, add the following to your `CMakeLists.txt`:

```
REQUIRES protocomm
```

or

```
PRIV_REQUIRES protocomm
```

Header File

- [components/protocomm/include/crypto/srp6a/esp_srp.h](#)
- This header file can be included with:

```
#include "esp_srp.h"
```

- This header file is a part of the API provided by the `protocomm` component. To declare that your component depends on `protocomm`, add the following to your `CMakeLists.txt`:

```
REQUIRES protocomm
```

or

```
PRIV_REQUIRES protocomm
```

Functions

esp_srp_handle_t ***esp_srp_init** (*esp_ng_type_t* ng)

Initialize srp context for given NG type.

Note: the handle gets freed with `esp_srp_free`

Parameters **ng** -- NG type given by `esp_ng_type_t`

Returns `esp_srp_handle_t`* srp handle

void **esp_srp_free** (*esp_srp_handle_t* *hd)

free `esp_srp_context`

Parameters **hd** -- handle to be free

esp_err_t **esp_srp_srv_pubkey** (*esp_srp_handle_t* *hd, const char *username, int username_len, const char *pass, int pass_len, int salt_len, char **bytes_B, int *len_B, char **bytes_salt)

Returns B (pub key) and salt. [Step2.b].

Note: `*bytes_B` MUST NOT BE FREED BY THE CALLER

Note: `*bytes_salt` MUST NOT BE FREE BY THE CALLER

Parameters

- **hd** -- `esp_srp` handle
- **username** -- Username not expected NULL terminated
- **username_len** -- Username length
- **pass** -- Password not expected to be NULL terminated
- **pass_len** -- Password length
- **salt_len** -- Salt length
- **bytes_B** -- Public Key returned
- **len_B** -- Length of the public key
- **bytes_salt** -- Salt bytes generated

Returns `esp_err_t` ESP_OK on success, appropriate error otherwise

esp_err_t **esp_srp_gen_salt_verifier** (const char *username, int username_len, const char *pass, int pass_len, char **bytes_salt, int salt_len, char **verifier, int *verifier_len)

Generate salt-verifier pair, given username, password and salt length.

Note: if API has returned ESP_OK, salt and verifier generated need to be freed by caller

Note: Usually, username and password are not saved on the device. Rather salt and verifier are generated outside the device and are embedded. this convenience API can be used to generate salt and verifier on the fly for development use case. OR for devices which intentionally want to generate different password each time and can send it to the client securely. e.g., a device has a display and it shows the pin

Parameters

- **username** -- [in] username
- **username_len** -- [in] length of the username
- **pass** -- [in] password
- **pass_len** -- [in] length of the password
- **bytes_salt** -- [out] generated salt on successful generation, or NULL
- **salt_len** -- [in] salt length
- **verifier** -- [out] generated verifier on successful generation, or NULL
- **verifier_len** -- [out] length of the generated verifier

Returns esp_err_t ESP_OK on success, appropriate error otherwise

esp_err_t **esp_srp_set_salt_verifier** (*esp_srp_handle_t* *hd, const char *salt, int salt_len, const char *verifier, int verifier_len)

Set the Salt and Verifier pre-generated for a given password. This should be used only if the actual password is not available. The public key can then be generated using *esp_srp_srv_pubkey_from_salt_verifier()* and not *esp_srp_srv_pubkey()*

Parameters

- **hd** -- esp_srp_handle
- **salt** -- pre-generated salt bytes
- **salt_len** -- length of the salt bytes
- **verifier** -- pre-generated verifier
- **verifier_len** -- length of the verifier bytes

Returns esp_err_t ESP_OK on success, appropriate error otherwise

esp_err_t **esp_srp_srv_pubkey_from_salt_verifier** (*esp_srp_handle_t* *hd, char **bytes_B, int *len_B)

Returns B (pub key)[Step2.b] when the salt and verifier are set using *esp_srp_set_salt_verifier()*

Note: *bytes_B MUST NOT BE FREED BY THE CALLER

Parameters

- **hd** -- esp_srp handle
- **bytes_B** -- Key returned to the called
- **len_B** -- Length of the key returned

Returns esp_err_t ESP_OK on success, appropriate error otherwise

esp_err_t **esp_srp_get_session_key** (*esp_srp_handle_t* *hd, char *bytes_A, int len_A, char **bytes_key, uint16_t *len_key)

Get session key in *bytes_key given by len in *len_key. [Step2.c].

This calculated session key is used for further communication given the proofs are exchanged/authenticated with *esp_srp_exchange_proofs*

Note: `*bytes_key` MUST NOT BE FREED BY THE CALLER

Parameters

- **hd** -- `esp_srp_handle`
- **bytes_A** -- Private Key
- **len_A** -- Private Key length
- **bytes_key** -- Key returned to the caller
- **len_key** -- length of the key in `*bytes_key`

Returns `esp_err_t` `ESP_OK` on success, appropriate error otherwise

`esp_err_t esp_srp_exchange_proofs(esp_srp_handle_t *hd, char *username, uint16_t username_len, char *bytes_user_proof, char *bytes_host_proof)`

Complete the authentication. If this step fails, the session_key exchanged should not be used.

This is the final authentication step in SRP algorithm [Step4.1, Step4.b, Step4.c]

Parameters

- **hd** -- `esp_srp_handle`
- **username** -- Username not expected NULL terminated
- **username_len** -- Username length
- **bytes_user_proof** -- param in
- **bytes_host_proof** -- parameter out (should be `SHA512_DIGEST_LENGTH`) bytes in size

Returns `esp_err_t` `ESP_OK` if user's proof is ok and subsequently `bytes_host_proof` is populated with our own proof.

Type Definitions

typedef struct `esp_srp_handle` **esp_srp_handle_t**

`esp_srp_handle` as the result of `esp_srp_init`

The handle is returned by `esp_srp_init` on successful init. It is then passed for subsequent API calls as an argument. `esp_srp_free` can be used to clean up the handle. After `esp_srp_free` the handle becomes invalid.

Enumerations

enum **esp_ng_type_t**

Large prime+generator to be used for the algorithm.

Values:

enumerator **ESP_NG_3072**

Header File

- `components/protocomm/include/transport/protocomm_httpd.h`
- This header file can be included with:

```
#include "protocomm_httpd.h"
```

- This header file is a part of the API provided by the `protocomm` component. To declare that your component depends on `protocomm`, add the following to your `CMakeLists.txt`:

```
REQUIRES protocomm
```

or

PRIV_REQUIRES protocomm

Functions

esp_err_t **protocomm_httpd_start** (*protocomm_t* *pc, const *protocomm_httpd_config_t* *config)

Start HTTPD protocomm transport.

This API internally creates a framework to allow endpoint registration and security configuration for the protocomm.

Note: This is a singleton. ie. Protocomm can have multiple instances, but only one instance can be bound to an HTTP transport layer.

Parameters

- **pc** -- [in] Protocomm instance pointer obtained from `protocomm_new()`
- **config** -- [in] Pointer to config structure for initializing HTTP server

Returns

- ESP_OK : Success
- ESP_ERR_INVALID_ARG : Null arguments
- ESP_ERR_NOT_SUPPORTED : Transport layer bound to another protocomm instance
- ESP_ERR_INVALID_STATE : Transport layer already bound to this protocomm instance
- ESP_ERR_NO_MEM : Memory allocation for server resource failed
- ESP_ERR_HTTPD_* : HTTP server error on start

esp_err_t **protocomm_httpd_stop** (*protocomm_t* *pc)

Stop HTTPD protocomm transport.

This API cleans up the HTTPD transport protocomm and frees all the handlers registered with the protocomm.

Parameters **pc** -- [in] Same protocomm instance that was passed to `protocomm_httpd_start()`

Returns

- ESP_OK : Success
- ESP_ERR_INVALID_ARG : Null / incorrect protocomm instance pointer

Unions

union **protocomm_httpd_config_data_t**

#include <protocomm_httpd.h> Protocomm HTTPD Configuration Data

Public Members

void ***handle**

HTTP Server Handle, if `ext_handle_provided` is set to true

protocomm_http_server_config_t **config**

HTTP Server Configuration, if a server is not already active

Structures

struct **protocomm_http_server_config_t**

Config parameters for protocomm HTTP server.

Public Members

uint16_t **port**

Port on which the HTTP server will listen

size_t **stack_size**

Stack size of server task, adjusted depending upon stack usage of endpoint handler

unsigned **task_priority**

Priority of server task

struct **protocomm_httpd_config_t**

Config parameters for protocomm HTTP server.

Public Members

bool **ext_handle_provided**

Flag to indicate if an external HTTP Server Handle has been provided. In such a case, protocomm will use the same HTTP Server and not start a new one internally.

protocomm_httpd_config_data_t **data**

Protocomm HTTPD Configuration Data

Macros

PROTOCOLM_HTTPD_DEFAULT_CONFIG()

Header File

- [components/protocomm/include/transport/protocomm_ble.h](#)
- This header file can be included with:

```
#include "protocomm_ble.h"
```

- This header file is a part of the API provided by the `protocomm` component. To declare that your component depends on `protocomm`, add the following to your `CMakeLists.txt`:

```
REQUIRES protocomm
```

or

```
PRIV_REQUIRES protocomm
```

Functions

esp_err_t **protocomm_ble_start** (*protocomm_t* *pc, const *protocomm_ble_config_t* *config)

Start Bluetooth Low Energy based transport layer for provisioning.

Initialize and start required BLE service for provisioning. This includes the initialization for characteristics/service for BLE.

Parameters

- **pc** -- [in] Protocomm instance pointer obtained from `protocomm_new()`
- **config** -- [in] Pointer to config structure for initializing BLE

Returns

- `ESP_OK` : Success

- ESP_FAIL : Simple BLE start error
- ESP_ERR_NO_MEM : Error allocating memory for internal resources
- ESP_ERR_INVALID_STATE : Error in ble config
- ESP_ERR_INVALID_ARG : Null arguments

esp_err_t **protocomm_ble_stop** (*protocomm_t* *pc)

Stop Bluetooth Low Energy based transport layer for provisioning.

Stops service/task responsible for BLE based interactions for provisioning

Note: You might want to optionally reclaim memory from Bluetooth. Refer to the documentation of `esp_bt_mem_release` in that case.

Parameters **pc** -- [in] Same protocomm instance that was passed to `protocomm_ble_start()`

Returns

- ESP_OK : Success
- ESP_FAIL : Simple BLE stop error
- ESP_ERR_INVALID_ARG : Null / incorrect protocomm instance

Structures

struct **name_uuid**

This structure maps handler required by protocomm layer to UUIDs which are used to uniquely identify BLE characteristics from a smartphone or a similar client device.

Public Members

const char ***name**

Name of the handler, which is passed to protocomm layer

uint16_t **uuid**

UUID to be assigned to the BLE characteristic which is mapped to the handler

struct **protocomm_ble_event_t**

Structure for BLE events in Protocomm.

Public Members

uint16_t **evt_type**

This field indicates the type of BLE event that occurred.

uint16_t **conn_handle**

The handle of the relevant connection.

uint16_t **conn_status**

The status of the connection attempt; 0: the connection was successfully established. o BLE host error code: the connection attempt failed for the specified reason.

uint16_t **disconnect_reason**

Return code indicating the reason for the disconnect.

struct **protocomm_ble_config**

Config parameters for protocomm BLE service.

Public Members

char **device_name**[MAX_BLE_DEVNAME_LEN + 1]

BLE device name being broadcast at the time of provisioning

uint8_t **service_uuid**[BLE_UUID128_VAL_LENGTH]

128 bit UUID of the provisioning service

uint8_t ***manufacturer_data**

BLE device manufacturer data pointer in advertisement

ssize_t **manufacturer_data_len**

BLE device manufacturer data length in advertisement

ssize_t **nu_lookup_count**

Number of entries in the Name-UUID lookup table

protocomm_ble_name_uuid_t ***nu_lookup**

Pointer to the Name-UUID lookup table

unsigned **ble_bonding**

BLE bonding

unsigned **ble_sm_sc**

BLE security flag

unsigned **ble_link_encryption**

BLE security flag

uint8_t ***ble_addr**

BLE address

unsigned **keep_ble_on**

Flag to keep BLE on

unsigned **ble_notify**

BLE characteristic notify flag

Macros

MAX_BLE_DEVNAME_LEN

BLE device name cannot be larger than this value 31 bytes (max scan response size) - 1 byte (length) - 1 byte (type) = 29 bytes

BLE_UUID128_VAL_LENGTH

MAX_BLE_MANUFACTURER_DATA_LEN

Theoretically, the limit for max manufacturer length remains same as BLE device name i.e. 31 bytes (max scan response size) - 1 byte (length) - 1 byte (type) = 29 bytes. However, manufacturer data goes along with BLE device name in scan response. So, it is important to understand the actual length should be smaller than (29 - (BLE device name length) - 2).

BLE_ADDR_LEN**Type Definitions**

```
typedef struct name_uuid protocomm_ble_name_uuid_t
```

This structure maps handler required by protocomm layer to UUIDs which are used to uniquely identify BLE characteristics from a smartphone or a similar client device.

```
typedef struct protocomm_ble_config protocomm_ble_config_t
```

Config parameters for protocomm BLE service.

Enumerations

```
enum protocomm_transport_ble_event_t
```

Events generated by BLE transport.

These events are generated when the BLE transport is paired and disconnected.

Values:

enumerator **PROTOCOLM_TRANSPORT_BLE_CONNECTED**

enumerator **PROTOCOLM_TRANSPORT_BLE_DISCONNECTED**

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.8 Storage API

This section contains reference of the high-level storage APIs. They are based on low-level drivers such as SPI flash, SD/MMC.

- *Partitions API* allow block based access to SPI flash according to the *Partition Tables*.
- *Non-Volatile Storage library (NVS)* implements a fault-tolerant wear-levelled key-value storage in SPI NOR flash.
- *Virtual File System (VFS)* library provides an interface for registration of file system drivers. SPIFFS, FAT and various other file system libraries are based on the VFS.
- *SPIFFS* is a wear-levelled file system optimized for SPI NOR flash, well suited for small partition sizes and low throughput
- *FAT* is a standard file system which can be used in SPI flash or on SD/MMC cards
- *Wear Levelling* library implements a flash translation layer (FTL) suitable for SPI NOR flash. It is used as a container for FAT partitions in flash.

Note: It is suggested to use high-level APIs (`esp_partition` or file system) instead of low-level driver APIs to access the SPI NOR flash.

Due to the restriction of NOR flash and ESP hardware, accessing the main flash will affect the performance of the whole system. See [SPI Flash API](#) to learn more about the limitations.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.8.1 FAT Filesystem Support

ESP-IDF uses the [FatFs](#) library to work with FAT filesystems. FatFs resides in the `fatfs` component. Although the library can be used directly, many of its features can be accessed via VFS using the C standard library and POSIX API functions.

Additionally, FatFs has been modified to support the runtime pluggable disk I/O layer. This allows mapping of FatFs drives to physical disks at runtime.

Using FatFs with VFS

The header file `fatfs/vfs/esp_vfs_fat.h` defines the functions for connecting FatFs and VFS.

The function `esp_vfs_fat_register()` allocates a FATFS structure and registers a given path prefix in VFS. Subsequent operations on files starting with this prefix are forwarded to FatFs APIs.

The function `esp_vfs_fat_unregister_path()` deletes the registration with VFS, and frees the FATFS structure.

Most applications use the following workflow when working with `esp_vfs_fat_` functions:

1. Call `esp_vfs_fat_register()` to specify:
 - Path prefix where to mount the filesystem (e.g., `"/sdcard"`, `"/spiflash"`)
 - FatFs drive number
 - A variable which receives the pointer to the FATFS structure
2. Call `ff_diskio_register()` to register the disk I/O driver for the drive number used in Step 1.
3. Call the FatFs function `f_mount()`, and optionally `f_fdisk()`, `f_mkfs()`, to mount the filesystem using the same drive number which was passed to `esp_vfs_fat_register()`. For more information, see [FatFs documentation](#).
4. Call the C standard library and POSIX API functions to perform such actions on files as open, read, write, erase, copy, etc. Use paths starting with the path prefix passed to `esp_vfs_fat_register()` (for example, `"/sdcard/hello.txt"`). The filesystem uses [8.3 filenames](#) format (SFN) by default. If you need to use long filenames (LFN), enable the `CONFIG_FATFS_LONG_FILENAMES` option. More details on the FatFs filenames are available [here](#).
5. Optionally, by enabling the option `CONFIG_FATFS_USE_FASTSEEK`, you can use the POSIX `lseek` function to perform it faster. The fast seek does not work for files in write mode, so to take advantage of fast seek, you should open (or close and then reopen) the file in read-only mode.
6. Optionally, by enabling the option `CONFIG_FATFS_IMMEDIATE_FSYNC`, you can enable automatic calling of `f_sync()` to flush recent file changes after each call of `vfs_fat_write()`, `vfs_fat_pwrite()`, `vfs_fat_link()`, `vfs_fat_truncate()` and `vfs_fat_ftruncate()` functions. This feature improves file-consistency and size reporting accuracy for the FatFs, at a price on decreased performance due to frequent disk operations.
7. Optionally, call the FatFs library functions directly. In this case, use paths without a VFS prefix, for example, `"/hello.txt"`.
8. Close all open files.
9. Call the FatFs function `f_mount()` for the same drive number with `NULL FATFS*` argument to unmount the filesystem.

10. Call the FatFs function `ff_diskio_register()` with NULL `ff_diskio_impl_t*` argument and the same drive number to unregister the disk I/O driver.
11. Call `esp_vfs_fat_unregister_path()` with the path where the file system is mounted to remove FatFs from VFS, and free the FATFS structure allocated in Step 1.

The convenience functions `esp_vfs_fat_sdmmc_mount()`, `esp_vfs_fat_sdspi_mount()`, and `esp_vfs_fat_sdcard_unmount()` wrap the steps described above and also handle SD card initialization. These functions are described in the next section.

Using FatFs with VFS and SD Cards

The header file `fatfs/vfs/esp_vfs_fat.h` defines convenience functions `esp_vfs_fat_sdmmc_mount()`, `esp_vfs_fat_sdspi_mount()`, and `esp_vfs_fat_sdcard_unmount()`. These functions perform Steps 1–3 and 7–9 respectively and handle SD card initialization, but provide only limited error handling. Developers are encouraged to check its source code and incorporate more advanced features into production applications.

The convenience function `esp_vfs_fat_sdmmc_unmount()` unmounts the filesystem and releases the resources acquired by `esp_vfs_fat_sdmmc_mount()`.

Using FatFs with VFS in Read-Only Mode

The header file `fatfs/vfs/esp_vfs_fat.h` also defines the convenience functions `esp_vfs_fat_spiflash_mount_ro()` and `esp_vfs_fat_spiflash_unmount_ro()`. These functions perform Steps 1-3 and 7-9 respectively for read-only FAT partitions. These are particularly helpful for data partitions written only once during factory provisioning, which will not be changed by production application throughout the lifetime of the hardware.

FatFS Disk IO Layer

FatFs has been extended with API functions that register the disk I/O driver at runtime.

These APIs provide implementation of disk I/O functions for SD/MMC cards and can be registered for the given FatFs drive number using the function `ff_diskio_register_sdmmc()`.

void **ff_diskio_register** (BYTE pdrv, const `ff_diskio_impl_t` *discio_impl)

Register or unregister diskio driver for given drive number.

When FATFS library calls one of disk_XXX functions for driver number pdrv, corresponding function in discio_impl for given pdrv will be called.

Parameters

- **pdrv** -- drive number
- **discio_impl** -- pointer to `ff_diskio_impl_t` structure with diskio functions or NULL to unregister and free previously registered drive

struct **ff_diskio_impl_t**

Structure of pointers to disk IO driver functions.

See FatFs documentation for details about these functions

Public Members

DSTATUS (***init**)(unsigned char pdrv)

disk initialization function

DSTATUS (***status**)(unsigned char pdrv)

disk status check function

DRESULT (***read**)(unsigned char pdrv, unsigned char *buff, uint32_t sector, unsigned count)

sector read function

DRESULT (***write**)(unsigned char pdrv, const unsigned char *buff, uint32_t sector, unsigned count)

sector write function

DRESULT (***ioctl**)(unsigned char pdrv, unsigned char cmd, void *buff)

function to get info about disk and do some misc operations

void **ff_diskio_register_sdmmc** (unsigned char pdrv, [sdmmc_card_t](#) *card)

Register SD/MMC diskio driver

Parameters

- **pdrv** -- drive number
- **card** -- pointer to [sdmmc_card_t](#) structure describing a card; card should be initialized before calling `f_mount`.

[esp_err_t](#) **ff_diskio_register_wl_partition** (unsigned char pdrv, [wl_handle_t](#) flash_handle)

Register spi flash partition

Parameters

- **pdrv** -- drive number
- **flash_handle** -- handle of the wear levelling partition.

[esp_err_t](#) **ff_diskio_register_raw_partition** (unsigned char pdrv, const [esp_partition_t](#) *part_handle)

Register spi flash partition

Parameters

- **pdrv** -- drive number
- **part_handle** -- pointer to raw flash partition.

FatFs Partition Generator

We provide a partition generator for FatFs ([wl_fatfsngen.py](#)) which is integrated into the build system and could be easily used in the user project.

The tool is used to create filesystem images on a host and populate it with content of the specified host folder.

The script is based on the partition generator ([fatfsngen.py](#)). Apart from generating partition, it can also initialize wear levelling.

The latest version supports both short and long file names, FAT12 and FAT16. The long file names are limited to 255 characters and can contain multiple periods (.) characters within the filename and additional characters +, , , ;, =, [and].

An in-depth description of the FatFs partition generator and analyzer can be found at [Generating and parsing FAT partition on host](#).

Build System Integration with FatFs Partition Generator It is possible to invoke FatFs generator directly from the CMake build system by calling `fatfs_create_spiflash_image`:

```
fatfs_create_spiflash_image(<partition> <base_dir> [FLASH_IN_PROJECT])
```

If you prefer generating partition without wear levelling support, you can use `fatfs_create_rawflash_image`:

```
fatfs_create_rawflash_image(<partition> <base_dir> [FLASH_IN_PROJECT])
```

`fatfs_create_spiflash_image` respectively `fatfs_create_rawflash_image` must be called from project's `CMakeLists.txt`.

If you decide for any reason to use `fatfs_create_rawflash_image` (without wear levelling support), beware that it supports mounting only in read-only mode in the device.

The arguments of the function are as follows:

1. `partition` - the name of the partition as defined in the partition table (e.g., [storage/fatfs/gen/partitions_example.csv](#)).
2. `base_dir` - the directory that will be encoded to FatFs partition and optionally flashed into the device. Beware that you have to specify the suitable size of the partition in the partition table.
3. flag `FLASH_IN_PROJECT` - optionally, users can have the image automatically flashed together with the app binaries, partition tables, etc. on `idf.py flash -p <PORT>` by specifying `FLASH_IN_PROJECT`.
4. flag `PRESERVE_TIME` - optionally, users can force preserving the timestamps from the source folder to the target image. Without preserving the time, every timestamp will be set to the FATFS default initial time (1st January 1980).

For example:

```
fatfs_create_spiflash_image(my_fatfs_partition my_folder FLASH_IN_PROJECT)
```

If `FLASH_IN_PROJECT` is not specified, the image will still be generated, but you will have to flash it manually using `esptool.py` or a custom build system target.

For an example, see [storage/fatfs/gen](#).

FatFs Partition Analyzer

([fatfsparse.py](#)) is a partition analyzing tool for FatFs.

It is a reverse tool of ([fatfs.gen.py](#)), i.e., it can generate the folder structure on the host based on the FatFs image.

Usage:

```
./fatfsparse.py [-h] [--wl-layer {detect,enabled,disabled}] fatfs_image.img
```

High-level API Reference

Header File

- [components/fatfs/vfs/esp_vfs_fat.h](#)
- This header file can be included with:

```
#include "esp_vfs_fat.h"
```

- This header file is a part of the API provided by the `fatfs` component. To declare that your component depends on `fatfs`, add the following to your `CMakeLists.txt`:

```
REQUIRES fatfs
```

or

```
PRIV_REQUIRES fatfs
```

Functions

esp_err_t **esp_vfs_fat_register** (const char *base_path, const char *fat_drive, size_t max_files, FATFS **out_fs)

Register FATFS with VFS component.

This function registers given FAT drive in VFS, at the specified base path. If only one drive is used, fat_drive argument can be an empty string. Refer to FATFS library documentation on how to specify FAT drive. This function also allocates FATFS structure which should be used for f_mount call.

Note: This function doesn't mount the drive into FATFS, it just connects POSIX and C standard library IO function with FATFS. You need to mount desired drive into FATFS separately.

Parameters

- **base_path** -- path prefix where FATFS should be registered
- **fat_drive** -- FATFS drive specification; if only one drive is used, can be an empty string
- **max_files** -- maximum number of files which can be open at the same time
- **out_fs** -- [out] pointer to FATFS structure which can be used for FATFS f_mount call is returned via this argument.

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE if esp_vfs_fat_register was already called
- ESP_ERR_NO_MEM if not enough memory or too many VFSes already registered

esp_err_t **esp_vfs_fat_unregister_path** (const char *base_path)

Un-register FATFS from VFS.

Note: FATFS structure returned by esp_vfs_fat_register is destroyed after this call. Make sure to call f_mount function to unmount it before calling esp_vfs_fat_unregister_ctx. Difference between this function and the one above is that this one will release the correct drive, while the one above will release the last registered one

Parameters **base_path** -- path prefix where FATFS is registered. This is the same used when esp_vfs_fat_register was called

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE if FATFS is not registered in VFS

esp_err_t **esp_vfs_fat_sdmmc_mount** (const char *base_path, const *sdmmc_host_t* *host_config, const void *slot_config, const *esp_vfs_fat_mount_config_t* *mount_config, *sdmmc_card_t* **out_card)

Convenience function to get FAT filesystem on SD card registered in VFS.

This is an all-in-one function which does the following:

- initializes SDMMC driver or SPI driver with configuration in host_config
- initializes SD card with configuration in slot_config
- mounts FAT partition on SD card using FATFS library, with configuration in mount_config
- registers FATFS library with VFS, with prefix given by base_prefix variable

This function is intended to make example code more compact. For real world applications, developers should implement the logic of probing SD card, locating and mounting partition, and registering FATFS in VFS, with proper error checking and handling of exceptional conditions.

Note: Use this API to mount a card through SDSPI is deprecated. Please call esp_vfs_fat_sdspi_mount () instead for that case.

Parameters

- **base_path** -- path where partition should be registered (e.g. "/sdcard")
- **host_config** -- Pointer to structure describing SDMMC host. When using SDMMC peripheral, this structure can be initialized using SDMMC_HOST_DEFAULT() macro. When using SPI peripheral, this structure can be initialized using SD-SPI_HOST_DEFAULT() macro.
- **slot_config** -- Pointer to structure with slot configuration. For SDMMC peripheral, pass a pointer to [sdmmc_slot_config_t](#) structure initialized using SDMMC_SLOT_CONFIG_DEFAULT.
- **mount_config** -- pointer to structure with extra parameters for mounting FATFS
- **out_card** -- [out] if not NULL, pointer to the card information structure will be returned via this argument

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE if esp_vfs_fat_sdmmc_mount was already called
- ESP_ERR_NO_MEM if memory can not be allocated
- ESP_FAIL if partition can not be mounted
- other error codes from SDMMC or SPI drivers, SDMMC protocol, or FATFS drivers

[esp_err_t](#) **esp_vfs_fat_sdspi_mount** (const char *base_path, const [sdmmc_host_t](#) *host_config_input, const [sdspi_device_config_t](#) *slot_config, const [esp_vfs_fat_mount_config_t](#) *mount_config, [sdmmc_card_t](#) **out_card)

Convenience function to get FAT filesystem on SD card registered in VFS.

This is an all-in-one function which does the following:

- initializes an SPI Master device based on the SPI Master driver with configuration in slot_config, and attach it to an initialized SPI bus.
- initializes SD card with configuration in host_config_input
- mounts FAT partition on SD card using FATFS library, with configuration in mount_config
- registers FATFS library with VFS, with prefix given by base_prefix variable

This function is intended to make example code more compact. For real world applications, developers should implement the logic of probing SD card, locating and mounting partition, and registering FATFS in VFS, with proper error checking and handling of exceptional conditions.

Note: This function try to attach the new SD SPI device to the bus specified in host_config. Make sure the SPI bus specified in host_config->slot have been initialized by spi_bus_initialize() before.

Parameters

- **base_path** -- path where partition should be registered (e.g. "/sdcard")
- **host_config_input** -- Pointer to structure describing SDMMC host. This structure can be initialized using SDSPI_HOST_DEFAULT() macro.
- **slot_config** -- Pointer to structure with slot configuration. For SPI peripheral, pass a pointer to [sdspi_device_config_t](#) structure initialized using SDSPI_DEVICE_CONFIG_DEFAULT().
- **mount_config** -- pointer to structure with extra parameters for mounting FATFS
- **out_card** -- [out] If not NULL, pointer to the card information structure will be returned via this argument. It is suggested to hold this handle and use it to unmount the card later if needed. Otherwise it's not suggested to use more than one card at the same time and unmount one of them in your application.

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE if esp_vfs_fat_sdmmc_mount was already called
- ESP_ERR_NO_MEM if memory can not be allocated
- ESP_FAIL if partition can not be mounted
- other error codes from SDMMC or SPI drivers, SDMMC protocol, or FATFS drivers

esp_err_t **esp_vfs_fat_sdmmc_unmount** (void)

Unmount FAT filesystem and release resources acquired using `esp_vfs_fat_sdmmc_mount`.

Deprecated:

Use `esp_vfs_fat_sdcard_unmount()` instead.

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE if `esp_vfs_fat_sdmmc_mount` hasn't been called

esp_err_t **esp_vfs_fat_sdcard_unmount** (const char *base_path, *sdmmc_card_t* *card)

Unmount an SD card from the FAT filesystem and release resources acquired using `esp_vfs_fat_sdmmc_mount()` or `esp_vfs_fat_sdspi_mount()`

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if the card argument is unregistered
- ESP_ERR_INVALID_STATE if `esp_vfs_fat_sdmmc_mount` hasn't been called

esp_err_t **esp_vfs_fat_sdcard_format** (const char *base_path, *sdmmc_card_t* *card)

Format FAT filesystem.

Note: This API should be only called when the FAT is already mounted.

Parameters

- **base_path** -- Path where partition should be registered (e.g. "/sdcard")
- **card** -- Pointer to the card handle, which should be initialised by calling `esp_vfs_fat_sdspi_mount` first

Returns

- ESP_OK
- ESP_ERR_INVALID_STATE: FAT partition isn't mounted, call `esp_vfs_fat_sdmmc_mount` or `esp_vfs_fat_sdspi_mount` first
- ESP_ERR_NO_MEM: if memory can not be allocated
- ESP_FAIL: fail to format it, or fail to mount back

esp_err_t **esp_vfs_fat_spiflash_mount_rw_wl** (const char *base_path, const char *partition_label, const *esp_vfs_fat_mount_config_t* *mount_config, *wl_handle_t* *wl_handle)

Convenience function to initialize FAT filesystem in SPI flash and register it in VFS.

This is an all-in-one function which does the following:

- finds the partition with defined partition_label. Partition label should be configured in the partition table.
- initializes flash wear levelling library on top of the given partition
- mounts FAT partition using FATFS library on top of flash wear levelling library
- registers FATFS library with VFS, with prefix given by base_prefix variable

This function is intended to make example code more compact.

Parameters

- **base_path** -- path where FATFS partition should be mounted (e.g. "/spiflash")
- **partition_label** -- label of the partition which should be used
- **mount_config** -- pointer to structure with extra parameters for mounting FATFS
- **wl_handle** -- [out] wear levelling driver handle

Returns

- ESP_OK on success

- `ESP_ERR_NOT_FOUND` if the partition table does not contain FATFS partition with given label
- `ESP_ERR_INVALID_STATE` if `esp_vfs_fat_spiflash_mount_rw_wl` was already called
- `ESP_ERR_NO_MEM` if memory can not be allocated
- `ESP_FAIL` if partition can not be mounted
- other error codes from wear levelling library, SPI flash driver, or FATFS drivers

esp_err_t **esp_vfs_fat_spiflash_unmount_rw_wl** (const char *base_path, *wl_handle_t* wl_handle)

Unmount FAT filesystem and release resources acquired using `esp_vfs_fat_spiflash_mount_rw_wl`.

Parameters

- **base_path** -- path where partition should be registered (e.g. `"/spiflash"`)
- **wl_handle** -- wear levelling driver handle returned by `esp_vfs_fat_spiflash_mount_rw_wl`

Returns

- `ESP_OK` on success
- `ESP_ERR_INVALID_STATE` if `esp_vfs_fat_spiflash_mount_rw_wl` hasn't been called

esp_err_t **esp_vfs_fat_spiflash_format_rw_wl** (const char *base_path, const char *partition_label)

Format FAT filesystem.

Note: This API can be called when the FAT is mounted / not mounted. If this API is called when the FAT isn't mounted (by calling `esp_vfs_fat_spiflash_mount_rw_wl`), this API will first mount the FAT then format it, then restore back to the original state.

Parameters

- **base_path** -- Path where partition should be registered (e.g. `"/spiflash"`)
- **partition_label** -- Label of the partition which should be used

Returns

- `ESP_OK`
- `ESP_ERR_NO_MEM`: if memory can not be allocated
- Other errors from `esp_vfs_fat_spiflash_mount_rw_wl`

esp_err_t **esp_vfs_fat_spiflash_mount_ro** (const char *base_path, const char *partition_label, const *esp_vfs_fat_mount_config_t* *mount_config)

Convenience function to initialize read-only FAT filesystem and register it in VFS.

This is an all-in-one function which does the following:

- finds the partition with defined `partition_label`. Partition label should be configured in the partition table.
- mounts FAT partition using FATFS library
- registers FATFS library with VFS, with prefix given by `base_prefix` variable

Note: Wear levelling is not used when FAT is mounted in read-only mode using this function.

Parameters

- **base_path** -- path where FATFS partition should be mounted (e.g. `"/spiflash"`)
- **partition_label** -- label of the partition which should be used
- **mount_config** -- pointer to structure with extra parameters for mounting FATFS

Returns

- `ESP_OK` on success
- `ESP_ERR_NOT_FOUND` if the partition table does not contain FATFS partition with given label
- `ESP_ERR_INVALID_STATE` if `esp_vfs_fat_spiflash_mount_ro` was already called for the same partition

- `ESP_ERR_NO_MEM` if memory can not be allocated
- `ESP_FAIL` if partition can not be mounted
- other error codes from SPI flash driver, or FATFS drivers

esp_err_t **esp_vfs_fat_spiflash_unmount_ro** (const char *base_path, const char *partition_label)

Unmount FAT filesystem and release resources acquired using `esp_vfs_fat_spiflash_mount_ro`.

Parameters

- **base_path** -- path where partition should be registered (e.g. `"/spiflash"`)
- **partition_label** -- label of partition to be unmounted

Returns

- `ESP_OK` on success
- `ESP_ERR_INVALID_STATE` if `esp_vfs_fat_spiflash_mount_ro` hasn't been called

esp_err_t **esp_vfs_fat_info** (const char *base_path, uint64_t *out_total_bytes, uint64_t *out_free_bytes)

Get information for FATFS partition.

Parameters

- **base_path** -- Base path of the partition examined (e.g. `"/spiflash"`)
- **out_total_bytes** -- [out] Size of the file system
- **out_free_bytes** -- [out] Free bytes available in the file system

Returns

- `ESP_OK` on success
- `ESP_ERR_INVALID_STATE` if partition not found
- `ESP_FAIL` if another FRESULT error (saved in `errno`)

Structures

struct **esp_vfs_fat_mount_config_t**

Configuration arguments for `esp_vfs_fat_sdmmc_mount` and `esp_vfs_fat_spiflash_mount_rw_wl` functions.

Public Members

bool **format_if_mount_failed**

If FAT partition can not be mounted, and this parameter is true, create partition table and format the filesystem.

int **max_files**

Max number of open files.

size_t **allocation_unit_size**

If `format_if_mount_failed` is set, and mount fails, format the card with given allocation unit size. Must be a power of 2, between sector size and `128 * sector size`. For SD cards, sector size is always 512 bytes. For wear_levelling, sector size is determined by `CONFIG_WL_SECTOR_SIZE` option.

Using larger allocation unit size will result in higher read/write performance and higher overhead when storing small files.

Setting this field to 0 will result in allocation unit set to the sector size.

bool **disk_status_check_enable**

Enables real `ff_disk_status` function implementation for SD cards (`ff_sdmmc_status`). Possibly slows down IO performance.

Try to enable if you need to handle situations when SD cards are not unmounted properly before physical removal or you are experiencing issues with SD cards.

Doesn't do anything for other memory storage media.

Type Definitions

```
typedef esp_vfs_fat_mount_config_t esp_vfs_fat_sdmmc_mount_config_t
```

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

orphan

2.8.2 Generating and Parsing FATFS on Host

This document is intended mainly for developers of Python tools [fatfsgen.py](#) and [fatfsparse.py](#), and people with special interest in these tools and implementation of the FAT file system in ESP-IDF. If you are interested in using these tools, please refer to the user guide at [FatFs Partition Generator](#).

The FAT file system is composed of various logical units. The units are used to store general information about the file system, allocations, content of files and directories, and file's metadata. The tools [fatfsgen.py](#) and [fatfsparse.py](#) are used to implement the FAT file system while considering all these logical units, and they also provide support for wear levelling.

FAT File System Generator and Parser Design

This section describes particular units of the FAT file system generator and parser design. The implementation aims to create a valid model of the FAT structure with a focus on macro operations, generating and parsing the whole partition without modifying it in the run (mounting).

Class FATFS This is the most general entity responsible for modeling the FAT file system. It is composed of **FATFSState** (holding metadata and boot sector), **FAT** (holding file allocation table), and **Directory** (representing the root directory required by FAT12 and FAT16). The class processes all the requirements for the partition, analyses the local folder dedicated to transforming it into a binary image, and generates an internal representation of the local folder. Then, the class can generate a binary image from the internal FAT file system model.

Class WLFATFS The class extends the functionality of the class **FATFS**. It implements an encapsulation of the file system into the wear levelling, by adding the "dummy" sector for balancing the load (a redundant sector, see the section [Wear Levelling](#)), configuration sector and state sector. This class generates a binary FATFS partition with initialized wear levelling layer. For further analysis, it also provides an option to remove the wear levelling completely. The class is instantiated and invoked by the `wl_fatfsgen.py` script.

Class BootSectorState The instance of this class contains the metadata required for building a boot sector and BPB (BIOS Parameter Block). Boot sector is basically implemented for the cross-platform compatibility, i.e., when ESP chipsets are connected with other platforms, it will always follow all the FAT file system standards. However, during partition generation, chip does not consume the data in this boot sector and all the other data needed, as the data is constant. In other words, changing the fields with the prefix "BS" is usually unnecessary and often does not work. If you want to add new features, please focus on fields with the prefix "BPB". Another critical role of this class is to share access to the metadata and binary image over the whole class system. Because of this, every class in the system can access this singleton.




Fig. 23: FAT File System Generator and Parser Design Architecture

Class FATFSState The class **FATFSState** might be obsolete in the future, so developers could transfer its functionality into the **BootSectorState**. The class contains a reference to the **BootSectorState** and extends the data with some unknown information when creating a boot sector or unnecessary for the boot sector, such as the information generated when the file system supports long file names.

Class FAT FAT represents the File Allocation Table. FAT is a sequence of bytes spread over one or more sectors. The number of sectors is determined by the number of clusters and is calculated by the function `get_fat_sectors_count` in `utils.py`. The aim is to have as few sectors for one FAT as possible when you refer to every physical cluster in the file system. The FAT works as follows: For every physical cluster at `i * some_constant` address, FAT contains an entry at the `i`-th location which represents next address of the clusters in the file chain. Each version of the FAT file system uses a different size for FAT entries. FAT12 uses 12 bits per entry, thus 2 entries span 3 bytes. FAT16 uses 16 bits per entry, thus 1 entry spans 2 bytes. FAT32 uses 32 bits per FAT entry, thus one entry spans 4 bytes. All entries are in little-endian byte order.

All zeros at the `i`-th entry indicates that corresponding cluster is free, while all ones at the `i`-th entry indicates that corresponding cluster is occupied and is the last cluster in the file chain. The other number at `i`-th entry determines the next cluster's address in the file chain. These clusters are not necessarily stored adjacent to one another in the memory but instead are often fragmented throughout the data region.


For partition generation, the file is divided into several parts to fit the cluster. Notice that the structure allocation of the files is a linked list. Every cluster in the file allocation chain has entry in the FAT which refers to the next cluster or the information about the last cluster in the file chain. As mentioned, FAT12 uses 12 bits per FAT entry, thus it can set a maximum number of 4096 clusters, as with 12 bits (one and a half bytes), it can enumerate 4096 clusters at most. However, because of other overhead, FAT12 can have 4085 clusters at most. Similarly, FAT16 can have 65525 clusters at most and for FAT32 can have 268435445 clusters at most (as practically only 28 bits out of 32 bits are used to denote each FAT entry). The current implementation doesn't allow forcibly redefining the FAT file system with less than 4085 clusters to FAT16, even though the documentation claims it is possible. Notice that it would be meaningless to define it vice versa, i.e., to FAT12 with more than 4085 clusters (which implies clusters on higher addresses being inaccessible).

Class Cluster The **Cluster** class is responsible for accessing the FAT entries and the corresponding physical cluster. The class **FAT** is an aggregation of a particular number of **Cluster** instances. Every cluster is uniquely defined by the ID, which also determines its position in FAT and corresponding sectors in the data region. When the cluster is allocated, it includes a reference to a file or directory. It happens only if the cluster is the first in the allocation chain. The cluster contains information about whether it is empty and the last in the file allocation linked list. If not, it includes a reference to the next cluster in the linked list. In practical usage, the cluster doesn't need to access its file, but the other way around. The **File** or **Directory** accesses its cluster, to retrieve the entire content that might be chained.

Class Directory This class represents the file system directory. An instance of **Directory** contains the reference to the corresponding instance of **Cluster**, which has the first cluster in the allocation chain for the directory given. The root directory is a special case with a different count of sectors and a slightly different instantiation process. However, the root directory is still an instance of this class and is the only **Directory** instance associated with the class **FATFS** and **WLFATFS** respectively. The class **Directory** (except for the root directory) has one-to-one association with the class **Entry** that defines its entry in the parent directory. It also has an aggregation associated with the class **Entry**, because every directory contains multiple entries that consist of the actual directory's content (for example, aliases, files, and directories).

Class File Similar to the class **Directory**, **File** represents single file in the file system. This class has one-to-one association with its first cluster in the allocation chain. Through this cluster, the **File** class may access the corresponding physical address and thus modifying its content. Every file also has one-to-one association with **Entry** instance belonging to its parent directory.

Class Entry **Entry** encapsulates information about the file/directory name in the data region of corresponding parent directory. Every file system entity (File/Directory) has an entry. In case of the symlink, the entity can have multiple entries. The directory uses entries to access its descendant files and sub-directories, and enables traversing



the tree structure. Except for that, **Entry** holds the name, extension, size, and information regarding the used file name size (long file names or file names 8.3), etc.

fatfsngen.py

`fatfs.py` generates FAT file systems on the host.

`fatfsgen.py` recursively traverses the given folder's directory structure and adds files and/or directories inside the binary partition. Users can set if the script generates the partition with wear levelling support, long file names support, and support for preserving the modification date and time from the original folder on the host.


The `./fatfsgen.py` `Espressif` command generates a simple binary partition with the default settings. Here `Espressif` is the local folder (containing files and/or sub-directories) from which binary image is generated.

There exist two scripts for that purpose, `fatfsgen.py` and `wl_fatfsgen.py`. The difference is that `wl_fatfsgen.py` firstly uses `fatfsgen.py` for generating the partition and then initializes wear leveling.

The script command line arguments are as follows:

```
fatfsgen.py [-h] [--output_file OUTPUT_FILE] [--partition_size PARTITION_SIZE] [--
↳sector_size {4096}] [--long_name_support] [--use_default_datetime] input_
↳directory

--output_file: path to the generated binary partition
--partition_size: defines the size of the binary partition (decimal, hexa or
↳binary number)
--sector_size: the size of the sector
--long_name_support: flag for supporting long file names
--use_default_datetime: this flag forces using default dates and times (date ==
↳0x2100, time == 0x0000), not using argument to preserve the original file system
↳metadata
input_directory: required argument, name of the directory being encoded to the
↳binary fat-compatible partition
```



fatfsparse.py

`fatfsparse.py` translates the binary image into the internal representation and generates the folder with equivalent content on the host. If user requires a parsing partition with initialized wear levelling, the `fatfsparse.py` will remove the wear levelling sectors using the function `remove_wl` provided by `wl_fatfsgen.py`. After the sectors are removed, parsing of the partition is the same as with no initial wear levelling.

`./fatfsparse.py fatfs_image.img` command yields the directory with the equivalent content as the binary data image `fatfs_image.img`.

The script command line arguments are as follows:

```
fatfsparse.py [-h] [--wl-layer {detect,enabled,disabled}] input_image
--wl-layer: indicates if wear leveling is enabled, disabled or should be detected.
↳ (detection is ambiguous)
input_image: path to binary image
```

The long file names can be detected automatically. However, the wear leveling cannot be 100% detected, because one partition can be valid either with or without wear leveling, according to the user's context. When the script finds wear leveling sectors (cfg and state), it assumes wear leveling is enabled, however it might be a false positive.

Features

FAT12/FAT16 The supported FAT types are FAT12 and FAT16. For smaller partitions, FAT12 is sufficient. The type is detected according to the count of clusters, and cannot be changed by the user. If there are less than 4085 clusters, the selected type is FAT12 (FAT's entries have 12 bits). For partitions with 4085 to 65526 clusters (with 4085 and 65526 excluded), the type is FAT16. Currently `fatfsgen.py` or `fatfsparse.py` cannot process file systems with more than 65525 clusters.

Wear Levelling There are two types of operations related to the wear levelling layer, initializing wear leveling records and removing wear leveling records during generation and parsing of the FAT file system image.

1. Initializing Wear Levelling

When a new image with wear leveling support is generated, the script initializes few extra sectors necessary for the wear leveling function.

- The dummy sector: This is an empty sector placed at the beginning of the partition and it will be ignored when file system is being mounted. The dummy sector copies the content of the next sector and then swaps its position with the next sector (or the first sector in case dummy sector was the last) after particular number of erase cycles. In this way, each FAT file system sector traverses across the whole range of flash partition, and thus the erase cycles corresponding to this sector gets distributed across the entire flash.
- **The state sector: State sector has 64 byte data stored.**
 - pos: position of the dummy sector
 - max_pos: number of sectors in the partition (excluding config and state sectors)
 - move_count: indicates how many times dummy sector traversed through the entire flash
 - access_count: count of sector erase cycles after which dummy sector will swap its position
 - max_count: equal to `wl_config_t::updaterate`
 - block_size: equal to `wl_config_t::page_size`
 - version: equal to `wl_config_t::version`
 - device_id: generated randomly when the state is first initialized
 - reserved: 7 x 32-bit words, set to 0
 - crc32: crc32 of all the previous fields, including reserved

Also, the state sector will be appended by 16-byte `pos update record` for every value of `pos`. Thus, this record will help us to determine the position of the dummy sector.

Since `erase + write` operation of the state sector is not atomic, we may lose the data if the power is cut off between "erase" and "write". However, two copies of the state are maintained to recover the state after the power outage. On each update, both copies are updated. Thus, after power outage, we can revert the original valid state.

- **The config sector: This sector contains the information about the partition used by the wear leveling layer.**
 - start_addr: start address of partition (always 0)
 - full_mem_size: size of the partition, including data, dummy, state x 2, config sectors. Value is in bytes
 - page_size: equal to sector size (generally 4096)
 - sector_size: always 4096 for the types of NOR flash supported by ESP-IDF
 - updaterate: ESP-IDF always sets this to 16. Could be made a config option at some point
 - wr_size: always set to 16
 - version: current version is 2
 - temp_buff_size: always set to 32 (This shouldn't actually have been stored in flash)
 - crc: crc32 of all the previous values

2. Removing Wear Levelling While removing wear leveling records, we have to find the position of the dummy sector, and the original and valid orders of the partition (because traversing the dummy sector shuffles the partition). The script can remove other wear leveling sectors from the partition. Steps to remove wear leveling records are given below:

- Find the `pos`, position of the dummy sector, which will be determined by the number of `pos update records` in the state sector.
- Create the new image by removing dummy sector and merging remaining sectors before and after dummy sector.
- Then remove the wear leveling state sectors and config sector which are placed at the end of the partition.
- Reorder the new image to get its original order. `move_count` helps us to find the beginning of the partition. The partition will start at the position `end_of_partition - move_count`. Thus the beginning of the partition after removing wear leveling sectors will be `partition[end_of_partition - (move_count*page_size)]`.

File Names Encoding The protocol FAT supports two types of file names.

Short File Names (SFN) The SFN is mandatory for the implementation of file names. SFN refer to the 8.3 file name convention, with 8 characters for the file name and 3 characters for the extension. This pattern is case-insensitive, however, all file names are changed to uppercase in the inner representation of the generator. The entry describing the short file names is 32 bytes long and its structure is as follows:

Offset:	00	01	02	03	04	05	06	07	08	09	0A	0B	0C	0D	0E	0F	
0x000000:	46	49	4C	45	4E	41	4D	45	45	58	54	20	18	00	00	00	FILENAMEEXT.....
0x000010:	21	00	21	00	00	00	00	00	21	00	02	00	1E	00	00	00	!!.....!.....

The entry denotes the file with 8.3 file name ("FILENAME.EXT") `__(0x00/00-0A)__` of size `0x1E = 30` bytes `__(0x10/0x0C)__,` with default times of modification and creation (`0x0021`) `__(0x10/00,02 and 08)__. The relevant cluster for the file is located at __0x02 (0x10/0A)__. Please notice that a character is encoded using one byte (e.g., __0x46 == 'F'__`

Long File Names (LFN) The LFN supports 255 characters excluding the trailing NULL. The LFN supports any character as short file names with an additional period `.` and the following special characters: `+ , ; = [ ]`. LFN uses UNICODE, so the character is encoded using 2 bytes.

The structure of one name encoded using LFN is as follows:

00003000:	42	65	00	2E	00	74	00	78	00	74	00	0F	00	43	FF	FF	Be...t.x.t...C..
00003010:	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	00	00	FF	FF	FF	FF	
00003020:	01	74	00	68	00	69	00	73	00	69	00	0F	00	43	73	00	.t.h.i.s.i...Cs.
00003030:	6C	00	6F	00	6E	00	67	00	66	00	00	00	69	00	6C	00	l.o.n.g.f...i.l.
00003040:	54	48	49	53	49	53	7E	31	54	58	54	20	00	00	D6	45	THISIS~1TXT...VE
00003050:	26	55	26	55	00	00	D6	45	26	55	02	00	1C	00	00	00	&U&U...VE&U.....

The above example encodes a file name `thisislongfile.txt`. The record is composed of multiple entries. The first entry contains metadata and is equivalent to the SFN entry. This entry might be final if the file name conforms to the 8.3 file name convention. In such scenarios, the SFN pattern is used. Otherwise, the generator adds various entries with the LFN structure above the SFN entry. These entries hold information about the file name and its checksum for consistency. Every LFN record can hold 13 characters (26 bytes). The file name is firstly cut into some amount of 13-character substrings and these are added above the SFN entry.

We add LFN entries in reversed order, so the first entry in the directory is the last part of the file name and the last is SFN entry. In the above example, we can see that the first entry contains text `e.txt`, while the others contain the beginning of the name `thisislongfil`. The first byte in LFN entries denotes an order or the sequence number (numbered from 1). To determine the first entry of the LFN, the first byte is masked with `0x40` (`first_byte = | 0x40`). The specification says that the last entry value will be Ored with `0x40` and it is the mark for the last entry. For example, when the record is the second and also the last in the LFN entry, its first byte is `0x42`.

The LFN entry is signed at field **DIR_Attr** with value `ATTR_READ_ONLY | ATTR_HIDDEN | ATTR_SYSTEM | ATTR_VOLUME_ID` (see the file `long_filename_utils.py`). The SFN entry (possibly also within LFN) contains either `ATTR_DIRECTORY` or `ATTR_ARCHIVE` in this field for directory or file respectively.

The LFN entry is tagged at the field **DIR_NTRes** with the value `0x00`. This is a sign of the SFN entry in the LFN record, if the entry is a whole SFN record, the value is `0x18`. As you can see in the first example, the value at this field is `0x18`, because the name **"FILENAME.EXT"** fits the SFN. However, the recent example showing **"thisislongfile.txt"** has value `0x00` at field **DIR_NTRes** in the last entry, since it is a LFN. The SFN needs to be unique. For that purpose, the `fatfsgen.py` uses the first 6 characters from the file name, concatenating with `~` and with ID denoting the order of the name with the same prefix. The ID is between 0 to 127, which is the maximal amount of files with the same prefix.

Calculation of the checksum is described and implemented in the `utils.py` by function `lfn_checksum`. The `fatfsparse.py` assumes that the LFN entries might not be right next to each other, but it assumes the relative order is preserved. The approach is first to find the SFN belonging to some LFN record (using **DIR_NTRes** field). From then, the script starts to search by moving upwards to the beginning of the respective sector, until it finds the last entry in the LFN record (the one with the first half byte equal to 4). The entries are distinguished by their checksums. When finished, the file name can be composed.

Date and Time in FAT File System The FAT file system protocol used by ESP-IDF does not preserve the date or time on the chips' media, so all the images extracted from the device have the same default timestamp for all the FAT-specified date-time fields (creation and the last modification timestamp as well as creation, last modification and last access dates).

There are a couple of fields in the SFN entry describing time, such as **DIR_CrtTime** and **DIR_WrtTime**. Some fields are ignored by the FAT implementation used by ESP-IDF (see the file `entry.py`). However, changes in the fields **DIR_WrtTime** and **DIR_WrtDate** are preserved in the chip. Both time and data entry are 16-bit, where the granularity of the time is 2 seconds.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.8.3 Manufacturing Utility

Introduction

This utility is designed to create instances of factory NVS partition images on a per-device basis for mass manufacturing purposes. The NVS partition images are created from CSV files containing user-provided configurations and values.

Please note that this utility only creates manufacturing binary images which then need to be flashed onto your devices using:

- [esptool.py](#)
- **Flash Download tool (available on Windows only)**
 - Download and unzip it, and follow the instructions inside the *doc* folder.
- Direct flash programming using custom production tools.

Prerequisites

This utility is dependent on ESP-IDF's NVS Partition Generator Utility.


- **Operating System requirements:**
 - Linux / MacOS / Windows (standard distributions)
- **The following packages are needed to use this utility:**
 - [Python](#)

Note:

Before using this utility, please make sure that:

- The path to Python is added to the PATH environment variable.
 - You have installed the packages from *requirement.txt*, the file in the root of the ESP-IDF directory.
-

Workflow



CSV Configuration File

This file contains the configuration of the device to be flashed.

The data in the configuration file has the following format (the *REPEAT* tag is optional):

```
name1,namespace,      <-- First entry should be of type "namespace"
key1,type1,encoding1
key2,type2,encoding2,REPEAT
name2,namespace,
key3,type3,encoding3
key4,type4,encoding4
```

Note: The first line in this file should always be the namespace entry.

Each line should have three parameters: *key*, *type*, *encoding*, separated by a comma. If the *REPEAT* tag is present, the value corresponding to this key in the master value CSV file will be the same for all devices.

Please refer to README of the NVS Partition Generator Utility for detailed description of each parameter.

Below is a sample example of such a configuration file:

```
app,namespace,
firmware_key,data,hex2bin
serial_no,data,string,REPEAT
device_no,data,i32
```

Note:

Make sure there are no spaces:

- before and after ','
 - at the end of each line in a CSV file
-

Master Value CSV File

This file contains details of the devices to be flashed. Each line in this file corresponds to a device instance.

The data in the master value CSV file has the following format:

```
key1,key2,key3,....
value1,value2,value3,....
```

Note: The first line in the file should always contain the *key* names. All the keys from the configuration file should be present here in the **same order**. This file can have additional columns (keys). The additional keys will be treated as metadata and would not be part of the final binary files.

Each line should contain the `value` of the corresponding keys, separated by a comma. If the key has the `REPEAT` tag, its corresponding value **must** be entered in the second line only. Keep the entry empty for this value in the following lines.

The description of this parameter is as follows:

value Data value

Data value is the value of data corresponding to the key.

Below is a sample example of a master value CSV file:

```
id,firmware_key,serial_no,device_no
1,1a2b3c4d5e6faabb,A1,101
2,1a2b3c4d5e6fccdd,,102
3,1a2b3c4d5e6feeff,,103
```

Note: If the 'REPEAT' tag is present, a new master value CSV file will be created in the same folder as the input Master CSV File with the values inserted at each line for the key with the 'REPEAT' tag.

This utility creates intermediate CSV files which are used as input for the NVS partition utility to generate the binary files.

The format of this intermediate CSV file is as follows:

```
key,type,encoding,value
key,namespace, ,
key1,type1,encoding1,value1
key2,type2,encoding2,value2
```

An instance of an intermediate CSV file will be created for each device on an individual basis.

Running the utility

Usage:

```
python mfg_gen.py [-h] {generate,generate-key} ...
```

Optional Arguments:

No.	Parameter	Description
1	-h / --help	Show the help message and exit

Commands:

Run `mfg_gen.py {command} -h` for additional help

No.	Parameter	Description
1	generate	Generate NVS partition
2	generate-key	Generate keys for encryption

To generate factory images for each device (Default):

Usage:

```
python mfg_gen.py generate [-h] [--fileid FILEID] [--version {1,2}] [--keygen]
                             [--inputkey INPUTKEY] [--outdir OUTDIR]
                             [--key_protect_hmac] [--kp_hmac_keygen]
```

(continues on next page)

(continued from previous page)

```

[--kp_hmac_keyfile KP_HMAC_KEYFILE] [--kp_hmac_
↪inputkey KP_HMAC_INPUTKEY]
conf values prefix size

```

Positional Arguments:

Parameter	Description
conf	Path to configuration csv file to parse
values	Path to values csv file to parse
prefix	Unique name for each output filename prefix
size	Size of NVS partition in bytes (must be multiple of 4096)

Optional Arguments:

Parameter	Description
-h / --help	Show the help message and exit
--fileid FILEID	Unique file identifier (any key in values file) for each filename suffix (Default: numeric value(1,2,3...))
--version {1,2}	Set multipage blob version. (Default: Version 2) Version 1 - Multipage blob support disabled. Version 2 - Multipage blob support enabled.
--keygen	Generates key for encrypting NVS partition
--inputkey IN- PUTKEY	File having key for encrypting NVS partition
--outdir OUTDIR	Output directory to store files created (Default: current directory)
--key_protect_hmac	If set, the NVS encryption key protection scheme based on HMAC peripheral is used; else the default scheme based on Flash Encryption is used
--kp_hmac_keygen	Generate the HMAC key for HMAC-based encryption scheme
--kp_hmac_keyfile KP_HMAC_KEYFILE	Path to output HMAC key file
--kp_hmac_inputkey KP_HMAC_INPUTKEY	File having the HMAC key for generating the NVS encryption keys

You can run the utility to generate factory images for each device using the command below. A sample CSV file is provided with the utility:

```

python mfg_gen.py generate samples/sample_config.csv samples/sample_values_
↪singlepage_blob.csv Sample 0x3000

```

The master value CSV file should have the path in the `file` type relative to the directory from which you are running the utility.

To generate encrypted factory images for each device:

You can run the utility to encrypt factory images for each device using the command below. A sample CSV file is provided with the utility:

- Encrypt by allowing the utility to generate encryption keys:

```

python mfg_gen.py generate samples/sample_config.csv samples/sample_values_
↪singlepage_blob.csv Sample 0x3000 --keygen

```

Note: Encryption key of the following format `<outdir>/keys/keys-<prefix>-<fileid>.bin` is created. This newly created file having encryption keys in `keys/` directory is compatible with NVS key-partition structure. Refer to [NVS Key Partition](#) for more details.

- To generate an encrypted image using the HMAC-based scheme, the above command can be used along with some additional parameters.

- Encrypt by allowing the utility to generate encryption keys and the HMAC-key:

```
python mfg_gen.py generate samples/sample_config.csv samples/
↪sample_values_singlepage_blob.csv Sample 0x3000 --keygen --key_
↪protect_hmac --kp_hmac_keygen
```

Note: Encryption key of the format <outdir>/keys/keys-<timestamp>.bin and HMAC key of the format <outdir>/keys/hmac-keys-<timestamp>.bin are created.

- Encrypt by allowing the utility to generate encryption keys with user-provided HMAC-key:

```
python mfg_gen.py generate samples/sample_config.csv samples/sample_values_
↪singlepage_blob.csv Sample 0x3000 --keygen --key_protect_hmac --kp_hmac_
↪inputkey testdata/sample_hmac_key.bin
```

Note: You can provide the custom filename for the HMAC key as well as the encryption key as a parameter.

- Encrypt by providing the encryption keys as input binary file:

```
python mfg_gen.py generate samples/sample_config.csv samples/sample_values_
↪singlepage_blob.csv Sample 0x3000 --inputkey keys/sample_keys.bin
```

To generate only encryption keys:

Usage:: python mfg_gen.py generate-key [-h] [--keyfile KEYFILE] [--outdir OUTDIR]

Optional Arguments:

Parameter	Description
-h / --help	Show the help message and exit
--keyfile KEYFILE	Path to output encryption keys file
--outdir OUTDIR	Output directory to store files created. (Default: current directory)
--key_protect_hmac	If set, the NVS encryption key protection scheme based on HMAC peripheral is used; else the default scheme based on Flash Encryption is used
--kp_hmac_keygen	Generate the HMAC key for HMAC-based encryption scheme
--kp_hmac_keyfile KP_HMAC_KEYFILE	Path to output HMAC key file
--kp_hmac_inputkey KP_HMAC_INPUTKEY	File having the HMAC key for generating the NVS encryption keys

You can run the utility to generate only encryption keys using the command below:

```
python mfg_gen.py generate-key
```

Note: Encryption key of the following format <outdir>/keys/keys-<timestamp>.bin is created. Timestamp format is: %m-%d_%H-%M. To provide custom target filename use the --keyfile argument.

For generating encryption key for the HMAC-based scheme, the following commands can be used:

- Generate the HMAC key and the NVS encryption keys:

```
python mfg_gen.py generate-key --key_protect_hmac --kp_hmac_keygen
```

Note: Encryption key of the format <outdir>/keys/keys-<timestamp>.bin and HMAC key of the format <outdir>/keys/hmac-keys-<timestamp>.bin are created.

- Generate the NVS encryption keys, given the HMAC-key:

```
python mfg_gen.py generate-key --key_protect_hmac --kp_hmac_inputkey testdata/  
↪sample_hmac_key.bin
```

Note: You can provide the custom filename for the HMAC key as well as the encryption key as a parameter.

Generated encryption key binary file can further be used to encrypt factory images created on the per device basis.

The default numeric value: 1,2,3... of the `fileid` argument corresponds to each line bearing device instance values in the master value CSV file.

While running the manufacturing utility, the following folders will be created in the specified `outdir` directory:

- `bin/` for storing the generated binary files
- `csv/` for storing the generated intermediate CSV files
- `keys/` for storing encryption keys (when generating encrypted factory images)

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.8.4 Non-Volatile Storage Library

Introduction

Non-volatile storage (NVS) library is designed to store key-value pairs in flash. This section introduces some concepts used by NVS.

Underlying Storage Currently, NVS uses a portion of main flash memory through the [esp_partition](#) API. The library uses all the partitions with `data` type and `nvs` subtype. The application can choose to use the partition with the label `nvs` through the [nvs_open\(\)](#) API function or any other partition by specifying its name using the [nvs_open_from_partition\(\)](#) API function.

Future versions of this library may have other storage backends to keep data in another flash chip (SPI or I2C), RTC, FRAM, etc.

Note: if an NVS partition is truncated (for example, when the partition table layout is changed), its contents should be erased. ESP-IDF build system provides a `idf.py erase-flash` target to erase all contents of the flash chip.

Note: NVS works best for storing many small values, rather than a few large values of the type 'string' and 'blob'. If you need to store large blobs or strings, consider using the facilities provided by the FAT filesystem on top of the wear levelling library.

Keys and Values NVS operates on key-value pairs. Keys are ASCII strings; the maximum key length is currently 15 characters. Values can have one of the following types:

- integer types: `uint8_t`, `int8_t`, `uint16_t`, `int16_t`, `uint32_t`, `int32_t`, `uint64_t`, `int64_t`
- zero-terminated string

- variable length binary data (blob)

Note: String values are currently limited to 4000 bytes. This includes the null terminator. Blob values are limited to 508,000 bytes or 97.6% of the partition size - 4000 bytes, whichever is lower.

Additional types, such as `float` and `double` might be added later.

Keys are required to be unique. Assigning a new value to an existing key replaces the old value and data type with the value and data type specified by a write operation.

A data type check is performed when reading a value. An error is returned if the data type expected by read operation does not match the data type of entry found for the key provided.

Namespaces To mitigate potential conflicts in key names between different components, NVS assigns each key-value pair to one of namespaces. Namespace names follow the same rules as key names, i.e., the maximum length is 15 characters. Furthermore, there can be no more than 254 different namespaces in one NVS partition. Namespace name is specified in the `nvs_open()` or `nvs_open_from_partition` call. This call returns an opaque handle, which is used in subsequent calls to the `nvs_get_*`, `nvs_set_*`, and `nvs_commit()` functions. This way, a handle is associated with a namespace, and key names will not collide with same names in other namespaces. Please note that the namespaces with the same name in different NVS partitions are considered as separate namespaces.

NVS Iterators Iterators allow to list key-value pairs stored in NVS, based on specified partition name, namespace, and data type.

There are the following functions available:

- `nvs_entry_find()` creates an opaque handle, which is used in subsequent calls to the `nvs_entry_next()` and `nvs_entry_info()` functions.
- `nvs_entry_next()` advances an iterator to the next key-value pair.
- `nvs_entry_info()` returns information about each key-value pair

In general, all iterators obtained via `nvs_entry_find()` have to be released using `nvs_release_iterator()`, which also tolerates NULL iterators.

`nvs_entry_find()` and `nvs_entry_next()` set the given iterator to NULL or a valid iterator in all cases except a parameter error occurred (i.e., return `ESP_ERR_NVS_NOT_FOUND`). In case of a parameter error, the given iterator will not be modified. Hence, it is best practice to initialize the iterator to NULL before calling `nvs_entry_find()` to avoid complicated error checking before releasing the iterator.

Security, Tampering, and Robustness NVS is not directly compatible with the ESP32-P4 flash encryption system. However, data can still be stored in encrypted form if NVS encryption is used together with ESP32-P4 flash encryption or with the help of the HMAC peripheral. Please refer to [NVS Encryption](#) for more details.

If NVS encryption is not used, it is possible for anyone with physical access to the flash chip to alter, erase, or add key-value pairs. With NVS encryption enabled, it is not possible to alter or add a key-value pair and get recognized as a valid pair without knowing corresponding NVS encryption keys. However, there is no tamper-resistance against the erase operation.

The library does try to recover from conditions when flash memory is in an inconsistent state. In particular, one should be able to power off the device at any point and time and then power it back on. This should not result in loss of data, except for the new key-value pair if it was being written at the moment of powering off. The library should also be able to initialize properly with any random data present in flash memory.

NVS Encryption

Please refer to the [NVS Encryption](#) guide for more details.

NVS Partition Generator Utility

This utility helps generate NVS partition binary files which can be flashed separately on a dedicated partition via a flashing utility. Key-value pairs to be flashed onto the partition can be provided via a CSV file. For more details, please refer to [NVS Partition Generator Utility](#).

Instead of calling the `nvs_partition_gen.py` tool manually, the creation of the partition binary files can also be done directly from CMake using the function `nvs_create_partition_image`:

```
nvs_create_partition_image(<partition> <csv> [FLASH_IN_PROJECT] [DEPENDS dep dep_
↪dep ...])
```

Positional Arguments:

Parameter	Description
<code>partition</code>	Name of the NVS partition
<code>csv</code>	Path to CSV file to parse

Optional Arguments:

Parameter	Description
<code>FLASH_IN_PROJECT</code>	Name of the NVS partition
<code>DEPENDS</code>	Specify files on which the command depends

If `FLASH_IN_PROJECT` is not specified, the image will still be generated, but you will have to flash it manually using `idf.py <partition>-flash` (e.g., if your partition name is `nvs`, then use `idf.py nvs-flash`).

`nvs_create_partition_image` must be called from one of the component `CMakeLists.txt` files. Currently, only non-encrypted partitions are supported.

Application Example

You can find code examples in the [storage](#) directory of ESP-IDF examples:

[storage/nvs_rw_value](#)

Demonstrates how to read a single integer value from, and write it to NVS.

The value checked in this example holds the number of the ESP32-P4 module restarts. The value's function as a counter is only possible due to its storing in NVS.

The example also shows how to check if a read/write operation was successful, or if a certain value has not been initialized in NVS. The diagnostic procedure is provided in plain text to help you track the program flow and capture any issues on the way.

[storage/nvs_rw_blob](#)

Demonstrates how to read a single integer value and a blob (binary large object), and write them to NVS to preserve this value between ESP32-P4 module restarts.

- `value` - tracks the number of the ESP32-P4 module soft and hard restarts.
- `blob` - contains a table with module run times. The table is read from NVS to dynamically allocated RAM. A new run time is added to the table on each manually triggered soft restart, and then the added run time is written to NVS. Triggering is done by pulling down GPIO0.

The example also shows how to implement the diagnostic procedure to check if the read/write operation was successful.

[storage/nvs_rw_value_cxx](#)

This example does exactly the same as [storage/nvs_rw_value](#), except that it uses the C++ NVS handle class.

storage/nvs_statistics

This example demonstrates how to obtain and interpret NVS usage statistics: free/used/available/total number of entries and number of namespaces in given NVS partition.

Default NVS partition is erased for a clean run of this example. Then mock data string values are written.

Usage statistics are obtained prior to and post writing, with the differences being compared to expected values of newly used entries.

storage/nvs_iteration

This example demonstrates how to iterate over entries of specific (or any) NVS data type and how to obtain info about such entries.

Default NVS partition is erased for a clean run of this example. Then mock data consisting of different NVS integer data types are written.

After that, the example iterates over each individual data type as well as the generic NVS_TYPE_ANY type, and logs the information obtained from each iteration.

Internals

Log of Key-Value Pairs NVS stores key-value pairs sequentially, with new key-value pairs being added at the end. When a value of any given key has to be updated, a new key-value pair is added at the end of the log and the old key-value pair is marked as erased.

Pages and Entries NVS library uses two main entities in its operation: pages and entries. Page is a logical structure which stores a portion of the overall log. Logical page corresponds to one physical sector of flash memory. Pages which are in use have a *sequence number* associated with them. Sequence numbers impose an ordering on pages. Higher sequence numbers correspond to pages which were created later. Each page can be in one of the following states:

Empty/uninitialized Flash storage for the page is empty (all bytes are 0xff). Page is not used to store any data at this point and does not have a sequence number.


Active Flash storage is initialized, page header has been written to flash, page has a valid sequence number. Page has some empty entries and data can be written there. No more than one page can be in this state at any given moment.

Full Flash storage is in a consistent state and is filled with key-value pairs. Writing new key-value pairs into this page is not possible. It is still possible to mark some key-value pairs as erased.

Erasing Non-erased key-value pairs are being moved into another page so that the current page can be erased. This is a transient state, i.e., page should never stay in this state at the time when any API call returns. In case of a sudden power off, the move-and-erase process will be completed upon the next power-on.

Corrupted Page header contains invalid data, and further parsing of page data was canceled. Any items previously written into this page will not be accessible. The corresponding flash sector will not be erased immediately and will be kept along with sectors in **uninitialized** state for later use. This may be useful for debugging.


Mapping from flash sectors to logical pages does not have any particular order. The library will inspect sequence numbers of pages found in each flash sector and organize pages in a list based on these numbers.



Structure of a Page For now, we assume that flash sector size is 4096 bytes and that ESP32-P4 flash encryption hardware operates on 32-byte blocks. It is possible to introduce some settings configurable at compile-time (e.g., via menuconfig) to accommodate flash chips with different sector sizes (although it is not clear if other components in the system, e.g., SPI flash driver and SPI flash cache can support these other sizes).

Page consists of three parts: header, entry state bitmap, and entries themselves. To be compatible with ESP32-P4 flash encryption, the entry size is 32 bytes. For integer types, an entry holds one key-value pair. For strings and blobs, an entry holds part of key-value pair (more on that in the entry structure description).

The following diagram illustrates the page structure. Numbers in parentheses indicate the size of each part in bytes.



Page header and entry state bitmap are always written to flash unencrypted. Entries are encrypted if flash encryption feature of ESP32-P4 is used.

Page state values are defined in such a way that changing state is possible by writing 0 into some of the bits. Therefore it is not necessary to erase the page to change its state unless that is a change to the *erased* state.

The version field in the header reflects the NVS format version used. For backward compatibility reasons, it is decremented for every version upgrade starting at 0xff (i.e., 0xff for version-1, 0xfe for version-2 and so on).

CRC32 value in the header is calculated over the part which does not include a state value (bytes 4 to 28). The unused part is currently filled with 0xff bytes.

The following sections describe the structure of entry state bitmap and entry itself.


Entry and Entry State Bitmap Each entry can be in one of the following three states represented with two bits in the entry state bitmap. The final four bits in the bitmap (256 - 2 * 126) are not used.

Empty (2'b11) Nothing is written into the specific entry yet. It is in an uninitialized state (all bytes are 0xff).

Written (2'b10) A key-value pair (or part of key-value pair which spans multiple entries) has been written into the entry.


Erased (2'b00) A key-value pair in this entry has been discarded. Contents of this entry will not be parsed anymore.

Structure of Entry For values of primitive types (currently integers from 1 to 8 bytes long), entry holds one key-value pair. For string and blob types, entry holds part of the whole key-value pair. For strings, in case when a key-value pair spans multiple entries, all entries are stored in the same page. Blobs are allowed to span over multiple pages by dividing them into smaller chunks. For tracking these chunks, an additional fixed length metadata entry is stored called "blob index". Earlier formats of blobs are still supported (can be read and modified). However, once the blobs are modified, they are stored using the new format.



(continues on next page)

(continued from previous page)



Individual fields in entry structure have the following meanings:

NS Namespace index for this entry. For more information on this value, see the section on namespaces implementation.

Type One byte indicating the value data type. See the `ItemType` enumeration in `nvs_flash/include/nvs_handle.hpp` for possible values.

Span Number of entries used by this key-value pair. For integer types, this is equal to 1. For strings and blobs, this depends on value length.

ChunkIndex Used to store the index of a blob-data chunk for blob types. For other types, this should be `0xff`.

CRC32 Checksum calculated over all the bytes in this entry, except for the CRC32 field itself.

Key Zero-terminated ASCII string containing a key name. Maximum string length is 15 bytes, excluding a zero terminator.

Data For integer types, this field contains the value itself. If the value itself is shorter than 8 bytes, it is padded to the right, with unused bytes filled with `0xff`.

For "blob index" entry, these 8 bytes hold the following information about data-chunks:

- **Size** (Only for blob index.) Size, in bytes, of complete blob data.
- **ChunkCount** (Only for blob index.) Total number of blob-data chunks into which the blob was divided during storage.
- **ChunkStart** (Only for blob index.) ChunkIndex of the first blob-data chunk of this blob. Subsequent chunks have chunkIndex incrementally allocated (step of 1).

For string and blob data chunks, these 8 bytes hold additional data about the value, which are described below:

- **Size** (Only for strings and blobs.) Size, in bytes, of actual data. For strings, this includes zero terminators.
- **CRC32** (Only for strings and blobs.) Checksum calculated over all bytes of data.

Variable length values (strings and blobs) are written into subsequent entries, 32 bytes per entry. The `Span` field of the first entry indicates how many entries are used.

Namespaces As mentioned above, each key-value pair belongs to one of the namespaces. Namespace identifiers (strings) are stored as keys of key-value pairs in namespace with index 0. Values corresponding to these keys are indexes of these namespaces.

NS=0 Type=uint8_t Key="wifi" Value=1	Entry describing namespace "wifi"
NS=1 Type=uint32_t Key="channel" Value=6	Key "channel" in namespace "wifi"
NS=0 Type=uint8_t Key="pwm" Value=2	Entry describing namespace "pwm"
NS=2 Type=uint16_t Key="channel" Value=20	Key "channel" in namespace "pwm"

Item Hash List To reduce the number of reads from flash memory, each member of the Page class maintains a list of pairs: item index; item hash. This list makes searches much quicker. Instead of iterating over all entries, reading them from flash one at a time, *Page::findItem* first performs a search for the item hash in the hash list. This gives the item index within the page if such an item exists. Due to a hash collision, it is possible that a different item is found. This is handled by falling back to iteration over items in flash.

Each node in the hash list contains a 24-bit hash and 8-bit item index. Hash is calculated based on item namespace, key name, and ChunkIndex. CRC32 is used for calculation; the result is truncated to 24 bits. To reduce the overhead for storing 32-bit entries in a linked list, the list is implemented as a double-linked list of arrays. Each array holds 29 entries, for the total size of 128 bytes, together with linked list pointers and a 32-bit count field. The minimum amount of extra RAM usage per page is therefore 128 bytes; maximum is 640 bytes.

API Reference

Header File

- [components/nvs_flash/include/nvs_flash.h](#)
- This header file can be included with:

```
#include "nvs_flash.h"
```

- This header file is a part of the API provided by the `nvs_flash` component. To declare that your component depends on `nvs_flash`, add the following to your CMakeLists.txt:

```
REQUIRES nvs_flash
```

or

```
PRIV_REQUIRES nvs_flash
```

Functions

esp_err_t **nvs_flash_init** (void)

Initialize the default NVS partition.

This API initialises the default NVS partition. The default NVS partition is the one that is labeled "nvs" in the partition table.

When "NVS_ENCRYPTION" is enabled in the menuconfig, this API enables the NVS encryption for the default NVS partition as follows

- a. Read security configurations from the first NVS key partition listed in the partition table. (NVS key partition is any "data" type partition which has the subtype value set to "nvs_keys")
- b. If the NVS key partition obtained in the previous step is empty, generate and store new keys in that NVS key partition.
- c. Internally call "nvs_flash_secure_init()" with the security configurations obtained/generated in the previous steps.

Post initialization NVS read/write APIs remain the same irrespective of NVS encryption.

Returns

- ESP_OK if storage was successfully initialized.
- ESP_ERR_NVS_NO_FREE_PAGES if the NVS storage contains no empty pages (which may happen if NVS partition was truncated)
- ESP_ERR_NOT_FOUND if no partition with label "nvs" is found in the partition table
- ESP_ERR_NO_MEM in case memory could not be allocated for the internal structures
- one of the error codes from the underlying flash storage driver
- error codes from `nvs_flash_read_security_cfg` API (when "NVS_ENCRYPTION" is enabled).
- error codes from `nvs_flash_generate_keys` API (when "NVS_ENCRYPTION" is enabled).

- error codes from `nvs_flash_secure_init_partition` API (when "NVS_ENCRYPTION" is enabled) .

esp_err_t **nvs_flash_init_partition** (const char *partition_label)

Initialize NVS flash storage for the specified partition.

Parameters `partition_label` -- [in] Label of the partition. Must be no longer than 16 characters.

Returns

- ESP_OK if storage was successfully initialized.
- ESP_ERR_NVS_NO_FREE_PAGES if the NVS storage contains no empty pages (which may happen if NVS partition was truncated)
- ESP_ERR_NOT_FOUND if specified partition is not found in the partition table
- ESP_ERR_NO_MEM in case memory could not be allocated for the internal structures
- one of the error codes from the underlying flash storage driver

esp_err_t **nvs_flash_init_partition_ptr** (const *esp_partition_t* *partition)

Initialize NVS flash storage for the partition specified by partition pointer.

Parameters `partition` -- [in] pointer to a partition obtained by the ESP partition API.

Returns

- ESP_OK if storage was successfully initialized
- ESP_ERR_NVS_NO_FREE_PAGES if the NVS storage contains no empty pages (which may happen if NVS partition was truncated)
- ESP_ERR_INVALID_ARG in case partition is NULL
- ESP_ERR_NO_MEM in case memory could not be allocated for the internal structures
- one of the error codes from the underlying flash storage driver

esp_err_t **nvs_flash_deinit** (void)

Deinitialize NVS storage for the default NVS partition.

Default NVS partition is the partition with "nvs" label in the partition table.

Returns

- ESP_OK on success (storage was deinitialized)
- ESP_ERR_NVS_NOT_INITIALIZED if the storage was not initialized prior to this call

esp_err_t **nvs_flash_deinit_partition** (const char *partition_label)

Deinitialize NVS storage for the given NVS partition.

Parameters `partition_label` -- [in] Label of the partition

Returns

- ESP_OK on success
- ESP_ERR_NVS_NOT_INITIALIZED if the storage for given partition was not initialized prior to this call

esp_err_t **nvs_flash_erase** (void)

Erase the default NVS partition.

Erases all contents of the default NVS partition (one with label "nvs").

Note: If the partition is initialized, this function first de-initializes it. Afterwards, the partition has to be initialized again to be used.

Returns

- ESP_OK on success
- ESP_ERR_NOT_FOUND if there is no NVS partition labeled "nvs" in the partition table
- different error in case de-initialization fails (shouldn't happen)

esp_err_t **nvs_flash_erase_partition** (const char *part_name)

Erase specified NVS partition.

Erase all content of a specified NVS partition

Note: If the partition is initialized, this function first de-initializes it. Afterwards, the partition has to be initialized again to be used.

Parameters **part_name** -- [in] Name (label) of the partition which should be erased

Returns

- ESP_OK on success
- ESP_ERR_NOT_FOUND if there is no NVS partition with the specified name in the partition table
- different error in case de-initialization fails (shouldn't happen)

esp_err_t **nvs_flash_erase_partition_ptr** (const *esp_partition_t* *partition)

Erase custom partition.

Erase all content of specified custom partition.

Note: If the partition is initialized, this function first de-initializes it. Afterwards, the partition has to be initialized again to be used.

Parameters **partition** -- [in] pointer to a partition obtained by the ESP partition API.

Returns

- ESP_OK on success
- ESP_ERR_NOT_FOUND if there is no partition with the specified parameters in the partition table
- ESP_ERR_INVALID_ARG in case partition is NULL
- one of the error codes from the underlying flash storage driver

esp_err_t **nvs_flash_secure_init** (*nvs_sec_cfg_t* *cfg)

Initialize the default NVS partition.

This API initialises the default NVS partition. The default NVS partition is the one that is labeled "nvs" in the partition table.

Parameters **cfg** -- [in] Security configuration (keys) to be used for NVS encryption/decryption. If cfg is NULL, no encryption is used.

Returns

- ESP_OK if storage has been initialized successfully.
- ESP_ERR_NVS_NO_FREE_PAGES if the NVS storage contains no empty pages (which may happen if NVS partition was truncated)
- ESP_ERR_NOT_FOUND if no partition with label "nvs" is found in the partition table
- ESP_ERR_NO_MEM in case memory could not be allocated for the internal structures
- one of the error codes from the underlying flash storage driver

esp_err_t **nvs_flash_secure_init_partition** (const char *partition_label, *nvs_sec_cfg_t* *cfg)

Initialize NVS flash storage for the specified partition.

Parameters

- **partition_label** -- [in] Label of the partition. Note that internally, a reference to passed value is kept and it should be accessible for future operations
- **cfg** -- [in] Security configuration (keys) to be used for NVS encryption/decryption. If cfg is null, no encryption/decryption is used.

Returns

- ESP_OK if storage has been initialized successfully.

- `ESP_ERR_NVS_NO_FREE_PAGES` if the NVS storage contains no empty pages (which may happen if NVS partition was truncated)
- `ESP_ERR_NOT_FOUND` if specified partition is not found in the partition table
- `ESP_ERR_NO_MEM` in case memory could not be allocated for the internal structures
- one of the error codes from the underlying flash storage driver

esp_err_t **nvs_flash_generate_keys** (const *esp_partition_t* *partition, *nvs_sec_cfg_t* *cfg)

Generate and store NVS keys in the provided esp partition.

Parameters

- **partition** -- [in] Pointer to partition structure obtained using `esp_partition_find_first` or `esp_partition_get`. Must be non-NULL.
- **cfg** -- [out] Pointer to nvs security configuration structure. Pointer must be non-NULL. Generated keys will be populated in this structure.

Returns

- `ESP_OK`, if cfg was read successfully;
- `ESP_ERR_INVALID_ARG`, if partition or cfg is NULL;
- or error codes from `esp_partition_write/erase` APIs.

esp_err_t **nvs_flash_read_security_cfg** (const *esp_partition_t* *partition, *nvs_sec_cfg_t* *cfg)

Read NVS security configuration from a partition.

Note: Provided partition is assumed to be marked 'encrypted'.

Parameters

- **partition** -- [in] Pointer to partition structure obtained using `esp_partition_find_first` or `esp_partition_get`. Must be non-NULL.
- **cfg** -- [out] Pointer to nvs security configuration structure. Pointer must be non-NULL.

Returns

- `ESP_OK`, if cfg was read successfully;
- `ESP_ERR_INVALID_ARG`, if partition or cfg is NULL
- `ESP_ERR_NVS_KEYS_NOT_INITIALIZED`, if the partition is not yet written with keys.
- `ESP_ERR_NVS_CORRUPT_KEY_PART`, if the partition containing keys is found to be corrupt
- or error codes from `esp_partition_read` API.

esp_err_t **nvs_flash_register_security_scheme** (*nvs_sec_scheme_t* *scheme_cfg)

Registers the given security scheme for NVS encryption The scheme registered with `sec_scheme_id` by this API be used as the default security scheme for the "nvs" partition. Users will have to call this API explicitly in their application.

Parameters **scheme_cfg** -- [in] Pointer to the security scheme configuration structure that the user (or the `nvs_key_provider`) wants to register.

Returns

- `ESP_OK`, if security scheme registration succeeds;
- `ESP_ERR_INVALID_ARG`, if `scheme_cfg` is NULL;
- `ESP_FAIL`, if security scheme registration fails

nvs_sec_scheme_t ***nvs_flash_get_default_security_scheme** (void)

Fetch the configuration structure for the default active security scheme for NVS encryption.

Returns Pointer to the default active security scheme configuration (NULL if no scheme is registered yet i.e. active)

esp_err_t **nvs_flash_generate_keys_v2** (*nvs_sec_scheme_t* *scheme_cfg, *nvs_sec_cfg_t* *cfg)

Generate (and store) the NVS keys using the specified key-protection scheme.

Parameters

- **scheme_cfg** -- [in] Security scheme specific configuration
- **cfg** -- [out] Security configuration (encryption keys)

Returns

- ESP_OK, if cfg was populated successfully with generated encryption keys;
- ESP_ERR_INVALID_ARG, if scheme_cfg or cfg is NULL;
- ESP_FAIL, if the key generation process fails

esp_err_t **nvs_flash_read_security_cfg_v2** (*nvs_sec_scheme_t* *scheme_cfg, *nvs_sec_cfg_t* *cfg)

Read NVS security configuration set by the specified security scheme.

Parameters

- **scheme_cfg** -- [in] Security scheme specific configuration
- **cfg** -- [out] Security configuration (encryption keys)

Returns

- ESP_OK, if cfg was read successfully;
- ESP_ERR_INVALID_ARG, if scheme_cfg or cfg is NULL;
- ESP_FAIL, if the key reading process fails

Structures

struct **nvs_sec_cfg_t**

Key for encryption and decryption.

Public Members

uint8_t **eky**[NVS_KEY_SIZE]

XTS encryption and decryption key

uint8_t **tky**[NVS_KEY_SIZE]

XTS tweak key

struct **nvs_sec_scheme_t**

NVS encryption: Security scheme configuration structure.

Public Members

int **scheme_id**

Security Scheme ID (E.g. HMAC)

void ***scheme_data**

Scheme-specific data (E.g. eFuse block for HMAC-based key generation)

nvs_flash_generate_keys_t **nvs_flash_key_gen**

Callback for the nvs_flash_key_gen implementation

nvs_flash_read_cfg_t **nvs_flash_read_cfg**

Callback for the nvs_flash_read_keys implementation

Macros

NVS_KEY_SIZE

Type Definitions

typedef *esp_err_t* (***nvs_flash_generate_keys_t**)(const void *scheme_data, *nvs_sec_cfg_t* *cfg)

Callback function prototype for generating the NVS encryption keys.

typedef *esp_err_t* (***nvs_flash_read_cfg_t**)(const void *scheme_data, *nvs_sec_cfg_t* *cfg)

Callback function prototype for reading the NVS encryption keys.

Header File

- `components/nvs_flash/include/nvs.h`
- This header file can be included with:

```
#include "nvs.h"
```

- This header file is a part of the API provided by the `nvs_flash` component. To declare that your component depends on `nvs_flash`, add the following to your `CMakeLists.txt`:

```
REQUIRES nvs_flash
```

or

```
PRIV_REQUIRES nvs_flash
```

Functions

esp_err_t **nvs_set_i8** (*nvs_handle_t* handle, const char *key, int8_t value)

set int8_t value for given key

Set value for the key, given its name. Note that the actual storage will not be updated until `nvs_commit` is called. Regardless whether key-value pair is created or updated, function always requires at least one `nvs` available entry. See `nvs_get_stats`. After create type of operation, the number of available entries is decreased by one. After update type of operation, the number of available entries remains the same.

Parameters

- **handle** -- [in] Handle obtained from `nvs_open` function. Handles that were opened read only cannot be used.
- **key** -- [in] Key name. Maximum length is `(NVS_KEY_NAME_MAX_SIZE-1)` characters. Shouldn't be empty.
- **value** -- [in] The value to set.

Returns

- `ESP_OK` if value was set successfully
- `ESP_FAIL` if there is an internal error; most likely due to corrupted NVS partition (only if NVS assertion checks are disabled)
- `ESP_ERR_NVS_INVALID_HANDLE` if handle has been closed or is NULL
- `ESP_ERR_NVS_READ_ONLY` if storage handle was opened as read only
- `ESP_ERR_NVS_INVALID_NAME` if key name doesn't satisfy constraints
- `ESP_ERR_NVS_NOT_ENOUGH_SPACE` if there is not enough space in the underlying storage to save the value
- `ESP_ERR_NVS_REMOVE_FAILED` if the value wasn't updated because flash write operation has failed. The value was written however, and update will be finished after re-initialization of `nvs`, provided that flash operation doesn't fail again.

esp_err_t **nvs_set_u8** (*nvs_handle_t* handle, const char *key, uint8_t value)

set uint8_t value for given key

This function is the same as `nvs_set_i8` except for the data type.

esp_err_t **nvs_set_i16** (*nvs_handle_t* handle, const char *key, int16_t value)

set int16_t value for given key

This function is the same as `nvs_set_i8` except for the data type.

esp_err_t **nvs_set_u16** (*nvs_handle_t* handle, const char *key, uint16_t value)
set uint16_t value for given key

This function is the same as `nvs_set_i8` except for the data type.

esp_err_t **nvs_set_i32** (*nvs_handle_t* handle, const char *key, int32_t value)
set int32_t value for given key

This function is the same as `nvs_set_i8` except for the data type.

esp_err_t **nvs_set_u32** (*nvs_handle_t* handle, const char *key, uint32_t value)
set uint32_t value for given key

This function is the same as `nvs_set_i8` except for the data type.

esp_err_t **nvs_set_i64** (*nvs_handle_t* handle, const char *key, int64_t value)
set int64_t value for given key

This function is the same as `nvs_set_i8` except for the data type.

esp_err_t **nvs_set_u64** (*nvs_handle_t* handle, const char *key, uint64_t value)
set uint64_t value for given key

This function is the same as `nvs_set_i8` except for the data type.

esp_err_t **nvs_set_str** (*nvs_handle_t* handle, const char *key, const char *value)
set string for given key

Sets string value for the key. Function requires whole space for new data to be available as contiguous entries in same nvs page. Operation consumes 1 overhead entry and 1 entry per each 32 characters of new string including zero character to be set. In case of value update for existing key, entries occupied by the previous value and overhead entry are returned to the pool of available entries. Note that storage of long string values can fail due to fragmentation of nvs pages even if `available_entries` returned by `nvs_get_stats` suggests enough overall space available. Note that the underlying storage will not be updated until `nvs_commit` is called.

Parameters

- **handle** -- [in] Handle obtained from `nvs_open` function. Handles that were opened read only cannot be used.
- **key** -- [in] Key name. Maximum length is (NVS_KEY_NAME_MAX_SIZE-1) characters. Shouldn't be empty.
- **value** -- [in] The value to set. For strings, the maximum length (including null character) is 4000 bytes, if there is one complete page free for writing. This decreases, however, if the free space is fragmented.

Returns

- `ESP_OK` if value was set successfully
- `ESP_ERR_NVS_INVALID_HANDLE` if handle has been closed or is NULL
- `ESP_ERR_NVS_READ_ONLY` if storage handle was opened as read only
- `ESP_ERR_NVS_INVALID_NAME` if key name doesn't satisfy constraints
- `ESP_ERR_NVS_NOT_ENOUGH_SPACE` if there is not enough space in the underlying storage to save the value
- `ESP_ERR_NVS_REMOVE_FAILED` if the value wasn't updated because flash write operation has failed. The value was written however, and update will be finished after re-initialization of nvs, provided that flash operation doesn't fail again.
- `ESP_ERR_NVS_VALUE_TOO_LONG` if the string value is too long

esp_err_t **nvs_get_i8** (*nvs_handle_t* handle, const char *key, int8_t *out_value)
get int8_t value for given key

These functions retrieve value for the key, given its name. If `key` does not exist, or the requested variable type doesn't match the type which was used when setting a value, an error is returned.

In case of any error, `out_value` is not modified.

`out_value` has to be a pointer to an already allocated variable of the given type.

```
// Example of using nvs_get_i32:
int32_t max_buffer_size = 4096; // default value
esp_err_t err = nvs_get_i32(my_handle, "max_buffer_size", &max_buffer_size);
assert(err == ESP_OK || err == ESP_ERR_NVS_NOT_FOUND);
// if ESP_ERR_NVS_NOT_FOUND was returned, max_buffer_size will still
// have its default value.
```

Parameters

- **handle** -- [in] Handle obtained from `nvs_open` function.
- **key** -- [in] Key name. Maximum length is (NVS_KEY_NAME_MAX_SIZE-1) characters. Shouldn't be empty.
- **out_value** -- Pointer to the output value. May be NULL for `nvs_get_str` and `nvs_get_blob`, in this case required length will be returned in length argument.

Returns

- ESP_OK if the value was retrieved successfully
- ESP_FAIL if there is an internal error; most likely due to corrupted NVS partition (only if NVS assertion checks are disabled)
- ESP_ERR_NVS_NOT_FOUND if the requested key doesn't exist
- ESP_ERR_NVS_INVALID_HANDLE if handle has been closed or is NULL
- ESP_ERR_NVS_INVALID_NAME if key name doesn't satisfy constraints
- ESP_ERR_NVS_INVALID_LENGTH if length is not sufficient to store data

esp_err_t **nvs_get_u8** (*nvs_handle_t* handle, const char *key, uint8_t *out_value)

get uint8_t value for given key

This function is the same as `nvs_get_i8` except for the data type.

esp_err_t **nvs_get_i16** (*nvs_handle_t* handle, const char *key, int16_t *out_value)

get int16_t value for given key

This function is the same as `nvs_get_i8` except for the data type.

esp_err_t **nvs_get_u16** (*nvs_handle_t* handle, const char *key, uint16_t *out_value)

get uint16_t value for given key

This function is the same as `nvs_get_i8` except for the data type.

esp_err_t **nvs_get_i32** (*nvs_handle_t* handle, const char *key, int32_t *out_value)

get int32_t value for given key

This function is the same as `nvs_get_i8` except for the data type.

esp_err_t **nvs_get_u32** (*nvs_handle_t* handle, const char *key, uint32_t *out_value)

get uint32_t value for given key

This function is the same as `nvs_get_i8` except for the data type.

esp_err_t **nvs_get_i64** (*nvs_handle_t* handle, const char *key, int64_t *out_value)

get int64_t value for given key

This function is the same as `nvs_get_i8` except for the data type.

esp_err_t **nvs_get_u64** (*nvs_handle_t* handle, const char *key, uint64_t *out_value)

get uint64_t value for given key

This function is the same as `nvs_get_i8` except for the data type.

esp_err_t **nvs_get_str** (*nvs_handle_t* handle, const char *key, char *out_value, size_t *length)

get string value for given key

These functions retrieve the data of an entry, given its key. If key does not exist, or the requested variable type doesn't match the type which was used when setting a value, an error is returned.

In case of any error, `out_value` is not modified.

All functions expect `out_value` to be a pointer to an already allocated variable of the given type.

`nvs_get_str` and `nvs_get_blob` functions support WinAPI-style length queries. To get the size necessary to store the value, call `nvs_get_str` or `nvs_get_blob` with zero `out_value` and non-zero pointer to length. Variable pointed to by length argument will be set to the required length. For `nvs_get_str`, this length includes the zero terminator. When calling `nvs_get_str` and `nvs_get_blob` with non-zero `out_value`, length has to be non-zero and has to point to the length available in `out_value`. It is suggested that `nvs_get/set_str` is used for zero-terminated C strings, and `nvs_get/set_blob` used for arbitrary data structures.

```
// Example (without error checking) of using nvs_get_str to get a string into
// dynamic array:
size_t required_size;
nvs_get_str(my_handle, "server_name", NULL, &required_size);
char* server_name = malloc(required_size);
nvs_get_str(my_handle, "server_name", server_name, &required_size);

// Example (without error checking) of using nvs_get_blob to get a binary data
// into a static array:
uint8_t mac_addr[6];
size_t size = sizeof(mac_addr);
nvs_get_blob(my_handle, "dst_mac_addr", mac_addr, &size);
```

Parameters

- **handle** -- [in] Handle obtained from `nvs_open` function.
- **key** -- [in] Key name. Maximum length is (NVS_KEY_NAME_MAX_SIZE-1) characters. Shouldn't be empty.
- **out_value** -- [out] Pointer to the output value. May be NULL for `nvs_get_str` and `nvs_get_blob`, in this case required length will be returned in length argument.
- **length** -- [inout] A non-zero pointer to the variable holding the length of `out_value`. In case `out_value` a zero, will be set to the length required to hold the value. In case `out_value` is not zero, will be set to the actual length of the value written. For `nvs_get_str` this includes zero terminator.

Returns

- ESP_OK if the value was retrieved successfully
- ESP_FAIL if there is an internal error; most likely due to corrupted NVS partition (only if NVS assertion checks are disabled)
- ESP_ERR_NVS_NOT_FOUND if the requested key doesn't exist
- ESP_ERR_NVS_INVALID_HANDLE if handle has been closed or is NULL
- ESP_ERR_NVS_INVALID_NAME if key name doesn't satisfy constraints
- ESP_ERR_NVS_INVALID_LENGTH if `length` is not sufficient to store data

esp_err_t **nvs_get_blob** (*nvs_handle_t* handle, const char *key, void *out_value, size_t *length)

get blob value for given key

This function behaves the same as `nvs_get_str`, except for the data type.

esp_err_t **nvs_open** (const char *namespace_name, *nvs_open_mode_t* open_mode, *nvs_handle_t* *out_handle)

Open non-volatile storage with a given namespace from the default NVS partition.

Multiple internal ESP-IDF and third party application modules can store their key-value pairs in the NVS module. In order to reduce possible conflicts on key names, each module can use its own namespace. The default NVS partition is the one that is labelled "nvs" in the partition table.

Parameters

- **namespace_name** -- [in] Namespace name. Maximum length is (NVS_KEY_NAME_MAX_SIZE-1) characters. Shouldn't be empty.
- **open_mode** -- [in] NVS_READWRITE or NVS_READONLY. If NVS_READONLY, will open a handle for reading only. All write requests will be rejected for this handle.
- **out_handle** -- [out] If successful (return code is zero), handle will be returned in this argument.

Returns

- ESP_OK if storage handle was opened successfully
- ESP_FAIL if there is an internal error; most likely due to corrupted NVS partition (only if NVS assertion checks are disabled)
- ESP_ERR_NVS_NOT_INITIALIZED if the storage driver is not initialized
- ESP_ERR_NVS_PART_NOT_FOUND if the partition with label "nvs" is not found
- ESP_ERR_NVS_NOT_FOUND id namespace doesn't exist yet and mode is NVS_READONLY
- ESP_ERR_NVS_INVALID_NAME if namespace name doesn't satisfy constraints
- ESP_ERR_NO_MEM in case memory could not be allocated for the internal structures
- ESP_ERR_NVS_NOT_ENOUGH_SPACE if there is no space for a new entry or there are too many different namespaces (maximum allowed different namespaces: 254)
- ESP_ERR_NOT_ALLOWED if the NVS partition is read-only and mode is NVS_READWRITE
- ESP_ERR_INVALID_ARG if out_handle is equal to NULL
- other error codes from the underlying storage driver

esp_err_t **nvs_open_from_partition** (const char *part_name, const char *namespace_name, *nvs_open_mode_t* open_mode, *nvs_handle_t* *out_handle)

Open non-volatile storage with a given namespace from specified partition.

The behaviour is same as `nvs_open()` API. However this API can operate on a specified NVS partition instead of default NVS partition. Note that the specified partition must be registered with NVS using `nvs_flash_init_partition()` API.

Parameters

- **part_name** -- [in] Label (name) of the partition of interest for object read/write/erase
- **namespace_name** -- [in] Namespace name. Maximum length is (NVS_KEY_NAME_MAX_SIZE-1) characters. Shouldn't be empty.
- **open_mode** -- [in] NVS_READWRITE or NVS_READONLY. If NVS_READONLY, will open a handle for reading only. All write requests will be rejected for this handle.
- **out_handle** -- [out] If successful (return code is zero), handle will be returned in this argument.

Returns

- ESP_OK if storage handle was opened successfully
- ESP_FAIL if there is an internal error; most likely due to corrupted NVS partition (only if NVS assertion checks are disabled)
- ESP_ERR_NVS_NOT_INITIALIZED if the storage driver is not initialized
- ESP_ERR_NVS_PART_NOT_FOUND if the partition with specified name is not found
- ESP_ERR_NVS_NOT_FOUND id namespace doesn't exist yet and mode is NVS_READONLY
- ESP_ERR_NVS_INVALID_NAME if namespace name doesn't satisfy constraints
- ESP_ERR_NO_MEM in case memory could not be allocated for the internal structures
- ESP_ERR_NVS_NOT_ENOUGH_SPACE if there is no space for a new entry or there are too many different namespaces (maximum allowed different namespaces: 254)
- ESP_ERR_NOT_ALLOWED if the NVS partition is read-only and mode is NVS_READWRITE
- ESP_ERR_INVALID_ARG if out_handle is equal to NULL
- other error codes from the underlying storage driver

esp_err_t **nvs_set_blob** (*nvs_handle_t* handle, const char *key, const void *value, size_t length)

set variable length binary value for given key

Sets variable length binary value for the key. Function uses 2 overhead and 1 entry per each 32 bytes of new data from the pool of available entries. See `nvs_get_stats`. In case of value update for existing key, space occupied by the existing value and 2 overhead entries are returned to the pool of available entries. Note that the underlying storage will not be updated until `nvs_commit` is called.

Parameters

- **handle** -- [in] Handle obtained from `nvs_open` function. Handles that were opened read only cannot be used.

- **key** -- **[in]** Key name. Maximum length is (NVS_KEY_NAME_MAX_SIZE-1) characters. Shouldn't be empty.
- **value** -- **[in]** The value to set.
- **length** -- **[in]** length of binary value to set, in bytes; Maximum length is 508000 bytes or (97.6% of the partition size - 4000) bytes whichever is lower.

Returns

- ESP_OK if value was set successfully
- ESP_FAIL if there is an internal error; most likely due to corrupted NVS partition (only if NVS assertion checks are disabled)
- ESP_ERR_NVS_INVALID_HANDLE if handle has been closed or is NULL
- ESP_ERR_NVS_READ_ONLY if storage handle was opened as read only
- ESP_ERR_NVS_INVALID_NAME if key name doesn't satisfy constraints
- ESP_ERR_NVS_NOT_ENOUGH_SPACE if there is not enough space in the underlying storage to save the value
- ESP_ERR_NVS_REMOVE_FAILED if the value wasn't updated because flash write operation has failed. The value was written however, and update will be finished after re-initialization of nvs, provided that flash operation doesn't fail again.
- ESP_ERR_NVS_VALUE_TOO_LONG if the value is too long

esp_err_t **nvs_find_key** (*nvs_handle_t* handle, const char *key, *nvs_type_t* *out_type)

Lookup key-value pair with given key name.

Note that function may indicate both existence of the key as well as the data type of NVS entry if it is found.

Parameters

- **handle** -- **[in]** Storage handle obtained with nvs_open.
- **key** -- **[in]** Key name. Maximum length is (NVS_KEY_NAME_MAX_SIZE-1) characters. Shouldn't be empty.
- **out_type** -- **[out]** Pointer to the output variable populated with data type of NVS entry in case key was found. May be NULL, respective data type is then not provided.

Returns

- ESP_OK if NVS entry for key provided was found
- ESP_ERR_NVS_NOT_FOUND if the requested key doesn't exist
- ESP_ERR_NVS_INVALID_HANDLE if handle has been closed or is NULL
- ESP_FAIL if there is an internal error; most likely due to corrupted NVS partition (only if NVS assertion checks are disabled)
- other error codes from the underlying storage driver

esp_err_t **nvs_erase_key** (*nvs_handle_t* handle, const char *key)

Erase key-value pair with given key name.

Note that actual storage may not be updated until nvs_commit function is called.

Parameters

- **handle** -- **[in]** Storage handle obtained with nvs_open. Handles that were opened read only cannot be used.
- **key** -- **[in]** Key name. Maximum length is (NVS_KEY_NAME_MAX_SIZE-1) characters. Shouldn't be empty.

Returns

- ESP_OK if erase operation was successful
- ESP_FAIL if there is an internal error; most likely due to corrupted NVS partition (only if NVS assertion checks are disabled)
- ESP_ERR_NVS_INVALID_HANDLE if handle has been closed or is NULL
- ESP_ERR_NVS_READ_ONLY if handle was opened as read only
- ESP_ERR_NVS_NOT_FOUND if the requested key doesn't exist
- other error codes from the underlying storage driver

esp_err_t **nvs_erase_all** (*nvs_handle_t* handle)

Erase all key-value pairs in a namespace.

Note that actual storage may not be updated until nvs_commit function is called.

Parameters **handle** -- [in] Storage handle obtained with `nvs_open`. Handles that were opened read only cannot be used.

Returns

- `ESP_OK` if erase operation was successful
- `ESP_FAIL` if there is an internal error; most likely due to corrupted NVS partition (only if NVS assertion checks are disabled)
- `ESP_ERR_NVS_INVALID_HANDLE` if handle has been closed or is `NULL`
- `ESP_ERR_NVS_READ_ONLY` if handle was opened as read only
- other error codes from the underlying storage driver

esp_err_t **nvs_commit** (*nvs_handle_t* handle)

Write any pending changes to non-volatile storage.

After setting any values, `nvs_commit()` must be called to ensure changes are written to non-volatile storage. Individual implementations may write to storage at other times, but this is not guaranteed.

Parameters **handle** -- [in] Storage handle obtained with `nvs_open`. Handles that were opened read only cannot be used.

Returns

- `ESP_OK` if the changes have been written successfully
- `ESP_ERR_NVS_INVALID_HANDLE` if handle has been closed or is `NULL`
- other error codes from the underlying storage driver

void **nvs_close** (*nvs_handle_t* handle)

Close the storage handle and free any allocated resources.

This function should be called for each handle opened with `nvs_open` once the handle is not in use any more. Closing the handle may not automatically write the changes to nonvolatile storage. This has to be done explicitly using `nvs_commit` function. Once this function is called on a handle, the handle should no longer be used.

Parameters **handle** -- [in] Storage handle to close

esp_err_t **nvs_get_stats** (const char *part_name, *nvs_stats_t* *nvs_stats)

Fill structure *nvs_stats_t*. It provides info about memory used by NVS.

This function calculates the number of used entries, free entries, available entries, total entries and number of namespaces in partition.

```
// Example of nvs_get_stats() to get overview of actual statistics of data_
→entries :
nvs_stats_t nvs_stats;
nvs_get_stats(NULL, &nvs_stats);
printf("Count: UsedEntries = (%lu), FreeEntries = (%lu), AvailableEntries = (
→%lu), AllEntries = (%lu)\n",
      nvs_stats.used_entries, nvs_stats.free_entries, nvs_stats.available_
→entries, nvs_stats.total_entries);
```

Parameters

- **part_name** -- [in] Partition name NVS in the partition table. If pass a `NULL` than will use `NVS_DEFAULT_PART_NAME` ("nvs").
- **nvs_stats** -- [out] Returns filled structure `nvs_states_t`. It provides info about used memory the partition.

Returns

- `ESP_OK` if the changes have been written successfully. Return param `nvs_stats` will be filled.
- `ESP_ERR_NVS_PART_NOT_FOUND` if the partition with label "name" is not found. Return param `nvs_stats` will be filled 0.
- `ESP_ERR_NVS_NOT_INITIALIZED` if the storage driver is not initialized. Return param `nvs_stats` will be filled 0.
- `ESP_ERR_INVALID_ARG` if `nvs_stats` is equal to `NULL`.

- **ESP_ERR_INVALID_STATE** if there is page with the status of **INVALID**. Return param **nvs_stats** will be filled not with correct values because not all pages will be counted. Counting will be interrupted at the first **INVALID** page.

esp_err_t nvs_get_used_entry_count (*nvs_handle_t* handle, size_t *used_entries)

Calculate all entries in a namespace.

An entry represents the smallest storage unit in NVS. Strings and blobs may occupy more than one entry. Note that to find out the total number of entries occupied by the namespace, add one to the returned value **used_entries** (if **err** is equal to **ESP_OK**). Because the name space entry takes one entry.

```
// Example of nvs_get_used_entry_count() to get amount of all key-value pairs_
↳in one namespace:
nvs_handle_t handle;
nvs_open("namespace1", NVS_READWRITE, &handle);
...
size_t used_entries;
size_t total_entries_namespace;
if (nvs_get_used_entry_count(handle, &used_entries) == ESP_OK) {
    // the total number of entries occupied by the namespace
    total_entries_namespace = used_entries + 1;
}
```

Parameters

- **handle** -- **[in]** Handle obtained from **nvs_open** function.
- **used_entries** -- **[out]** Returns amount of used entries from a namespace.

Returns

- **ESP_OK** if the changes have been written successfully. Return param **used_entries** will be filled valid value.
- **ESP_ERR_NVS_NOT_INITIALIZED** if the storage driver is not initialized. Return param **used_entries** will be filled 0.
- **ESP_ERR_NVS_INVALID_HANDLE** if **handle** has been closed or is **NULL**. Return param **used_entries** will be filled 0.
- **ESP_ERR_INVALID_ARG** if **used_entries** is equal to **NULL**.
- Other error codes from the underlying storage driver. Return param **used_entries** will be filled 0.

esp_err_t nvs_entry_find (const char *part_name, const char *namespace_name, *nvs_type_t* type, *nvs_iterator_t* *output_iterator)

Create an iterator to enumerate NVS entries based on one or more parameters.

```
// Example of listing all the key-value pairs of any type under specified_
↳partition and namespace
nvs_iterator_t it = NULL;
esp_err_t res = nvs_entry_find(<nvs_partition_name>, <namespace>, NVS_TYPE_
↳ANY, &it);
while (res == ESP_OK) {
    nvs_entry_info_t info;
    nvs_entry_info(it, &info); // Can omit error check if parameters are_
↳guaranteed to be non-NULL
    printf("key '%s', type '%d' \n", info.key, info.type);
    res = nvs_entry_next(&it);
}
nvs_release_iterator(it);
```

Parameters

- **part_name** -- **[in]** Partition name
- **namespace_name** -- **[in]** Set this value if looking for entries with a specific namespace. Pass **NULL** otherwise.

- **type** -- [in] One of `nvs_type_t` values.
- **output_iterator** -- [out] Set to a valid iterator to enumerate all the entries found. Set to NULL if no entry for specified criteria was found. If any other error except `ESP_ERR_INVALID_ARG` occurs, `output_iterator` is NULL, too. If `ESP_ERR_INVALID_ARG` occurs, `output_iterator` is not changed. If a valid iterator is obtained through this function, it has to be released using `nvs_release_iterator` when not used any more, unless `ESP_ERR_INVALID_ARG` is returned.

Returns

- `ESP_OK` if no internal error or programming error occurred.
- `ESP_ERR_NVS_NOT_FOUND` if no element of specified criteria has been found.
- `ESP_ERR_NO_MEM` if memory has been exhausted during allocation of internal structures.
- `ESP_ERR_INVALID_ARG` if any of the parameters is NULL. Note: don't release `output_iterator` in case `ESP_ERR_INVALID_ARG` has been returned

`esp_err_t nvs_entry_find_in_handle` (*nvs_handle_t* handle, *nvs_type_t* type, *nvs_iterator_t* *output_iterator)

Create an iterator to enumerate NVS entries based on a handle and type.

```
// Example of listing all the key-value pairs of any type under specified
// handle (which defines a partition and namespace)
nvs_iterator_t it = NULL;
esp_err_t res = nvs_entry_find_in_handle(<nvs_handle>, NVS_TYPE_ANY, &it);
while(res == ESP_OK) {
    nvs_entry_info_t info;
    nvs_entry_info(it, &info); // Can omit error check if parameters are
    // guaranteed to be non-NULL
    printf("key '%s', type '%d' \n", info.key, info.type);
    res = nvs_entry_next(&it);
}
nvs_release_iterator(it);
```

Parameters

- **handle** -- [in] Handle obtained from `nvs_open` function.
- **type** -- [in] One of `nvs_type_t` values.
- **output_iterator** -- [out] Set to a valid iterator to enumerate all the entries found. Set to NULL if no entry for specified criteria was found. If any other error except `ESP_ERR_INVALID_ARG` occurs, `output_iterator` is NULL, too. If `ESP_ERR_INVALID_ARG` occurs, `output_iterator` is not changed. If a valid iterator is obtained through this function, it has to be released using `nvs_release_iterator` when not used any more, unless `ESP_ERR_INVALID_ARG` is returned.

Returns

- `ESP_OK` if no internal error or programming error occurred.
- `ESP_ERR_NVS_NOT_FOUND` if no element of specified criteria has been found.
- `ESP_ERR_NO_MEM` if memory has been exhausted during allocation of internal structures.
- `ESP_ERR_NVS_INVALID_HANDLE` if unknown handle was specified.
- `ESP_ERR_INVALID_ARG` if `output_iterator` parameter is NULL. Note: don't release `output_iterator` in case `ESP_ERR_INVALID_ARG` has been returned

`esp_err_t nvs_entry_next` (*nvs_iterator_t* *iterator)

Advances the iterator to next item matching the iterator criteria.

Note that any copies of the iterator will be invalid after this call.

Parameters **iterator** -- [inout] Iterator obtained from `nvs_entry_find` or `nvs_entry_find_in_handle` function. Must be non-NULL. If any error ex-

cept `ESP_ERR_INVALID_ARG` occurs, `iterator` is set to `NULL`. If `ESP_ERR_INVALID_ARG` occurs, `iterator` is not changed.

Returns

- `ESP_OK` if no internal error or programming error occurred.
- `ESP_ERR_NVS_NOT_FOUND` if no next element matching the iterator criteria.
- `ESP_ERR_INVALID_ARG` if `iterator` is `NULL`.
- Possibly other errors in the future for internal programming or flash errors.

esp_err_t **nvs_entry_info** (const *nvs_iterator_t* iterator, *nvs_entry_info_t* *out_info)

Fills *nvs_entry_info_t* structure with information about entry pointed to by the iterator.

Parameters

- **iterator** -- [in] Iterator obtained from `nvs_entry_find` or `nvs_entry_find_in_handle` function. Must be non-`NULL`.
- **out_info** -- [out] Structure to which entry information is copied.

Returns

- `ESP_OK` if all parameters are valid; current iterator data has been written to `out_info`
- `ESP_ERR_INVALID_ARG` if one of the parameters is `NULL`.

void **nvs_release_iterator** (*nvs_iterator_t* iterator)

Release iterator.

Parameters **iterator** -- [in] Release iterator obtained from `nvs_entry_find` or `nvs_entry_find_in_handle` or `nvs_entry_next` function. `NULL` argument is allowed.

Structures

struct **nvs_entry_info_t**

information about entry obtained from `nvs_entry_info` function

Public Members

char **namespace_name**[`NVS_NS_NAME_MAX_SIZE`]

Namespace to which key-value belong

char **key**[`NVS_KEY_NAME_MAX_SIZE`]

Key of stored key-value pair

nvs_type_t **type**

Type of stored key-value pair

struct **nvs_stats_t**

Note: Info about storage space NVS.

Public Members

size_t **used_entries**

Number of used entries.

size_t **free_entries**

Number of free entries. It includes also reserved entries.

size_t **available_entries**

Number of entries available for data storage.

size_t **total_entries**

Number of all entries.

size_t **namespace_count**

Number of namespaces.

Macros

ESP_ERR_NVS_BASE

Starting number of error codes

ESP_ERR_NVS_NOT_INITIALIZED

The storage driver is not initialized

ESP_ERR_NVS_NOT_FOUND

A requested entry couldn't be found or namespace doesn't exist yet and mode is NVS_READONLY

ESP_ERR_NVS_TYPE_MISMATCH

The type of set or get operation doesn't match the type of value stored in NVS

ESP_ERR_NVS_READ_ONLY

Storage handle was opened as read only

ESP_ERR_NVS_NOT_ENOUGH_SPACE

There is not enough space in the underlying storage to save the value

ESP_ERR_NVS_INVALID_NAME

Namespace name doesn't satisfy constraints

ESP_ERR_NVS_INVALID_HANDLE

Handle has been closed or is NULL

ESP_ERR_NVS_REMOVE_FAILED

The value wasn't updated because flash write operation has failed. The value was written however, and update will be finished after re-initialization of nvs, provided that flash operation doesn't fail again.

ESP_ERR_NVS_KEY_TOO_LONG

Key name is too long

ESP_ERR_NVS_PAGE_FULL

Internal error; never returned by nvs API functions

ESP_ERR_NVS_INVALID_STATE

NVS is in an inconsistent state due to a previous error. Call `nvs_flash_init` and `nvs_open` again, then retry.

ESP_ERR_NVS_INVALID_LENGTH

String or blob length is not sufficient to store data

ESP_ERR_NVS_NO_FREE_PAGES

NVS partition doesn't contain any empty pages. This may happen if NVS partition was truncated. Erase the whole partition and call `nvs_flash_init` again.

ESP_ERR_NVS_VALUE_TOO_LONG

Value doesn't fit into the entry or string or blob length is longer than supported by the implementation

ESP_ERR_NVS_PART_NOT_FOUND

Partition with specified name is not found in the partition table

ESP_ERR_NVS_NEW_VERSION_FOUND

NVS partition contains data in new format and cannot be recognized by this version of code

ESP_ERR_NVS_XTS_ENCR_FAILED

XTS encryption failed while writing NVS entry

ESP_ERR_NVS_XTS_DECR_FAILED

XTS decryption failed while reading NVS entry

ESP_ERR_NVS_XTS_CFG_FAILED

XTS configuration setting failed

ESP_ERR_NVS_XTS_CFG_NOT_FOUND

XTS configuration not found

ESP_ERR_NVS_ENCR_NOT_SUPPORTED

NVS encryption is not supported in this version

ESP_ERR_NVS_KEYS_NOT_INITIALIZED

NVS key partition is uninitialized

ESP_ERR_NVS_CORRUPT_KEY_PART

NVS key partition is corrupt

ESP_ERR_NVS_WRONG_ENCRYPTION

NVS partition is marked as encrypted with generic flash encryption. This is forbidden since the NVS encryption works differently.

ESP_ERR_NVS_CONTENT_DIFFERS

Internal error; never returned by `nvs` API functions. NVS key is different in comparison

NVS_DEFAULT_PART_NAME

Default partition name of the NVS partition in the partition table

NVS_PART_NAME_MAX_SIZE

maximum length of partition name (excluding null terminator)

NVS_KEY_NAME_MAX_SIZE

Maximum length of NVS key name (including null terminator)

NVS_NS_NAME_MAX_SIZE

Maximum length of NVS namespace name (including null terminator)

Type Definitions

typedef uint32_t **nvs_handle_t**

Opaque pointer type representing non-volatile storage handle

typedef *nvs_handle_t* **nvs_handle**

typedef *nvs_open_mode_t* **nvs_open_mode**

typedef struct nvs_opaque_iterator_t ***nvs_iterator_t**

Opaque pointer type representing iterator to nvs entries

Enumerations

enum **nvs_open_mode_t**

Mode of opening the non-volatile storage.

Values:

enumerator **NVS_READONLY**

Read only

enumerator **NVS_READWRITE**

Read and write

enum **nvs_type_t**

Types of variables.

Values:

enumerator **NVS_TYPE_U8**

Type uint8_t

enumerator **NVS_TYPE_I8**

Type int8_t

enumerator **NVS_TYPE_U16**

Type uint16_t

enumerator **NVS_TYPE_I16**

Type int16_t

enumerator **NVS_TYPE_U32**

Type uint32_t

enumerator **NVS_TYPE_I32**

Type int32_t

enumerator **NVS_TYPE_U64**

Type uint64_t

enumerator **NVS_TYPE_I64**

Type int64_t

enumerator **NVS_TYPE_STR**

Type string

enumerator **NVS_TYPE_BLOB**

Type blob

enumerator **NVS_TYPE_ANY**

Must be last

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.8.5 NVS Encryption

Overview

This guide provides an overview of the NVS encryption feature. NVS encryption helps to achieve secure storage on the device flash memory.

Data stored in NVS partitions can be encrypted using XTS-AES in the manner similar to the one mentioned in disk encryption standard IEEE P1619. For the purpose of encryption, each entry is treated as one `sector` and relative address of the entry (w.r.t., partition-start) is fed to the encryption algorithm as `sector-number`.

NVS encryption can be facilitated by enabling [CONFIG_NVS_ENCRYPTION](#) and [CONFIG_NVS_SEC_KEY_PROTECTION_SCHEME](#) > `CONFIG_NVS_SEC_KEY_PROTECT_USING_FLASH_ENC` or `CONFIG_NVS_SEC_KEY_PROTECT_USING_HMAC` depending on the scheme to be used.

NVS Encryption: Flash Encryption-Based Scheme

In this scheme, the keys required for NVS encryption are stored in yet another partition, which is protected using [Flash Encryption](#). Therefore, enabling [Flash Encryption](#) becomes a prerequisite for NVS encryption here.

NVS encryption should be enabled when [Flash Encryption](#) is enabled because the Wi-Fi driver stores credentials (like SSID and passphrase) in the default NVS partition. It is important to encrypt them if platform-level encryption is already enabled.

For using NVS encryption using this scheme, the partition table must contain the [NVS Key Partition](#). Two partition tables containing the [NVS Key Partition](#) are provided for NVS encryption under the partition table option (`menuconfig`>Partition Table). They can be selected with the project configuration menu (`idf.py menuconfig`). Please refer to the example [security/flash_encryption](#) for how to configure and use the NVS encryption feature.

NVS Key Partition An application requiring NVS encryption support (using the Flash Encryption-based scheme) needs to be compiled with a key-partition of the type `data` and subtype `nvs_keys`. This partition should be marked as `encrypted` and its size should be the minimum partition size (4 KB). Refer to [Partition Tables](#) for more details. Two additional partition tables which contain the *NVS Key Partition* are provided under the partition table option (`menuconfig > Partition Table`). They can be directly used for NVS encryption. The structure of these partitions is depicted below:

+	-----+	-----+	-----+	-----+	+
	XTS encryption key (32)				
+	-----+				+
	XTS tweak key (32)				
+	-----+				+
	CRC32 (4)				
+	-----+				+

The XTS encryption keys in the *NVS Key Partition* can be generated in one of the following two ways.

Generate the keys on ESP32-P4 chip itself

- When NVS encryption is enabled, the `nvs_flash_init()` API function can be used to initialize the encrypted default NVS partition. The API function internally generates the XTS encryption keys on the ESP chip. The API function finds the first *NVS Key Partition*.
- Then the API function automatically generates and stores the NVS keys in that partition by making use of the `nvs_flash_generate_keys()` API function provided by `nvs_flash/include/nvs_flash.h`. New keys are generated and stored only when the respective key partition is empty. The same key partition can then be used to read the security configurations for initializing a custom encrypted NVS partition with help of `nvs_flash_secure_init_partition()`.
- The API functions `nvs_flash_secure_init()` and `nvs_flash_secure_init_partition()` do not generate the keys internally. When these API functions are used for initializing encrypted NVS partitions, the keys can be generated after startup using the `nvs_flash_generate_keys()` API function provided by `nvs_flash.h`. The API function then writes those keys onto the key-partition in encrypted form.

Note: Please note that `nvs_keys` partition must be completely erased before you start the application in this approach. Otherwise the application may generate the `ESP_ERR_NVS_CORRUPT_KEY_PART` error code assuming that `nvs_keys` partition is not empty and contains malformed data. You can use the following command for this:

```
parttool.py --port PORT --partition-table-file=PARTITION_TABLE_FILE --
partition-table-offset PARTITION_TABLE_OFFSET erase_partition --
partition-type=data --partition-subtype=nvs keys
```

Use a pre-generated NVS key partition

This option will be required by the user when keys in the *NVS Key Partition* are not generated by the application. The *NVS Key Partition* containing the XTS encryption keys can be generated with the help of *NVS Partition Generator Utility*. Then the user can store the pre-generated key partition on the flash with help of the following two commands:

- ## 1. Build and flash the partition table

```
idf.py partition-table partition-table-flash
```

2. Store the keys in the *NVS Key Partition* (on the flash) with the help of `parttool.py` (see Partition Tool section in *partition-tables* for more details)

```
parttool.py --port PORT --partition-table-offset PARTITION_TABLE_OFFSET --  
write_partition --partition-name="name of nvs_key partition" --input   
NVS_KEY_PARTITION_FILE (continues on next page)
```

(continues on next page)

(continued from previous page)

Note: If the device is encrypted in flash encryption development mode and you want to renew the NVS key partition, you need to tell `parttool.py` to encrypt the NVS key partition and you also need to give it a pointer to the unencrypted partition table in your build directory (`build/partition_table`) since the partition table on the device is encrypted, too. You can use the following command:

```
parttool.py --esptool-write-args encrypt --port PORT --partition-table-
→file=PARTITION_TABLE_FILE --partition-table-offset PARTITION_TABLE_
→OFFSET write_partition --partition-name="name of nvs_key partition" --
→input NVS_KEY_PARTITION_FILE
```

Since the key partition is marked as `encrypted` and *Flash Encryption* is enabled, the bootloader will encrypt this partition using flash encryption key on the first boot.

It is possible for an application to use different keys for different NVS partitions and thereby have multiple key-partitions. However, it is a responsibility of the application to provide the correct key-partition and keys for encryption or decryption.

NVS Encryption: HMAC Peripheral-Based Scheme

In this scheme, the XTS keys required for NVS encryption are derived from an HMAC key programmed in eFuse with the purpose `esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_HMAC_UP`. Since the encryption keys are derived at runtime, they are not stored anywhere in the flash. Thus, this feature does not require a separate *NVS Key Partition*.

Note: This scheme enables us to achieve secure storage on ESP32-P4 **without enabling flash encryption**.

Important: Please take note that this scheme uses one eFuse block for storing the HMAC key required for deriving the encryption keys.

- When NVS encryption is enabled, the `nvs_flash_init()` API function can be used to initialize the encrypted default NVS partition. The API function first checks whether an HMAC key is present at `CONFIG_NVS_SEC_HMAC_EFUSE_KEY_ID`.

Note: The valid range for the config `CONFIG_NVS_SEC_HMAC_EFUSE_KEY_ID` is from 0 (`hmac_key_id_t::HMAC_KEY0`) to 5 (`hmac_key_id_t::HMAC_KEY5`). By default, the config is set to 6 (`hmac_key_id_t::HMAC_KEY_MAX`), which have to be configured before building the user application.

- If no key is found, a key is generated internally and stored at the eFuse block specified at `CONFIG_NVS_SEC_HMAC_EFUSE_KEY_ID`.
- If a key is found with the purpose `esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_HMAC_UP`, the same is used for the derivation of the XTS encryption keys.
- If the specified eFuse block is found to be occupied with a key with a purpose other than `esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_HMAC_UP`, an error is thrown.
- The API `nvs_flash_init()` then automatically generates the NVS keys on demand by using the `nvs_flash_generate_keys_v2()` API function provided by the `nvs_flash/include/nvs_flash.h`. The same keys can also be used to read the security configurations (see `nvs_flash_read_security_cfg_v2()`) for initializing a custom encrypted NVS partition with help of `nvs_flash_secure_init_partition()`.

- The API functions `nvs_flash_secure_init()` and `nvs_flash_secure_init_partition()` do not generate the keys internally. When these API functions are used for initializing encrypted NVS partitions, the keys can be generated after startup using the `nvs_flash_generate_keys_v2()` API function or take and populate the NVS security configuration structure `nvs_sec_cfg_t` with `nvs_flash_read_security_cfg_v2()` and feed them into the above APIs.

Note: Users can program their own HMAC key in eFuse block beforehand by using the following command:

```
espefuse.py -p PORT burn_key <BLOCK_KEYN> <hmac_key_file.bin> HMAC_UP
```

Encrypted Read/Write

The same NVS API functions `nvs_get_*` or `nvs_set_*` can be used for reading of, and writing to an encrypted NVS partition as well.

Encrypt the default NVS partition

- To enable encryption for the default NVS partition, no additional step is necessary. When `CONFIG_NVS_ENCRYPTION` is enabled, the `nvs_flash_init()` API function internally performs some additional steps to enable encryption for the default NVS partition depending on the scheme being used (set by `CONFIG_NVS_SEC_KEY_PROTECTION_SCHEME`).
- For the flash encryption-based scheme, the first *NVS Key Partition* found is used to generate the encryption keys while for the HMAC one, keys are generated using the HMAC key burnt in eFuse at `CONFIG_NVS_SEC_HMAC_EFUSE_KEY_ID` (refer to the API documentation for more details).

Alternatively, `nvs_flash_secure_init()` API function can also be used to enable encryption for the default NVS partition.

Encrypt a custom NVS partition

- To enable encryption for a custom NVS partition, `nvs_flash_secure_init_partition()` API function is used instead of `nvs_flash_init_partition()`.
- When `nvs_flash_secure_init()` and `nvs_flash_secure_init_partition()` API functions are used, the applications are expected to follow the steps below in order to perform NVS read/write operations with encryption enabled:
 1. Populate the NVS security configuration structure `nvs_sec_cfg_t`
 - For the Flash Encryption-based scheme
 - * Find key partition and NVS data partition using `esp_partition_find*` API functions.
 - * Populate the `nvs_sec_cfg_t` struct using the `nvs_flash_read_security_cfg()` or `nvs_flash_generate_keys()` API functions.
 - For the HMAC-based scheme
 - * Set the scheme-specific config data with `nvs_sec_config_hmac_t` and register the HMAC-based scheme with the API `nvs_sec_provider_register_hmac()` which will also populate the scheme-specific handle (see `nvs_sec_scheme_t`).
 - * Populate the `nvs_sec_cfg_t` struct using the `nvs_flash_read_security_cfg_v2()` or `nvs_flash_generate_keys_v2()` API functions.

```
nvs_sec_cfg_t cfg = {};
nvs_sec_scheme_t *sec_scheme_handle = NULL;

nvs_sec_config_hmac_t sec_scheme_cfg = {};
hmac_key_id_t hmac_key = HMAC_KEY0;
sec_scheme_cfg.hmac_key_id = hmac_key;

ret = nvs_sec_provider_register_hmac(&sec_scheme_cfg, &sec_scheme_
↪handle);
```

(continues on next page)

(continued from previous page)

```

if (ret != ESP_OK) {
    return ret;
}

ret = nvs_flash_read_security_cfg_v2(sec_scheme_handle, &cfg);
if (ret != ESP_OK) {
    if (ret == ESP_ERR_NVS_SEC_HMAC_KEY_NOT_FOUND) {
        ret = nvs_flash_generate_keys_v2(&sec_scheme_handle, &cfg);
        if (ret != ESP_OK) {
            ESP_LOGE(TAG, "Failed to generate NVS encr-keys!");
            return ret;
        }
    }
    ESP_LOGE(TAG, "Failed to read NVS security cfg!");
    return ret;
}

```

2. Initialise NVS flash partition using the `nvs_flash_secure_init()` or `nvs_flash_secure_init_partition()` API functions.
3. Open a namespace using the `nvs_open()` or `nvs_open_from_partition()` API functions.
4. Perform NVS read/write operations using `nvs_get_*` or `nvs_set_*`.
5. Deinitialise an NVS partition using `nvs_flash_deinit()`.

Note: While using the HMAC-based scheme, the above workflow can be used without enabling any of the config options for NVS encryption - `CONFIG_NVS_ENCRYPTION`, `CONFIG_NVS_SEC_KEY_PROTECTION_SCHEME` -> `CONFIG_NVS_SEC_KEY_PROTECT_USING_HMAC` and `CONFIG_NVS_SEC_HMAC_EFUSE_KEY_ID` to encrypt the default as well as custom NVS partitions with `nvs_flash_secure_init()` API.

NVS Security Provider

The component `nvs_sec_provider` stores all the implementation-specific code for the NVS encryption schemes and would also accommodate any future schemes. This component acts as an interface to the `nvs_flash` component for the handling of encryption keys. `nvs_sec_provider` has a configuration menu of its own, based on which the selected security scheme and the corresponding settings are registered for the `nvs_flash` component.

This component offers factory functions with which a particular security scheme can be registered without having to worry about the APIs to generate and read the encryption keys (e.g., `nvs_sec_provider_register_hmac()`). Refer to the `security/nvs_encryption_hmac` example for API usage.

API Reference

Header File

- `components/nvs_sec_provider/include/nvs_sec_provider.h`
- This header file can be included with:

```
#include "nvs_sec_provider.h"
```

- This header file is a part of the API provided by the `nvs_sec_provider` component. To declare that your component depends on `nvs_sec_provider`, add the following to your `CMakeLists.txt`:

```
REQUIRES nvs_sec_provider
```

or

```
PRIV_REQUIRES nvs_sec_provider
```

Functions

esp_err_t **nvs_sec_provider_register_flash_enc** (const *nvs_sec_config_flash_enc_t* *sec_scheme_cfg, *nvs_sec_scheme_t* **sec_scheme_handle_out)

Register the Flash-Encryption based scheme for NVS Encryption.

Parameters

- **sec_scheme_cfg** -- [in] Security scheme specific configuration data
- **sec_scheme_handle_out** -- [out] Security scheme specific configuration handle

Returns

- ESP_OK, if sec_scheme_handle_out was populated successfully with the scheme configuration;
- ESP_ERR_INVALID_ARG, if scheme_cfg_hmac is NULL;
- ESP_ERR_NO_MEM, No memory for the scheme-specific handle sec_scheme_handle_out
- ESP_ERR_NOT_FOUND, if no nvs_keys partition is found

esp_err_t **nvs_sec_provider_register_hmac** (const *nvs_sec_config_hmac_t* *sec_scheme_cfg, *nvs_sec_scheme_t* **sec_scheme_handle_out)

Register the HMAC-based scheme for NVS Encryption.

Parameters

- **sec_scheme_cfg** -- [in] Security scheme specific configuration data
- **sec_scheme_handle_out** -- [out] Security scheme specific configuration handle

Returns

- ESP_OK, if sec_scheme_handle_out was populated successfully with the scheme configuration;
- ESP_ERR_INVALID_ARG, if scheme_cfg_hmac is NULL;
- ESP_ERR_NO_MEM, No memory for the scheme-specific handle sec_scheme_handle_out

esp_err_t **nvs_sec_provider_deregister** (*nvs_sec_scheme_t* *sec_scheme_handle)

Deregister the NVS encryption scheme registered with the given handle.

Parameters **sec_scheme_handle** -- [in] Security scheme specific configuration handle

Returns

- ESP_OK, if the scheme registered with sec_scheme_handle was deregistered successfully
- ESP_ERR_INVALID_ARG, if sec_scheme_handle is NULL;

Structures

struct **nvs_sec_config_flash_enc_t**

Flash encryption-based scheme specific configuration data.

Public Members

const *esp_partition_t* ***nvs_keys_part**

Partition of subtype *nvs_keys* holding the NVS encryption keys

struct **nvs_sec_config_hmac_t**

HMAC-based scheme specific configuration data.

Public Members

hmac_key_id_t **hmac_key_id**

HMAC Key ID used for generating the NVS encryption keys

Macros

ESP_ERR_NVS_SEC_BASE

Starting number of error codes

ESP_ERR_NVS_SEC_HMAC_KEY_NOT_FOUND

HMAC Key required to generate the NVS encryption keys not found

ESP_ERR_NVS_SEC_HMAC_KEY_BLK_ALREADY_USED

Provided eFuse block for HMAC key generation is already in use

ESP_ERR_NVS_SEC_HMAC_KEY_GENERATION_FAILED

Failed to generate/write the HMAC key to eFuse

ESP_ERR_NVS_SEC_HMAC_XTS_KEYS_DERIV_FAILED

Failed to derive the NVS encryption keys based on the HMAC-based scheme

NVS_SEC_PROVIDER_CFG_FLASH_ENC_DEFAULT ()

Helper for populating the Flash encryption-based scheme specific configuration data.

NVS_SEC_PROVIDER_CFG_HMAC_DEFAULT ()

Helper for populating the HMAC-based scheme specific configuration data.

Enumerations

enum **nvs_sec_scheme_id_t**

NVS Encryption Keys Protection Scheme.

Values:

enumerator **NVS_SEC_SCHEME_FLASH_ENC**

Protect NVS encryption keys using Flash Encryption

enumerator **NVS_SEC_SCHEME_HMAC**

Protect NVS encryption keys using HMAC peripheral

enumerator **NVS_SEC_SCHEME_MAX**

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.8.6 NVS Partition Generator Utility

Introduction

The utility `nvs_flash/nvs_partition_generator/nvs_partition_gen.py` creates a binary file, compatible with the NVS architecture defined in *Non-Volatile Storage Library*, based on the key-value pairs provided in a CSV file.

This utility is ideally suited for generating a binary blob, containing data specific to ODM/OEM, which can be flashed externally at the time of device manufacturing. This allows manufacturers to generate many instances of the same application firmware with customized parameters for each device, such as a serial number.

Prerequisites

To use this utility in encryption mode, install the following packages:

- `cryptography`

All the required packages are included in *requirements.txt* in the root of the ESP-IDF directory.

CSV File Format Each line of a CSV file should contain 4 parameters, separated by a comma. The table below describes each of these parameters.

No.	Parameter	Description	Notes
1	Key	Key of the data. The data can be accessed later from an application using this key.	
2	Type	Supported values are <code>file</code> , <code>data</code> , and <code>namespace</code> .	
3	Encoding	Supported values are: <code>u8</code> , <code>i8</code> , <code>u16</code> , <code>i16</code> , <code>u32</code> , <code>i32</code> , <code>u64</code> , <code>i64</code> , <code>string</code> , <code>hex2bin</code> , <code>base64</code> , and <code>binary</code> . This specifies how actual data values are encoded in the resulting binary file. The difference between the <code>string</code> and <code>binary</code> encoding is that <code>string</code> data is terminated with a NULL character, whereas <code>binary</code> data is not.	As of now, for the <code>file</code> type, only <code>hex2bin</code> , <code>base64</code> , <code>string</code> , and <code>binary</code> encoding is supported.
4	Value	Data value	Encoding and Value cells for the <code>namespace</code> field type should be empty. Encoding and Value of <code>namespace</code> are fixed and are not configurable. Any values in these cells are ignored.

Note: The first line of the CSV file should always be the column header and it is not configurable.

Below is an example dump of such a CSV file:

```
key,type,encoding,value      <-- column header
namespace_name,namespace,,   <-- First entry should be of type "namespace"
key1,data,u8,1
key2,file,string,/path/to/file
```

Note:

Make sure there are no spaces:

- before and after `'`

- at the end of each line in a CSV file

NVS Entry and Namespace Association

When a namespace entry is encountered in a CSV file, each following entry will be treated as part of that namespace until the next namespace entry is found. At this point, all the following entries will be treated as part of the new namespace.

Note: First entry in a CSV file should always be a namespace entry.

Multipage Blob Support

By default, binary blobs are allowed to span over multiple pages and are written in the format mentioned in Section [Structure of Entry](#). If you intend to use the older format, the utility provides an option to disable this feature.

Encryption-Decryption Support

The NVS Partition Generator utility also allows you to create an encrypted binary file and decrypt an encrypted one. The utility uses the XTS-AES encryption. Please refer to [NVS Encryption](#) for more details.

Running the Utility

Usage:

```
python nvs_partition_gen.py [-h] {generate,generate-key,encrypt,decrypt} ...
```

Optional Arguments:

No.	Parameter	Description
1	-h / --help	Show the help message and exit

Commands:

Run `nvs_partition_gen.py {command} -h` for additional help

No.	Parameter	Description
1	generate	Generate NVS partition
2	generate-key	Generate keys for encryption
3	encrypt	Generate NVS encrypted partition
4	decrypt	Decrypt NVS encrypted partition

Generate NVS Partition (Default) Usage:

```
python nvs_partition_gen.py generate [-h] [--version {1,2}] [--outdir OUTDIR]   
↪ input output size
```

Positional Arguments:

Parameter	Description
input	Path to CSV file to parse
output	Path to output NVS binary file
size	Size of NVS partition in bytes (must be multiple of 4096)

Optional Arguments:

Parameter	Description
-h / --help	Show the help message and exit
--version {1,2}	Set multipage blob version (Default: Version 2) Version 1 - Multipage blob support disabled Version 2 - Multipage blob support enabled
--outdir OUTDIR	Output directory to store file created (Default: current directory)

You can run the utility to generate NVS partition using the command below. A sample CSV file is provided with the utility:

```
python nvs_partition_gen.py generate sample_singlepage_blob.csv sample.bin 0x3000
```

Generate Encryption Keys Partition Usage:

```
python nvs_partition_gen.py generate-key [-h] [--key_protect_hmac] [--kp_hmac_
↪keygen]
                                     [--kp_hmac_keyfile KP_HMAC_KEYFILE] [--
↪kp_hmac_inputkey KP_HMAC_INPUTKEY]
                                     [--keyfile KEYFILE] [--outdir OUTDIR]
```

Optional Arguments:

Parameter	Description
-h / --help	Show the help message and exit
--keyfile KEYFILE	Path to output encryption keys file
--outdir OUTDIR	Output directory to store files created. (Default: current directory)

Optional Arguments (HMAC scheme-specific):

Parameter	Description
--key_protect_hmac	If set, the NVS encryption key protection scheme based on HMAC peripheral is used; else the default scheme based on flash encryption is used
--kp_hmac_keygen	Generate the HMAC key for HMAC-based encryption scheme
--kp_hmac_keyfile KP_HMAC_KEYFILE	Path to output the HMAC key file
--kp_hmac_inputkey KP_HMAC_INPUTKEY	File having the HMAC key for generating the NVS encryption keys

You can run the utility to generate only the encryption key partition using the command below:

```
python nvs_partition_gen.py generate-key
```

For generating encryption key for the HMAC-based scheme, the following commands can be used:

- Generate the HMAC key and the NVS encryption keys:

```
python nvs_partition_gen.py generate-key --key_protect_hmac --kp_hmac_keygen
```

Note: Encryption key of the format <outdir>/keys/keys-<timestamp>.bin and HMAC key of the format <outdir>/keys/hmac-keys-<timestamp>.bin are created.

- Generate the NVS encryption keys, given the HMAC key:

```
python nvs_partition_gen.py generate-key --key_protect_hmac --kp_hmac_inputkey_
↳testdata/sample_hmac_key.bin
```

Note: You can provide the custom filename for the HMAC key as well as the encryption key as a parameter.

Generate Encrypted NVS Partition Usage:

```
python nvs_partition_gen.py encrypt [-h] [--version {1,2}] [--keygen]
                                  [--keyfile KEYFILE] [--inputkey INPUTKEY] [--
↳outdir OUTDIR]
                                  [--key_protect_hmac] [--kp_hmac_keygen]
                                  [--kp_hmac_keyfile KP_HMAC_KEYFILE] [--kp_hmac_
↳inputkey KP_HMAC_INPUTKEY]
                                  input output size
```

Positional Arguments:

Parameter	Description
input	Path to CSV file to parse
output	Path to output NVS binary file
size	Size of NVS partition in bytes (must be multiple of 4096)

Optional Arguments:

Parameter	Description
-h / --help	Show the help message and exit
--version {1,2}	Set multipage blob version (Default: Version 2) Version 1 - Multipage blob support disabled Version 2 - Multipage blob support enabled
--keygen	Generates key for encrypting NVS partition
--keyfile KEYFILE	Path to output encryption keys file
--inputkey INPUTKEY	File having key for encrypting NVS partition
--outdir OUTDIR	Output directory to store file created (Default: current directory)

Optional Arguments (HMAC scheme-specific):

Parameter	Description
--key_protect_hmac	If set, the NVS encryption key protection scheme based on HMAC peripheral is used; else the default scheme based on flash encryption is used
--kp_hmac_keygen	Generate the HMAC key for HMAC-based encryption scheme
--kp_hmac_keyfile KP_HMAC_KEYFILE	Path to output HMAC key file
--kp_hmac_inputkey KP_HMAC_INPUTKEY	File having the HMAC key for generating the NVS encryption keys

You can run the utility to encrypt NVS partition using the command below. A sample CSV file is provided with the utility:

- Encrypt by allowing the utility to generate encryption keys:

```
python nvs_partition_gen.py encrypt sample_singlepage_blob.csv sample_encr.bin
↪ 0x3000 --keygen
```

Note: Encryption key of the format <outdir>/keys/keys-<timestamp>.bin is created.

- To generate an encrypted partition using the HMAC-based scheme, the above command can be used along with some additional parameters.
 - Encrypt by allowing the utility to generate encryption keys and the HMAC-key:

```
python nvs_partition_gen.py encrypt sample_singlepage_blob.csv sample_encr.
↪ bin 0x3000 --keygen --key_protect_hmac --kp_hmac_keygen
```

Note: Encryption key of the format <outdir>/keys/keys-<timestamp>.bin and HMAC key of the format <outdir>/keys/hmac-keys-<timestamp>.bin are created.

- Encrypt by allowing the utility to generate encryption keys with user-provided HMAC-key:

```
python nvs_partition_gen.py encrypt sample_singlepage_blob.csv sample_encr.
↪ bin 0x3000 --keygen --key_protect_hmac --kp_hmac_inputkey testdata/
↪ sample_hmac_key.bin
```

Note: You can provide the custom filename for the HMAC key as well as the encryption key as a parameter.

- Encrypt by allowing the utility to generate encryption keys and store it in provided custom filename:

```
python nvs_partition_gen.py encrypt sample_singlepage_blob.csv sample_encr.bin
↪ 0x3000 --keygen --keyfile sample_keys.bin
```

Note:

- Encryption key of the format <outdir>/keys/sample_keys.bin is created.
- This newly created file having encryption keys in keys/ directory is compatible with NVS key-partition structure. Refer to [NVS Key Partition](#) for more details.

- Encrypt by providing the encryption keys as input binary file:

```
python nvs_partition_gen.py encrypt sample_singlepage_blob.csv sample_encr.bin
↪ 0x3000 --inputkey sample_keys.bin
```

Decrypt Encrypted NVS Partition Usage:

```
python nvs_partition_gen.py decrypt [-h] [--outdir OUTDIR] input key output
```

Positional Arguments:

Parameter	Description
input	Path to encrypted NVS partition file to parse
key	Path to file having keys for decryption
output	Path to output decrypted binary file

Optional Arguments:

Parameter	Description
<code>-h / --help</code>	Show the help message and exit
<code>--outdir OUTDIR</code>	Output directory to store files created. (Default: current directory)

You can run the utility to decrypt encrypted NVS partition using the command below:

```
python nvs_partition_gen.py decrypt sample_encr.bin sample_keys.bin sample_decr.bin
```

You can also provide the format version number:

- Multipage blob support disabled (Version 1)
- Multipage blob support enabled (Version 2)

Multipage Blob Support Disabled (Version 1) You can run the utility in this format by setting the version parameter to 1, as shown below. A sample CSV file for the same is provided with the utility:

```
python nvs_partition_gen.py generate sample_singlepage_blob.csv sample.bin 0x3000 -
↪--version 1
```

Multipage Blob Support Enabled (Version 2) You can run the utility in this format by setting the version parameter to 2, as shown below. A sample CSV file for the same is provided with the utility:

```
python nvs_partition_gen.py generate sample_multipage_blob.csv sample.bin 0x4000 --
↪version 2
```

Note:

- Minimum NVS Partition Size needed is 0x3000 bytes.
- When flashing the binary onto the device, make sure it is consistent with the application's sdkconfig.

Caveats

- Utility does not check for duplicate keys and will write data pertaining to both keys. You need to make sure that the keys are distinct.
- Once a new page is created, no data will be written in the space left on the previous page. Fields in the CSV file need to be ordered in such a way as to optimize memory.
- 64-bit datatype is not yet supported.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.8.7 NVS Partition Parser Utility

Introduction

The utility `nvs_flash/nvs_partition_tool/nvs_tool.py` loads and parses an NVS storage partition for easier debugging and data extraction. The utility also features integrity check which scans the partition for potential errors. Data blobs are encoded in `base64` format.

Encrypted Partitions

This utility does not support decryption. To decrypt the NVS partition, please use the [NVS Partition Generator Utility](#) which does support NVS partition encryption and decryption.

Usage

There are two output format styles available with the `-f` or `--format` option:

- *json* - All of the output is printed as a JSON.
- *text* - The output is printed as a human-readable text with different selectable output styles mentioned below.

For the *text* output format, the utility provides six different output styles with the `-d` or `--dump` option:

- *all* (default) - Prints all entries with metadata.
- *written* - Prints only written entries with metadata.
- *minimal* - Prints written *namespace:key = value* pairs.
- *namespaces* - Prints all written namespaces
- *blobs* - Prints all blobs and strings (reconstructs them if they are chunked).
- *storage_info* - Prints entry states count for every page.

Note: There is also a *none* option which will not print anything. This can be used with the integrity check option if the NVS partition contents are irrelevant.

The utility also provides an integrity check feature via the `-i` or `--integrity-check` option (available only with the *text* format as it would invalidate the *json* output). This feature scans through the entire partition and prints potential errors. It can be used with the `-d none` option which will print only the potential errors.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.8.8 SD/SDIO/MMC Driver

Overview

The SD/SDIO/MMC driver currently supports SD memory, SDIO cards, and eMMC chips. This is a protocol level driver built on top of SDMMC and SD SPI host drivers.

SDMMC and SD SPI host drivers ([driver/sdmmc/include/driver/sdmmc_host.h](#) and [driver/spi/include/driver/sdsp_i_host.h](#)) provide API functions for:

- Sending commands to slave devices
- Sending and receiving data
- Handling error conditions within the bus

For functions used to initialize and configure:

- SDMMC host, see [SDMMC Host API](#)
- SD SPI host, see [SD SPI Host API](#)

The SDMMC protocol layer described in this document handles the specifics of the SD protocol, such as the card initialization and data transfer commands.

The protocol layer works with the host via the `sdmmc_host_t` structure. This structure contains pointers to various functions of the host.

Pin Configurations

SDMMC have two slots:

- slot 0 pins are dedicated for UHS-I mode. This is not yet supported in the driver.
- slot 1 pins are routed via GPIO Matrix, and it's for non UHS-I usage. You will need to configure the pins in `sdmmc_slot_config_t` to use the slot 1.

Application Example

An example which combines the SDMMC driver with the FATFS library is provided in the `storage/sd_card` directory of ESP-IDF examples. This example initializes the card, then writes and reads data from it using POSIX and C library APIs. See README.md file in the example directory for more information.

Combo (Memory + IO) Cards The driver does not support SD combo cards. Combo cards are treated as IO cards.

Thread Safety Most applications need to use the protocol layer only in one task. For this reason, the protocol layer does not implement any kind of locking on the `sdmmc_card_t` structure, or when accessing SDMMC or SD SPI host drivers. Such locking is usually implemented on a higher layer, e.g., in the filesystem driver.

Protocol Layer API

The protocol layer is given the `sdmmc_host_t` structure. This structure describes the SD/MMC host driver, lists its capabilities, and provides pointers to functions of the driver. The protocol layer stores card-specific information in the `sdmmc_card_t` structure. When sending commands to the SD/MMC host driver, the protocol layer uses the `sdmmc_command_t` structure to describe the command, arguments, expected return values, and data to transfer if there is any.

Using API with SD Memory Cards

1. To initialize the host, call the host driver functions, e.g., `sdmmc_host_init()`, `sdmmc_host_init_slot()`.
2. To initialize the card, call `sdmmc_card_init()` and pass to it the parameters `host` - the host driver information, and `card` - a pointer to the structure `sdmmc_card_t` which will be filled with information about the card when the function completes.
3. To read and write sectors of the card, use `sdmmc_read_sectors()` and `sdmmc_write_sectors()` respectively and pass to it the parameter `card` - a pointer to the card information structure.
4. If the card is not used anymore, call the host driver function - e.g., `sdmmc_host_deinit()` - to disable the host peripheral and free the resources allocated by the driver.

Using API with eMMC Chips From the protocol layer's perspective, eMMC memory chips behave exactly like SD memory cards. Even though eMMCs are chips and do not have a card form factor, the terminology for SD cards can still be applied to eMMC due to the similarity of the protocol (`sdmmc_card_t`, `sdmmc_card_init`). Note that eMMC chips cannot be used over SPI, which makes them incompatible with the SD SPI host driver.

To initialize eMMC memory and perform read/write operations, follow the steps listed for SD cards in the previous section.

Using API with SDIO Cards Initialization and the probing process are the same as with SD memory cards. The only difference is in data transfer commands in SDIO mode.

During the card initialization and probing, performed with `sdmmc_card_init()`, the driver only configures the following registers of the IO card:

1. The IO portion of the card is reset by setting RES bit in the I/O Abort (0x06) register.
2. If 4-line mode is enabled in host and slot configuration, the driver attempts to set the Bus width field in the Bus Interface Control (0x07) register. If setting the field is successful, which means that the slave supports 4-line mode, the host is also switched to 4-line mode.
3. If high-speed mode is enabled in the host configuration, the SHS bit is set in the High Speed (0x13) register.

In particular, the driver does not set any bits in (1) I/O Enable and Int Enable registers, (2) I/O block sizes, etc. Applications can set them by calling `sdmmc_io_write_byte()`.

For card configuration and data transfer, choose the pair of functions relevant to your case from the table below.

Action	Read Function	Write Function
Read and write a single byte using IO_RW_DIRECT (CMD52)	<code>sd-mmc_io_read_byte()</code>	<code>sd-mmc_io_write_byte()</code>
Read and write multiple bytes using IO_RW_EXTENDED (CMD53) in byte mode	<code>sd-mmc_io_read_bytes()</code>	<code>sd-mmc_io_write_bytes()</code>
Read and write blocks of data using IO_RW_EXTENDED (CMD53) in block mode	<code>sd-mmc_io_read_blocks()</code>	<code>sd-mmc_io_write_blocks()</code>

SDIO interrupts can be enabled by the application using the function `sdmmc_io_enable_int()`. When using SDIO in 1-line mode, the D1 line also needs to be connected to use SDIO interrupts.

If you want the application to wait until the SDIO interrupt occurs, use `sdmmc_io_wait_int()`.

API Reference

Header File

- `components/sdmmc/include/sdmmc_cmd.h`
- This header file can be included with:

```
#include "sdmmc_cmd.h"
```

- This header file is a part of the API provided by the `sdmmc` component. To declare that your component depends on `sdmmc`, add the following to your `CMakeLists.txt`:

```
REQUIRES sdmmc
```

or

```
PRIV_REQUIRES sdmmc
```

Functions

`esp_err_t sdmmc_card_init` (const `sdmmc_host_t` *host, `sdmmc_card_t` *out_card)

Probe and initialize SD/MMC card using given host

Note: Only SD cards (SDSC and SDHC/SDXC) are supported now. Support for MMC/eMMC cards will be added later.

Parameters

- **host** -- pointer to structure defining host controller
- **out_card** -- pointer to structure which will receive information about the card when the function completes

Returns

- ESP_OK on success
- One of the error codes from SDMMC host controller

void **sdmmc_card_print_info** (FILE *stream, const [sdmmc_card_t](#) *card)

Print information about the card to a stream.

Parameters

- **stream** -- stream obtained using fopen or fdopen
- **card** -- card information structure initialized using sdmmc_card_init

[esp_err_t](#) **sdmmc_get_status** ([sdmmc_card_t](#) *card)

Get status of SD/MMC card

Parameters **card** -- pointer to card information structure previously initialized using sdmmc_card_init

Returns

- ESP_OK on success
- One of the error codes from SDMMC host controller

[esp_err_t](#) **sdmmc_write_sectors** ([sdmmc_card_t](#) *card, const void *src, size_t start_sector, size_t sector_count)

Write given number of sectors to SD/MMC card

Parameters

- **card** -- pointer to card information structure previously initialized using sdmmc_card_init
- **src** -- pointer to data buffer to read data from; data size must be equal to sector_count * card->csd.sector_size
- **start_sector** -- sector where to start writing
- **sector_count** -- number of sectors to write

Returns

- ESP_OK on success or sector_count equal to 0
- One of the error codes from SDMMC host controller

[esp_err_t](#) **sdmmc_read_sectors** ([sdmmc_card_t](#) *card, void *dst, size_t start_sector, size_t sector_count)

Read given number of sectors from the SD/MMC card

Parameters

- **card** -- pointer to card information structure previously initialized using sdmmc_card_init
- **dst** -- pointer to data buffer to write into; buffer size must be at least sector_count * card->csd.sector_size
- **start_sector** -- sector where to start reading
- **sector_count** -- number of sectors to read

Returns

- ESP_OK on success or sector_count equal to 0
- One of the error codes from SDMMC host controller

[esp_err_t](#) **sdmmc_erase_sectors** ([sdmmc_card_t](#) *card, size_t start_sector, size_t sector_count, [sdmmc_erase_arg_t](#) arg)

Erase given number of sectors from the SD/MMC card

Note: When sdmmc_erase_sectors used with cards in SDSPI mode, it was observed that card requires re-init after erase operation.

Parameters

- **card** -- pointer to card information structure previously initialized using sdmmc_card_init
- **start_sector** -- sector where to start erase

- **sector_count** -- number of sectors to erase
- **arg** -- erase command (CMD38) argument

Returns

- ESP_OK on success or sector_count equal to 0
- One of the error codes from SDMMC host controller

esp_err_t **sdmmc_can_discard**(*sdmmc_card_t* *card)

Check if SD/MMC card supports discard

Parameters **card** -- pointer to card information structure previously initialized using sdmmc_card_init

Returns

- ESP_OK if supported by the card/device
- ESP_FAIL if not supported by the card/device

esp_err_t **sdmmc_can_trim**(*sdmmc_card_t* *card)

Check if SD/MMC card supports trim

Parameters **card** -- pointer to card information structure previously initialized using sdmmc_card_init

Returns

- ESP_OK if supported by the card/device
- ESP_FAIL if not supported by the card/device

esp_err_t **sdmmc_mmc_can_sanitize**(*sdmmc_card_t* *card)

Check if SD/MMC card supports sanitize

Parameters **card** -- pointer to card information structure previously initialized using sdmmc_card_init

Returns

- ESP_OK if supported by the card/device
- ESP_FAIL if not supported by the card/device

esp_err_t **sdmmc_mmc_sanitize**(*sdmmc_card_t* *card, uint32_t timeout_ms)

Sanitize the data that was unmapped by a Discard command

Note: Discard command has to precede sanitize operation. To discard, use MMC_DICARD_ARG with sdmmc_erase_sectors argument

Parameters

- **card** -- pointer to card information structure previously initialized using sdmmc_card_init
- **timeout_ms** -- timeout value in milliseconds required to sanitize the selected range of sectors.

Returns

- ESP_OK on success
- One of the error codes from SDMMC host controller

esp_err_t **sdmmc_full_erase**(*sdmmc_card_t* *card)

Erase complete SD/MMC card

Parameters **card** -- pointer to card information structure previously initialized using sdmmc_card_init

Returns

- ESP_OK on success
- One of the error codes from SDMMC host controller

esp_err_t **sdmmc_io_read_byte**(*sdmmc_card_t* *card, uint32_t function, uint32_t reg, uint8_t *out_byte)

Read one byte from an SDIO card using IO_RW_DIRECT (CMD52)

Parameters

- **card** -- pointer to card information structure previously initialized using `sdmmc_card_init`
- **function** -- IO function number
- **reg** -- byte address within IO function
- **out_byte** -- **[out]** output, receives the value read from the card

Returns

- ESP_OK on success
- One of the error codes from SDMMC host controller

esp_err_t **sdmmc_io_write_byte** (*sdmmc_card_t* *card, uint32_t function, uint32_t reg, uint8_t in_byte, uint8_t *out_byte)

Write one byte to an SDIO card using IO_RW_DIRECT (CMD52)

Parameters

- **card** -- pointer to card information structure previously initialized using `sdmmc_card_init`
- **function** -- IO function number
- **reg** -- byte address within IO function
- **in_byte** -- value to be written
- **out_byte** -- **[out]** if not NULL, receives new byte value read from the card (read-after-write).

Returns

- ESP_OK on success
- One of the error codes from SDMMC host controller

esp_err_t **sdmmc_io_read_bytes** (*sdmmc_card_t* *card, uint32_t function, uint32_t addr, void *dst, size_t size)

Read multiple bytes from an SDIO card using IO_RW_EXTENDED (CMD53)

This function performs read operation using CMD53 in byte mode. For block mode, see `sdmmc_io_read_blocks`.

Parameters

- **card** -- pointer to card information structure previously initialized using `sdmmc_card_init`
- **function** -- IO function number
- **addr** -- byte address within IO function where reading starts
- **dst** -- buffer which receives the data read from card
- **size** -- number of bytes to read

Returns

- ESP_OK on success
- ESP_ERR_INVALID_SIZE if size exceeds 512 bytes
- One of the error codes from SDMMC host controller

esp_err_t **sdmmc_io_write_bytes** (*sdmmc_card_t* *card, uint32_t function, uint32_t addr, const void *src, size_t size)

Write multiple bytes to an SDIO card using IO_RW_EXTENDED (CMD53)

This function performs write operation using CMD53 in byte mode. For block mode, see `sdmmc_io_write_blocks`.

Parameters

- **card** -- pointer to card information structure previously initialized using `sdmmc_card_init`
- **function** -- IO function number
- **addr** -- byte address within IO function where writing starts
- **src** -- data to be written
- **size** -- number of bytes to write

Returns

- ESP_OK on success

- ESP_ERR_INVALID_SIZE if size exceeds 512 bytes
- One of the error codes from SDMMC host controller

esp_err_t **sdmmc_io_read_blocks** (*sdmmc_card_t* *card, uint32_t function, uint32_t addr, void *dst, size_t size)

Read blocks of data from an SDIO card using IO_RW_EXTENDED (CMD53)

This function performs read operation using CMD53 in block mode. For byte mode, see sdmmc_io_read_bytes.

Parameters

- **card** -- pointer to card information structure previously initialized using sdmmc_card_init
- **function** -- IO function number
- **addr** -- byte address within IO function where writing starts
- **dst** -- buffer which receives the data read from card
- **size** -- number of bytes to read, must be divisible by the card block size.

Returns

- ESP_OK on success
- ESP_ERR_INVALID_SIZE if size is not divisible by 512 bytes
- One of the error codes from SDMMC host controller

esp_err_t **sdmmc_io_write_blocks** (*sdmmc_card_t* *card, uint32_t function, uint32_t addr, const void *src, size_t size)

Write blocks of data to an SDIO card using IO_RW_EXTENDED (CMD53)

This function performs write operation using CMD53 in block mode. For byte mode, see sdmmc_io_write_bytes.

Parameters

- **card** -- pointer to card information structure previously initialized using sdmmc_card_init
- **function** -- IO function number
- **addr** -- byte address within IO function where writing starts
- **src** -- data to be written
- **size** -- number of bytes to read, must be divisible by the card block size.

Returns

- ESP_OK on success
- ESP_ERR_INVALID_SIZE if size is not divisible by 512 bytes
- One of the error codes from SDMMC host controller

esp_err_t **sdmmc_io_enable_int** (*sdmmc_card_t* *card)

Enable SDIO interrupt in the SDMMC host

Parameters **card** -- pointer to card information structure previously initialized using sdmmc_card_init

Returns

- ESP_OK on success
- ESP_ERR_NOT_SUPPORTED if the host controller does not support IO interrupts

esp_err_t **sdmmc_io_wait_int** (*sdmmc_card_t* *card, TickType_t timeout_ticks)

Block until an SDIO interrupt is received

Slave uses D1 line to signal interrupt condition to the host. This function can be used to wait for the interrupt.

Parameters

- **card** -- pointer to card information structure previously initialized using sdmmc_card_init
- **timeout_ticks** -- time to wait for the interrupt, in RTOS ticks

Returns

- ESP_OK if the interrupt is received
- ESP_ERR_NOT_SUPPORTED if the host controller does not support IO interrupts
- ESP_ERR_TIMEOUT if the interrupt does not happen in timeout_ticks

esp_err_t **sdmmc_io_get_cis_data** (*sdmmc_card_t* *card, uint8_t *out_buffer, size_t buffer_size, size_t *inout_cis_size)

Get the data of CIS region of an SDIO card.

You may provide a buffer not sufficient to store all the CIS data. In this case, this function stores as much data into your buffer as possible. Also, this function will try to get and return the size required for you.

Parameters

- **card** -- pointer to card information structure previously initialized using `sdmmc_card_init`
- **out_buffer** -- Output buffer of the CIS data
- **buffer_size** -- Size of the buffer.
- **inout_cis_size** -- Mandatory, pointer to a size, input and output.
 - input: Limitation of maximum searching range, should be 0 or larger than `buffer_size`. The function searches for `CIS_CODE_END` until this range. Set to 0 to search infinitely.
 - output: The size required to store all the CIS data, if `CIS_CODE_END` is found.

Returns

- `ESP_OK`: on success
- `ESP_ERR_INVALID_RESPONSE`: if the card does not (correctly) support CIS.
- `ESP_ERR_INVALID_SIZE`: `CIS_CODE_END` found, but `buffer_size` is less than required size, which is stored in the `inout_cis_size` then.
- `ESP_ERR_NOT_FOUND`: if the `CIS_CODE_END` not found. Increase input value of `inout_cis_size` or set it to 0, if you still want to search for the end; output value of `inout_cis_size` is invalid in this case.
- and other error code return from `sdmmc_io_read_bytes`

esp_err_t **sdmmc_io_print_cis_info** (uint8_t *buffer, size_t buffer_size, FILE *fp)

Parse and print the CIS information of an SDIO card.

Note: Not all the CIS codes and all kinds of tuples are supported. If you see some unresolved code, you can add the parsing of these code in `sdmmc_io.c` and contribute to the IDF through the Github repository.

`using sdmmc_card_init`

Parameters

- **buffer** -- Buffer to parse
- **buffer_size** -- Size of the buffer.
- **fp** -- File pointer to print to, set to NULL to print to stdout.

Returns

- `ESP_OK`: on success
- `ESP_ERR_NOT_SUPPORTED`: if the value from the card is not supported to be parsed.
- `ESP_ERR_INVALID_SIZE`: if the CIS size fields are not correct.

Header File

- [components/driver/sdmmc/include/driver/sdmmc_types.h](#)
- This header file can be included with:

`#include "driver/sdmmc_types.h"`

- This header file is a part of the API provided by the `driver` component. To declare that your component depends on `driver`, add the following to your `CMakeLists.txt`:

`REQUIRES driver`

or

PRIV_REQUIRES driver

Structures

struct **sdmmc_csd_t**

Decoded values from SD card Card Specific Data register

Public Members

int **csd_ver**

CSD structure format

int **mmc_ver**

MMC version (for CID format)

int **capacity**

total number of sectors

int **sector_size**

sector size in bytes

int **read_block_len**

block length for reads

int **card_command_class**

Card Command Class for SD

int **tr_speed**

Max transfer speed

struct **sdmmc_cid_t**

Decoded values from SD card Card IDentification register

Public Members

int **mfg_id**

manufacturer identification number

int **oem_id**

OEM/product identification number

char **name**[8]

product name (MMC v1 has the longest)

int **revision**

product revision

int **serial**

product serial number

int **date**

manufacturing date

struct **sdmmc_scr_t**

Decoded values from SD Configuration Register Note: When new member is added, update reserved bits accordingly

Public Members

uint32_t **sd_spec**

SD Physical layer specification version, reported by card

uint32_t **erase_mem_state**

data state on card after erase whether 0 or 1 (card vendor dependent)

uint32_t **bus_width**

bus widths supported by card: BIT(0) —1-bit bus, BIT(2) —4-bit bus

uint32_t **reserved**

reserved for future expansion

uint32_t **rsvd_mnf**

reserved for manufacturer usage

struct **sdmmc_ssr_t**

Decoded values from SD Status Register Note: When new member is added, update reserved bits accordingly

Public Members

uint32_t **alloc_unit_kb**

Allocation unit of the card, in multiples of kB (1024 bytes)

uint32_t **erase_size_au**

Erase size for the purpose of timeout calculation, in multiples of allocation unit

uint32_t **cur_bus_width**

SD current bus width

uint32_t **discard_support**

SD discard feature support

uint32_t **fule_support**

SD FULE (Full User Area Logical Erase) feature support

uint32_t **erase_timeout**

Timeout (in seconds) for erase of a single allocation unit

uint32_t **erase_offset**

Constant timeout offset (in seconds) for any erase operation

uint32_t **reserved**

reserved for future expansion

struct **sdmmc_ext_csd_t**

Decoded values of Extended Card Specific Data

Public Members

uint8_t **rev**

Extended CSD Revision

uint8_t **power_class**

Power class used by the card

uint8_t **erase_mem_state**

data state on card after erase whether 0 or 1 (card vendor dependent)

uint8_t **sec_feature**

secure data management features supported by the card

struct **sdmmc_switch_func_rsp_t**

SD SWITCH_FUNC response buffer

Public Members

uint32_t **data**[512 / 8 / sizeof(uint32_t)]

response data

struct **sdmmc_command_t**

SD/MMC command information

Public Members

uint32_t **opcode**

SD or MMC command index

uint32_t **arg**

SD/MMC command argument

sdmmc_response_t **response**

response buffer

void ***data**

buffer to send or read into

size_t **datalen**

length of data in the buffer

size_t **buflen**

length of the buffer

size_t **blklen**

block length

int **flags**

see below

esp_err_t **error**

error returned from transfer

uint32_t **timeout_ms**

response timeout, in milliseconds

struct **sdmmc_host_t**

SD/MMC Host description

This structure defines properties of SD/MMC host and functions of SD/MMC host which can be used by upper layers.

Public Members

uint32_t **flags**

flags defining host properties

int **slot**

slot number, to be passed to host functions

int **max_freq_khz**

max frequency supported by the host

float **io_voltage**

I/O voltage used by the controller (voltage switching is not supported)

esp_err_t (***init**)(void)

Host function to initialize the driver

esp_err_t (***set_bus_width**)(int slot, size_t width)

host function to set bus width

`size_t (*get_bus_width)(int slot)`

host function to get bus width

`esp_err_t (*set_bus_ddr_mode)(int slot, bool ddr_enable)`

host function to set DDR mode

`esp_err_t (*set_card_clk)(int slot, uint32_t freq_khz)`

host function to set card clock frequency

`esp_err_t (*set_cclk_always_on)(int slot, bool cclk_always_on)`

host function to set whether the clock is always enabled

`esp_err_t (*do_transaction)(int slot, sdmmc_command_t *cmdinfo)`

host function to do a transaction

`esp_err_t (*deinit)(void)`

host function to deinitialize the driver

`esp_err_t (*deinit_p)(int slot)`

host function to deinitialize the driver, called with the `slot`

`esp_err_t (*io_int_enable)(int slot)`

Host function to enable SDIO interrupt line

`esp_err_t (*io_int_wait)(int slot, TickType_t timeout_ticks)`

Host function to wait for SDIO interrupt line to be active

`int command_timeout_ms`

timeout, in milliseconds, of a single command. Set to 0 to use the default value.

`esp_err_t (*get_real_freq)(int slot, int *real_freq)`

Host function to provide real working freq, based on SDMMC controller setup

`sdmmc_delay_phase_t input_delay_phase`

input delay phase, this will only take into effect when the host works in SDMMC_FREQ_HIGHSPEED or SDMMC_FREQ_52M. Driver will print out how long the delay is

`esp_err_t (*set_input_delay)(int slot, sdmmc_delay_phase_t delay_phase)`

set input delay phase

struct `sdmmc_card_t`

SD/MMC card information structure

Public Members

`sdmmc_host_t host`

Host with which the card is associated

uint32_t **ocr**

OCR (Operation Conditions Register) value

sdmmc_cid_t **cid**

decoded CID (Card Identification) register value

sdmmc_response_t **raw_cid**

raw CID of MMC card to be decoded after the CSD is fetched in the data transfer mode

sdmmc_csd_t **csd**

decoded CSD (Card-Specific Data) register value

sdmmc_scr_t **scr**

decoded SCR (SD card Configuration Register) value

sdmmc_ssr_t **ssr**

decoded SSR (SD Status Register) value

sdmmc_ext_csd_t **ext_csd**

decoded EXT_CSD (Extended Card Specific Data) register value

uint16_t **rca**

RCA (Relative Card Address)

uint16_t **max_freq_khz**

Maximum frequency, in kHz, supported by the card

int **real_freq_khz**

Real working frequency, in kHz, configured on the host controller

uint32_t **is_mem**

Bit indicates if the card is a memory card

uint32_t **is_sdio**

Bit indicates if the card is an IO card

uint32_t **is_mmc**

Bit indicates if the card is MMC

uint32_t **num_io_functions**

If is_sdio is 1, contains the number of IO functions on the card

uint32_t **log_bus_width**

log2(bus width supported by card)

uint32_t **is_ddr**

Card supports DDR mode

uint32_t **reserved**

Reserved for future expansion

Macros

SDMMC_HOST_FLAG_1BIT

host supports 1-line SD and MMC protocol

SDMMC_HOST_FLAG_4BIT

host supports 4-line SD and MMC protocol

SDMMC_HOST_FLAG_8BIT

host supports 8-line MMC protocol

SDMMC_HOST_FLAG_SPI

host supports SPI protocol

SDMMC_HOST_FLAG_DDR

host supports DDR mode for SD/MMC

SDMMC_HOST_FLAG_DEINIT_ARG

host `deinit` function called with the slot argument

SDMMC_FREQ_DEFAULT

SD/MMC Default speed (limited by clock divider)

SDMMC_FREQ_HIGHSPEED

SD High speed (limited by clock divider)

SDMMC_FREQ_PROBING

SD/MMC probing speed

SDMMC_FREQ_52M

MMC 52MHz speed

SDMMC_FREQ_26M

MMC 26MHz speed

Type Definitions

```
typedef uint32_t sdmmc_response_t[4]
```

SD/MMC command response buffer

Enumerations

```
enum sdmmc_delay_phase_t
```

SD/MMC Host clock timing delay phases

This will only take effect when the host works in `SDMMC_FREQ_HIGHSPEED` or `SDMMC_FREQ_52M`. Driver will print out how long the delay is, in picosecond (ps).

Values:

enumerator **SDMMC_DELAY_PHASE_0**

Delay phase 0

enumerator **SDMMC_DELAY_PHASE_1**

Delay phase 1

enumerator **SDMMC_DELAY_PHASE_2**

Delay phase 2

enumerator **SDMMC_DELAY_PHASE_3**

Delay phase 3

enum **sdmmc_erase_arg_t**

SD/MMC erase command(38) arguments SD: ERASE: Erase the write blocks, physical/hard erase.

DISCARD: Card may deallocate the discarded blocks partially or completely. After discard operation the previously written data may be partially or fully read by the host depending on card implementation.

MMC: ERASE: Does TRIM, applies erase operation to write blocks instead of Erase Group.

DISCARD: The Discard function allows the host to identify data that is no longer required so that the device can erase the data if necessary during background erase events. Applies to write blocks instead of Erase Group. After discard operation, the original data may be remained partially or fully accessible to the host dependent on device.

Values:

enumerator **SDMMC_ERASE_ARG**

Erase operation on SD, Trim operation on MMC

enumerator **SDMMC_DISCARD_ARG**

Discard operation for SD/MMC

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.8.9 Partitions API

Overview

The `esp_partition` component has higher-level API functions which work with partitions defined in the [Partition Tables](#). These APIs are based on lower level API provided by [SPI Flash API](#).

Partition Table API

ESP-IDF projects use a partition table to maintain information about various regions of SPI flash memory (bootloader, various application binaries, data, filesystems). More information can be found in [Partition Tables](#).

This component provides API functions to enumerate partitions found in the partition table and perform operations on them. These functions are declared in `esp_partition.h`:

- [esp_partition_find\(\)](#) checks a partition table for entries with specific type, returns an opaque iterator.
- [esp_partition_get\(\)](#) returns a structure describing the partition for a given iterator.
- [esp_partition_next\(\)](#) shifts the iterator to the next found partition.
- [esp_partition_iterator_release\(\)](#) releases iterator returned by [esp_partition_find\(\)](#).

- `esp_partition_find_first()` is a convenience function which returns the structure describing the first partition found by `esp_partition_find()`.
- `esp_partition_read()`, `esp_partition_write()`, `esp_partition_erase_range()` are equivalent to `esp_flash_read()`, `esp_flash_write()`, `esp_flash_erase_region()`, but operate within partition boundaries.

See Also

- [Partition Tables](#)
- [Over The Air Updates \(OTA\)](#) provides high-level API for updating applications stored in flash.
- [Non-Volatile Storage Library](#) provides a structured API for storing small pieces of data in SPI flash.

API Reference - Partition Table

Header File

- `components/esp_partition/include/esp_partition.h`
- This header file can be included with:

```
#include "esp_partition.h"
```

- This header file is a part of the API provided by the `esp_partition` component. To declare that your component depends on `esp_partition`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_partition
```

or

```
PRIV_REQUIRES esp_partition
```

Functions

`esp_partition_iterator_t` **esp_partition_find** (`esp_partition_type_t` type, `esp_partition_subtype_t` subtype, `const char *label`)

Find partition based on one or more parameters.

Parameters

- **type** -- Partition type, one of `esp_partition_type_t` values or an 8-bit unsigned integer. To find all partitions, no matter the type, use `ESP_PARTITION_TYPE_ANY`, and set subtype argument to `ESP_PARTITION_SUBTYPE_ANY`.
- **subtype** -- Partition subtype, one of `esp_partition_subtype_t` values or an 8-bit unsigned integer. To find all partitions of given type, use `ESP_PARTITION_SUBTYPE_ANY`.
- **label** -- (optional) Partition label. Set this value if looking for partition with a specific name. Pass `NULL` otherwise.

Returns iterator which can be used to enumerate all the partitions found, or `NULL` if no partitions were found. Iterator obtained through this function has to be released using `esp_partition_iterator_release` when not used any more.

`const esp_partition_t *`**esp_partition_find_first** (`esp_partition_type_t` type, `esp_partition_subtype_t` subtype, `const char *label`)

Find first partition based on one or more parameters.

Parameters

- **type** -- Partition type, one of `esp_partition_type_t` values or an 8-bit unsigned integer. To find all partitions, no matter the type, use `ESP_PARTITION_TYPE_ANY`, and set subtype argument to `ESP_PARTITION_SUBTYPE_ANY`.
- **subtype** -- Partition subtype, one of `esp_partition_subtype_t` values or an 8-bit unsigned integer. To find all partitions of given type, use `ESP_PARTITION_SUBTYPE_ANY`.
- **label** -- (optional) Partition label. Set this value if looking for partition with a specific name. Pass `NULL` otherwise.

Returns pointer to [esp_partition_t](#) structure, or NULL if no partition is found. This pointer is valid for the lifetime of the application.

const [esp_partition_t](#) ***esp_partition_get** ([esp_partition_iterator_t](#) iterator)

Get [esp_partition_t](#) structure for given partition.

Parameters **iterator** -- Iterator obtained using `esp_partition_find`. Must be non-NULL.

Returns pointer to [esp_partition_t](#) structure. This pointer is valid for the lifetime of the application.

[esp_partition_iterator_t](#) **esp_partition_next** ([esp_partition_iterator_t](#) iterator)

Move partition iterator to the next partition found.

Any copies of the iterator will be invalid after this call.

Parameters **iterator** -- Iterator obtained using `esp_partition_find`. Must be non-NULL.

Returns NULL if no partition was found, valid `esp_partition_iterator_t` otherwise.

void **esp_partition_iterator_release** ([esp_partition_iterator_t](#) iterator)

Release partition iterator.

Parameters **iterator** -- Iterator obtained using `esp_partition_find`. The iterator is allowed to be NULL, so it is not necessary to check its value before calling this function.

const [esp_partition_t](#) ***esp_partition_verify** (const [esp_partition_t](#) *partition)

Verify partition data.

Given a pointer to partition data, verify this partition exists in the partition table (all fields match.)

This function is also useful to take partition data which may be in a RAM buffer and convert it to a pointer to the permanent partition data stored in flash.

Pointers returned from this function can be compared directly to the address of any pointer returned from [esp_partition_get\(\)](#), as a test for equality.

Parameters **partition** -- Pointer to partition data to verify. Must be non-NULL. All fields of this structure must match the partition table entry in flash for this function to return a successful match.

Returns

- If partition not found, returns NULL.
- If found, returns a pointer to the [esp_partition_t](#) structure in flash. This pointer is always valid for the lifetime of the application.

[esp_err_t](#) **esp_partition_read** (const [esp_partition_t](#) *partition, size_t src_offset, void *dst, size_t size)

Read data from the partition.

Partitions marked with an encryption flag will automatically be read and decrypted via a cache mapping.

Parameters

- **partition** -- Pointer to partition structure obtained using `esp_partition_find_first` or `esp_partition_get`. Must be non-NULL.
- **dst** -- Pointer to the buffer where data should be stored. Pointer must be non-NULL and buffer must be at least 'size' bytes long.
- **src_offset** -- Address of the data to be read, relative to the beginning of the partition.
- **size** -- Size of data to be read, in bytes.

Returns ESP_OK, if data was read successfully; ESP_ERR_INVALID_ARG, if `src_offset` exceeds partition size; ESP_ERR_INVALID_SIZE, if read would go out of bounds of the partition; or one of error codes from lower-level flash driver.

[esp_err_t](#) **esp_partition_write** (const [esp_partition_t](#) *partition, size_t dst_offset, const void *src, size_t size)

Write data to the partition.

Before writing data to flash, corresponding region of flash needs to be erased. This can be done using `esp_partition_erase_range` function.

Partitions marked with an encryption flag will automatically be written via the `esp_flash_write_encrypted()` function. If writing to an encrypted partition, all write offsets and lengths must be multiples of 16 bytes. See the `esp_flash_write_encrypted()` function for more details. Unencrypted partitions do not have this restriction.

Note: Prior to writing to flash memory, make sure it has been erased with `esp_partition_erase_range` call.

Parameters

- **partition** -- Pointer to partition structure obtained using `esp_partition_find_first` or `esp_partition_get`. Must be non-NULL.
- **dst_offset** -- Address where the data should be written, relative to the beginning of the partition.
- **src** -- Pointer to the source buffer. Pointer must be non-NULL and buffer must be at least 'size' bytes long.
- **size** -- Size of data to be written, in bytes.

Returns `ESP_OK`, if data was written successfully; `ESP_ERR_INVALID_ARG`, if `dst_offset` exceeds partition size; `ESP_ERR_INVALID_SIZE`, if write would go out of bounds of the partition; `ESP_ERR_NOT_ALLOWED`, if partition is read-only; or one of error codes from lower-level flash driver.

esp_err_t **esp_partition_read_raw** (const *esp_partition_t* *partition, size_t src_offset, void *dst, size_t size)

Read data from the partition without any transformation/decryption.

Note: This function is essentially the same as `esp_partition_read()` above. It just never decrypts data but returns it as is.

Parameters

- **partition** -- Pointer to partition structure obtained using `esp_partition_find_first` or `esp_partition_get`. Must be non-NULL.
- **dst** -- Pointer to the buffer where data should be stored. Pointer must be non-NULL and buffer must be at least 'size' bytes long.
- **src_offset** -- Address of the data to be read, relative to the beginning of the partition.
- **size** -- Size of data to be read, in bytes.

Returns `ESP_OK`, if data was read successfully; `ESP_ERR_INVALID_ARG`, if `src_offset` exceeds partition size; `ESP_ERR_INVALID_SIZE`, if read would go out of bounds of the partition; or one of error codes from lower-level flash driver.

esp_err_t **esp_partition_write_raw** (const *esp_partition_t* *partition, size_t dst_offset, const void *src, size_t size)

Write data to the partition without any transformation/encryption.

Before writing data to flash, corresponding region of flash needs to be erased. This can be done using `esp_partition_erase_range` function.

Note: This function is essentially the same as `esp_partition_write()` above. It just never encrypts data but writes it as is.

Note: Prior to writing to flash memory, make sure it has been erased with `esp_partition_erase_range` call.

Parameters

- **partition** -- Pointer to partition structure obtained using `esp_partition_find_first` or `esp_partition_get`. Must be non-NULL.
- **dst_offset** -- Address where the data should be written, relative to the beginning of the partition.
- **src** -- Pointer to the source buffer. Pointer must be non-NULL and buffer must be at least 'size' bytes long.
- **size** -- Size of data to be written, in bytes.

Returns `ESP_OK`, if data was written successfully; `ESP_ERR_INVALID_ARG`, if `dst_offset` exceeds partition size; `ESP_ERR_INVALID_SIZE`, if write would go out of bounds of the partition; `ESP_ERR_NOT_ALLOWED`, if partition is read-only; or one of the error codes from lower-level flash driver.

esp_err_t **esp_partition_erase_range** (const *esp_partition_t* *partition, size_t offset, size_t size)

Erase part of the partition.

Parameters

- **partition** -- Pointer to partition structure obtained using `esp_partition_find_first` or `esp_partition_get`. Must be non-NULL.
- **offset** -- Offset from the beginning of partition where erase operation should start. Must be aligned to `partition->erase_size`.
- **size** -- Size of the range which should be erased, in bytes. Must be divisible by `partition->erase_size`.

Returns `ESP_OK`, if the range was erased successfully; `ESP_ERR_INVALID_ARG`, if iterator or `dst` are NULL; `ESP_ERR_INVALID_SIZE`, if erase would go out of bounds of the partition; `ESP_ERR_NOT_ALLOWED`, if partition is read-only; or one of error codes from lower-level flash driver.

esp_err_t **esp_partition_mmap** (const *esp_partition_t* *partition, size_t offset, size_t size, *esp_partition_mmap_memory_t* memory, const void **out_ptr, *esp_partition_mmap_handle_t* *out_handle)

Configure MMU to map partition into data memory.

Unlike `spi_flash_mmap` function, which requires a 64kB aligned base address, this function doesn't impose such a requirement. If offset results in a flash address which is not aligned to 64kB boundary, address will be rounded to the lower 64kB boundary, so that mapped region includes requested range. Pointer returned via `out_ptr` argument will be adjusted to point to the requested offset (not necessarily to the beginning of mmap-ed region).

To release mapped memory, pass handle returned via `out_handle` argument to `esp_partition_munmap` function.

Parameters

- **partition** -- Pointer to partition structure obtained using `esp_partition_find_first` or `esp_partition_get`. Must be non-NULL.
- **offset** -- Offset from the beginning of partition where mapping should start.
- **size** -- Size of the area to be mapped.
- **memory** -- Memory space where the region should be mapped
- **out_ptr** -- Output, pointer to the mapped memory region
- **out_handle** -- Output, handle which should be used for `esp_partition_munmap` call

Returns `ESP_OK`, if successful

void **esp_partition_munmap** (*esp_partition_mmap_handle_t* handle)

Release region previously obtained using `esp_partition_mmap`.

Note: Calling this function will not necessarily unmap memory region. Region will only be unmapped when there are no other handles which reference this region. In case of partially overlapping regions it is possible that memory will be unmapped partially.

Parameters **handle** -- Handle obtained from `spi_flash_mmap`

esp_err_t **esp_partition_get_sha256** (const *esp_partition_t* *partition, uint8_t *sha_256)

Get SHA-256 digest for required partition.

For apps with SHA-256 appended to the app image, the result is the appended SHA-256 value for the app image content. The hash is verified before returning, if app content is invalid then the function returns ESP_ERR_IMAGE_INVALID. For apps without SHA-256 appended to the image, the result is the SHA-256 of all bytes in the app image. For other partition types, the result is the SHA-256 of the entire partition.

Parameters

- **partition** -- [in] Pointer to info for partition containing app or data. (fields: address, size and type, are required to be filled).
- **sha_256** -- [out] Returned SHA-256 digest for a given partition.

Returns

- ESP_OK: In case of successful operation.
- ESP_ERR_INVALID_ARG: The size was 0 or the sha_256 was NULL.
- ESP_ERR_NO_MEM: Cannot allocate memory for sha256 operation.
- ESP_ERR_IMAGE_INVALID: App partition doesn't contain a valid app image.
- ESP_FAIL: An allocation error occurred.

bool **esp_partition_check_identity** (const *esp_partition_t* *partition_1, const *esp_partition_t* *partition_2)

Check for the identity of two partitions by SHA-256 digest.

Parameters

- **partition_1** -- [in] Pointer to info for partition 1 containing app or data. (fields: address, size and type, are required to be filled).
- **partition_2** -- [in] Pointer to info for partition 2 containing app or data. (fields: address, size and type, are required to be filled).

Returns

- True: In case of the two firmware is equal.
- False: Otherwise

esp_err_t **esp_partition_register_external** (*esp_flash_t* *flash_chip, size_t offset, size_t size, const char *label, *esp_partition_type_t* type, *esp_partition_subtype_t* subtype, const *esp_partition_t* **out_partition)

Register a partition on an external flash chip.

This API allows designating certain areas of external flash chips (identified by the *esp_flash_t* structure) as partitions. This allows using them with components which access SPI flash through the esp_partition API.

Parameters

- **flash_chip** -- Pointer to the structure identifying the flash chip
- **offset** -- Address in bytes, where the partition starts
- **size** -- Size of the partition in bytes
- **label** -- Partition name
- **type** -- One of the partition types (ESP_PARTITION_TYPE_*), or an integer. Note that applications can not be booted from external flash chips, so using ESP_PARTITION_TYPE_APP is not supported.
- **subtype** -- One of the partition subtypes (ESP_PARTITION_SUBTYPE_*), or an integer.
- **out_partition** -- [out] Output, if non-NULL, receives the pointer to the resulting *esp_partition_t* structure

Returns

- ESP_OK on success
- ESP_ERR_NO_MEM if memory allocation has failed
- ESP_ERR_INVALID_ARG if the new partition overlaps another partition on the same flash chip
- ESP_ERR_INVALID_SIZE if the partition doesn't fit into the flash chip size

esp_err_t **esp_partition_deregister_external** (const *esp_partition_t* *partition)

Deregister the partition previously registered using `esp_partition_register_external`.

Parameters *partition* -- pointer to the partition structure obtained from `esp_partition_register_external`,

Returns

- ESP_OK on success
- ESP_ERR_NOT_FOUND if the partition pointer is not found
- ESP_ERR_INVALID_ARG if the partition comes from the partition table
- ESP_ERR_INVALID_ARG if the partition was not registered using `esp_partition_register_external` function.

void **esp_partition_unload_all** (void)

Unload partitions and free space allocated by them.

Structures

struct **esp_partition_t**

partition information structure

This is not the format in flash, that format is `esp_partition_info_t`.

However, this is the format used by this API.

Public Members

esp_flash_t ***flash_chip**

SPI flash chip on which the partition resides

esp_partition_type_t **type**

partition type (app/data)

esp_partition_subtype_t **subtype**

partition subtype

uint32_t **address**

starting address of the partition in flash

uint32_t **size**

size of the partition, in bytes

uint32_t **erase_size**

size the erase operation should be aligned to

char **label**[17]

partition label, zero-terminated ASCII string

bool **encrypted**

flag is set to true if partition is encrypted

bool **readonly**

flag is set to true if partition is read-only

Macros

ESP_PARTITION_SUBTYPE_OTA (i)

Convenience macro to get esp_partition_subtype_t value for the i-th OTA partition.

Type Definitions

typedef uint32_t **esp_partition_mmap_handle_t**

Opaque handle for memory region obtained from esp_partition_mmap.

typedef struct esp_partition_iterator_opaque_ ***esp_partition_iterator_t**

Opaque partition iterator type.

Enumerations

enum **esp_partition_mmap_memory_t**

Enumeration which specifies memory space requested in an mmap call.

Values:

enumerator **ESP_PARTITION_MMAP_DATA**

map to data memory (Vaddr0), allows byte-aligned access, 4 MB total

enumerator **ESP_PARTITION_MMAP_INST**

map to instruction memory (Vaddr1-3), allows only 4-byte-aligned access, 11 MB total

enum **esp_partition_type_t**

Partition type.

Note: Partition types with integer value 0x00-0x3F are reserved for partition types defined by ESP-IDF. Any other integer value 0x40-0xFE can be used by individual applications, without restriction.

Values:

enumerator **ESP_PARTITION_TYPE_APP**

Application partition type.

enumerator **ESP_PARTITION_TYPE_DATA**

Data partition type.

enumerator **ESP_PARTITION_TYPE_ANY**

Used to search for partitions with any type.

enum **esp_partition_subtype_t**

Partition subtype.

Application-defined partition types (0x40-0xFE) can set any numeric subtype value.

Note: These ESP-IDF-defined partition subtypes apply to partitions of type ESP_PARTITION_TYPE_APP and ESP_PARTITION_TYPE_DATA.

Values:

enumerator **ESP_PARTITION_SUBTYPE_APP_FACTORY**

Factory application partition.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_MIN**

Base for OTA partition subtypes.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_0**

OTA partition 0.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_1**

OTA partition 1.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_2**

OTA partition 2.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_3**

OTA partition 3.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_4**

OTA partition 4.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_5**

OTA partition 5.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_6**

OTA partition 6.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_7**

OTA partition 7.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_8**

OTA partition 8.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_9**

OTA partition 9.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_10**

OTA partition 10.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_11**

OTA partition 11.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_12**

OTA partition 12.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_13**

OTA partition 13.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_14**

OTA partition 14.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_15**

OTA partition 15.

enumerator **ESP_PARTITION_SUBTYPE_APP_OTA_MAX**

Max subtype of OTA partition.

enumerator **ESP_PARTITION_SUBTYPE_APP_TEST**

Test application partition.

enumerator **ESP_PARTITION_SUBTYPE_DATA_OTA**

OTA selection partition.

enumerator **ESP_PARTITION_SUBTYPE_DATA_PHY**

PHY init data partition.

enumerator **ESP_PARTITION_SUBTYPE_DATA_NVS**

NVS partition.

enumerator **ESP_PARTITION_SUBTYPE_DATA_COREDUMP**

COREDUMP partition.

enumerator **ESP_PARTITION_SUBTYPE_DATA_NVS_KEYS**

Partition for NVS keys.

enumerator **ESP_PARTITION_SUBTYPE_DATA_EFUSE_EM**

Partition for emulate eFuse bits.

enumerator **ESP_PARTITION_SUBTYPE_DATA_UNDEFINED**

Undefined (or unspecified) data partition.

enumerator **ESP_PARTITION_SUBTYPE_DATA_ESPHTTPD**

ESPHTTPD partition.

enumerator **ESP_PARTITION_SUBTYPE_DATA_FAT**

FAT partition.

enumerator **ESP_PARTITION_SUBTYPE_DATA_SPIFFS**

SPIFFS partition.

enumerator **ESP_PARTITION_SUBTYPE_DATA_LITTLEFS**

LITTLEFS partition.

enumerator **ESP_PARTITION_SUBTYPE_ANY**

Used to search for partitions with any subtype.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.8.10 SPIFFS Filesystem

Overview

SPIFFS is a file system intended for SPI NOR flash devices on embedded targets. It supports wear levelling, file system consistency checks, and more.

Notes

- Currently, SPIFFS does not support directories, it produces a flat structure. If SPIFFS is mounted under `/spiffs`, then creating a file with the path `/spiffs/tmp/myfile.txt` will create a file called `/tmp/myfile.txt` in SPIFFS, instead of `myfile.txt` in the directory `/spiffs/tmp`.
- It is not a real-time stack. One write operation might take much longer than another.
- For now, it does not detect or handle bad blocks.
- SPIFFS is able to reliably utilize only around 75% of assigned partition space.
- When the filesystem is running out of space, the garbage collector is trying to find free space by scanning the filesystem multiple times, which can take up to several seconds per write function call, depending on required space. This is caused by the SPIFFS design and the issue has been reported multiple times (e.g., [here](#)) and in the official [SPIFFS github repository](#). The issue can be partially mitigated by the [SPIFFS configuration](#).
- Deleting a file does not always remove the whole file, which leaves unusable sections throughout the filesystem.
- When the chip experiences a power loss during a file system operation it could result in SPIFFS corruption. However the file system still might be recovered via `esp_spiffs_check` function. More details in the official SPIFFS [FAQ](#).

Tools

spiffsgen.py `spiffsgen.py` is a write-only Python SPIFFS implementation used to create filesystem images from the contents of a host folder. To use `spiffsgen.py`, open Terminal and run:

```
python spiffsgen.py <image_size> <base_dir> <output_file>
```

The required arguments are as follows:

- **image_size**: size of the partition onto which the created SPIFFS image will be flashed.
- **base_dir**: directory for which the SPIFFS image needs to be created.
- **output_file**: SPIFFS image output file.

There are also other arguments that control image generation. Documentation on these arguments can be found in the tool's help:

```
python spiffsgen.py --help
```

These optional arguments correspond to a possible SPIFFS build configuration. To generate the right image, please make sure that you use the same arguments/configuration as were used to build SPIFFS. As a guide, the help output indicates the SPIFFS build configuration to which the argument corresponds. In cases when these arguments are not specified, the default values shown in the help output will be used.

When the image is created, it can be flashed using `esptool.py` or `parttool.py`.

Aside from invoking the `spiffsgen.py` standalone by manually running it from the command line or a script, it is also possible to invoke `spiffsgen.py` directly from the build system by calling `spiffs_create_partition_image`:

```
spiffs_create_partition_image(<partition> <base_dir> [FLASH_IN_PROJECT] [DEPENDS_
↳dep dep dep...])
```

This is more convenient as the build configuration is automatically passed to the tool, ensuring that the generated image is valid for that build. An example of this is while the **image_size** is required for the standalone invocation, only the **partition** name is required when using `spiffs_create_partition_image` -- the image size is automatically obtained from the project's partition table.

`spiffs_create_partition_image` must be called from one of the component `CMakeLists.txt` files.

Optionally, users can opt to have the image automatically flashed together with the app binaries, partition tables, etc. on `idf.py flash` by specifying `FLASH_IN_PROJECT`. For example:

```
spiffs_create_partition_image(my_spiffs_partition my_folder FLASH_IN_PROJECT)
```

If `FLASH_IN_PROJECT/SPIFFS_IMAGE_FLASH_IN_PROJECT` is not specified, the image will still be generated, but you will have to flash it manually using `esptool.py`, `parttool.py`, or a custom build system target.

There are cases where the contents of the base directory itself is generated at build time. Users can use `DEPENDS/SPIFFS_IMAGE_DEPENDS` to specify targets that should be executed before generating the image:

```
add_custom_target(dep COMMAND ...)

spiffs_create_partition_image(my_spiffs_partition my_folder DEPENDS dep)
```

For an example, see [storage/spiffsgen](#).

mkspiffs Another tool for creating SPIFFS partition images is `mkspiffs`. Similar to `spiffsgen.py`, it can be used to create an image from a given folder and then flash that image using `esptool.py`

For that, you need to obtain the following parameters:

- **Block Size:** 4096 (standard for SPI Flash)
- **Page Size:** 256 (standard for SPI Flash)
- **Image Size:** Size of the partition in bytes (can be obtained from a partition table)
- **Partition Offset:** Starting address of the partition (can be obtained from a partition table)

To pack a folder into a 1-Megabyte image, run:

```
mkspiffs -c [src_folder] -b 4096 -p 256 -s 0x100000 spiffs.bin
```

To flash the image onto ESP32-P4 at offset 0x110000, run:

```
python esptool.py --chip esp32p4 --port [port] --baud [baud] write_flash -z_
↳0x110000 spiffs.bin
```

Note: You can configure the `write_flash` command of `esptool.py` to [write the spiffs data to an external SPI flash chip](#) using the `--spi-connection <CLK>, <Q>, <D>, <HD>, <CS>` option. Just specify the GPIO pins assigned to the external flash, e.g. `python esptool.py write_flash --spi-connection 6, 7, 8, 9, 11 -z 0x110000 spiffs.bin`.

Notes on Which SPIFFS Tool to Use The two tools presented above offer very similar functionality. However, there are reasons to prefer one over the other, depending on the use case.

Use `spiffsgen.py` in the following cases:

1. If you want to simply generate a SPIFFS image during the build. `spiffsgen.py` makes it very convenient by providing functions/commands from the build system itself.
2. If the host has no C/C++ compiler available, because `spiffsgen.py` does not require compilation.

Use `mkspiffs` in the following cases:

1. If you need to unpack SPIFFS images in addition to image generation. For now, it is not possible with `spiffsgen.py`.
2. If you have an environment where a Python interpreter is not available, but a host compiler is available. Otherwise, a pre-compiled `mkspiffs` binary can do the job. However, there is no build system integration for `mkspiffs` and the user has to do the corresponding work: compiling `mkspiffs` during build (if a pre-compiled binary is not used), creating build rules/targets for the output files, passing proper parameters to the tool, etc.

See Also

- [Partition Table documentation](#)

Application Example

An example of using SPIFFS is provided in the [storage/spiffs](#) directory. This example initializes and mounts a SPIFFS partition, then writes and reads data from it using POSIX and C library APIs. See the `README.md` file in the example directory for more information.

High-level API Reference

Header File

- [components/spiffs/include/esp_spiffs.h](#)
- This header file can be included with:

```
#include "esp_spiffs.h"
```

- This header file is a part of the API provided by the `spiffs` component. To declare that your component depends on `spiffs`, add the following to your `CMakeLists.txt`:

```
REQUIRES spiffs
```

or

```
PRIV_REQUIRES spiffs
```

Functions

`esp_err_t esp_vfs_spiffs_register` (const `esp_vfs_spiffs_conf_t` *conf)

Register and mount SPIFFS to VFS with given path prefix.

Parameters `conf` -- Pointer to `esp_vfs_spiffs_conf_t` configuration structure

Returns

- `ESP_OK` if success
- `ESP_ERR_NO_MEM` if objects could not be allocated
- `ESP_ERR_INVALID_STATE` if already mounted or partition is encrypted
- `ESP_ERR_NOT_FOUND` if partition for SPIFFS was not found
- `ESP_FAIL` if mount or format fails

`esp_err_t esp_vfs_spiffs_unregister` (const char *partition_label)

Unregister and unmount SPIFFS from VFS

Parameters `partition_label` -- Same label as passed to `esp_vfs_spiffs_register`.

Returns

- ESP_OK if successful
- ESP_ERR_INVALID_STATE already unregistered

bool **esp_spiffs_mounted** (const char *partition_label)

Check if SPIFFS is mounted

Parameters **partition_label** -- Optional, label of the partition to check. If not specified, first partition with subtype=spiffs is used.

Returns

- true if mounted
- false if not mounted

esp_err_t **esp_spiffs_format** (const char *partition_label)

Format the SPIFFS partition

Parameters **partition_label** -- Same label as passed to esp_vfs_spiffs_register.

Returns

- ESP_OK if successful
- ESP_FAIL on error

esp_err_t **esp_spiffs_info** (const char *partition_label, size_t *total_bytes, size_t *used_bytes)

Get information for SPIFFS

Parameters

- **partition_label** -- Same label as passed to esp_vfs_spiffs_register
- **total_bytes** -- [out] Size of the file system
- **used_bytes** -- [out] Current used bytes in the file system

Returns

- ESP_OK if success
- ESP_ERR_INVALID_STATE if not mounted

esp_err_t **esp_spiffs_check** (const char *partition_label)

Check integrity of SPIFFS

Parameters **partition_label** -- Same label as passed to esp_vfs_spiffs_register

Returns

- ESP_OK if successful
- ESP_ERR_INVALID_STATE if not mounted
- ESP_FAIL on error

esp_err_t **esp_spiffs_gc** (const char *partition_label, size_t size_to_gc)

Perform garbage collection in SPIFFS partition.

Call this function to run GC and ensure that at least the given amount of space is available in the partition. This function will fail with ESP_ERR_NOT_FINISHED if it is not possible to reclaim the requested space (that is, not enough free or deleted pages in the filesystem). This function will also fail if it fails to reclaim the requested space after CONFIG_SPIFFS_GC_MAX_RUNS number of GC iterations. On one GC iteration, SPIFFS will erase one logical block (4kB). Therefore the value of CONFIG_SPIFFS_GC_MAX_RUNS should be set at least to the maximum expected size_to_gc, divided by 4096. For example, if the application expects to make room for a 1MB file and calls esp_spiffs_gc(label, 1024 * 1024), CONFIG_SPIFFS_GC_MAX_RUNS should be set to at least 256. On the other hand, increasing CONFIG_SPIFFS_GC_MAX_RUNS value increases the maximum amount of time for which any SPIFFS GC or write operation may potentially block.

Parameters

- **partition_label** -- Label of the partition to be garbage-collected. The partition must be already mounted.
- **size_to_gc** -- The number of bytes that the GC process should attempt to make available.

Returns

- ESP_OK on success
- ESP_ERR_NOT_FINISHED if GC fails to reclaim the size given by size_to_gc
- ESP_ERR_INVALID_STATE if the partition is not mounted
- ESP_FAIL on all other errors

Structures

struct **esp_vfs_spiffs_conf_t**

Configuration structure for esp_vfs_spiffs_register.

Public Members

const char ***base_path**

File path prefix associated with the filesystem.

const char ***partition_label**

Optional, label of SPIFFS partition to use. If set to NULL, first partition with subtype=spiffs will be used.

size_t **max_files**

Maximum files that could be open at the same time.

bool **format_if_mount_failed**

If true, it will format the file system if it fails to mount.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.8.11 Virtual Filesystem Component

Overview

Virtual filesystem (VFS) component provides a unified interface for drivers which can perform operations on file-like objects. These can be real filesystems (FAT, SPIFFS, etc.) or device drivers which provide a file-like interface.

This component allows C library functions, such as fopen and fprintf, to work with FS drivers. At a high level, each FS driver is associated with some path prefix. When one of C library functions needs to open a file, the VFS component searches for the FS driver associated with the file path and forwards the call to that driver. VFS also forwards read, write, and other calls for the given file to the same FS driver.

For example, one can register a FAT filesystem driver with the /fat prefix and call fopen("/fat/file.txt", "w"). Then the VFS component calls the function open of the FAT driver and pass the argument /file.txt to it together with appropriate mode flags. All subsequent calls to C library functions for the returned FILE* stream will also be forwarded to the FAT driver.

FS Registration

To register an FS driver, an application needs to define an instance of the *esp_vfs_t* structure and populate it with function pointers to FS APIs:

```
esp_vfs_t myfs = {  
    .flags = ESP_VFS_FLAG_DEFAULT,  
    .write = &myfs_write,  
    .open = &myfs_open,  
    .fstat = &myfs_fstat,  
};
```

(continues on next page)

(continued from previous page)

```

        .close = &myfs_close,
        .read = &myfs_read,
    };

ESP_ERROR_CHECK(esp_vfs_register("/data", &myfs, NULL));

```

Depending on the way how the FS driver declares its API functions, either `read`, `write`, etc., or `read_p`, `write_p`, etc., should be used.

Case 1: API functions are declared without an extra context pointer (the FS driver is a singleton):

```

ssize_t myfs_write(int fd, const void * data, size_t size);

// In definition of esp_vfs_t:
    .flags = ESP_VFS_FLAG_DEFAULT,
    .write = &myfs_write,
// ... other members initialized

// When registering FS, context pointer (third argument) is NULL:
ESP_ERROR_CHECK(esp_vfs_register("/data", &myfs, NULL));

```

Case 2: API functions are declared with an extra context pointer (the FS driver supports multiple instances):

```

ssize_t myfs_write(myfs_t* fs, int fd, const void * data, size_t size);

// In definition of esp_vfs_t:
    .flags = ESP_VFS_FLAG_CONTEXT_PTR,
    .write_p = &myfs_write,
// ... other members initialized

// When registering FS, pass the FS context pointer into the third argument
// (hypothetical myfs_mount function is used for illustrative purposes)
myfs_t* myfs_inst1 = myfs_mount(partition1->offset, partition1->size);
ESP_ERROR_CHECK(esp_vfs_register("/data1", &myfs, myfs_inst1));

// Can register another instance:
myfs_t* myfs_inst2 = myfs_mount(partition2->offset, partition2->size);
ESP_ERROR_CHECK(esp_vfs_register("/data2", &myfs, myfs_inst2));

```

Synchronous Input/Output Multiplexing Synchronous input/output multiplexing by `select()` is supported in the VFS component. The implementation works in the following way.

1. `select()` is called with file descriptors which could belong to various VFS drivers.
2. The file descriptors are divided into groups each belonging to one VFS driver.
3. The file descriptors belonging to non-socket VFS drivers are handed over to the given VFS drivers by `start_select()`, described later on this page. This function represents the driver-specific implementation of `select()` for the given driver. This should be a non-blocking call which means the function should immediately return after setting up the environment for checking events related to the given file descriptors.
4. The file descriptors belonging to the socket VFS driver are handed over to the socket driver by `socket_select()` described later on this page. This is a blocking call which means that it will return only if there is an event related to socket file descriptors or a non-socket driver signals `socket_select()` to exit.
5. Results are collected from each VFS driver and all drivers are stopped by de-initialization of the environment for checking events.
6. The `select()` call ends and returns the appropriate results.

Non-Socket VFS Drivers If you want to use `select()` with a file descriptor belonging to a non-socket VFS driver, then you need to register the driver with functions `start_select()` and `end_select()` similarly to the following example:

```
// In definition of esp_vfs_t:
    .start_select = &uart_start_select,
    .end_select = &uart_end_select,
// ... other members initialized
```

`start_select()` is called for setting up the environment for detection of read/write/error conditions on file descriptors belonging to the given VFS driver.

`end_select()` is called to stop/deinitialize/free the environment which was setup by `start_select()`.

Note: `end_select()` might be called without a previous `start_select()` call in some rare circumstances. `end_select()` should fail gracefully if this is the case (i.e., should not crash but return an error instead).

Please refer to the reference implementation for the UART peripheral in [vfs/vfs_uart.c](#) and most particularly to the functions `esp_vfs_dev_uart_register()`, `uart_start_select()`, and `uart_end_select()` for more information.

Please check the following examples that demonstrate the use of `select()` with VFS file descriptors:

- [peripherals/uart/uart_select](#)
- [system/select](#)

Socket VFS Drivers A socket VFS driver is using its own internal implementation of `select()` and non-socket VFS drivers notify it upon read/write/error conditions.

A socket VFS driver needs to be registered with the following functions defined:

```
// In definition of esp_vfs_t:
    .socket_select = &lwip_select,
    .get_socket_select_semaphore = &lwip_get_socket_select_semaphore,
    .stop_socket_select = &lwip_stop_socket_select,
    .stop_socket_select_isr = &lwip_stop_socket_select_isr,
// ... other members initialized
```

`socket_select()` is the internal implementation of `select()` for the socket driver. It works only with file descriptors belonging to the socket VFS.

`get_socket_select_semaphore()` returns the signalization object (semaphore) which is used in non-socket drivers to stop the waiting in `socket_select()`.

`stop_socket_select()` call is used to stop the waiting in `socket_select()` by passing the object returned by `get_socket_select_semaphore()`.

`stop_socket_select_isr()` has the same functionality as `stop_socket_select()` but it can be used from ISR.

Please see [lwip/port/esp32xx/vfs_lwip.c](#) for a reference socket driver implementation using LWIP.

Note: If you use `select()` for socket file descriptors only then you can disable the `CONFIG_VFS_SUPPORT_SELECT` option to reduce the code size and improve performance. You should not change the socket driver during an active `select()` call or you might experience some undefined behavior.

Paths

Each registered FS has a path prefix associated with it. This prefix can be considered as a "mount point" of this partition.

In case when mount points are nested, the mount point with the longest matching path prefix is used when opening the file. For instance, suppose that the following filesystems are registered in VFS:

- FS 1 on /data
- FS 2 on /data/static

Then:

- FS 1 will be used when opening a file called /data/log.txt
- FS 2 will be used when opening a file called /data/static/index.html
- Even if /index.html " does not exist in FS 2, FS 1 will **not** be searched for /static/index.html.

As a general rule, mount point names must start with the path separator (/) and must contain at least one character after path separator. However, an empty mount point name is also supported and might be used in cases when an application needs to provide a "fallback" filesystem or to override VFS functionality altogether. Such filesystem will be used if no prefix matches the path given.

VFS does not handle dots (.) in path names in any special way. VFS does not treat .. as a reference to the parent directory. In the above example, using a path /data/static/../../log.txt will not result in a call to FS 1 to open /log.txt. Specific FS drivers (such as FATFS) might handle dots in file names differently.

When opening files, the FS driver receives only relative paths to files. For example:

1. The myfs driver is registered with /data as a path prefix.
2. The application calls `fopen("/data/config.json", ...)`.
3. The VFS component calls `myfs_open("/config.json", ...)`.
4. The myfs driver opens the /config.json file.

VFS does not impose any limit on total file path length, but it does limit the FS path prefix to `ESP_VFS_PATH_MAX` characters. Individual FS drivers may have their own filename length limitations.

File Descriptors

File descriptors are small positive integers from 0 to `FD_SETSIZE - 1`, where `FD_SETSIZE` is defined in `sys/select.h`. The largest file descriptors (configured by `CONFIG_LWIP_MAX_SOCKETS`) are reserved for sockets. The VFS component contains a lookup-table called `s_fd_table` for mapping global file descriptors to VFS driver indexes registered in the `s_vfs` array.

Standard IO Streams (`stdin`, `stdout`, `stderr`)

If the menuconfig option `UART` for console output is not set to `None`, then `stdin`, `stdout`, and `stderr` are configured to read from, and write to, a UART. It is possible to use `UART0` or `UART1` for standard IO. By default, `UART0` is used with 115200 baud rate; TX pin is `GPIO1`; RX pin is `GPIO3`. These parameters can be changed in menuconfig.

Writing to `stdout` or `stderr` sends characters to the UART transmit FIFO. Reading from `stdin` retrieves characters from the UART receive FIFO.

By default, VFS uses simple functions for reading from and writing to UART. Writes busy-wait until all data is put into UART FIFO, and reads are non-blocking, returning only the data present in the FIFO. Due to this non-blocking read behavior, higher level C library calls, such as `fscanf("%d\n", &var);`, might not have desired results.

Applications which use the UART driver can instruct VFS to use the driver's interrupt driven, blocking read and write functions instead. This can be done using a call to the `esp_vfs_dev_uart_use_driver` function. It is also possible to revert to the basic non-blocking functions using a call to `esp_vfs_dev_uart_use_nonblocking`.

VFS also provides an optional newline conversion feature for input and output. Internally, most applications send and receive lines terminated by the LF ("n") character. Different terminal programs may require different line termination, such as CR or CRLF. Applications can configure this separately for input and output either via menuconfig, or by calls to the functions `esp_vfs_dev_uart_port_set_rx_line_endings` and `esp_vfs_dev_uart_port_set_tx_line_endings`.

Standard Streams and FreeRTOS Tasks FILE objects for `stdin`, `stdout`, and `stderr` are shared between all FreeRTOS tasks, but the pointers to these objects are stored in per-task `struct _reent`.

The following code is transferred to `fprintf(__getreent()->_stderr, "42\n");` by the preprocessor:

```
fprintf(stderr, "42\n");
```

The `__getreent()` function returns a per-task pointer to `struct _reent` in `newlib` libc. This structure is allocated on the TCB of each task. When a task is initialized, `_stdin`, `_stdout`, and `_stderr` members of `struct _reent` are set to the values of `_stdin`, `_stdout`, and `_stderr` of `_GLOBAL_REENT` (i.e., the structure which is used before FreeRTOS is started).

Such a design has the following consequences:

- It is possible to set `stdin`, `stdout`, and `stderr` for any given task without affecting other tasks, e.g., by doing `stdin = fopen("/dev/uart/1", "r")`.
- Closing default `stdin`, `stdout`, or `stderr` using `fclose` closes the FILE stream object, which will affect all other tasks.
- To change the default `stdin`, `stdout`, `stderr` streams for new tasks, modify `_GLOBAL_REENT->_stdin(_stdout, _stderr)` before creating the task.

eventfd()

`eventfd()` call is a powerful tool to notify a `select()` based loop of custom events. The `eventfd()` implementation in ESP-IDF is generally the same as described in [man\(2\) eventfd](#) except for:

- `esp_vfs_eventfd_register()` has to be called before calling `eventfd()`
- Options `EFD_CLOEXEC`, `EFD_NONBLOCK` and `EFD_SEMAPHORE` are not supported in flags.
- Option `EFD_SUPPORT_ISR` has been added in flags. This flag is required to read and write the `eventfd` in an interrupt handler.

Note that creating an `eventfd` with `EFD_SUPPORT_ISR` will cause interrupts to be temporarily disabled when reading, writing the file and during the beginning and the ending of the `select()` when this file is set.

API Reference

Header File

- [components/vfs/include/esp_vfs.h](#)
- This header file can be included with:

```
#include "esp_vfs.h"
```

- This header file is a part of the API provided by the `vfs` component. To declare that your component depends on `vfs`, add the following to your `CMakeLists.txt`:

```
REQUIRES vfs
```

or

```
PRIV_REQUIRES vfs
```

Functions

`ssize_t esp_vfs_write` (`struct _reent *r`, `int fd`, `const void *data`, `size_t size`)

These functions are to be used in `newlib` syscall table. They will be called by `newlib` when it needs to use any of the syscalls.

`off_t esp_vfs_lseek` (`struct _reent *r`, `int fd`, `off_t size`, `int mode`)

`ssize_t esp_vfs_read` (`struct _reent *r`, `int fd`, `void *dst`, `size_t size`)

int **esp_vfs_open** (struct _reent *r, const char *path, int flags, int mode)

int **esp_vfs_close** (struct _reent *r, int fd)

int **esp_vfs_fstat** (struct _reent *r, int fd, struct stat *st)

int **esp_vfs_stat** (struct _reent *r, const char *path, struct stat *st)

int **esp_vfs_link** (struct _reent *r, const char *n1, const char *n2)

int **esp_vfs_unlink** (struct _reent *r, const char *path)

int **esp_vfs_rename** (struct _reent *r, const char *src, const char *dst)

int **esp_vfs_utime** (const char *path, const struct utimbuf *times)

esp_err_t **esp_vfs_register** (const char *base_path, const *esp_vfs_t* *vfs, void *ctx)

Register a virtual filesystem for given path prefix.

Parameters

- **base_path** -- file path prefix associated with the filesystem. Must be a zero-terminated C string, may be empty. If not empty, must be up to ESP_VFS_PATH_MAX characters long, and at least 2 characters long. Name must start with a "/" and must not end with "/". For example, "/data" or "/dev/spi" are valid. These VFSes would then be called to handle file paths such as "/data/myfile.txt" or "/dev/spi/0". In the special case of an empty base_path, a "fallback" VFS is registered. Such VFS will handle paths which are not matched by any other registered VFS.
- **vfs** -- Pointer to *esp_vfs_t*, a structure which maps syscalls to the filesystem driver functions. VFS component doesn't assume ownership of this pointer.
- **ctx** -- If vfs->flags has ESP_VFS_FLAG_CONTEXT_PTR set, a pointer which should be passed to VFS functions. Otherwise, NULL.

Returns ESP_OK if successful, ESP_ERR_NO_MEM if too many VFSes are registered.

esp_err_t **esp_vfs_register_fd_range** (const *esp_vfs_t* *vfs, void *ctx, int min_fd, int max_fd)

Special case function for registering a VFS that uses a method other than open() to open new file descriptors from the interval <min_fd; max_fd).

This is a special-purpose function intended for registering LWIP sockets to VFS.

Parameters

- **vfs** -- Pointer to *esp_vfs_t*. Meaning is the same as for esp_vfs_register().
- **ctx** -- Pointer to context structure. Meaning is the same as for esp_vfs_register().
- **min_fd** -- The smallest file descriptor this VFS will use.
- **max_fd** -- Upper boundary for file descriptors this VFS will use (the biggest file descriptor plus one).

Returns ESP_OK if successful, ESP_ERR_NO_MEM if too many VFSes are registered, ESP_ERR_INVALID_ARG if the file descriptor boundaries are incorrect.

esp_err_t **esp_vfs_register_with_id** (const *esp_vfs_t* *vfs, void *ctx, *esp_vfs_id_t* *vfs_id)

Special case function for registering a VFS that uses a method other than open() to open new file descriptors. In comparison with esp_vfs_register_fd_range, this function doesn't pre-registers an interval of file descriptors. File descriptors can be registered later, by using esp_vfs_register_fd.

Parameters

- **vfs** -- Pointer to *esp_vfs_t*. Meaning is the same as for esp_vfs_register().
- **ctx** -- Pointer to context structure. Meaning is the same as for esp_vfs_register().
- **vfs_id** -- Here will be written the VFS ID which can be passed to esp_vfs_register_fd for registering file descriptors.

Returns ESP_OK if successful, ESP_ERR_NO_MEM if too many VFSes are registered, ESP_ERR_INVALID_ARG if the file descriptor boundaries are incorrect.

esp_err_t **esp_vfs_unregister** (const char *base_path)

Unregister a virtual filesystem for given path prefix

Parameters **base_path** -- file prefix previously used in `esp_vfs_register` call

Returns `ESP_OK` if successful, `ESP_ERR_INVALID_STATE` if VFS for given prefix hasn't been registered

esp_err_t **esp_vfs_unregister_with_id** (*esp_vfs_id_t* vfs_id)

Unregister a virtual filesystem with the given index

Parameters **vfs_id** -- The VFS ID returned by `esp_vfs_register_with_id`

Returns `ESP_OK` if successful, `ESP_ERR_INVALID_STATE` if VFS for the given index hasn't been registered

esp_err_t **esp_vfs_register_fd** (*esp_vfs_id_t* vfs_id, int *fd)

Special function for registering another file descriptor for a VFS registered by `esp_vfs_register_with_id`. This function should only be used to register permanent file descriptors (socket fd) that are not removed after being closed.

Parameters

- **vfs_id** -- VFS identifier returned by `esp_vfs_register_with_id`.
- **fd** -- The registered file descriptor will be written to this address.

Returns `ESP_OK` if the registration is successful, `ESP_ERR_NO_MEM` if too many file descriptors are registered, `ESP_ERR_INVALID_ARG` if the arguments are incorrect.

esp_err_t **esp_vfs_register_fd_with_local_fd** (*esp_vfs_id_t* vfs_id, int local_fd, bool permanent, int *fd)

Special function for registering another file descriptor with given `local_fd` for a VFS registered by `esp_vfs_register_with_id`.

Parameters

- **vfs_id** -- VFS identifier returned by `esp_vfs_register_with_id`.
- **local_fd** -- The fd in the local vfs. Passing -1 will set the local fd as the (*fd) value.
- **permanent** -- Whether the fd should be treated as permanent (not removed after close())
- **fd** -- The registered file descriptor will be written to this address.

Returns `ESP_OK` if the registration is successful, `ESP_ERR_NO_MEM` if too many file descriptors are registered, `ESP_ERR_INVALID_ARG` if the arguments are incorrect.

esp_err_t **esp_vfs_unregister_fd** (*esp_vfs_id_t* vfs_id, int fd)

Special function for unregistering a file descriptor belonging to a VFS registered by `esp_vfs_register_with_id`.

Parameters

- **vfs_id** -- VFS identifier returned by `esp_vfs_register_with_id`.
- **fd** -- File descriptor which should be unregistered.

Returns `ESP_OK` if the registration is successful, `ESP_ERR_INVALID_ARG` if the arguments are incorrect.

int **esp_vfs_select** (int nfds, fd_set *readfds, fd_set *writefds, fd_set *errorfds, struct timeval *timeout)

Synchronous I/O multiplexing which implements the functionality of POSIX `select()` for VFS.

Parameters

- **nfds** -- Specifies the range of descriptors which should be checked. The first `nfds` descriptors will be checked in each set.
- **readfds** -- If not NULL, then points to a descriptor set that on input specifies which descriptors should be checked for being ready to read, and on output indicates which descriptors are ready to read.
- **writefds** -- If not NULL, then points to a descriptor set that on input specifies which descriptors should be checked for being ready to write, and on output indicates which descriptors are ready to write.
- **errorfds** -- If not NULL, then points to a descriptor set that on input specifies which descriptors should be checked for error conditions, and on output indicates which descriptors have error conditions.

- **timeout** -- If not NULL, then points to timeval structure which specifies the time period after which the functions should time-out and return. If it is NULL, then the function will not time-out. Note that the timeout period is rounded up to the system tick and incremented by one.

Returns The number of descriptors set in the descriptor sets, or -1 when an error (specified by `errno`) have occurred.

void **esp_vfs_select_triggered** (*esp_vfs_select_sem_t* sem)

Notification from a VFS driver about a read/write/error condition.

This function is called when the VFS driver detects a read/write/error condition as it was requested by the previous call to `start_select`.

Parameters **sem** -- semaphore structure which was passed to the driver by the `start_select` call

void **esp_vfs_select_triggered_isr** (*esp_vfs_select_sem_t* sem, BaseType_t *woken)

Notification from a VFS driver about a read/write/error condition (ISR version)

This function is called when the VFS driver detects a read/write/error condition as it was requested by the previous call to `start_select`.

Parameters

- **sem** -- semaphore structure which was passed to the driver by the `start_select` call
- **woken** -- is set to `pdTRUE` if the function wakes up a task with higher priority

ssize_t **esp_vfs_pread** (int fd, void *dst, size_t size, off_t offset)

Implements the VFS layer of POSIX `pread()`

Parameters

- **fd** -- File descriptor used for read
- **dst** -- Pointer to the buffer where the output will be written
- **size** -- Number of bytes to be read
- **offset** -- Starting offset of the read

Returns A positive return value indicates the number of bytes read. -1 is return on failure and `errno` is set accordingly.

ssize_t **esp_vfs_pwrite** (int fd, const void *src, size_t size, off_t offset)

Implements the VFS layer of POSIX `pwrite()`

Parameters

- **fd** -- File descriptor used for write
- **src** -- Pointer to the buffer from where the output will be read
- **size** -- Number of bytes to write
- **offset** -- Starting offset of the write

Returns A positive return value indicates the number of bytes written. -1 is return on failure and `errno` is set accordingly.

Structures

struct **esp_vfs_select_sem_t**

VFS semaphore type for `select()`

Public Members

bool **is_sem_local**

type of "sem" is `SemaphoreHandle_t` when true, defined by socket driver otherwise

void ***sem**

semaphore instance

struct **esp_vfs_t**

VFS definition structure.

This structure should be filled with pointers to corresponding FS driver functions.

VFS component will translate all FDs so that the filesystem implementation sees them starting at zero. The caller sees a global FD which is prefixed with an pre-filesystem-implementation.

Some FS implementations expect some state (e.g. pointer to some structure) to be passed in as a first argument. For these implementations, populate the members of this structure which have `_p` suffix, set flags member to `ESP_VFS_FLAG_CONTEXT_PTR` and provide the context pointer to `esp_vfs_register` function. If the implementation doesn't use this extra argument, populate the members without `_p` suffix and set flags member to `ESP_VFS_FLAG_DEFAULT`.

If the FS driver doesn't provide some of the functions, set corresponding members to `NULL`.

Public Members

int **flags**

`ESP_VFS_FLAG_CONTEXT_PTR` and/or `ESP_VFS_FLAG_READONLY_FS` or
`ESP_VFS_FLAG_DEFAULT`

`ssize_t (*write_p)(void *p, int fd, const void *data, size_t size)`

Write with context pointer

`ssize_t (*write)(int fd, const void *data, size_t size)`

Write without context pointer

`off_t (*lseek_p)(void *p, int fd, off_t size, int mode)`

Seek with context pointer

`off_t (*lseek)(int fd, off_t size, int mode)`

Seek without context pointer

`ssize_t (*read_p)(void *ctx, int fd, void *dst, size_t size)`

Read with context pointer

`ssize_t (*read)(int fd, void *dst, size_t size)`

Read without context pointer

`ssize_t (*pread_p)(void *ctx, int fd, void *dst, size_t size, off_t offset)`

pread with context pointer

`ssize_t (*pread)(int fd, void *dst, size_t size, off_t offset)`

pread without context pointer

`ssize_t (*pwrite_p)(void *ctx, int fd, const void *src, size_t size, off_t offset)`

pwrite with context pointer

`ssize_t (*pwrite)(int fd, const void *src, size_t size, off_t offset)`

pwrite without context pointer

int (***open_p**)(void *ctx, const char *path, int flags, int mode)

open with context pointer

int (***open**)(const char *path, int flags, int mode)

open without context pointer

int (***close_p**)(void *ctx, int fd)

close with context pointer

int (***close**)(int fd)

close without context pointer

int (***fstat_p**)(void *ctx, int fd, struct *stat* *st)

fstat with context pointer

int (***fstat**)(int fd, struct *stat* *st)

fstat without context pointer

int (***stat_p**)(void *ctx, const char *path, struct *stat* *st)

stat with context pointer

int (***stat**)(const char *path, struct *stat* *st)

stat without context pointer

int (***link_p**)(void *ctx, const char *n1, const char *n2)

link with context pointer

int (***link**)(const char *n1, const char *n2)

link without context pointer

int (***unlink_p**)(void *ctx, const char *path)

unlink with context pointer

int (***unlink**)(const char *path)

unlink without context pointer

int (***rename_p**)(void *ctx, const char *src, const char *dst)

rename with context pointer

int (***rename**)(const char *src, const char *dst)

rename without context pointer

DIR *(***opendir_p**)(void *ctx, const char *name)

opendir with context pointer

DIR *(***opendir**)(const char *name)

opendir without context pointer

struct dirent *(***readdir_p**)(void *ctx, DIR *pdir)

readdir with context pointer

struct dirent *(***readdir**)(DIR *pdir)

readdir without context pointer

int (***readdir_r_p**)(void *ctx, DIR *pdir, struct dirent *entry, struct dirent **out_dirent)

readdir_r with context pointer

int (***readdir_r**)(DIR *pdir, struct dirent *entry, struct dirent **out_dirent)

readdir_r without context pointer

long (***telldir_p**)(void *ctx, DIR *pdir)

telldir with context pointer

long (***telldir**)(DIR *pdir)

telldir without context pointer

void (***seekdir_p**)(void *ctx, DIR *pdir, long offset)

seekdir with context pointer

void (***seekdir**)(DIR *pdir, long offset)

seekdir without context pointer

int (***closedir_p**)(void *ctx, DIR *pdir)

closedir with context pointer

int (***closedir**)(DIR *pdir)

closedir without context pointer

int (***mkdir_p**)(void *ctx, const char *name, mode_t mode)

mkdir with context pointer

int (***mkdir**)(const char *name, mode_t mode)

mkdir without context pointer

int (***rmdir_p**)(void *ctx, const char *name)

rmdir with context pointer

int (***rmdir**)(const char *name)

rmdir without context pointer

int (***fcntl_p**)(void *ctx, int fd, int cmd, int arg)

fcntl with context pointer

int (***fcntl**)(int fd, int cmd, int arg)

fcntl without context pointer

int (***ioctl_p**)(void *ctx, int fd, int cmd, va_list args)

ioctl with context pointer

int (***ioctl**)(int fd, int cmd, va_list args)

ioctl without context pointer

int (***fsync_p**)(void *ctx, int fd)

fsync with context pointer

int (***fsync**)(int fd)

fsync without context pointer

int (***access_p**)(void *ctx, const char *path, int amode)

access with context pointer

int (***access**)(const char *path, int amode)

access without context pointer

int (***truncate_p**)(void *ctx, const char *path, off_t length)

truncate with context pointer

int (***truncate**)(const char *path, off_t length)

truncate without context pointer

int (***ftruncate_p**)(void *ctx, int fd, off_t length)

ftruncate with context pointer

int (***ftruncate**)(int fd, off_t length)

ftruncate without context pointer

int (***utime_p**)(void *ctx, const char *path, const struct utimbuf *times)

utime with context pointer

int (***utime**)(const char *path, const struct utimbuf *times)

utime without context pointer

int (***tcsetattr_p**)(void *ctx, int fd, int optional_actions, const struct termios *p)

tcsetattr with context pointer

int (***tcsetattr**)(int fd, int optional_actions, const struct termios *p)

tcsetattr without context pointer

int (***tcgetattr_p**)(void *ctx, int fd, struct termios *p)

tcgetattr with context pointer

int (***tcgetattr**)(int fd, struct termios *p)

tcgetattr without context pointer

int (***tcdrain_p**)(void *ctx, int fd)

tcdrain with context pointer

int (***tcdrain**)(int fd)

tcdrain without context pointer

int (***tcflush_p**)(void *ctx, int fd, int select)

tcflush with context pointer

int (***tcflush**)(int fd, int select)

tcflush without context pointer

int (***tcflow_p**)(void *ctx, int fd, int action)

tcflow with context pointer

int (***tcflow**)(int fd, int action)

tcflow without context pointer

pid_t (***tcgetsid_p**)(void *ctx, int fd)

tcgetsid with context pointer

pid_t (***tcgetsid**)(int fd)

tcgetsid without context pointer

int (***tcsendbreak_p**)(void *ctx, int fd, int duration)

tcsendbreak with context pointer

int (***tcsendbreak**)(int fd, int duration)

tcsendbreak without context pointer

esp_err_t (***start_select**)(int nfds, fd_set *readfds, fd_set *writefds, fd_set *exceptfds,

esp_vfs_select_sem_t sem, void **end_select_args)

start_select is called for setting up synchronous I/O multiplexing of the desired file descriptors in the given VFS

int (***socket_select**)(int nfds, fd_set *readfds, fd_set *writefds, fd_set *errorfds, struct timeval *timeout)

socket select function for socket FDs with the functionality of POSIX select(); this should be set only for the socket VFS

void (***stop_socket_select**)(void *sem)

called by VFS to interrupt the socket_select call when select is activated from a non-socket VFS driver; set only for the socket driver

void (***stop_socket_select_isr**)(void *sem, BaseType_t *woken)

stop_socket_select which can be called from ISR; set only for the socket driver

void (***get_socket_select_semaphore**)(void)

end_select is called to stop the I/O multiplexing and deinitialize the environment created by start_select for the given VFS

esp_err_t (***end_select**)(void *end_select_args)

get_socket_select_semaphore returns semaphore allocated in the socket driver; set only for the socket driver

Macros

MAX_FDS

Maximum number of (global) file descriptors.

ESP_VFS_PATH_MAX

Maximum length of path prefix (not including zero terminator)

ESP_VFS_FLAG_DEFAULT

Default value of flags member in *esp_vfs_t* structure.

ESP_VFS_FLAG_CONTEXT_PTR

Flag which indicates that FS needs extra context pointer in syscalls.

ESP_VFS_FLAG_READONLY_FS

Flag which indicates that FS is located on read-only partition.

Type Definitions

typedef int **esp_vfs_id_t**

Header File

- [components/vfs/include/esp_vfs_dev.h](#)
- This header file can be included with:

```
#include "esp_vfs_dev.h"
```

- This header file is a part of the API provided by the `vfs` component. To declare that your component depends on `vfs`, add the following to your CMakeLists.txt:

```
REQUIRES vfs
```

or

```
PRIV_REQUIRES vfs
```

Functions

void **esp_vfs_dev_uart_register** (void)

add /dev/uart virtual filesystem driver

This function is called from startup code to enable serial output

void **esp_vfs_dev_uart_set_rx_line_endings** (esp_line_endings_t mode)

Set the line endings expected to be received on UART.

This specifies the conversion between line endings received on UART and newlines ('', LF) passed into stdin:

- `ESP_LINE_ENDINGS_CRLF`: convert CRLF to LF
- `ESP_LINE_ENDINGS_CR`: convert CR to LF
- `ESP_LINE_ENDINGS_LF`: no modification

Note: this function is not thread safe w.r.t. reading from UART

Parameters `mode` -- line endings expected on UART

void **`esp_vfs_dev_uart_set_tx_line_endings`** (`esp_line_endings_t mode`)

Set the line endings to sent to UART.

This specifies the conversion between newlines ('
, LF) on stdout and line endings sent over UART:

- `ESP_LINE_ENDINGS_CRLF`: convert LF to CRLF
- `ESP_LINE_ENDINGS_CR`: convert LF to CR
- `ESP_LINE_ENDINGS_LF`: no modification

Note: this function is not thread safe w.r.t. writing to UART

Parameters `mode` -- line endings to send to UART

int **`esp_vfs_dev_uart_port_set_rx_line_endings`** (`int uart_num, esp_line_endings_t mode`)

Set the line endings expected to be received on specified UART.

This specifies the conversion between line endings received on UART and newlines ('
, LF) passed into stdin:

- `ESP_LINE_ENDINGS_CRLF`: convert CRLF to LF
- `ESP_LINE_ENDINGS_CR`: convert CR to LF
- `ESP_LINE_ENDINGS_LF`: no modification

Note: this function is not thread safe w.r.t. reading from UART

Parameters

- `uart_num` -- the UART number
- `mode` -- line endings to send to UART

Returns 0 if succeeded, or -1 when an error (specified by `errno`) have occurred.

int **`esp_vfs_dev_uart_port_set_tx_line_endings`** (`int uart_num, esp_line_endings_t mode`)

Set the line endings to sent to specified UART.

This specifies the conversion between newlines ('
, LF) on stdout and line endings sent over UART:

- `ESP_LINE_ENDINGS_CRLF`: convert LF to CRLF
- `ESP_LINE_ENDINGS_CR`: convert LF to CR
- `ESP_LINE_ENDINGS_LF`: no modification

Note: this function is not thread safe w.r.t. writing to UART

Parameters

- **uart_num** -- the UART number
- **mode** -- line endings to send to UART

Returns 0 if succeeded, or -1 when an error (specified by `errno`) have occurred.

void **esp_vfs_dev_uart_use_nonblocking** (int uart_num)

set VFS to use simple functions for reading and writing UART Read is non-blocking, write is busy waiting until TX FIFO has enough space. These functions are used by default.

Parameters **uart_num** -- UART peripheral number

void **esp_vfs_dev_uart_use_driver** (int uart_num)

set VFS to use UART driver for reading and writing

Note: application must configure UART driver before calling these functions With these functions, read and write are blocking and interrupt-driven.

Parameters **uart_num** -- UART peripheral number

void **esp_vfs_usb_serial_jtag_use_driver** (void)

set VFS to use USB-SERIAL-JTAG driver for reading and writing

Note: application must configure USB-SERIAL-JTAG driver before calling these functions With these functions, read and write are blocking and interrupt-driven.

void **esp_vfs_usb_serial_jtag_use_nonblocking** (void)

set VFS to use simple functions for reading and writing UART Read is non-blocking, write is busy waiting until TX FIFO has enough space. These functions are used by default.

Header File

- [components/vfs/include/esp_vfs_eventfd.h](#)
- This header file can be included with:

```
#include "esp_vfs_eventfd.h"
```

- This header file is a part of the API provided by the `vfs` component. To declare that your component depends on `vfs`, add the following to your `CMakeLists.txt`:

```
REQUIRES vfs
```

or

```
PRIV_REQUIRES vfs
```

Functions

esp_err_t **esp_vfs_eventfd_register** (const *esp_vfs_eventfd_config_t* *config)

Registers the event vfs.

Returns ESP_OK if successful, ESP_ERR_NO_MEM if too many VFSes are registered.

esp_err_t **esp_vfs_eventfd_unregister** (void)

Unregisters the event vfs.

Returns ESP_OK if successful, ESP_ERR_INVALID_STATE if VFS for given prefix hasn't been registered

int **eventfd** (unsigned int initval, int flags)

Structures

struct **esp_vfs_eventfd_config_t**

Eventfd vfs initialization settings.

Public Members

size_t **max_fds**

The maximum number of eventfds supported

Macros

EFD_SUPPORT_ISR

ESP_VFS_EVENTFD_CONFIG_DEFAULT ()

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.8.12 Wear Levelling API

Overview

Most of flash memory and especially SPI flash that is used in ESP32-P4 has a sector-based organization and also has a limited number of erase/modification cycles per memory sector. The wear levelling component helps to distribute wear and tear among sectors more evenly without requiring any attention from the user.

The wear levelling component provides API functions related to reading, writing, erasing, and memory mapping of data in external SPI flash through the partition component. The component also has higher-level API functions which work with the FAT filesystem defined in *FAT filesystem*.

The wear levelling component, together with the FAT FS component, uses FAT FS sectors of 4096 bytes, which is a standard size for flash memory. With this size, the component shows the best performance but needs additional memory in RAM.

To save internal memory, the component has two additional modes which both use sectors of 512 bytes:

- **Performance mode.** Erase sector operation data is stored in RAM, the sector is erased, and then data is copied back to flash memory. However, if a device is powered off for any reason, all 4096 bytes of data is lost.
- **Safety mode.** The data is first saved to flash memory, and after the sector is erased, the data is saved back. If a device is powered off, the data can be recovered as soon as the device boots up.

The default settings are as follows:

- Sector size is 512 bytes
- Performance mode

You can change the settings through the configuration menu.

The wear levelling component does not cache data in RAM. The write and erase functions modify flash directly, and flash contents are consistent when the function returns.

Wear Levelling access API functions

This is the set of API functions for working with data in flash:

- `wl_mount` - initializes the wear levelling module and mounts the specified partition
- `wl_unmount` - unmounts the partition and deinitializes the wear levelling module
- `wl_erase_range` - erases a range of addresses in flash
- `wl_write` - writes data to a partition
- `wl_read` - reads data from a partition
- `wl_size` - returns the size of available memory in bytes
- `wl_sector_size` - returns the size of one sector

As a rule, try to avoid using raw wear levelling functions and use filesystem-specific functions instead.

Memory Size

The memory size is calculated in the wear levelling module based on partition parameters. The module uses some sectors of flash for internal data.

See Also

- [FAT Filesystem Support](#)
- [Partition Tables](#)

Application Example

An example that combines the wear levelling driver with the FATFS library is provided in the [storage/wear_levelling](#) directory. This example initializes the wear levelling driver, mounts FatFs partition, as well as writes and reads data from it using POSIX and C library APIs. See [storage/wear_levelling/README.md](#) for more information.

High-level API Reference

Header Files

- [fatfs/vfs/esp_vfs_fat.h](#)

High-level wear levelling functions `esp_vfs_fat_spiflash_mount_rw_wl()`, `esp_vfs_fat_spiflash_unmount_rw_wl()` and struct `esp_vfs_fat_mount_config_t` are described in [FAT Filesystem Support](#).

Mid-level API Reference

Header File

- [components/wear_levelling/include/wear_levelling.h](#)
- This header file can be included with:

```
#include "wear_levelling.h"
```

- This header file is a part of the API provided by the `wear_levelling` component. To declare that your component depends on `wear_levelling`, add the following to your `CMakeLists.txt`:

```
REQUIRES wear_levelling
```

or

```
PRIV_REQUIRES wear_levelling
```

Functions

esp_err_t **wl_mount** (const *esp_partition_t* *partition, *wl_handle_t* *out_handle)

Mount WL for defined partition.

Parameters

- **partition** -- that will be used for access
- **out_handle** -- handle of the WL instance

Returns

- ESP_OK, if the WL allocation is successful;
- ESP_ERR_INVALID_ARG, if the arguments for WL configuration are not valid;
- ESP_ERR_NO_MEM, if the WL allocation fails because of insufficient memory;

esp_err_t **wl_unmount** (*wl_handle_t* handle)

Unmount WL for defined partition.

Parameters **handle** -- WL partition handle

Returns

- ESP_OK, if the operation is successful;
- or one of error codes from lower-level flash driver.

esp_err_t **wl_erase_range** (*wl_handle_t* handle, size_t start_addr, size_t size)

Erase part of the WL storage.

Parameters

- **handle** -- WL handle that are related to the partition
- **start_addr** -- Address from where erase operation should start. Must be aligned to the result of function `wl_sector_size(...)`.
- **size** -- Size of the range which should be erased, in bytes. Must be divisible by the result of function `wl_sector_size(...)`.

Returns

- ESP_OK, if the given range was erased successfully;
- ESP_ERR_INVALID_ARG, if iterator or dst are NULL;
- ESP_ERR_INVALID_SIZE, if erase would go out of bounds of the partition;
- or one of error codes from lower-level flash driver.

esp_err_t **wl_write** (*wl_handle_t* handle, size_t dest_addr, const void *src, size_t size)

Write data to the WL storage.

Before writing data to flash, corresponding region of flash needs to be erased. This can be done using `wl_erase_range` function.

Note: Prior to writing to WL storage, make sure it has been erased with `wl_erase_range` call.

Parameters

- **handle** -- WL handle corresponding to the WL partition
- **dest_addr** -- Address where the data should be written, relative to the beginning of the partition.
- **src** -- Pointer to the source buffer. Pointer must be non-NULL and buffer must be at least 'size' bytes long.

- **size** -- Size of data to be written, in bytes.

Returns

- ESP_OK, if data was written successfully;
- ESP_ERR_INVALID_ARG, if dst_offset exceeds partition size;
- ESP_ERR_INVALID_SIZE, if write would go out of bounds of the partition;
- or one of error codes from lower-level flash driver.

esp_err_t **wl_read** (*wl_handle_t* handle, size_t src_addr, void *dest, size_t size)

Read data from the WL storage.

Parameters

- **handle** -- WL module instance that was initialized before
- **dest** -- Pointer to the buffer where data should be stored. The Pointer must be non-NULL and the buffer must be at least 'size' bytes long.
- **src_addr** -- Address of the data to be read, relative to the beginning of the partition.
- **size** -- Size of data to be read, in bytes.

Returns

- ESP_OK, if data was read successfully;
- ESP_ERR_INVALID_ARG, if src_offset exceeds partition size;
- ESP_ERR_INVALID_SIZE, if read would go out of bounds of the partition;
- or one of error codes from lower-level flash driver.

size_t **wl_size** (*wl_handle_t* handle)

Get the actual flash size in use for the WL storage partition.

Parameters **handle** -- WL module handle that was initialized before

Returns usable size, in bytes

size_t **wl_sector_size** (*wl_handle_t* handle)

Get sector size of the WL instance.

Parameters **handle** -- WL module handle that was initialized before

Returns sector size, in bytes

Macros

WL_INVALID_HANDLE

Type Definitions

typedef int32_t **wl_handle_t**

wear levelling handle

Code examples for this API section are provided in the [storage](#) directory of ESP-IDF examples.

2.9 System API

2.9.1 App Image Format

Application Image Structures

An application image consists of the following:

1. The `esp_image_header_t` structure describes the mode of SPI flash and the count of memory segments.
2. The `esp_image_segment_header_t` structure describes each segment, its length, and its location in ESP32-P4's memory, followed by the data with a length of `data_len`. The data offset for each segment in the image is calculated in the following way:
 - offset for 0 Segment = `sizeof(esp_image_header_t) + sizeof(esp_image_segment_header_t)`
 - offset for 1 Segment = offset for 0 Segment + length of 0 Segment + `sizeof(esp_image_segment_header_t)`
 - offset for 2 Segment = offset for 1 Segment + length of 1 Segment + `sizeof(esp_image_segment_header_t)`
 - ...

The count of each segment is defined in the `segment_count` field that is stored in `esp_image_header_t`. The count cannot be more than `ESP_IMAGE_MAX_SEGMENTS`.

To get the list of your image segments, please run the following command:

```
esptool.py --chip esp32p4 image_info build/app.bin
```

```
esptool.py v2.3.1
Image version: 1
Entry point: 40080ea4
13 segments

Segment 1: len 0x13ce0 load 0x3f400020 file_offs 0x00000018 SOC_DROM
Segment 2: len 0x00000 load 0x3ff80000 file_offs 0x00013d00 SOC_RTC_DRAM
Segment 3: len 0x00000 load 0x3ff80000 file_offs 0x00013d08 SOC_RTC_DRAM
Segment 4: len 0x028e0 load 0x3ffb0000 file_offs 0x00013d10 DRAM
Segment 5: len 0x00000 load 0x3ffb28e0 file_offs 0x000165f8 DRAM
Segment 6: len 0x00400 load 0x40080000 file_offs 0x00016600 SOC_IRAM
Segment 7: len 0x09600 load 0x40080400 file_offs 0x00016a08 SOC_IRAM
Segment 8: len 0x62e4c load 0x400d0018 file_offs 0x00020010 SOC_IROM
Segment 9: len 0x06cec load 0x40089a00 file_offs 0x00082e64 SOC_IROM
Segment 10: len 0x00000 load 0x400c0000 file_offs 0x00089b58 SOC_RTC_IRAM
Segment 11: len 0x00004 load 0x50000000 file_offs 0x00089b60 SOC_RTC_DATA
Segment 12: len 0x00000 load 0x50000004 file_offs 0x00089b6c SOC_RTC_DATA
Segment 13: len 0x00000 load 0x50000004 file_offs 0x00089b74 SOC_RTC_DATA
Checksum: e8 (valid)
Validation Hash: 407089ca0eae2bbf83b4120979d3354b1c938a49cb7a0c997f240474ef2ec76b_
↳ (valid)
```

You can also see the information on segments in the ESP-IDF logs while your application is booting:

```
I (443) esp_image: segment 0: paddr=0x00020020 vaddr=0x3f400020 size=0x13ce0 ( 81120) map
I (489) esp_image: segment 1: paddr=0x00033d08 vaddr=0x3ff80000 size=0x00000 ( 0) load
I (530) esp_image: segment 2: paddr=0x00033d10 vaddr=0x3ff80000 size=0x00000 ( 0) load
I (571) esp_image: segment 3: paddr=0x00033d18 vaddr=0x3ffb0000 size=0x028e0 ( 10464) load
I (612) esp_image: segment 4: paddr=0x00033600 vaddr=0x3ffb28e0 size=0x00000 ( 0) load
I (654) esp_image: segment 5: paddr=0x00033608 vaddr=0x40080000 size=0x00400 ( 1024) load
I (695) esp_image: segment 6: paddr=0x000336a10 vaddr=0x40080400 size=0x09600 ( 38400) load
I (737) esp_image: segment 7: paddr=0x00040018 vaddr=0x400d0018 size=0x62e4c ( 405068) map
```

(continues on next page)

(continued from previous page)

```

I (847) esp_image: segment 8: paddr=0x000a2e6c vaddr=0x40089a00 size=0x06cec ( 0) load
I (888) esp_image: segment 9: paddr=0x000a9b60 vaddr=0x400c0000 size=0x00000 ( 0) load
I (929) esp_image: segment 10: paddr=0x000a9b68 vaddr=0x50000000 size=0x00004 ( 4) load
I (971) esp_image: segment 11: paddr=0x000a9b74 vaddr=0x50000004 size=0x00000 ( 0) load
I (1012) esp_image: segment 12: paddr=0x000a9b7c vaddr=0x50000004 size=0x00000 ( 0) load

```

For more details on the type of memory segments and their address ranges, see **ESP32-P4 Technical Reference Manual > System and Memory > Internal Memory** [PDF].

3. The image has a single checksum byte after the last segment. This byte is written on a sixteen byte padded boundary, so the application image might need padding.
4. If the `hash_appended` field from `esp_image_header_t` is set then a SHA256 checksum will be appended. The value of the SHA256 hash is calculated on the range from the first byte and up to this field. The length of this field is 32 bytes.
5. If the option `CONFIG_SECURE_SIGNED_APPS_SCHEME` is set to ECDSA then the application image will have an additional 68 bytes for an ECDSA signature, which includes:
 - version word (4 bytes)
 - signature data (64 bytes)
6. If the option `CONFIG_SECURE_SIGNED_APPS_SCHEME` is set to RSA or ECDSA (V2) then the application image will have an additional signature sector of 4 KB in size. For more details on the format of this signature sector, please refer to *Signature Block Format*.

Application Description

The DROM segment of the application binary starts with the `esp_app_desc_t` structure which carries specific fields describing the application:

- `magic_word`: the magic word for the `esp_app_desc_t` structure
- `secure_version`: see *Anti-rollback*
- `version`: see *App version*¹
- `project_name`: filled from `PROJECT_NAME`¹
- `time and date`: compile time and date
- `idf_ver`: version of ESP-IDF¹
- `app_elf_sha256`: contains SHA256 hash for the application ELF file

This structure is useful for identification of images uploaded via Over-the-Air (OTA) updates because it has a fixed offset = `sizeof(esp_image_header_t) + sizeof(esp_image_segment_header_t)`. As soon as a device receives the first fragment containing this structure, it has all the information to determine whether the update should be continued with or not.

To obtain the `esp_app_desc_t` structure for the currently running application, use `esp_app_get_description()`.

To obtain the `esp_app_desc_t` structure for another OTA partition, use `esp_ota_get_partition_description()`.

Adding a Custom Structure to an Application

Users also have the opportunity to have similar structure with a fixed offset relative to the beginning of the image.

The following pattern can be used to add a custom structure to your image:

¹ The maximum length is 32 characters, including null-termination character. For example, if the length of `PROJECT_NAME` exceeds 31 characters, the excess characters will be disregarded.

```
const __attribute__((section(".rodata_custom_desc"))) esp_custom_app_desc_t custom_
↪app_desc = { ... }
```

Offset for custom structure is `sizeof(esp_image_header_t) + sizeof(esp_image_segment_header_t) + sizeof(esp_app_desc_t)`.

To guarantee that the custom structure is located in the image even if it is not used, you need to add `target_link_libraries(${COMPONENT_TARGET} "-u custom_app_desc")` into `CMakeLists.txt`.

API Reference

Header File

- `components/bootloader_support/include/esp_app_format.h`
- This header file can be included with:

```
#include "esp_app_format.h"
```

- This header file is a part of the API provided by the `bootloader_support` component. To declare that your component depends on `bootloader_support`, add the following to your `CMakeLists.txt`:

```
REQUIRES bootloader_support
```

or

```
PRIV_REQUIRES bootloader_support
```

Structures

struct **esp_image_header_t**

Main header of binary image.

Public Members

uint8_t **magic**

Magic word `ESP_IMAGE_HEADER_MAGIC`

uint8_t **segment_count**

Count of memory segments

uint8_t **spi_mode**

flash read mode (`esp_image_spi_mode_t` as `uint8_t`)

uint8_t **spi_speed**

flash frequency (`esp_image_spi_freq_t` as `uint8_t`)

uint8_t **spi_size**

flash chip size (`esp_image_flash_size_t` as `uint8_t`)

uint32_t **entry_addr**

Entry address

uint8_t **wp_pin**

WP pin when SPI pins set via efuse (read by ROM bootloader, the IDF bootloader uses software to configure the WP pin and sets this field to 0xEE=disabled)

uint8_t **spi_pin_drv**[3]

Drive settings for the SPI flash pins (read by ROM bootloader)

esp_chip_id_t **chip_id**

Chip identification number

uint8_t **min_chip_rev**

Minimal chip revision supported by image After the Major and Minor revision eFuses were introduced into the chips, this field is no longer used. But for compatibility reasons, we keep this field and the data in it. Use min_chip_rev_full instead. The software interprets this as a Major version for most of the chips and as a Minor version for the ESP32-C3.

uint16_t **min_chip_rev_full**

Minimal chip revision supported by image, in format: major * 100 + minor

uint16_t **max_chip_rev_full**

Maximal chip revision supported by image, in format: major * 100 + minor

uint8_t **reserved**[4]

Reserved bytes in additional header space, currently unused

uint8_t **hash_appended**

If 1, a SHA256 digest "simple hash" (of the entire image) is appended after the checksum. Included in image length. This digest is separate to secure boot and only used for detecting corruption. For secure boot signed images, the signature is appended after this (and the simple hash is included in the signed data).

struct **esp_image_segment_header_t**

Header of binary image segment.

Public Members

uint32_t **load_addr**

Address of segment

uint32_t **data_len**

Length of data

Macros

ESP_IMAGE_HEADER_MAGIC

The magic word for the *esp_image_header_t* structure.

ESP_IMAGE_MAX_SEGMENTS

Max count of segments in the image.

Enumerations

enum **esp_chip_id_t**

ESP chip ID.

Values:

enumerator **ESP_CHIP_ID_ESP32**

chip ID: ESP32

enumerator **ESP_CHIP_ID_ESP32S2**

chip ID: ESP32-S2

enumerator **ESP_CHIP_ID_ESP32C3**

chip ID: ESP32-C3

enumerator **ESP_CHIP_ID_ESP32S3**

chip ID: ESP32-S3

enumerator **ESP_CHIP_ID_ESP32C2**

chip ID: ESP32-C2

enumerator **ESP_CHIP_ID_ESP32C6**

chip ID: ESP32-C6

enumerator **ESP_CHIP_ID_ESP32H2**

chip ID: ESP32-H2

enumerator **ESP_CHIP_ID_ESP32P4**

chip ID: ESP32-P4

enumerator **ESP_CHIP_ID_INVALID**

Invalid chip ID (we defined it to make sure the esp_chip_id_t is 2 bytes size)

enum **esp_image_spi_mode_t**

SPI flash mode, used in [esp_image_header_t](#).

Values:

enumerator **ESP_IMAGE_SPI_MODE_QIO**

SPI mode QIO

enumerator **ESP_IMAGE_SPI_MODE_QOUT**

SPI mode QOUT

enumerator **ESP_IMAGE_SPI_MODE_DIO**

SPI mode DIO

enumerator **ESP_IMAGE_SPI_MODE_DOUT**

SPI mode DOUT

enumerator **ESP_IMAGE_SPI_MODE_FAST_READ**
SPI mode FAST_READ

enumerator **ESP_IMAGE_SPI_MODE_SLOW_READ**
SPI mode SLOW_READ

enum **esp_image_spi_freq_t**
SPI flash clock division factor.

Values:

enumerator **ESP_IMAGE_SPI_SPEED_DIV_2**
The SPI flash clock frequency is divided by 2 of the clock source

enumerator **ESP_IMAGE_SPI_SPEED_DIV_3**
The SPI flash clock frequency is divided by 3 of the clock source

enumerator **ESP_IMAGE_SPI_SPEED_DIV_4**
The SPI flash clock frequency is divided by 4 of the clock source

enumerator **ESP_IMAGE_SPI_SPEED_DIV_1**
The SPI flash clock frequency equals to the clock source

enum **esp_image_flash_size_t**
Supported SPI flash sizes.

Values:

enumerator **ESP_IMAGE_FLASH_SIZE_1MB**
SPI flash size 1 MB

enumerator **ESP_IMAGE_FLASH_SIZE_2MB**
SPI flash size 2 MB

enumerator **ESP_IMAGE_FLASH_SIZE_4MB**
SPI flash size 4 MB

enumerator **ESP_IMAGE_FLASH_SIZE_8MB**
SPI flash size 8 MB

enumerator **ESP_IMAGE_FLASH_SIZE_16MB**
SPI flash size 16 MB

enumerator **ESP_IMAGE_FLASH_SIZE_32MB**
SPI flash size 32 MB

enumerator **ESP_IMAGE_FLASH_SIZE_64MB**
SPI flash size 64 MB

enumerator **ESP_IMAGE_FLASH_SIZE_128MB**

SPI flash size 128 MB

enumerator **ESP_IMAGE_FLASH_SIZE_MAX**

SPI flash size MAX

2.9.2 Bootloader Image Format

The bootloader image consists of the same structures as the application image, see [Application Image Structures](#). The only difference is in the [Bootloader Description](#) structure.

To get information about the bootloader image, please run the following command:

```
esptool.py --chip esp32p4 image_info build/bootloader/bootloader.bin --version 2
```

```
File size: 26576 (bytes)

ESP32 image header
=====
Image version: 1
Entry point: 0x40080658
Segments: 4
Flash size: 2MB
Flash freq: 40m
Flash mode: DIO

ESP32 extended image header
=====
WP pin: 0xee
Flash pins drive settings: clk_drv: 0x0, q_drv: 0x0, d_drv: 0x0, cs0_drv: 0x0, hd_
↳drv: 0x0, wp_drv: 0x0
Chip ID: 0
Minimal chip revision: v0.0, (legacy min_rev = 0)
Maximal chip revision: v3.99

Segments information
=====
Segment   Length   Load addr   File offs   Memory types
-----
1  0x01bb0  0x3fff0030  0x00000018  BYTE_ACCESSIBLE, DRAM, DIRAM_DRAM
2  0x03c90  0x40078000  0x00001bd0  CACHE_APP
3  0x00004  0x40080400  0x00005868  IRAM
4  0x00f2c  0x40080404  0x00005874  IRAM

ESP32 image footer
=====
Checksum: 0x65 (valid)
Validation hash: 6f31a7f8512f26f6bce7c3b270f93bf6cf1ee4602c322998ca8ce27433527e92_
↳ (valid)

Bootloader information
=====
Bootloader version: 1
ESP-IDF: v5.1-dev-4304-gcb51a3b-dirty
Compile time: Mar 30 2023 19:14:17
```

Bootloader Description

The DRAM0 segment of the bootloader binary starts with the `esp_bootloader_desc_t` structure which carries specific fields describing the bootloader. This structure is located at a fixed offset = `sizeof(esp_image_header_t) + sizeof(esp_image_segment_header_t)`.

- `magic_byte` - the magic byte for the `esp_bootloader_desc` structure.
- `reserved` - reserved for the future IDF use.
- `version` - bootloader version, see [CONFIG_BOOTLOADER_PROJECT_VER](#)
- `idf_ver` - ESP-IDF version. *
- `date and time` - compile date and time.
- `reserved2` - reserved for the future IDF use.

* - The maximum length is 32 characters, including null-termination character.

To get the `esp_bootloader_desc_t` structure from the running bootloader, use `esp_bootloader_get_description()`.

To get the `esp_bootloader_desc_t` structure from a running application, use `esp_ota_get_bootloader_description()`.

API Reference

Header File

- [components/esp_bootloader_format/include/esp_bootloader_desc.h](#)
- This header file can be included with:

```
#include "esp_bootloader_desc.h"
```

- This header file is a part of the API provided by the `esp_bootloader_format` component. To declare that your component depends on `esp_bootloader_format`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_bootloader_format
```

or

```
PRIV_REQUIRES esp_bootloader_format
```

Functions

const `esp_bootloader_desc_t` *`esp_bootloader_get_description` (void)

Return `esp_bootloader_desc` structure.

Intended for use by the bootloader.

Returns Pointer to `esp_bootloader_desc` structure.

Structures

struct `esp_bootloader_desc_t`

Bootloader description structure.

Public Members

uint8_t `magic_byte`

Magic byte `ESP_BOOTLOADER_DESC_MAGIC_BYTE`

uint8_t **reserved**[3]

reserved for IDF

uint32_t **version**

Bootloader version

char **idf_ver**[32]

Version IDF

char **date_time**[24]

Compile date and time

uint8_t **reserved2**[16]

reserved for IDF

Macros

ESP_BOOTLOADER_DESC_MAGIC_BYTE

The magic byte for the esp_bootloader_desc structure that is in DRAM.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.9.3 Application Level Tracing

Overview

ESP-IDF provides a useful feature for application behavior analysis called **Application Level Tracing**. The feature can be enabled in menuconfig and allows transfer of arbitrary data between the host and ESP32-P4 via JTAG interface with minimal overhead on program execution.

Developers can use this library to send application specific state of execution to the host, and receive commands or other types of information in the opposite direction at runtime. The main use cases of this library are:

1. Collecting application specific data, see [Application Specific Tracing](#).
2. Lightweight logging to the host, see [Logging to Host](#).
3. System behaviour analysis, see [System Behavior Analysis with SEGGER SystemView](#).

API Reference

Header File

- [components/app_trace/include/esp_app_trace.h](#)
- This header file can be included with:

```
#include "esp_app_trace.h"
```

- This header file is a part of the API provided by the app_trace component. To declare that your component depends on app_trace, add the following to your CMakeLists.txt:

REQUIRES app_trace

or

PRIV_REQUIRES app_trace

Functions

esp_err_t **esp_appttrace_init** (void)

Initializes application tracing module.

Note: Should be called before any esp_appttrace_xxx call.

Returns ESP_OK on success, otherwise see esp_err_t

void **esp_appttrace_down_buffer_config** (uint8_t *buf, uint32_t size)

Configures down buffer.

Note: Needs to be called before attempting to receive any data using esp_appttrace_down_buffer_get and esp_appttrace_read. This function does not protect internal data by lock.

Parameters

- **buf** -- Address of buffer to use for down channel (host to target) data.
- **size** -- Size of the buffer.

uint8_t ***esp_appttrace_buffer_get** (*esp_appttrace_dest_t* dest, uint32_t size, uint32_t tmo)

Allocates buffer for trace data. Once the data in the buffer is ready to be sent, esp_appttrace_buffer_put must be called to indicate it.

Parameters

- **dest** -- Indicates HW interface to send data.
- **size** -- Size of data to write to trace buffer.
- **tmo** -- Timeout for operation (in us). Use ESP_APPTRACE_TMO_INFINITE to wait indefinitely.

Returns non-NULL on success, otherwise NULL.

esp_err_t **esp_appttrace_buffer_put** (*esp_appttrace_dest_t* dest, uint8_t *ptr, uint32_t tmo)

Indicates that the data in the buffer is ready to be sent. This function is a counterpart of and must be preceded by esp_appttrace_buffer_get.

Parameters

- **dest** -- Indicates HW interface to send data. Should be identical to the same parameter in call to esp_appttrace_buffer_get.
- **ptr** -- Address of trace buffer to release. Should be the value returned by call to esp_appttrace_buffer_get.
- **tmo** -- Timeout for operation (in us). Use ESP_APPTRACE_TMO_INFINITE to wait indefinitely.

Returns ESP_OK on success, otherwise see esp_err_t

esp_err_t **esp_appttrace_write** (*esp_appttrace_dest_t* dest, const void *data, uint32_t size, uint32_t tmo)

Writes data to trace buffer.

Parameters

- **dest** -- Indicates HW interface to send data.
- **data** -- Address of data to write to trace buffer.
- **size** -- Size of data to write to trace buffer.

- **tmo** -- Timeout for operation (in us). Use ESP_APPTRACE_TMO_INFINITE to wait indefinitely.

Returns ESP_OK on success, otherwise see esp_err_t

int **esp_apptrace_vprintf_to** (*esp_apptrace_dest_t* dest, uint32_t tmo, const char *fmt, va_list ap)
vprintf-like function to send log messages to host via specified HW interface.

Parameters

- **dest** -- Indicates HW interface to send data.
- **tmo** -- Timeout for operation (in us). Use ESP_APPTRACE_TMO_INFINITE to wait indefinitely.
- **fmt** -- Address of format string.
- **ap** -- List of arguments.

Returns Number of bytes written.

int **esp_apptrace_vprintf** (const char *fmt, va_list ap)
vprintf-like function to send log messages to host.

Parameters

- **fmt** -- Address of format string.
- **ap** -- List of arguments.

Returns Number of bytes written.

esp_err_t **esp_apptrace_flush** (*esp_apptrace_dest_t* dest, uint32_t tmo)

Flushes remaining data in trace buffer to host.

Parameters

- **dest** -- Indicates HW interface to flush data on.
- **tmo** -- Timeout for operation (in us). Use ESP_APPTRACE_TMO_INFINITE to wait indefinitely.

Returns ESP_OK on success, otherwise see esp_err_t

esp_err_t **esp_apptrace_flush_nolock** (*esp_apptrace_dest_t* dest, uint32_t min_sz, uint32_t tmo)

Flushes remaining data in trace buffer to host without locking internal data. This is a special version of esp_apptrace_flush which should be called from panic handler.

Parameters

- **dest** -- Indicates HW interface to flush data on.
- **min_sz** -- Threshold for flushing data. If current filling level is above this value, data will be flushed. TRAX destinations only.
- **tmo** -- Timeout for operation (in us). Use ESP_APPTRACE_TMO_INFINITE to wait indefinitely.

Returns ESP_OK on success, otherwise see esp_err_t

esp_err_t **esp_apptrace_read** (*esp_apptrace_dest_t* dest, void *data, uint32_t *size, uint32_t tmo)

Reads host data from trace buffer.

Parameters

- **dest** -- Indicates HW interface to read the data on.
- **data** -- Address of buffer to put data from trace buffer.
- **size** -- Pointer to store size of read data. Before call to this function pointed memory must hold requested size of data
- **tmo** -- Timeout for operation (in us). Use ESP_APPTRACE_TMO_INFINITE to wait indefinitely.

Returns ESP_OK on success, otherwise see esp_err_t

uint8_t ***esp_apptrace_down_buffer_get** (*esp_apptrace_dest_t* dest, uint32_t *size, uint32_t tmo)

Retrieves incoming data buffer if any. Once data in the buffer is processed, esp_apptrace_down_buffer_put must be called to indicate it.

Parameters

- **dest** -- Indicates HW interface to receive data.

- **size** -- Address to store size of available data in down buffer. Must be initialized with requested value.
- **tmo** -- Timeout for operation (in us). Use ESP_APPTRACE_TMO_INFINITE to wait indefinitely.

Returns non-NULL on success, otherwise NULL.

[*esp_err_t*](#) **esp_appttrace_down_buffer_put** ([*esp_appttrace_dest_t*](#) dest, uint8_t *ptr, uint32_t tmo)

Indicates that the data in the down buffer is processed. This function is a counterpart of and must be preceded by `esp_appttrace_down_buffer_get`.

Parameters

- **dest** -- Indicates HW interface to receive data. Should be identical to the same parameter in call to `esp_appttrace_down_buffer_get`.
- **ptr** -- Address of trace buffer to release. Should be the value returned by call to `esp_appttrace_down_buffer_get`.
- **tmo** -- Timeout for operation (in us). Use ESP_APPTRACE_TMO_INFINITE to wait indefinitely.

Returns ESP_OK on success, otherwise see `esp_err_t`

bool **esp_appttrace_host_is_connected** ([*esp_appttrace_dest_t*](#) dest)

Checks whether host is connected.

Parameters **dest** -- Indicates HW interface to use.

Returns true if host is connected, otherwise false

void ***esp_appttrace_fopen** ([*esp_appttrace_dest_t*](#) dest, const char *path, const char *mode)

Opens file on host. This function has the same semantic as 'fopen' except for the first argument.

Parameters

- **dest** -- Indicates HW interface to use.
- **path** -- Path to file.
- **mode** -- Mode string. See fopen for details.

Returns non zero file handle on success, otherwise 0

int **esp_appttrace_fclose** ([*esp_appttrace_dest_t*](#) dest, void *stream)

Closes file on host. This function has the same semantic as 'fclose' except for the first argument.

Parameters

- **dest** -- Indicates HW interface to use.
- **stream** -- File handle returned by `esp_appttrace_fopen`.

Returns Zero on success, otherwise non-zero. See fclose for details.

size_t **esp_appttrace_fwrite** ([*esp_appttrace_dest_t*](#) dest, const void *ptr, size_t size, size_t nmemb, void *stream)

Writes to file on host. This function has the same semantic as 'fwrite' except for the first argument.

Parameters

- **dest** -- Indicates HW interface to use.
- **ptr** -- Address of data to write.
- **size** -- Size of an item.
- **nmemb** -- Number of items to write.
- **stream** -- File handle returned by `esp_appttrace_fopen`.

Returns Number of written items. See fwrite for details.

size_t **esp_appttrace_fread** ([*esp_appttrace_dest_t*](#) dest, void *ptr, size_t size, size_t nmemb, void *stream)

Read file on host. This function has the same semantic as 'fread' except for the first argument.

Parameters

- **dest** -- Indicates HW interface to use.
- **ptr** -- Address to store read data.
- **size** -- Size of an item.
- **nmemb** -- Number of items to read.
- **stream** -- File handle returned by `esp_appttrace_fopen`.

Returns Number of read items. See fread for details.

int **esp_apptrace_fseek** (*esp_apptrace_dest_t* dest, void *stream, long offset, int whence)

Set position indicator in file on host. This function has the same semantic as 'fseek' except for the first argument.

Parameters

- **dest** -- Indicates HW interface to use.
- **stream** -- File handle returned by esp_apptrace_fopen.
- **offset** -- Offset. See fseek for details.
- **whence** -- Position in file. See fseek for details.

Returns Zero on success, otherwise non-zero. See fseek for details.

int **esp_apptrace_ftell** (*esp_apptrace_dest_t* dest, void *stream)

Get current position indicator for file on host. This function has the same semantic as 'ftell' except for the first argument.

Parameters

- **dest** -- Indicates HW interface to use.
- **stream** -- File handle returned by esp_apptrace_fopen.

Returns Current position in file. See ftell for details.

int **esp_apptrace_fstop** (*esp_apptrace_dest_t* dest)

Indicates to the host that all file operations are complete. This function should be called after all file operations are finished and indicate to the host that it can perform cleanup operations (close open files etc.).

Parameters **dest** -- Indicates HW interface to use.

Returns ESP_OK on success, otherwise see esp_err_t

int **esp_apptrace_feof** (*esp_apptrace_dest_t* dest, void *stream)

Test end-of-file indicator on a stream. This function has the same semantic as 'feof' except for the first argument.

Parameters

- **dest** -- Indicates HW interface to use.
- **stream** -- File handle returned by esp_apptrace_fopen.

Returns Non-Zero if end-of-file indicator is set for stream. See feof for details.

void **esp_gcov_dump** (void)

Triggers gcov info dump. This function waits for the host to connect to target before dumping data.

Enumerations

enum **esp_apptrace_dest_t**

Application trace data destinations bits.

Values:

enumerator **ESP_APPTRACE_DEST_JTAG**

JTAG destination.

enumerator **ESP_APPTRACE_DEST_TRAX**

xxx_TRAX name is obsolete, use more common xxx_JTAG

enumerator **ESP_APPTRACE_DEST_UART**

UART destination.

enumerator **ESP_APPTRACE_DEST_MAX**

enumerator **ESP_APPTRACE_DEST_NUM**

Header File

- [components/app_trace/include/esp_sysview_trace.h](#)
- This header file can be included with:

```
#include "esp_sysview_trace.h"
```

- This header file is a part of the API provided by the `app_trace` component. To declare that your component depends on `app_trace`, add the following to your `CMakeLists.txt`:

```
REQUIRES app_trace
```

or

```
PRIV_REQUIRES app_trace
```

Functions

static inline [esp_err_t](#) **esp_sysview_flush** (uint32_t tmo)

Flushes remaining data in SystemView trace buffer to host.

Parameters `tmo` -- Timeout for operation (in us). Use `ESP_APPTRACE_TMO_INFINITE` to wait indefinitely.

Returns `ESP_OK`.

int **esp_sysview_vprintf** (const char *format, va_list args)

vprintf-like function to send log messages to the host.

Parameters

- **format** -- Address of format string.
- **args** -- List of arguments.

Returns Number of bytes written.

[esp_err_t](#) **esp_sysview_heap_trace_start** (uint32_t tmo)

Starts SystemView heap tracing.

Parameters `tmo` -- Timeout (in us) to wait for the host to be connected. Use -1 to wait forever.

Returns `ESP_OK` on success, `ESP_ERR_TIMEOUT` if operation has been timed out.

[esp_err_t](#) **esp_sysview_heap_trace_stop** (void)

Stops SystemView heap tracing.

Returns `ESP_OK`.

void **esp_sysview_heap_trace_alloc** (void *addr, uint32_t size, const void *callers)

Sends heap allocation event to the host.

Parameters

- **addr** -- Address of allocated block.
- **size** -- Size of allocated block.
- **callers** -- Pointer to array with callstack addresses. Array size must be `CONFIG_HEAP_TRACING_STACK_DEPTH`.

void **esp_sysview_heap_trace_free** (void *addr, const void *callers)

Sends heap de-allocation event to the host.

Parameters

- **addr** -- Address of de-allocated block.
- **callers** -- Pointer to array with callstack addresses. Array size must be `CONFIG_HEAP_TRACING_STACK_DEPTH`.

2.9.4 Call Function with External Stack

Overview

A given function can be executed with a user-allocated stack space which is independent of current task stack. This mechanism can be used to save stack space wasted by tasks which call a common function with intensive stack usage such as `printf`. The given function can be called inside the shared stack space, which is a callback function deferred by calling `esp_execute_shared_stack_function()`, passing that function as a parameter.

Usage

`esp_execute_shared_stack_function()` takes four arguments:

- a mutex object allocated by the caller, which is used to protect if the same function shares its allocated stack
- a pointer to the top of stack used for that function
- the size of stack in bytes
- a pointer to the shared stack function

The user-defined function is deferred as a callback and can be called using the user-allocated space without taking space from current task stack.

The usage may look like the code below:

```
void external_stack_function(void)
{
    printf("Executing this printf from external stack! \n");
}

//Let us suppose we want to call printf using a separated stack space
//allowing the app to reduce its stack size.
void app_main()
{
    //Allocate a stack buffer, from heap or as a static form:
    StackType_t *shared_stack = malloc(8192 * sizeof(StackType_t));
    assert(shared_stack != NULL);

    //Allocate a mutex to protect its usage:
    SemaphoreHandle_t printf_lock = xSemaphoreCreateMutex();
    assert(printf_lock != NULL);

    //Call the desired function using the macro helper:
    esp_execute_shared_stack_function(printf_lock,
                                     shared_stack,
                                     8192,
                                     external_stack_function);

    vSemaphoreDelete(printf_lock);
    free(shared_stack);
}
```

API Reference

Header File

- [components/esp_system/include/esp_expression_with_stack.h](#)
- This header file can be included with:

```
#include "esp_expression_with_stack.h"
```

Functions

void **esp_execute_shared_stack_function** (*SemaphoreHandle_t* lock, void *stack, size_t stack_size, *shared_stack_function* function)

Calls user defined shared stack space function.

Note: if either lock, stack or stack size is invalid, the expression will be called using the current stack.

Parameters

- **lock** -- Mutex object to protect in case of shared stack
- **stack** -- Pointer to user allocated stack
- **stack_size** -- Size of current stack in bytes
- **function** -- pointer to the shared stack function to be executed

Macros

ESP_EXECUTE_EXPRESSION_WITH_STACK (lock, stack, stack_size, expression)

Type Definitions

typedef void (***shared_stack_function**)(void)

2.9.5 Chip Revision

Overview

ESP32-P4 may have different revisions. These revisions mainly fix some issues, and sometimes also bring new features to the chip. *Versioning Scheme* describes the versioning of these chip revisions, and the APIs to read the versions at runtime.

There are some considerations of compatibility among application, ESP-IDF version, and chip revisions:

- Applications may depend on some fixes/features provided by a chip revision.
- When using updated version of hardware, the hardware may be incompatible with earlier versions of ESP-IDF.

Compatibility Checks of ESP-IDF describes how the application can specify its chip revision requirements, and the way ESP-IDF checks the compatibility. After that, there is troubleshooting information for this mechanism.

Versioning Scheme

A chip's revision number is typically expressed as vX.Y, where:

- X means a **Major** wafer version. If it is changed, it means that the current software version is not compatible with this released chip and the software must be updated to use this chip.
- Y means a **Minor** wafer version. If it is changed that means the current software version is compatible with the released chip, and there is no need to update the software.

If a newly released chip does not contain breaking changes, the chip can run the same software as the previous chip. As such, the new chip's revision number will only increment the minor version while keeping the major version the same (e.g., v1.1 to v1.2).

Conversely, if a newly released chip contains breaking changes, the chip **cannot** run the same software as the previous chip. As such, the new chip's revision number will increment the major version and set the minor version to 0 (e.g., v1.1 to v2.0).

This versioning scheme was selected to indicate the derivation relationship of chip revisions, and clearly distinguish changes in chips between breaking changes and non-breaking changes.

ESP-IDF is designed to execute seamlessly on future chip minor revisions with the same logic as the chip's nearest previous minor revision. Thus, users can directly port their compiled binaries to newer MINOR chip revisions without upgrading their ESP-IDF version and re-compile the whole project.

When a binary is executed on a chip revision of unexpected MAJOR revision, the software is also able to report issues according to the MAJOR revision. The major and minor versioning scheme also allows hardware changes to be branchable.

Note: The current chip revision scheme using major and minor versions was introduced from ESP-IDF v5.0 onwards. Thus bootloaders built using earlier versions of ESP-IDF will still use the legacy chip revision scheme of wafer versions.

EFuse Bits for Chip Revisions Chips have several eFuse version fields:

- Major wafer version (WAFER_VERSION_MAJOR eFuse)
- Minor wafer version (WAFER_VERSION_MINOR eFuse)
- Ignore maximum wafer revision (DISABLE_WAFER_VERSION_MAJOR eFuse). See [Compatibility Checks of ESP-IDF](#) on how this is used.

Note: The previous versioning logic was based on a single eFuse version field (WAFER_VERSION). This approach makes it impossible to mark chips as breaking or non-breaking changes, and the versioning logic becomes linear.

EFuse Bits for eFuse Block Revisions EFuse block has version fields:

- Major efuse block version (BLK_VERSION_MAJOR eFuse)
- Minor efuse block version (BLK_VERSION_MINOR eFuse)
- Ignore maximum efuse block revision (DISABLE_BLK_VERSION_MAJOR eFuse). See [Compatibility Checks of ESP-IDF](#) on how this is used.

Chip Revision APIs These APIs helps to get chip revision from eFuses:

- [efuse_hal_chip_revision\(\)](#). It returns revision in the `major * 100 + minor` format.
- [efuse_hal_get_major_chip_version\(\)](#). It returns Major revision of wafer.
- [efuse_hal_get_minor_chip_version\(\)](#). It returns Minor revision of wafer.

The following Kconfig definitions (in `major * 100 + minor` format) that can help add the chip revision dependency to the code:

- `CONFIG_ESP32P4_REV_MIN_FULL`
- `CONFIG_ESP_REV_MIN_FULL`
- `CONFIG_ESP32P4_REV_MAX_FULL`
- `CONFIG_ESP_REV_MAX_FULL`

EFuse Block Revision APIs These APIs helps to get eFuse block revision from eFuses:

- [efuse_hal_blk_version\(\)](#). It returns revision in the `major * 100 + minor` format.
- [efuse_ll_get_blk_version_major\(\)](#). It returns Major revision of eFuse block.
- [efuse_ll_get_blk_version_minor\(\)](#). It returns Minor revision of eFuse block.

The following Kconfig definitions (in `major * 100 + minor` format) that can help add the eFuse block revision dependency to the code:

- `CONFIG_ESP_EFUSE_BLOCK_REV_MIN_FULL`
- `CONFIG_ESP_EFUSE_BLOCK_REV_MAX_FULL`

Compatibility Checks of ESP-IDF

When building an application that needs to support multiple revisions of a particular chip, the minimum and maximum chip revision numbers supported by the build are specified via Kconfig.

The minimum chip revision can be configured via the `CONFIG_ESP32P4_REV_MIN` option. Specifying the minimum chip revision will limit the software to only run on a chip revisions that are high enough to support some features or bugfixes.

The maximum chip revision cannot be configured and is automatically determined by the current ESP-IDF version being used. ESP-IDF will refuse to boot any chip revision exceeding the maximum chip revision. Given that it is impossible for a particular ESP-IDF version to foresee all future chip revisions, the maximum chip revision is usually set to `maximum supported MAJOR version + 99`. The "Ignore Maximum Revision" eFuse can be set to bypass the maximum revision limitation. However, the software is not guaranteed to work if the maximum revision is ignored.

The eFuse block revision is similar to the chip revision, but it mainly affects the coefficients that are specified in the eFuse (e.g. ADC calibration coefficients).

Below is the information about troubleshooting when the chip revision fails the compatibility check. Then there are technical details of the checking and software behavior on earlier version of ESP-IDF.

Troubleshooting

1. If the 2nd stage bootloader is run on a chip revision smaller than minimum revision specified in the image (i.e., the application), a reboot occurs. The following message will be printed:

```
Image requires chip rev >= v3.0, but chip is v1.0
```

To resolve this issue,

- Use a chip with the required minimum revision or higher.
 - Lower the `CONFIG_ESP32P4_REV_MIN` value and rebuild the image so that it is compatible with the chip revision being used.
2. If application does not match minimum and maximum chip revisions, a reboot occurs. The following message will be printed:

```
Image requires chip rev <= v2.99, but chip is v3.0
```

To resolve this issue, update ESP-IDF to a newer version that supports the chip's revision (`CONFIG_ESP32P4_REV_MAX_FULL`). Alternatively, set the Ignore maximal revision bit in eFuse or use a chip revision that is compatible with the current version of ESP-IDF.

Representing Revision Requirements of a Binary Image For the chip revision, the 2nd stage bootloader and the application binary images contain the `esp_image_header_t` header, which stores information specifying the chip revisions that the image is permitted to run on. This header has 3 fields related to the chip revisions:

- `min_chip_rev` - Minimum chip MAJOR revision required by image (but for ESP32-C3 it is MINOR revision). Its value is determined by `CONFIG_ESP32P4_REV_MIN`.
- `min_chip_rev_full` - Minimum chip MINOR revision required by image in format: `major * 100 + minor`. Its value is determined by `CONFIG_ESP32P4_REV_MIN`.
- `max_chip_rev_full` - Maximum chip revision required by image in format: `major * 100 + minor`. Its value is determined by `CONFIG_ESP32P4_REV_MAX_FULL`. It can not be changed by user. Only Espressif can change it when a new version will be supported in ESP-IDF.

For the eFuse revision, the requirements are stored in `esp_app_desc_t`, which is contained in the application binary image. We only check the application image because the eFuse block revision mostly affects the ADC calibration, which does not really matter in the bootloader. There are 2 fields related to eFuse block revisions:

- `min_efuse_blk_rev_full` - Minimum eFuse block MINOR revision required by image in format: `major * 100 + minor`. Its value is determined by `CONFIG_ESP_EFUSE_BLOCK_REV_MIN_FULL`.

- `max_efuse_blk_rev_full` - Maximum eFuse block MINOR revision required by image in format: `major * 100 + minor`. Its value is determined by `CONFIG_ESP_EFUSE_BLOCK_REV_MAX_FULL`. It reflects whether the current IDF version supports this efuse block format or not, and should not be changed by the user.

Maximum And Minimum Revision Restrictions The order for checking the minimum and maximum revisions during application boot up is as follows:

1. The 1st stage bootloader (ROM bootloader) does not check minimum and maximum revision fields from `esp_image_header_t` before running the 2nd stage bootloader.
2. The initialization phase of the 2nd stage bootloader checks that the 2nd stage bootloader itself can be launched on the chip of this revision. It extracts the minimum revision from the header of the bootloader image and checks against the chip revision from eFuses. If the chip revision is less than the minimum revision, the bootloader refuses to boot up and aborts. The maximum revision is not checked at this phase.
3. Then the 2nd stage bootloader checks the revision requirements of the application. It extracts the minimum and maximum revisions of the chip from the application image header, and the eFuse block from the segment header. Then the bootloader checks these versions against the chip and eFuse block revision from eFuses. If the these revisions are less than their minimum revision or higher than the maximum revision, the bootloader refuses to boot up and aborts. However, if the ignore maximum revision bit is set, the maximum revision constraint can be ignored. The ignore bits are set by the customer themselves when there is confirmation that the software is able to work with this chip revision or eFuse block revision.
4. Furthermore, at the OTA update stage, the running application checks if the new software matches the chip revision and eFuse block revision. It extracts the minimum and maximum chip revisions from the header of the new application image and the eFuse block constraints from the application description to check against the these revisions from eFuses. It checks for revisions matching in the same way that the bootloader does, so that the chip and eFuse block revisions are between their min and max revisions (logic of ignoring max revision also applies).

Backward Compatibility with Bootloaders Built by Older ESP-IDF Versions Please check the chip version using `esptool chip_id` command.

References

- [Compatibility Advisory for Chip Revision Numbering Scheme](#)
- [Compatibility Between ESP-IDF Releases and Revisions of Espressif SoCs](#)
- [SoC Errata](#)
- [ESP-IDF Versions](#)

API Reference

Header File

- `components/hal/include/hal/efuse_hal.h`
- This header file can be included with:

```
#include "hal/efuse_hal.h"
```

Functions

void **efuse_hal_get_mac** (uint8_t *mac)

get factory mac address

uint32_t **efuse_hal_chip_revision** (void)

Returns chip version.

Returns Chip version in format: Major * 100 + Minor

uint32_t **efuse_hal_blk_version** (void)

Return block version.

Returns Block version in format: Major * 100 + Minor

bool **efuse_hal_flash_encryption_enabled** (void)

Is flash encryption currently enabled in hardware?

Flash encryption is enabled if the FLASH_CRYPT_CNT efuse has an odd number of bits set.

Returns true if flash encryption is enabled.

bool **efuse_hal_get_disable_wafer_version_major** (void)

Returns the status of whether the bootloader (and OTA) will check the maximum chip version or not.

Returns true - Skip the maximum chip version check.

bool **efuse_hal_get_disable_blk_version_major** (void)

Returns the status of whether the app start-up (and OTA) will check the efuse block version or not.

Returns true - Skip the efuse block version check.

uint32_t **efuse_hal_get_major_chip_version** (void)

Returns major chip version.

uint32_t **efuse_hal_get_minor_chip_version** (void)

Returns minor chip version.

void **efuse_hal_set_ecdsa_key** (int efuse_key_blk)

Set the efuse block that should be used as ECDSA private key.

Note: The efuse block must be burnt with key purpose ECDSA_KEY

Parameters **efuse_key_blk** -- Efuse key block number (Must be in [EFUSE_BLK_KEY0...EFUSE_BLK_KEY_MAX - 1] range)

2.9.6 Console

ESP-IDF provides `console` component, which includes building blocks needed to develop an interactive console over serial port. This component includes the following features:

- Line editing, provided by [linenoise](#) library. This includes handling of backspace and arrow keys, scrolling through command history, command auto-completion, and argument hints.
- Splitting of command line into arguments.
- Argument parsing, provided by [argtable3](#) library. This library includes APIs used for parsing GNU style command line arguments.
- Functions for registration and dispatching of commands.
- Functions to establish a basic REPL (Read-Evaluate-Print-Loop) environment.

Note: These features can be used together or independently. For example, it is possible to use line editing and command registration features, but use `getopt` or custom code for argument parsing, instead of [argtable3](#). Likewise, it is possible to use simpler means of command input (such as `fgets`) together with the rest of the means for command splitting and argument parsing.

Note: When using a console application on a chip that supports a hardware USB serial interface, we suggest to disable the secondary serial console output. The secondary output will be output-only and consequently does not make sense in an interactive application.

Line Editing

Line editing feature lets users compose commands by typing them, erasing symbols using the `backspace` key, navigating within the command using the left/right keys, navigating to previously typed commands using the up/down keys, and performing autocompletion using the `tab` key.

Note: This feature relies on ANSI escape sequence support in the terminal application. As such, serial monitors which display raw UART data can not be used together with the line editing library. If you see `[6n` or similar escape sequence when running [system/console](#) example instead of a command prompt (e.g., `esp>`), it means that the serial monitor does not support escape sequences. Programs which are known to work are GNU screen, minicom, and `esp-idf-monitor` (which can be invoked using `idf.py monitor` from project directory).

Here is an overview of functions provided by [linenoise](#) library.

Configuration Linenoise library does not need explicit initialization. However, some configuration defaults may need to be changed before invoking the main line editing function.

- `linenoiseClearScreen()`
Clear terminal screen using an escape sequence and position the cursor at the top left corner.
- `linenoiseSetMultiLine()`
Switch between single line and multi line editing modes. In single line mode, if the length of the command exceeds the width of the terminal, the command text is scrolled within the line to show the end of the text. In this case the beginning of the text is hidden. Single line mode needs less data to be sent to refresh screen on each key press, so exhibits less glitching compared to the multi line mode. On the flip side, editing commands and copying command text from terminal in single line mode is harder. Default is single line mode.
- `linenoiseAllowEmpty()`
Set whether linenoise library returns a zero-length string (if `true`) or `NULL` (if `false`) for empty lines. By default, zero-length strings are returned.
- `linenoiseSetMaxLineLen()`
Set maximum length of the line for linenoise library. Default length is 4096 bytes. The default value can be updated to optimize RAM memory usage.

Main Loop

- `linenoise()`
In most cases, console applications have some form of read/eval loop. `linenoise()` is the single function which handles user's key presses and returns the completed line once the `enter` key is pressed. As such, it handles the `read` part of the loop.
- `linenoiseFree()`
This function must be called to release the command line buffer obtained from `linenoise()` function.

Hints and Completions

- `linenoiseSetCompletionCallback()`
When the user presses the `tab` key, linenoise library invokes the completion callback. The callback should inspect the contents of the command typed so far and provide a list of possible completions using calls to `linenoiseAddCompletion()` function. `linenoiseSetCompletionCallback()` function should be called to register this completion callback, if completion feature is desired.
`console` component provides a ready made function to provide completions for registered commands, [esp_console_get_completion\(\)](#) (see below).

- `linenoiseAddCompletion()`
Function to be called by completion callback to inform the library about possible completions of the currently typed command.
- `linenoiseSetHintsCallback()`
Whenever user input changes, linenoise invokes the hints callback. This callback can inspect the command line typed so far, and provide a string with hints (which can include list of command arguments, for example). The library then displays the hint text on the same line where editing happens, possibly with a different color.
- `linenoiseSetFreeHintsCallback()`
If the hint string returned by hints callback is dynamically allocated or needs to be otherwise recycled, the function which performs such cleanup should be registered via `linenoiseSetFreeHintsCallback()`.

History

- `linenoiseHistorySetMaxLen()`
This function sets the number of most recently typed commands to be kept in memory. Users can navigate the history using the up/down arrows keys.
- `linenoiseHistoryAdd()`
Linenoise does not automatically add commands to history. Instead, applications need to call this function to add command strings to the history.
- `linenoiseHistorySave()`
Function saves command history from RAM to a text file, for example on an SD card or on a filesystem in flash memory.
- `linenoiseHistoryLoad()`
Counterpart to `linenoiseHistorySave()`, loads history from a file.
- `linenoiseHistoryFree()`
Releases memory used to store command history. Call this function when done working with linenoise library.

Splitting of Command Line into Arguments

`console` component provides `esp_console_split_argv()` function to split command line string into arguments. The function returns the number of arguments found (`argc`) and fills an array of pointers which can be passed as `argv` argument to any function which accepts arguments in `argc, argv` format.

The command line is split into arguments according to the following rules:

- Arguments are separated by spaces
- If spaces within arguments are required, they can be escaped using `\` (backslash) character.
- Other escape sequences which are recognized are `\\` (which produces literal backslash) and `\"`, which produces a double quote.
- Arguments can be quoted using double quotes. Quotes may appear only in the beginning and at the end of the argument. Quotes within the argument must be escaped as mentioned above. Quotes surrounding the argument are stripped by `esp_console_split_argv` function.

Examples:

- `abc def 1 20 .3 > [ abc, def, 1, 20, .3 ]`
- `abc "123 456" def > [ abc, 123 456, def ]`
- ``a\ b\\c\" > [a b\\c"]`

Argument Parsing

For argument parsing, `console` component includes `argtable3` library. Please see [tutorial](#) for an introduction to `argtable3`. Github repository also includes [examples](#).

Command Registration and Dispatching

`console` component includes utility functions which handle registration of commands, matching commands typed by the user to registered ones, and calling these commands with the arguments given on the command line.

Application first initializes command registration module using a call to `esp_console_init()`, and calls `esp_console_cmd_register()` function to register command handlers.

For each command, application provides the following information (in the form of `esp_console_cmd_t` structure):

- Command name (string without spaces)
- Help text explaining what the command does
- Optional hint text listing the arguments of the command. If application uses Argtable3 for argument parsing, hint text can be generated automatically by providing a pointer to argtable argument definitions structure instead.
- The command handler function.

A few other functions are provided by the command registration module:

- `esp_console_run()`
This function takes the command line string, splits it into argc/argv argument list using `esp_console_split_argv()`, looks up the command in the list of registered components, and if it is found, executes its handler.
- `esp_console_register_help_command()`
Adds help command to the list of registered commands. This command prints the list of all the registered commands, along with their arguments and help texts.
- `esp_console_get_completion()`
Callback function to be used with `linenoiseSetCompletionCallback()` from linenoise library. Provides completions to linenoise based on the list of registered commands.
- `esp_console_get_hint()`
Callback function to be used with `linenoiseSetHintsCallback()` from linenoise library. Provides argument hints for registered commands to linenoise.

Initialize Console REPL Environment

To establish a basic REPL environment, `console` component provides several useful APIs, combining those functions described above.

In a typical application, you only need to call `esp_console_new_repl_uart()` to initialize the REPL environment based on UART device, including driver install, basic console configuration, spawning a thread to do REPL task and register several useful commands (e.g., `help`).

After that, you can register your own commands with `esp_console_cmd_register()`. The REPL environment keeps in init state until you call `esp_console_start_repl()`.

Application Example

Example application illustrating usage of the `console` component is available in `system/console` directory. This example shows how to initialize UART and VFS functions, set up linenoise library, read and handle commands from UART, and store command history in Flash. See README.md in the example directory for more details.

Besides that, ESP-IDF contains several useful examples which are based on the `console` component and can be treated as "tools" when developing applications. For example, [peripherals/i2c/i2c_tools](#), [wifi/iperf](#).

API Reference

Header File

- `components/console/esp_console.h`
- This header file can be included with:

```
#include "esp_console.h"
```

- This header file is a part of the API provided by the `console` component. To declare that your component depends on `console`, add the following to your CMakeLists.txt:

REQUIRES console

or

PRIV_REQUIRES console

Functions

esp_err_t **esp_console_init** (const *esp_console_config_t* *config)

initialize console module

Note: Call this once before using other console module features

Parameters **config** -- console configuration

Returns

- ESP_OK on success
- ESP_ERR_NO_MEM if out of memory
- ESP_ERR_INVALID_STATE if already initialized
- ESP_ERR_INVALID_ARG if the configuration is invalid

esp_err_t **esp_console_deinit** (void)

de-initialize console module

Note: Call this once when done using console module functions

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE if not initialized yet

esp_err_t **esp_console_cmd_register** (const *esp_console_cmd_t* *cmd)

Register console command.

Parameters **cmd** -- pointer to the command description; can point to a temporary value

Returns

- ESP_OK on success
- ESP_ERR_NO_MEM if out of memory
- ESP_ERR_INVALID_ARG if command description includes invalid arguments

esp_err_t **esp_console_run** (const char *cmdline, int *cmd_ret)

Run command line.

Parameters

- **cmdline** -- command line (command name followed by a number of arguments)
- **cmd_ret** -- [out] return code from the command (set if command was run)

Returns

- ESP_OK, if command was run
- ESP_ERR_INVALID_ARG, if the command line is empty, or only contained whitespace
- ESP_ERR_NOT_FOUND, if command with given name wasn't registered
- ESP_ERR_INVALID_STATE, if esp_console_init wasn't called

size_t **esp_console_split_argv** (char *line, char **argv, size_t argv_size)

Split command line into arguments in place.


```
* - This function finds whitespace-separated arguments in the given input line.
*
*   'abc def 1 20 .3' -> [ 'abc', 'def', '1', '20', '.3' ]
*
* - Argument which include spaces may be surrounded with quotes. In this case
*   spaces are preserved and quotes are stripped.
*
*   'abc "123 456" def' -> [ 'abc', '123 456', 'def' ]
*
* - Escape sequences may be used to produce backslash, double quote, and space:
*
*   'a\ b\\c\"' -> [ 'a b\c"' ]
*
```

Note: Pointers to at most `argv_size - 1` arguments are returned in `argv` array. The pointer after the last one (i.e. `argv[argc]`) is set to `NULL`.

Parameters

- **line** -- pointer to buffer to parse; it is modified in place
- **argv** -- array where the pointers to arguments are written
- **argv_size** -- number of elements in `argv_array` (max. number of arguments)

Returns number of arguments found (`argc`)

void **esp_console_get_completion** (const char *buf, [linenoiseCompletions](#) *lc)

Callback which provides command completion for linenoise library.

When using linenoise for line editing, command completion support can be enabled like this:

```
linenoiseSetCompletionCallback(&esp_console_get_completion);
```

Parameters

- **buf** -- the string typed by the user
- **lc** -- `linenoiseCompletions` to be filled in

const char ***esp_console_get_hint** (const char *buf, int *color, int *bold)

Callback which provides command hints for linenoise library.

When using linenoise for line editing, hints support can be enabled as follows:

```
linenoiseSetHintsCallback((linenoiseHintsCallback*) &esp_console_get_hint);
```

The extra cast is needed because `linenoiseHintsCallback` is defined as returning a `char*` instead of `const char*`.

Parameters

- **buf** -- line typed by the user
- **color** -- **[out]** ANSI color code to be used when displaying the hint
- **bold** -- **[out]** set to 1 if hint has to be displayed in bold

Returns string containing the hint text. This string is persistent and should not be freed (i.e. `linenoiseSetFreeHintsCallback` should not be used).

[esp_err_t](#) **esp_console_register_help_command** (void)

Register a 'help' command.

Default 'help' command prints the list of registered commands along with hints and help strings if no additional argument is given. If an additional argument is given, the help command will look for a command with the same name and only print the hints and help strings of that command.

Returns

- `ESP_OK` on success
- `ESP_ERR_INVALID_STATE`, if `esp_console_init` wasn't called


```
esp_err_t esp_console_new_repl_uart (const esp_console_dev_uart_config_t *dev_config, const  
                                     esp_console_repl_config_t *repl_config, esp_console_repl_t  
                                     **ret_repl)
```

Establish a console REPL environment over UART driver.

Attention This function is meant to be used in the examples to make the code more compact. Applications which use console functionality should be based on the underlying linenoise and esp_console functions.

Note: This is an all-in-one function to establish the environment needed for REPL, includes:

- Install the UART driver on the console UART (8n1, 115200, REF_TICK clock source)
 - Configures the stdin/stdout to go through the UART driver
 - Initializes linenoise
 - Spawn new thread to run REPL in the background
-

Parameters

- **dev_config** -- [in] UART device configuration
- **repl_config** -- [in] REPL configuration
- **ret_repl** -- [out] return REPL handle after initialization succeed, return NULL otherwise

Returns

- ESP_OK on success
- ESP_FAIL Parameter error

```
esp_err_t esp_console_start_repl (esp_console_repl_t *repl)
```

Start REPL environment.

Note: Once the REPL gets started, it won't be stopped until the user calls repl->del(repl) to destroy the REPL environment.

Parameters **repl** -- [in] REPL handle returned from esp_console_new_repl_XXX

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE, if repl has started already

Structures

```
struct esp_console_config_t
```

Parameters for console initialization.

Public Members

```
size_t max_cmdline_length
```

length of command line buffer, in bytes

```
size_t max_cmdline_args
```

maximum number of command line arguments to parse

```
uint32_t heap_alloc_caps
```

where to (e.g. MALLOC_CAP_SPIRAM) allocate heap objects such as cmds used by esp_console

int **hint_color**

ASCII color code of hint text.

int **hint_bold**

Set to 1 to print hint text in bold.

struct **esp_console_repl_config_t**

Parameters for console REPL (Read Eval Print Loop)

Public Members

uint32_t **max_history_len**

maximum length for the history

const char ***history_save_path**

file path used to save history commands, set to NULL won't save to file system

uint32_t **task_stack_size**

repl task stack size

uint32_t **task_priority**

repl task priority

const char ***prompt**

prompt (NULL represents default: "esp> ")

size_t **max_cmdline_length**

maximum length of a command line. If 0, default value will be used

struct **esp_console_dev_uart_config_t**

Parameters for console device: UART.

Public Members

int **channel**

UART channel number (count from zero)

int **baud_rate**

Communication baud rate.

int **tx_gpio_num**

GPIO number for TX path, -1 means using default one.

int **rx_gpio_num**

GPIO number for RX path, -1 means using default one.

struct **esp_console_cmd_t**

Console command description.

Public Members

const char ***command**

Command name. Must not be NULL, must not contain spaces. The pointer must be valid until the call to `esp_console_deinit`.

const char ***help**

Help text for the command, shown by help command. If set, the pointer must be valid until the call to `esp_console_deinit`. If not set, the command will not be listed in 'help' output.

const char ***hint**

Hint text, usually lists possible arguments. If set to NULL, and 'argtable' field is non-NULL, hint will be generated automatically

esp_console_cmd_func_t **func**

Pointer to a function which implements the command.

void ***argtable**

Array or structure of pointers to `arg_xxx` structures, may be NULL. Used to generate hint text if 'hint' is set to NULL. Array/structure which this field points to must end with an `arg_end`. Only used for the duration of `esp_console_cmd_register` call.

struct **esp_console_repl_s**

Console REPL base structure.

Public Members

esp_err_t (***del**)(*esp_console_repl_t* *repl)

Delete console REPL environment.

Param repl [in] REPL handle returned from `esp_console_new_repl_xxx`

Return

- ESP_OK on success
- ESP_FAIL on errors

Macros

ESP_CONSOLE_CONFIG_DEFAULT ()

Default console configuration value.

ESP_CONSOLE_REPL_CONFIG_DEFAULT ()

Default console repl configuration value.

ESP_CONSOLE_DEV_UART_CONFIG_DEFAULT ()

Type Definitions

typedef struct *linenoiseCompletions* **linenoiseCompletions**

typedef int (***esp_console_cmd_func_t**)(int argc, char **argv)

Console command main function.

Param argc number of arguments

Param argv array with argc entries, each pointing to a zero-terminated string argument

Return console command return code, 0 indicates "success"

typedef struct *esp_console_repl_s* **esp_console_repl_t**

Type defined for console REPL.

2.9.7 eFuse Manager

Introduction

The eFuse Manager library is designed to structure access to eFuse bits and make using these easy. This library operates eFuse bits by a structure name which is assigned in eFuse table. This sections introduces some concepts used by eFuse Manager.

Hardware Description

The ESP32-P4 has a number of eFuses which can store system and user parameters. Each eFuse is a one-bit field which can be programmed to 1 after which it cannot be reverted back to 0. Some of system parameters are using these eFuse bits directly by hardware modules and have special place (for example EFUSE_BLK0).

For more details, see **ESP32-P4 Technical Reference Manual > eFuse Controller (eFuse)** [\[PDF\]](#). Some eFuse bits are available for user applications.

ESP32-P4 has 11 eFuse blocks each of the size of 256 bits (not all bits are available):

- EFUSE_BLK0 is used entirely for system purposes;
- EFUSE_BLK1 is used entirely for system purposes;
- EFUSE_BLK2 is used entirely for system purposes;
- EFUSE_BLK3 (also named EFUSE_BLK_USER_DATA) can be used for user purposes;
- EFUSE_BLK4 (also named EFUSE_BLK_KEY0) can be used as key (for secure_boot or flash_encryption) or for user purposes;
- EFUSE_BLK5 (also named EFUSE_BLK_KEY1) can be used as key (for secure_boot or flash_encryption) or for user purposes;
- EFUSE_BLK6 (also named EFUSE_BLK_KEY2) can be used as key (for secure_boot or flash_encryption) or for user purposes;
- EFUSE_BLK7 (also named EFUSE_BLK_KEY3) can be used as key (for secure_boot or flash_encryption) or for user purposes;
- EFUSE_BLK8 (also named EFUSE_BLK_KEY4) can be used as key (for secure_boot or flash_encryption) or for user purposes;
- EFUSE_BLK9 (also named EFUSE_BLK_KEY5) can be used as key (for secure_boot or flash_encryption) or for user purposes;
- EFUSE_BLK10 (also named EFUSE_BLK_SYS_DATA_PART2) is reserved for system purposes.

Each block is divided into 8 32-bits registers.

eFuse Manager Component

The component has API functions for reading and writing fields. Access to the fields is carried out through the structures that describe the location of the eFuse bits in the blocks. The component provides the ability to form fields of any length and from any number of individual bits. The description of the fields is made in a CSV file in a table form. To generate from a tabular form (CSV file) in the C-source uses the tool `efuse_table_gen.py`. The tool checks the CSV file for uniqueness of field names and bit intersection, in case of using a *custom* file from the user's project directory, the utility checks with the *common* CSV file.

CSV files:

- common (*esp_efuse_table.csv*) - contains eFuse fields which are used inside the ESP-IDF. C-source generation should be done manually when changing this file (run command `idf.py efuse-common-table`). Note that changes in this file can lead to incorrect operation.
- custom - (optional and can be enabled by [CONFIG_EFUSE_CUSTOM_TABLE](#)) contains eFuse fields that are used by the user in their application. C-source generation should be done manually when changing this file and running `idf.py efuse-custom-table`.

Description CSV File

The CSV file contains a description of the eFuse fields. In the simple case, one field has one line of description. Table header:

```
# field_name, efuse_block(EFUSE_BLK0..EFUSE_BLK10), bit_start(0..255), bit_
↪count(1..256), comment
```

Individual params in CSV file the following meanings:

field_name

Name of field. The prefix *ESP_EFUSE_* is added to the name, and this field name is available in the code. This name is used to access the fields. The name must be unique for all fields. If the line has an empty name, then this line is combined with the previous field. This allows you to set an arbitrary order of bits in the field, and expand the field as well (see *MAC_FACTORY* field in the common table). The field_name supports structured format using `.` to show that the field belongs to another field (see *WR_DIS* and *RD_DIS* in the common table).

efuse_block

Block number. It determines where the eFuse bits are placed for this field. Available *EFUSE_BLK0..EFUSE_BLK10*.

bit_start

Start bit number (0..255). The *bit_start* field can be omitted. In this case, it is set to *bit_start* + *bit_count* from the previous record, if it has the same *efuse_block*. Otherwise (if *efuse_block* is different, or this is the first entry), an error will be generated.

bit_count

The number of bits to use in this field (1..-). This parameter cannot be omitted. This field also may be *MAX_BLK_LEN* in this case, the field length has the maximum block length.

comment

This param is using for comment field, it also move to C-header file. The comment field can be omitted.

If a non-sequential bit order is required to describe a field, then the field description in the following lines should be continued without specifying a name, indicating that it belongs to one field. For example two fields *MAC_FACTORY* and *MAC_FACTORY_CRC*:

```
# Factory MAC address #
#####
MAC_FACTORY, EFUSE_BLK0, 72, 8, Factory MAC addr [0]
, EFUSE_BLK0, 64, 8, Factory MAC addr [1]
, EFUSE_BLK0, 56, 8, Factory MAC addr [2]
, EFUSE_BLK0, 48, 8, Factory MAC addr [3]
, EFUSE_BLK0, 40, 8, Factory MAC addr [4]
, EFUSE_BLK0, 32, 8, Factory MAC addr [5]
MAC_FACTORY_CRC, EFUSE_BLK0, 80, 8, CRC8 for factory MAC address
```

This field is available in code as *ESP_EFUSE_MAC_FACTORY* and *ESP_EFUSE_MAC_FACTORY_CRC*.

Structured eFuse Fields

WR_DIS,	EFUSE_BLK0,	0,	32,	Write protection
WR_DIS.RD_DIS,	EFUSE_BLK0,	0,	1,	Write protection for
↪RD_DIS				
WR_DIS.FIELD_1,	EFUSE_BLK0,	1,	1,	Write protection for
↪FIELD_1				
WR_DIS.FIELD_2,	EFUSE_BLK0,	2,	4,	Write protection for
↪FIELD_2 (includes B1 and B2)				
WR_DIS.FIELD_2.B1,	EFUSE_BLK0,	2,	2,	Write protection for
↪FIELD_2.B1				
WR_DIS.FIELD_2.B2,	EFUSE_BLK0,	4,	2,	Write protection for
↪FIELD_2.B2				
WR_DIS.FIELD_3,	EFUSE_BLK0,	5,	1,	Write protection for
↪FIELD_3				
WR_DIS.FIELD_3.ALIAS,	EFUSE_BLK0,	5,	1,	Write protection for
↪FIELD_3 (just a alias for WR_DIS.FIELD_3)				
WR_DIS.FIELD_4,	EFUSE_BLK0,	7,	1,	Write protection for
↪FIELD_4				

The structured eFuse field looks like `WR_DIS.RD_DIS` where the dot points that this field belongs to the parent field - `WR_DIS` and cannot be out of the parent's range.

It is possible to use some levels of structured fields as `WR_DIS.FIELD_2.B1` and `B2`. These fields should not be crossed each other and should be in the range of two fields: `WR_DIS` and `WR_DIS.FIELD_2`.

It is possible to create aliases for fields with the same range, see `WR_DIS.FIELD_3` and `WR_DIS.FIELD_3.ALIAS`.

The ESP-IDF names for structured eFuse fields should be unique. The `efuse_table_gen` tool generates the final names where the dot is replaced by `_`. The names for using in ESP-IDF are `ESP_EFUSE_WR_DIS`, `ESP_EFUSE_WR_DIS_RD_DIS`, `ESP_EFUSE_WR_DIS_FIELD_2_B1`, etc.

The `efuse_table_gen` tool checks that the fields do not overlap each other and must be within the range of a field if there is a violation, then throws the following error:

```
Field at USER_DATA, EFUSE_BLK3, 0, 256 intersected with SERIAL_NUMBER, EFUSE_
↪BLK3, 0, 32
```

Solution: Describe `SERIAL_NUMBER` to be included in `USER_DATA`. (`USER_DATA.SERIAL_NUMBER`).

```
Field at FEILD, EFUSE_BLK3, 0, 50 out of range FEILD.MAJOR_NUMBER, EFUSE_BLK3,
↪60, 32
```

Solution: Change `bit_start` for `FIELD.MAJOR_NUMBER` from 60 to 0, so `MAJOR_NUMBER` is in the `FEILD` range.

efuse_table_gen.py Tool

The tool is designed to generate C-source files from CSV file and validate fields. First of all, the check is carried out on the uniqueness of the names and overlaps of the field bits. If an additional *custom* file is used, it will be checked with the existing *common* file (`esp_efuse_table.csv`). In case of errors, a message will be displayed and the string that caused the error. C-source files contain structures of type `esp_efuse_desc_t`.

To generate a *common* files, use the following command `idf.py efuse-common-table` or:

```
cd $IDF_PATH/components/efuse/
./efuse_table_gen.py --idf_target esp32p4 esp32p4/esp_efuse_table.csv
```

After generation in the folder `$IDF_PATH/components/efuse/esp32p4` create:

- `esp_efuse_table.c` file.

- In *include* folder *esp_efuse_table.c* file.

To generate a *custom* files, use the following command `idf.py efuse-custom-table` or:

```
cd $IDF_PATH/components/efuse/  
./efuse_table_gen.py --idf_target esp32p4 esp32p4/esp_efuse_table.csv PROJECT_PATH/  
↪main/esp_efuse_custom_table.csv
```

After generation in the folder `PROJECT_PATH/main` create:

- *esp_efuse_custom_table.c* file.
- In *include* folder *esp_efuse_custom_table.c* file.

To use the generated fields, you need to include two files:

```
#include "esp_efuse.h"  
#include "esp_efuse_table.h" // or "esp_efuse_custom_table.h"
```

Supported Coding Scheme

Coding schemes are used to protect against data corruption. ESP32-P4 supports two coding schemes:

- **None.** EFUSE_BLK0 is stored with four backups, meaning each bit is stored four times. This backup scheme is automatically applied by the hardware and is not visible to software. EFUSE_BLK0 can be written many times.
- **RS.** EFUSE_BLK1 - EFUSE_BLK10 use Reed-Solomon coding scheme that supports up to 5 bytes of automatic error correction. Software encodes the 32-byte EFUSE_BLKx using RS (44, 32) to generate a 12-byte check code, and then burn the EFUSE_BLKx and the check code into eFuse at the same time. The eFuse Controller automatically decodes the RS encoding and applies error correction when reading back the eFuse block. Because the RS check codes are generated across the entire 256-bit eFuse block, each block can only be written to one time.

To write some fields into one block, or different blocks in one time, you need to use the `batch writing` mode. Firstly set this mode through `esp_efuse_batch_write_begin()` function then write some fields as usual using the `esp_efuse_write_...` functions. At the end to burn them, call the `esp_efuse_batch_write_commit()` function. It burns prepared data to the eFuse blocks and disables the batch recording mode.

Note: If there is already pre-written data in the eFuse block using the Reed-Solomon encoding scheme, then it is not possible to write anything extra (even if the required bits are empty) without breaking the previous encoding data. This encoding data will be overwritten with new encoding data and completely destroyed (however, the payload eFuses are not damaged). It can be related to: `CUSTOM_MAC`, `SPI_PAD_CONFIG_HD`, `SPI_PAD_CONFIG_CS`, etc. Please contact Espressif to order the required pre-burnt eFuses.

FOR TESTING ONLY (NOT RECOMMENDED): You can ignore or suppress errors that violate encoding scheme data in order to burn the necessary bits in the eFuse block.

eFuse API

Access to the fields is via a pointer to the description structure. API functions have some basic operation:

- `esp_efuse_read_field_blob()` - returns an array of read eFuse bits.
- `esp_efuse_read_field_cnt()` - returns the number of bits programmed as "1".
- `esp_efuse_write_field_blob()` - writes an array.
- `esp_efuse_write_field_cnt()` - writes a required count of bits as "1".
- `esp_efuse_get_field_size()` - returns the number of bits by the field name.
- `esp_efuse_read_reg()` - returns value of eFuse register.
- `esp_efuse_write_reg()` - writes value to eFuse register.
- `esp_efuse_get_coding_scheme()` - returns eFuse coding scheme for blocks.

- `esp_efuse_read_block()` - reads key to eFuse block starting at the offset and the required size.
- `esp_efuse_write_block()` - writes key to eFuse block starting at the offset and the required size.
- `esp_efuse_batch_write_begin()` - set the batch mode of writing fields.
- `esp_efuse_batch_write_commit()` - writes all prepared data for batch writing mode and reset the batch writing mode.
- `esp_efuse_batch_write_cancel()` - reset the batch writing mode and prepared data.
- `esp_efuse_get_key_dis_read()` - Returns a read protection for the key block.
- `esp_efuse_set_key_dis_read()` - Sets a read protection for the key block.
- `esp_efuse_get_key_dis_write()` - Returns a write protection for the key block.
- `esp_efuse_set_key_dis_write()` - Sets a write protection for the key block.
- `esp_efuse_get_key_purpose()` - Returns the current purpose set for an eFuse key block.
- `esp_efuse_write_key()` - Programs a block of key data to an eFuse block
- `esp_efuse_write_keys()` - Programs keys to unused eFuse blocks
- `esp_efuse_find_purpose()` - Finds a key block with the particular purpose set.
- `esp_efuse_get_keypurpose_dis_write()` - Returns a write protection of the key purpose field for an eFuse key block (for esp32 always true).
- `esp_efuse_key_block_unused()` - Returns true if the key block is unused, false otherwise.

For frequently used fields, special functions are made, like this `esp_efuse_get_pkg_ver()`.

eFuse API for Keys

EFUSE_BLK_KEY0 - EFUSE_BLK_KEY5 are intended to keep up to 6 keys with a length of 256-bits. Each key has an ESP_EFUSE_KEY_PURPOSE_x field which defines the purpose of these keys. The purpose field is described in `esp_efuse_purpose_t`.

The purposes like ESP_EFUSE_KEY_PURPOSE_XTS_AES... are used for flash encryption.

The purposes like ESP_EFUSE_KEY_PURPOSE_SECURE_BOOT_DIGEST... are used for secure boot.

There are some eFuse APIs useful to work with states of keys.

- `esp_efuse_get_purpose_field()` - Returns a pointer to a key purpose for an eFuse key block.
- `esp_efuse_get_key()` - Returns a pointer to a key block.
- `esp_efuse_set_key_purpose()` - Sets a key purpose for an eFuse key block.
- `esp_efuse_set_keypurpose_dis_write()` - Sets a write protection of the key purpose field for an eFuse key block.
- `esp_efuse_find_unused_key_block()` - Search for an unused key block and return the first one found.
- `esp_efuse_count_unused_key_blocks()` - Returns the number of unused eFuse key blocks in the range EFUSE_BLK_KEY0..EFUSE_BLK_KEY_MAX
- `esp_efuse_get_digest_revoke()` - Returns the status of the Secure Boot public key digest revocation bit.
- `esp_efuse_set_digest_revoke()` - Sets the Secure Boot public key digest revocation bit.
- `esp_efuse_get_write_protect_of_digest_revoke()` - Returns a write protection of the Secure Boot public key digest revocation bit.
- `esp_efuse_set_write_protect_of_digest_revoke()` - Sets a write protection of the Secure Boot public key digest revocation bit.

How to Add a New Field

1. Find a free bits for field. Show `esp_efuse_table.csv` file or run `idf.py show-efuse-table` or the next command:

```
$ ./efuse_table_gen.py -t IDF_TARGET_PATH_NAME esp32p4/esp_efuse_table.csv --info
Max number of bits in BLK 256
Parsing efuse CSV input file ../esp32p4/esp_efuse_table.csv ...
```

(continues on next page)

(continued from previous page)

Verifying efuse table...

Sorted efuse table:

#	field_name	efuse_block	bit_start	bit_count
1	WR_DIS	EFUSE_BLK0	0	32
2	WR_DIS.RD_DIS	EFUSE_BLK0	0	1
3	WR_DIS.SPI_BOOT_CRYPT_CNT	EFUSE_BLK0	4	1
4	WR_DIS.SECURE_BOOT_KEY_REVOKE0	EFUSE_BLK0	5	1
5	WR_DIS.SECURE_BOOT_KEY_REVOKE1	EFUSE_BLK0	6	1
6	WR_DIS.SECURE_BOOT_KEY_REVOKE2	EFUSE_BLK0	7	1
7	WR_DIS.KEY_PURPOSE_0	EFUSE_BLK0	8	1
8	WR_DIS.KEY_PURPOSE_1	EFUSE_BLK0	9	1
9	WR_DIS.KEY_PURPOSE_2	EFUSE_BLK0	10	1
10	WR_DIS.KEY_PURPOSE_3	EFUSE_BLK0	11	1
11	WR_DIS.KEY_PURPOSE_4	EFUSE_BLK0	12	1
12	WR_DIS.KEY_PURPOSE_5	EFUSE_BLK0	13	1
13	WR_DIS.SECURE_BOOT_EN	EFUSE_BLK0	15	1
14	WR_DIS.BLK1	EFUSE_BLK0	20	1
15	WR_DIS.MAC	EFUSE_BLK0	20	1
16	WR_DIS.MAC_EXT	EFUSE_BLK0	20	1
17	WR_DIS.BLOCK_SYS_DATA1	EFUSE_BLK0	21	1
18	WR_DIS.BLOCK_USR_DATA	EFUSE_BLK0	22	1
19	WR_DIS.BLOCK_KEY0	EFUSE_BLK0	23	1
20	WR_DIS.BLOCK_KEY1	EFUSE_BLK0	24	1
21	WR_DIS.BLOCK_KEY2	EFUSE_BLK0	25	1
22	WR_DIS.BLOCK_KEY3	EFUSE_BLK0	26	1
23	WR_DIS.BLOCK_KEY4	EFUSE_BLK0	27	1
24	WR_DIS.BLOCK_KEY5	EFUSE_BLK0	28	1
25	WR_DIS.BLOCK_SYS_DATA2	EFUSE_BLK0	29	1
26	RD_DIS	EFUSE_BLK0	32	7
27	RD_DIS.BLOCK_KEY0	EFUSE_BLK0	32	1
28	RD_DIS.BLOCK_KEY1	EFUSE_BLK0	33	1
29	RD_DIS.BLOCK_KEY2	EFUSE_BLK0	34	1
30	RD_DIS.BLOCK_KEY3	EFUSE_BLK0	35	1
31	RD_DIS.BLOCK_KEY4	EFUSE_BLK0	36	1
32	RD_DIS.BLOCK_KEY5	EFUSE_BLK0	37	1
33	RD_DIS.BLOCK_SYS_DATA2	EFUSE_BLK0	38	1
34	USB_DEVICE_EXCHG_PINS	EFUSE_BLK0	39	1
35	USB_OTG11_EXCHG_PINS	EFUSE_BLK0	40	1
36	DIS_USB_JTAG	EFUSE_BLK0	41	1
37	POWERGLITCH_EN	EFUSE_BLK0	42	1
38	DIS_FORCE_DOWNLOAD	EFUSE_BLK0	44	1
39	SPI_DOWNLOAD_MSPI_DIS	EFUSE_BLK0	45	1
40	DIS_TWAI	EFUSE_BLK0	46	1
41	JTAG_SEL_ENABLE	EFUSE_BLK0	47	1
42	SOFT_DIS_JTAG	EFUSE_BLK0	48	3
43	DIS_PAD_JTAG	EFUSE_BLK0	51	1
44	DIS_DOWNLOAD_MANUAL_ENCRYPT	EFUSE_BLK0	52	1
45	USB_PHY_SEL	EFUSE_BLK0	57	1
46	KM_HUK_GEN_STATE_LOW	EFUSE_BLK0	58	6
47	KM_HUK_GEN_STATE_HIGH	EFUSE_BLK0	64	3
48	KM_RND_SWITCH_CYCLE	EFUSE_BLK0	67	2
49	KM_DEPLOY_ONLY_ONCE	EFUSE_BLK0	69	4
50	FORCE_USE_KEY_MANAGER_KEY	EFUSE_BLK0	73	4
51	FORCE_DISABLE_SW_INIT_KEY	EFUSE_BLK0	77	1
52	XTS_KEY_LENGTH_256	EFUSE_BLK0	78	1
53	WDT_DELAY_SEL	EFUSE_BLK0	80	2
54	SPI_BOOT_CRYPT_CNT	EFUSE_BLK0	82	3
55	SECURE_BOOT_KEY_REVOKE0	EFUSE_BLK0	85	1
56	SECURE_BOOT_KEY_REVOKE1	EFUSE_BLK0	86	1
57	SECURE_BOOT_KEY_REVOKE2	EFUSE_BLK0	87	1
58	KEY_PURPOSE_0	EFUSE_BLK0	88	4

(continues on next page)

(continued from previous page)

59	KEY_PURPOSE_1	EFUSE_BLK0	92	4
60	KEY_PURPOSE_2	EFUSE_BLK0	96	4
61	KEY_PURPOSE_3	EFUSE_BLK0	100	4
62	KEY_PURPOSE_4	EFUSE_BLK0	104	4
63	KEY_PURPOSE_5	EFUSE_BLK0	108	4
64	SEC_DPA_LEVEL	EFUSE_BLK0	112	2
65	ECDSA_ENABLE_SOFT_K	EFUSE_BLK0	114	1
66	CRYPT_DPA_ENABLE	EFUSE_BLK0	115	1
67	SECURE_BOOT_EN	EFUSE_BLK0	116	1
68	SECURE_BOOT_AGGRESSIVE_REVOKE	EFUSE_BLK0	117	1
69	FLASH_TYPE	EFUSE_BLK0	119	1
70	FLASH_PAGE_SIZE	EFUSE_BLK0	120	2
71	FLASH_ECC_EN	EFUSE_BLK0	122	1
72	DIS_USB_OTG_DOWNLOAD_MODE	EFUSE_BLK0	123	1
73	FLASH_TPUW	EFUSE_BLK0	124	4
74	DIS_DOWNLOAD_MODE	EFUSE_BLK0	128	1
75	DIS_DIRECT_BOOT	EFUSE_BLK0	129	1
76	DIS_USB_SERIAL_JTAG_ROM_PRINT	EFUSE_BLK0	130	1
77	LOCK_KM_KEY	EFUSE_BLK0	131	1
78	DIS_USB_SERIAL_JTAG_DOWNLOAD_MODE	EFUSE_BLK0	132	1
↪1				
79	ENABLE_SECURITY_DOWNLOAD	EFUSE_BLK0	133	1
80	UART_PRINT_CONTROL	EFUSE_BLK0	134	2
81	FORCE_SEND_RESUME	EFUSE_BLK0	136	1
82	SECURE_VERSION	EFUSE_BLK0	137	16
83	SECURE_BOOT_DISABLE_FAST_WAKE	EFUSE_BLK0	153	1
84	HYS_EN_PAD	EFUSE_BLK0	154	1
85	DCDC_VSET	EFUSE_BLK0	155	5
86	PXA0_TIEH_SEL_0	EFUSE_BLK0	160	2
87	PXA0_TIEH_SEL_1	EFUSE_BLK0	162	2
88	PXA0_TIEH_SEL_2	EFUSE_BLK0	164	2
89	PXA0_TIEH_SEL_3	EFUSE_BLK0	166	2
90	KM_DISABLE_DEPLOY_MODE	EFUSE_BLK0	168	4
91	HP_PWR_SRC_SEL	EFUSE_BLK0	178	1
92	DCDC_VSET_EN	EFUSE_BLK0	179	1
93	DIS_WDT	EFUSE_BLK0	180	1
94	DIS_SWD	EFUSE_BLK0	181	1
95	MAC	EFUSE_BLK1	0	8
96	MAC	EFUSE_BLK1	8	8
97	MAC	EFUSE_BLK1	16	8
98	MAC	EFUSE_BLK1	24	8
99	MAC	EFUSE_BLK1	32	8
100	MAC	EFUSE_BLK1	40	8
101	MAC_EXT	EFUSE_BLK1	48	8
102	MAC_EXT	EFUSE_BLK1	56	8
103	SYS_DATA_PART2	EFUSE_BLK10	0	256
104	BLOCK_SYS_DATA1	EFUSE_BLK2	0	256
105	USER_DATA	EFUSE_BLK3	0	256
106	USER_DATA.MAC_CUSTOM	EFUSE_BLK3	200	48
107	KEY0	EFUSE_BLK4	0	256
108	KEY1	EFUSE_BLK5	0	256
109	KEY2	EFUSE_BLK6	0	256
110	KEY3	EFUSE_BLK7	0	256
111	KEY4	EFUSE_BLK8	0	256
112	KEY5	EFUSE_BLK9	0	256

Used bits in efuse table:

EFUSE_BLK0

[0 31] [0 0] [4 13] [15 15] [20 20] [20 20] [20 29] [32 38] [32 42] [44 52] [57-
↪78] [80 117] [119 171] [178 181]

(continues on next page)

(continued from previous page)

```

EFUSE_BLK1
[0 63]

EFUSE_BLK10
[0 255]

EFUSE_BLK2
[0 255]

EFUSE_BLK3
[0 255] [200 247]

EFUSE_BLK4
[0 255]

EFUSE_BLK5
[0 255]

EFUSE_BLK6
[0 255]

EFUSE_BLK7
[0 255]

EFUSE_BLK8
[0 255]

EFUSE_BLK9
[0 255]
Note: Not printed ranges are free for using. (bits in EFUSE_BLK0 are reserved for
↳Espressif)

```

The number of bits not included in square brackets is free (some bits are reserved for Espressif). All fields are checked for overlapping.

To add fields to an existing field, use the *Structured efuse fields* technique. For example, adding the fields: SERIAL_NUMBER, MODEL_NUMBER and HARDWARE REV to an existing USER_DATA field. Use . (dot) to show an attachment in a field.

USER_DATA.SERIAL_NUMBER,	EFUSE_BLK3,	0,	32,
USER_DATA.MODEL_NUMBER,	EFUSE_BLK3,	32,	10,
USER_DATA.HARDWARE_REV,	EFUSE_BLK3,	42,	10,

2. Fill a line for field: field_name, efuse_block, bit_start, bit_count, comment.
3. Run a show_efuse_table command to check eFuse table. To generate source files run efuse_common_table or efuse_custom_table command.

You may get errors such as intersects with or out of range. Please see how to solve them in the *Structured efuse fields* article.

Bit Order

The eFuses bit order is little endian (see the example below), it means that eFuse bits are read and written from LSB to MSB:

```

$ espefuse.py dump

USER_DATA      (BLOCK3      ) [3 ] read_regs: 03020100 07060504 0B0A0908
↳0F0E0D0C 13121111 17161514 1B1A1918 1F1E1D1C
BLOCK4         (BLOCK4      ) [4 ] read_regs: 03020100 07060504 0B0A0908
↳0F0E0D0C 13121111 17161514 1B1A1918 1F1E1D1C

```

(continues on next page)

(continued from previous page)

where is the register representation:

```
EFUSE_RD_USR_DATA0_REG = 0x03020100
EFUSE_RD_USR_DATA1_REG = 0x07060504
EFUSE_RD_USR_DATA2_REG = 0x0B0A0908
EFUSE_RD_USR_DATA3_REG = 0x0F0E0D0C
EFUSE_RD_USR_DATA4_REG = 0x13121111
EFUSE_RD_USR_DATA5_REG = 0x17161514
EFUSE_RD_USR_DATA6_REG = 0x1B1A1918
EFUSE_RD_USR_DATA7_REG = 0x1F1E1D1C
```

where is the byte representation:

```
byte[0] = 0x00, byte[1] = 0x01, ... byte[3] = 0x03, byte[4] = 0x04, ..., byte[31] =
↪ 0x1F
```

For example, csv file describes the USER_DATA field, which occupies all 256 bits (a whole block).

USER_DATA,	EFUSE_BLK3,	0,	256,	User data
USER_DATA.FIELD1,	EFUSE_BLK3,	16,	16,	Field1
ID,	EFUSE_BLK4,	8,	3,	ID bit[0..2]
,	EFUSE_BLK4,	16,	2,	ID bit[3..4]
,	EFUSE_BLK4,	32,	3,	ID bit[5..7]

Thus, reading the eFuse USER_DATA block written as above gives the following results:

```
uint8_t buf[32] = { 0 };
esp_efuse_read_field_blob(ESP_EFUSE_USER_DATA, &buf, sizeof(buf) * 8);
// buf[0] = 0x00, buf[1] = 0x01, ... buf[31] = 0x1F

uint32_t field1 = 0;
size_t field1_size = ESP_EFUSE_USER_DATA[0]->bit_count; // can be used for this_
↪ case because it only consists of one entry
esp_efuse_read_field_blob(ESP_EFUSE_USER_DATA, &field1, field1_size);
// field1 = 0x0302

uint32_t field1_1 = 0;
esp_efuse_read_field_blob(ESP_EFUSE_USER_DATA, &field1_1, 2); // reads only first_
↪ 2 bits
// field1 = 0x0002

uint8_t id = 0;
size_t id_size = esp_efuse_get_field_size(ESP_EFUSE_ID); // returns 6
// size_t id_size = ESP_EFUSE_USER_DATA[0]->bit_count; // cannot be used because_
↪ it consists of 3 entries. It returns 3 not 6.
esp_efuse_read_field_blob(ESP_EFUSE_ID, &id, id_size);
// id = 0x91
// b'100 10 001
//   [3] [2] [3]

uint8_t id_1 = 0;
esp_efuse_read_field_blob(ESP_EFUSE_ID, &id_1, 3);
// id = 0x01
// b'001
```

Get eFuses During Build

There is a way to get the state of eFuses at the build stage of the project. There are two cmake functions for this:

- `espefuse_get_json_summary()` - It calls the `espefuse.py summary --format json` command and returns a json string (it is not stored in a file).
- `espefuse_get_efuse()` - It finds a given eFuse name in the json string and returns its property.

The json string has the following properties:

```
{
  "MAC": {
    "bit_len": 48,
    "block": 0,
    "category": "identity",
    "description": "Factory MAC Address",
    "efuse_type": "bytes:6",
    "name": "MAC",
    "pos": 0,
    "readable": true,
    "value": "94:b9:7e:5a:6e:58 (CRC 0xe2 OK)",
    "word": 1,
    "writeable": true
  },
}
```

These functions can be used from a top-level project `CMakeLists.txt` ([get-started/hello_world/CMakeLists.txt](#)):

```
# ...
project(hello_world)

espefuse_get_json_summary(efuse_json)
espefuse_get_efuse(ret_data ${efuse_json} "MAC" "value")
message("MAC:" ${ret_data})
```

The format of the value property is the same as shown in `espefuse.py summary`.

```
MAC:94:b9:7e:5a:6e:58 (CRC 0xe2 OK)
```

There is an example test [system/efuse/CMakeLists.txt](#) which adds a custom target `efuse-summary`. This allows you to run the `idf.py efuse-summary` command to read the required eFuses (specified in the `efuse_names` list) at any time, not just at project build time.

Debug eFuse & Unit Tests

Virtual eFuses The Kconfig option `CONFIG_EFUSE_VIRTUAL` virtualizes eFuse values inside the eFuse Manager, so writes are emulated and no eFuse values are permanently changed. This can be useful for debugging app and unit tests. During startup, the eFuses are copied to RAM. All eFuse operations (read and write) are performed with RAM instead of the real eFuse registers.

In addition to the `CONFIG_EFUSE_VIRTUAL` option there is `CONFIG_EFUSE_VIRTUAL_KEEP_IN_FLASH` option that adds a feature to keep eFuses in flash memory. To use this mode the `partition_table` should have the `efuse` partition. `partition.csv`: `"efuse_em, data, efuse, , 0x2000, "`. During startup, the eFuses are copied from flash or, in case if flash is empty, from real eFuse to RAM and then update flash. This option allows keeping eFuses after reboots (possible to test `secure_boot` and `flash_encryption` features with this option).

Flash Encryption Testing Flash Encryption (FE) is a hardware feature that requires the physical burning of eFuses: `key` and `FLASH_CRYPT_CNT`. If FE is not actually enabled then enabling the `CONFIG_EFUSE_VIRTUAL_KEEP_IN_FLASH` option just gives testing possibilities and does not encrypt anything in the flash, even though the logs say encryption happens. The `bootloader_flash_write()` is adapted for this purpose. But if FE is already enabled on the chip and you run an application or bootloader created with the `CONFIG_EFUSE_VIRTUAL_KEEP_IN_FLASH` option then the flash encryption/decryption operations will work properly (data are encrypted as it is written into an encrypted flash partition and decrypted when they are read from an encrypted partition).

espefuse.py esptool includes a useful tool for reading/writing ESP32-P4 eFuse bits - [espefuse.py](#).

```
espefuse.py -p PORT summary

espefuse.py v4.7.dev1
Connecting....
Detecting chip type... ESP32-P4

=== Run "summary" command ===
EFUSE_NAME (Block) Description = [Meaningful Value] [Readable/Writeable] (Hex↵
↵Value)
-----
↵-----
Config fuses:
WR_DIS (BLOCK0)                               Disable programming of↵
↵individual eFuses                        = 0 R/W (0x00000000)
RD_DIS (BLOCK0)                               Disable reading from BLOCK4-10 ↵
↵                                = 0 R/W (0b00000000)
POWERGLITCH_EN (BLOCK0)                     Represents whether power glitch↵
↵function is enable = False R/W (0b0)
DIS_TWAI (BLOCK0)                           d. 1: enabled. 0: disabled
Represents whether TWAI↵
KM_HUK_GEN_STATE_LOW (BLOCK0)               abled. 1: disabled. 0: enabled
Set this bit to control↵
↵validation of HUK generate = 0 R/W (0b0000000)
mode. Odd of 1 is invalid;↵
↵even of 1 is valid
KM_HUK_GEN_STATE_HIGH (BLOCK0)             Set this bit to control↵
↵validation of HUK generate = 0 R/W (0b000)
mode. Odd of 1 is invalid;↵
↵even of 1 is valid
KM_RND_SWITCH_CYCLE (BLOCK0)              Set bits to control key manager↵
↵random number swit = 0 R/W (0b00)
ch cycle. 0: control by↵
↵register. 1: 8 km clk cycl
es. 2: 16 km cycles. 3: 32 km↵
↵cycles
KM_DEPLOY_ONLY_ONCE (BLOCK0)              Set each bit to control whether↵
↵corresponding key = 0 R/W (0x0)
can only be deployed once. 1 is↵
↵true; 0 is false.
Bit0: ecdsa. Bit1: xts. Bit2:↵
↵hmac. Bit3: ds
DIS_DIRECT_BOOT (BLOCK0)                  Represents whether direct boot↵
↵mode is disabled or = False R/W (0b0)
enabled. 1: disabled. 0:↵
↵enabled
UART_PRINT_CONTROL (BLOCK0)               Represents the type of UART↵
↵printing. 00: force en = 0 R/W (0b00)
able printing. 01: enable↵
↵printing when GPIO8 is r
eset at low level. 10: enable↵
↵printing when GPIO8
is reset at high level. 11:↵
↵force disable printing
HYS_EN_PAD (BLOCK0)                      Represents whether the↵
↵hysteresis function of corr = False R/W (0b0)
esponding PAD is enabled. 1:↵
↵enabled. 0:disabled
DCDC_VSET (BLOCK0)                       Set the dcdc voltage default ↵
↵                                = 0 R/W (0b000000)
PXA0_TIEH_SEL_0 (BLOCK0)                 TBD ↵
↵                                = 0 R/W (0b00)
(continues on next page)
```

(continued from previous page)

```

PXA0_TIEH_SEL_1 (BLOCK0)                                TBD
↪ = 0 R/W (0b00)
PXA0_TIEH_SEL_2 (BLOCK0)                                TBD
↪ = 0 R/W (0b00)
PXA0_TIEH_SEL_3 (BLOCK0)                                TBD
↪ = 0 R/W (0b00)
KM_DISABLE_DEPLOY_MODE (BLOCK0)                          TBD
↪ = 0 R/W (0x0)
HP_PWR_SRC_SEL (BLOCK0)                                  HP system power source select.
↪0:LDO. 1: DCDC = False R/W (0b0)
DCDC_VSET_EN (BLOCK0)                                    Select dcdc vset use efuse_dcdc
↪vset = False R/W (0b0)
DIS_SWD (BLOCK0)                                         Set this bit to disable super-
↪watchdog = False R/W (0b0)
BLOCK_SYS_DATA1 (BLOCK2)                                System data part 1
= 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
↪00 00 00 00 00 00 00 R/W
BLOCK_USR_DATA (BLOCK3)                                User data
= 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
↪00 00 00 00 00 00 00 R/W
BLOCK_SYS_DATA2 (BLOCK10)                               System data part 2 (reserved)
= 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
↪00 00 00 00 00 00 00 R/W

Flash fuses:
FLASH_TYPE (BLOCK0)                                    The type of interfaced flash.
↪0: four data lines; = False R/W (0b0)
1: eight data lines
FLASH_PAGE_SIZE (BLOCK0)                                Set flash page size
↪ = 0 R/W (0b00)
FLASH_ECC_EN (BLOCK0)                                    Set this bit to enable ecc for
↪flash boot = False R/W (0b0)
FLASH_TPUW (BLOCK0)                                     Represents the flash waiting
↪time after power-up; = 0 R/W (0x0)
in unit of ms. When the value
↪less than 15; the wa
iting time is the programmed
↪value. Otherwise; the
waiting time is 2 times the
↪programmed value
FORCE_SEND_RESUME (BLOCK0)                               Represents whether ROM code is
↪forced to send a re = False R/W (0b0)
sume command during SPI boot.
↪1: forced. 0: not for
ced

Jtag fuses:
JTAG_SEL_ENABLE (BLOCK0)                                Represents whether the
↪selection between usb_to_jt = False R/W (0b0)
ag and pad_to_jtag through
↪strapping gpio15 when b
oth EFUSE_DIS_PAD_JTAG and
↪EFUSE_DIS_USB_JTAG are
equal to 0 is enabled or
↪disabled. 1: enabled. 0:
disabled
SOFT_DIS_JTAG (BLOCK0)                                   Represents whether JTAG is
↪disabled in soft way. 0 = 0 R/W (0b000)
dd number: disabled. Even
↪number: enabled
DIS_PAD_JTAG (BLOCK0)                                    Represents whether JTAG is
↪disabled in the hard wa = False R/W (0b0)

```

(continues on next page)

(continued from previous page)

<code>↪enabled</code>	y(permanently). 1: disabled. 0: <code>↪</code>
Mac fuses:	
MAC (BLOCK1)	MAC address
= 00:00:00:00:00:00 (OK) R/W	
MAC_EXT (BLOCK1)	Stores the extended bits of MAC <code>↪</code>
<code>↪address</code> = 00:00 (OK) R/W	
MAC_EUI64 (BLOCK1)	calc MAC_EUI64 = <code>↪</code>
<code>↪MAC[0]:MAC[1]:MAC[2]:MAC_EXT[0]:M</code>	
= 00:00:00:00:00:00:00:00 (OK) R/W	AC_EXT[1]:MAC[3]:MAC[4]:MAC[5]
Security fuses:	
DIS_FORCE_DOWNLOAD (BLOCK0)	Represents whether the function <code>↪</code>
<code>↪that forces chip i = False R/W (0b0)</code>	nto download mode is disabled <code>↪</code>
<code>↪or enabled. 1: disab</code>	led. 0: enabled
SPI_DOWNLOAD_MSPI_DIS (BLOCK0)	Set this bit to disable <code>↪</code>
<code>↪accessing MSPI flash/MSPI = False R/W (0b0)</code>	ram by SYS AXI matrix during <code>↪</code>
<code>↪boot_mode_download</code>	
DIS_DOWNLOAD_MANUAL_ENCRYPT (BLOCK0)	Represents whether flash <code>↪</code>
<code>↪encrypt function is disab = False R/W (0b0)</code>	led or enabled(except in SPI <code>↪</code>
<code>↪boot mode). 1: disabl</code>	ed. 0: enabled
FORCE_USE_KEY_MANAGER_KEY (BLOCK0)	Set each bit to control whether <code>↪</code>
<code>↪corresponding key = 0 R/W (0x0)</code>	must come from key manager.. 1 <code>↪</code>
<code>↪is true; 0 is false</code>	. Bit0: ecdsa. Bit1: xts. Bit2: <code>↪</code>
<code>↪hmac. Bit3: ds</code>	
FORCE_DISABLE_SW_INIT_KEY (BLOCK0)	Set this bit to disable <code>↪</code>
<code>↪software written init key; = False R/W (0b0)</code>	and force use efuse_init_key
XTS_KEY_LENGTH_256 (BLOCK0)	Set this bit to configure flash <code>↪</code>
<code>↪encryption use xts = False R/W (0b0)</code>	-128 key; else use xts-256 key
SPI_BOOT_CRYPT_CNT (BLOCK0)	Enables flash encryption when 1 <code>↪</code>
<code>↪or 3 bits are set = Disable R/W (0b000)</code>	and disables otherwise
SECURE_BOOT_KEY_REVOKE0 (BLOCK0)	Revoke 1st secure boot key <code>↪</code>
<code>↪</code> = False R/W (0b0)	
SECURE_BOOT_KEY_REVOKE1 (BLOCK0)	Revoke 2nd secure boot key <code>↪</code>
<code>↪</code> = False R/W (0b0)	
SECURE_BOOT_KEY_REVOKE2 (BLOCK0)	Revoke 3rd secure boot key <code>↪</code>
<code>↪</code> = False R/W (0b0)	
KEY_PURPOSE_0 (BLOCK0)	Represents the purpose of Key0 <code>↪</code>
<code>↪</code> = USER R/W (0x0)	
KEY_PURPOSE_1 (BLOCK0)	Represents the purpose of Key1 <code>↪</code>
<code>↪</code> = USER R/W (0x0)	
KEY_PURPOSE_2 (BLOCK0)	Represents the purpose of Key2 <code>↪</code>
<code>↪</code> = USER R/W (0x0)	
KEY_PURPOSE_3 (BLOCK0)	Represents the purpose of Key3 <code>↪</code>
<code>↪</code> = USER R/W (0x0)	
KEY_PURPOSE_4 (BLOCK0)	Represents the purpose of Key4 <code>↪</code>
<code>↪</code> = USER R/W (0x0)	
KEY_PURPOSE_5 (BLOCK0)	Represents the purpose of Key5 <code>↪</code>
<code>↪</code> = USER R/W (0x0)	

(continues on next page)

[illegible]

(continued from previous page)

```
Purpose: USER
           Key5 or user data
= 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 R/W
```

```
Usb fuses:
```

```
USB_DEVICE_EXCHG_PINS (BLOCK0) Enable usb device exchange pins
↳of D+ and D- = False R/W (0b0)
```

```
USB_OTG11_EXCHG_PINS (BLOCK0) Enable usb otg11 exchange pins
↳of D+ and D- = False R/W (0b0)
```

```
DIS_USB_JTAG (BLOCK0) Represents whether the function
↳of usb switch to j = False R/W (0b0)
tag is disabled or enabled. 1: disabled. 0: enable
```

```
USB_PHY_SEL (BLOCK0) TBD
↳ = False R/W (0b0)
```

```
DIS_USB_OTG_DOWNLOAD_MODE (BLOCK0) Set this bit to disable
↳download via USB-OTG = False R/W (0b0)
```

```
DIS_USB_SERIAL_JTAG_ROM_PRINT (BLOCK0) Represents whether print from
↳USB-Serial-JTAG is d = False R/W (0b0)
isabled or enabled. 1: disabled.
```

```
DIS_USB_SERIAL_JTAG_DOWNLOAD_MODE (BLOCK0) Represents whether the USB-
↳Serial-JTAG download fu = False R/W (0b0)
nction is disabled or enabled. 1: disabled. 0: enabled
```

```
Wdt fuses:
```

```
WDT_DELAY_SEL (BLOCK0) Represents whether RTC watchdog
↳timeout threshold = 0 R/W (0b00)
is selected at startup. 1: disabled. 0: not select
```

```
DIS_WDT (BLOCK0) ed
Set this bit to disable watchdoq = False R/W (0b0)
```

To get a dump for all eFuse registers.

```

espefuse.py v4.7.dev1
Connecting....
Detecting chip type... ESP32-P4
BLOCK0          (          ) [0 ] read_regs: 00000000 00000000 00000000
↳00000000 00000000 00000000
MAC_SPI_8M_0    (BLOCK1    ) [1 ] read_regs: 00000000 00000000 00000000
↳00000000 00000000 00000000
BLOCK_SYS_DATA  (BLOCK2    ) [2 ] read_regs: 00000000 00000000 00000000
↳00000000 00000000 00000000 00000000 00000000
BLOCK_USR_DATA  (BLOCK3    ) [3 ] read_regs: 00000000 00000000 00000000
↳00000000 00000000 00000000 00000000 00000000
BLOCK_KEY0      (BLOCK4    ) [4 ] read_regs: 00000000 00000000 00000000
↳00000000 00000000 00000000 00000000 00000000
BLOCK_KEY1      (BLOCK5    ) [5 ] read_regs: 00000000 00000000 00000000
↳00000000 00000000 00000000 00000000 00000000
BLOCK_KEY2      (BLOCK6    ) [6 ] read_regs: 00000000 00000000 00000000
↳00000000 00000000 00000000 00000000 00000000
BLOCK_KEY3      (BLOCK7    ) [7 ] read_regs: 00000000 00000000 00000000
↳00000000 00000000 00000000 00000000 00000000
BLOCK_KEY4      (BLOCK8    ) [8 ] read_regs: 00000000 00000000 00000000
↳00000000 00000000 00000000 00000000 00000000

```

(continues on next page)

(continued from previous page)

```

BLOCK_KEY5      (BLOCK9      ) [9 ] read_regs: 00000000 00000000 00000000_
↪00000000 00000000 00000000 00000000 00000000
BLOCK_SYS_DATA2 (BLOCK10     ) [10] read_regs: 00000000 00000000 00000000_
↪00000000 00000000 00000000 00000000 00000000
BLOCK0          (            ) [0 ] err__regs: 00000000 00000000 00000000_
↪00000000 00000000 00000000
EFUSE_RD_RS_ERR0_REG      0x00000000
EFUSE_RD_RS_ERR1_REG      0x00000000
=== Run "dump" command ===

```

Header File

- [components/efuse/esp32p4/include/esp_efuse_chip.h](#)
- This header file can be included with:

```
#include "esp_efuse_chip.h"
```

- This header file is a part of the API provided by the efuse component. To declare that your component depends on efuse, add the following to your CMakeLists.txt:

```
REQUIRES efuse
```

or

```
PRIV_REQUIRES efuse
```

Enumerations

enum **esp_efuse_block_t**

Type of eFuse blocks ESP32P4.

Values:

enumerator **EFUSE_BLK0**

Number of eFuse BLOCK0. REPEAT_DATA

enumerator **EFUSE_BLK1**

Number of eFuse BLOCK1. MAC_SPI_8M_SYS

enumerator **EFUSE_BLK2**

Number of eFuse BLOCK2. SYS_DATA_PART1

enumerator **EFUSE_BLK_SYS_DATA_PART1**

Number of eFuse BLOCK2. SYS_DATA_PART1

enumerator **EFUSE_BLK3**

Number of eFuse BLOCK3. USER_DATA

enumerator **EFUSE_BLK_USER_DATA**

Number of eFuse BLOCK3. USER_DATA

enumerator **EFUSE_BLK4**

Number of eFuse BLOCK4. KEY0

enumerator **EFUSE_BLK_KEY0**

Number of eFuse BLOCK4. KEY0

enumerator **EFUSE_BLK5**

Number of eFuse BLOCK5. KEY1

enumerator **EFUSE_BLK_KEY1**

Number of eFuse BLOCK5. KEY1

enumerator **EFUSE_BLK6**

Number of eFuse BLOCK6. KEY2

enumerator **EFUSE_BLK_KEY2**

Number of eFuse BLOCK6. KEY2

enumerator **EFUSE_BLK7**

Number of eFuse BLOCK7. KEY3

enumerator **EFUSE_BLK_KEY3**

Number of eFuse BLOCK7. KEY3

enumerator **EFUSE_BLK8**

Number of eFuse BLOCK8. KEY4

enumerator **EFUSE_BLK_KEY4**

Number of eFuse BLOCK8. KEY4

enumerator **EFUSE_BLK9**

Number of eFuse BLOCK9. KEY5

enumerator **EFUSE_BLK_KEY5**

Number of eFuse BLOCK9. KEY5

enumerator **EFUSE_BLK_KEY_MAX**

enumerator **EFUSE_BLK10**

Number of eFuse BLOCK10. SYS_DATA_PART2

enumerator **EFUSE_BLK_SYS_DATA_PART2**

Number of eFuse BLOCK10. SYS_DATA_PART2

enumerator **EFUSE_BLK_MAX**

enum **esp_efuse_coding_scheme_t**

Type of coding scheme.

Values:

enumerator **EFUSE_CODING_SCHEME_NONE**

None

enumerator **EFUSE_CODING_SCHEME_RS**

Reed-Solomon coding

enum **esp_efuse_purpose_t**

Type of key purpose.

Values:

enumerator **ESP_EFUSE_KEY_PURPOSE_USER**

User purposes (software-only use)

enumerator **ESP_EFUSE_KEY_PURPOSE_ECDSA_KEY**

ECDSA private key (Expected in little endian order)

enumerator **ESP_EFUSE_KEY_PURPOSE_XTS_AES_256_KEY_1**

XTS_AES_256_KEY_1 (flash/PSRAM encryption)

enumerator **ESP_EFUSE_KEY_PURPOSE_XTS_AES_256_KEY_2**

XTS_AES_256_KEY_2 (flash/PSRAM encryption)

enumerator **ESP_EFUSE_KEY_PURPOSE_XTS_AES_128_KEY**

XTS_AES_128_KEY (flash/PSRAM encryption)

enumerator **ESP_EFUSE_KEY_PURPOSE_HMAC_DOWN_ALL**

HMAC Downstream mode

enumerator **ESP_EFUSE_KEY_PURPOSE_HMAC_DOWN_JTAG**

JTAG soft enable key (uses HMAC Downstream mode)

enumerator **ESP_EFUSE_KEY_PURPOSE_HMAC_DOWN_DIGITAL_SIGNATURE**

Digital Signature peripheral key (uses HMAC Downstream mode)

enumerator **ESP_EFUSE_KEY_PURPOSE_HMAC_UP**

HMAC Upstream mode

enumerator **ESP_EFUSE_KEY_PURPOSE_SECURE_BOOT_DIGEST0**

SECURE_BOOT_DIGEST0 (Secure Boot key digest)

enumerator **ESP_EFUSE_KEY_PURPOSE_SECURE_BOOT_DIGEST1**

SECURE_BOOT_DIGEST1 (Secure Boot key digest)

enumerator **ESP_EFUSE_KEY_PURPOSE_SECURE_BOOT_DIGEST2**

SECURE_BOOT_DIGEST2 (Secure Boot key digest)

enumerator **ESP_EFUSE_KEY_PURPOSE_KM_INIT_KEY**

KM_INIT_KEY

enumerator **ESP_EFUSE_KEY_PURPOSE_MAX**

MAX PURPOSE

Header File

- [components/efuse/include/esp_efuse.h](#)
- This header file can be included with:

```
#include "esp_efuse.h"
```

- This header file is a part of the API provided by the `efuse` component. To declare that your component depends on `efuse`, add the following to your `CMakeLists.txt`:

```
REQUIRES efuse
```

or

```
PRIV_REQUIRES efuse
```

Functions

esp_err_t **esp_efuse_read_field_blob** (const *esp_efuse_desc_t* *field[], void *dst, size_t dst_size_bits)

Reads bits from EFUSE field and writes it into an array.

The number of read bits will be limited to the minimum value from the description of the bits in "field" structure or "dst_size_bits" required size. Use "esp_efuse_get_field_size()" function to determine the length of the field.

Note: Please note that reading in the batch mode does not show uncommitted changes.

Parameters

- **field** -- **[in]** A pointer to the structure describing the fields of efuse.
- **dst** -- **[out]** A pointer to array that will contain the result of reading.
- **dst_size_bits** -- **[in]** The number of bits required to read. If the requested number of bits is greater than the field, the number will be limited to the field size.

Returns

- `ESP_OK`: The operation was successfully completed.
- `ESP_ERR_INVALID_ARG`: Error in the passed arguments.

bool **esp_efuse_read_field_bit** (const *esp_efuse_desc_t* *field[])

Read a single bit eFuse field as a boolean value.

Note: The value must exist and must be a single bit wide. If there is any possibility of an error in the provided arguments, call `esp_efuse_read_field_blob()` and check the returned value instead.

Note: If assertions are enabled and the parameter is invalid, execution will abort

Note: Please note that reading in the batch mode does not show uncommitted changes.

Parameters **field** -- **[in]** A pointer to the structure describing the fields of efuse.

Returns

- `true`: The field parameter is valid and the bit is set.
- `false`: The bit is not set, or the parameter is invalid and assertions are disabled.

esp_err_t **esp_efuse_read_field_cnt** (const *esp_efuse_desc_t* *field[], size_t *out_cnt)

Reads bits from EFUSE field and returns number of bits programmed as "1".

If the bits are set not sequentially, they will still be counted.

Note: Please note that reading in the batch mode does not show uncommitted changes.

Parameters

- **field** -- [in] A pointer to the structure describing the fields of efuse.
- **out_cnt** -- [out] A pointer that will contain the number of programmed as "1" bits.

Returns

- ESP_OK: The operation was successfully completed.
- ESP_ERR_INVALID_ARG: Error in the passed arguments.

esp_err_t **esp_efuse_write_field_blob** (const *esp_efuse_desc_t* *field[], const void *src, size_t src_size_bits)

Writes array to EFUSE field.

The number of write bits will be limited to the minimum value from the description of the bits in "field" structure or "src_size_bits" required size. Use "esp_efuse_get_field_size()" function to determine the length of the field. After the function is completed, the writing registers are cleared.

Parameters

- **field** -- [in] A pointer to the structure describing the fields of efuse.
- **src** -- [in] A pointer to array that contains the data for writing.
- **src_size_bits** -- [in] The number of bits required to write.

Returns

- ESP_OK: The operation was successfully completed.
- ESP_ERR_INVALID_ARG: Error in the passed arguments.
- ESP_ERR_EFUSE_REPEATED_PROG: Error repeated programming of programmed bits is strictly forbidden.
- ESP_ERR_CODING: Error range of data does not match the coding scheme.

esp_err_t **esp_efuse_write_field_cnt** (const *esp_efuse_desc_t* *field[], size_t cnt)

Writes a required count of bits as "1" to EFUSE field.

If there are no free bits in the field to set the required number of bits to "1", ESP_ERR_EFUSE_CNT_IS_FULL error is returned, the field will not be partially recorded. After the function is completed, the writing registers are cleared.

Parameters

- **field** -- [in] A pointer to the structure describing the fields of efuse.
- **cnt** -- [in] Required number of programmed as "1" bits.

Returns

- ESP_OK: The operation was successfully completed.
- ESP_ERR_INVALID_ARG: Error in the passed arguments.
- ESP_ERR_EFUSE_CNT_IS_FULL: Not all requested cnt bits is set.

esp_err_t **esp_efuse_write_field_bit** (const *esp_efuse_desc_t* *field[])

Write a single bit eFuse field to 1.

For use with eFuse fields that are a single bit. This function will write the bit to value 1 if it is not already set, or does nothing if the bit is already set.

This is equivalent to calling esp_efuse_write_field_cnt() with the cnt parameter equal to 1, except that it will return ESP_OK if the field is already set to 1.

Parameters **field** -- [in] Pointer to the structure describing the efuse field.

Returns

- ESP_OK: The operation was successfully completed, or the bit was already set to value 1.
- ESP_ERR_INVALID_ARG: Error in the passed arguments, including if the efuse field is not 1 bit wide.

esp_err_t **esp_efuse_set_write_protect** (*esp_efuse_block_t* blk)

Sets a write protection for the whole block.

After that, it is impossible to write to this block. The write protection does not apply to block 0.

Parameters **blk** -- [in] Block number of eFuse. (EFUSE_BLK1, EFUSE_BLK2 and EFUSE_BLK3)

Returns

- ESP_OK: The operation was successfully completed.
- ESP_ERR_INVALID_ARG: Error in the passed arguments.
- ESP_ERR_EFUSE_CNT_IS_FULL: Not all requested cnt bits is set.
- ESP_ERR_NOT_SUPPORTED: The block does not support this command.

esp_err_t **esp_efuse_set_read_protect** (*esp_efuse_block_t* blk)

Sets a read protection for the whole block.

After that, it is impossible to read from this block. The read protection does not apply to block 0.

Parameters **blk** -- [in] Block number of eFuse. (EFUSE_BLK1, EFUSE_BLK2 and EFUSE_BLK3)

Returns

- ESP_OK: The operation was successfully completed.
- ESP_ERR_INVALID_ARG: Error in the passed arguments.
- ESP_ERR_EFUSE_CNT_IS_FULL: Not all requested cnt bits is set.
- ESP_ERR_NOT_SUPPORTED: The block does not support this command.

int **esp_efuse_get_field_size** (const *esp_efuse_desc_t* *field[])

Returns the number of bits used by field.

Parameters **field** -- [in] A pointer to the structure describing the fields of efuse.

Returns Returns the number of bits used by field.

uint32_t **esp_efuse_read_reg** (*esp_efuse_block_t* blk, unsigned int num_reg)

Returns value of efuse register.

This is a thread-safe implementation. Example: EFUSE_BLK2_RDATA3_REG where (blk=2, num_reg=3)

Note: Please note that reading in the batch mode does not show uncommitted changes.

Parameters

- **blk** -- [in] Block number of eFuse.
- **num_reg** -- [in] The register number in the block.

Returns Value of register

esp_err_t **esp_efuse_write_reg** (*esp_efuse_block_t* blk, unsigned int num_reg, uint32_t val)

Write value to efuse register.

Apply a coding scheme if necessary. This is a thread-safe implementation. Example: EFUSE_BLK3_WDATA0_REG where (blk=3, num_reg=0)

Parameters

- **blk** -- [in] Block number of eFuse.
- **num_reg** -- [in] The register number in the block.
- **val** -- [in] Value to write.

Returns

- ESP_OK: The operation was successfully completed.
- ESP_ERR_EFUSE_REPEATED_PROG: Error repeated programming of programmed bits is strictly forbidden.

esp_efuse_coding_scheme_t **esp_efuse_get_coding_scheme** (*esp_efuse_block_t* blk)

Return efuse coding scheme for blocks.

Note: The coding scheme is applicable only to 1, 2 and 3 blocks. For 0 block, the coding scheme is always NONE.

Parameters **blk** -- [in] Block number of eFuse.

Returns Return efuse coding scheme for blocks

esp_err_t **esp_efuse_read_block** (*esp_efuse_block_t* blk, void *dst_key, size_t offset_in_bits, size_t size_bits)

Read key to efuse block starting at the offset and the required size.

Note: Please note that reading in the batch mode does not show uncommitted changes.

Parameters

- **blk** -- [in] Block number of eFuse.
- **dst_key** -- [in] A pointer to array that will contain the result of reading.
- **offset_in_bits** -- [in] Start bit in block.
- **size_bits** -- [in] The number of bits required to read.

Returns

- ESP_OK: The operation was successfully completed.
- ESP_ERR_INVALID_ARG: Error in the passed arguments.
- ESP_ERR_CODING: Error range of data does not match the coding scheme.

esp_err_t **esp_efuse_write_block** (*esp_efuse_block_t* blk, const void *src_key, size_t offset_in_bits, size_t size_bits)

Write key to efuse block starting at the offset and the required size.

Parameters

- **blk** -- [in] Block number of eFuse.
- **src_key** -- [in] A pointer to array that contains the key for writing.
- **offset_in_bits** -- [in] Start bit in block.
- **size_bits** -- [in] The number of bits required to write.

Returns

- ESP_OK: The operation was successfully completed.
- ESP_ERR_INVALID_ARG: Error in the passed arguments.
- ESP_ERR_CODING: Error range of data does not match the coding scheme.
- ESP_ERR_EFUSE_REPEATED_PROG: Error repeated programming of programmed bits

uint32_t **esp_efuse_get_pkg_ver** (void)

Returns chip package from efuse.

Returns chip package

void **esp_efuse_reset** (void)

Reset efuse write registers.

Efuse write registers are written to zero, to negate any changes that have been staged here.

Note: This function is not threadsafe, if calling code updates efuse values from multiple tasks then this is caller's responsibility to serialise.

esp_err_t **esp_efuse_disable_rom_download_mode** (void)

Disable ROM Download Mode via eFuse.

Permanently disables the ROM Download Mode feature. Once disabled, if the SoC is booted with strapping pins set for ROM Download Mode then an error is printed instead.

Note: Not all SoCs support this option. An error will be returned if called on an ESP32 with a silicon revision lower than 3, as these revisions do not support this option.

Note: If ROM Download Mode is already disabled, this function does nothing and returns success.

Returns

- ESP_OK If the eFuse was successfully burned, or had already been burned.
- ESP_ERR_NOT_SUPPORTED (ESP32 only) This SoC is not capable of disabling UART download mode
- ESP_ERR_INVALID_STATE (ESP32 only) This eFuse is write protected and cannot be written

esp_err_t **esp_efuse_set_rom_log_scheme** (*esp_efuse_rom_log_scheme_t* log_scheme)

Set boot ROM log scheme via eFuse.

Note: By default, the boot ROM will always print to console. This API can be called to set the log scheme only once per chip, once the value is changed from the default it can't be changed again.

Parameters **log_scheme** -- Supported ROM log scheme

Returns

- ESP_OK If the eFuse was successfully burned, or had already been burned.
- ESP_ERR_NOT_SUPPORTED (ESP32 only) This SoC is not capable of setting ROM log scheme
- ESP_ERR_INVALID_STATE This eFuse is write protected or has been burned already

esp_err_t **esp_efuse_enable_rom_secure_download_mode** (void)

Switch ROM Download Mode to Secure Download mode via eFuse.

Permanently enables Secure Download mode. This mode limits the use of ROM Download Mode functions to simple flash read, write and erase operations, plus a command to return a summary of currently enabled security features.

Note: If Secure Download mode is already enabled, this function does nothing and returns success.

Note: If UART DL mode is completely disabled then Secure Download mode can not be enabled and this API simply returns success.

Returns

- ESP_OK If the eFuse was successfully burned, or had already been burned, or UART DL mode is already disabled.
- Other errors If an error occurred while burning ESP_EFUSE_ENABLE_SECURITY_DOWNLOAD.

uint32_t **esp_efuse_read_secure_version** (void)

Return secure_version from efuse field.

Returns Secure version from efuse field

bool **esp_efuse_check_secure_version** (uint32_t secure_version)

Check secure_version from app and secure_version and from efuse field.

Parameters **secure_version** -- Secure version from app.

Returns

- True: If version of app is equal or more then secure_version from efuse.

esp_err_t **esp_efuse_update_secure_version** (uint32_t secure_version)

Write efuse field by secure_version value.

Update the secure_version value is available if the coding scheme is None. Note: Do not use this function in your applications. This function is called as part of the other API.

Parameters **secure_version** -- [in] Secure version from app.

Returns

- ESP_OK: Successful.
- ESP_FAIL: secure version of app cannot be set to efuse field.
- ESP_ERR_NOT_SUPPORTED: Anti rollback is not supported with the 3/4 and Repeat coding scheme.

esp_err_t **esp_efuse_batch_write_begin** (void)

Set the batch mode of writing fields.

This mode allows you to write the fields in the batch mode when need to burn several efuses at one time. To enable batch mode call begin() then perform as usually the necessary operations read and write and at the end call commit() to actually burn all written efuses. The batch mode can be used nested. The commit will be done by the last commit() function. The number of begin() functions should be equal to the number of commit() functions.

Note: If batch mode is enabled by the first task, at this time the second task cannot write/read efuses. The second task will wait for the first task to complete the batch operation.

```
// Example of using the batch writing mode.

// set the batch writing mode
esp_efuse_batch_write_begin();

// use any writing functions as usual
esp_efuse_write_field_blob(ESP_EFUSE_...);
esp_efuse_write_field_cnt(ESP_EFUSE_...);
esp_efuse_set_write_protect(EFUSE_BLKx);
esp_efuse_write_reg(EFUSE_BLKx, ...);
esp_efuse_write_block(EFUSE_BLKx, ...);
esp_efuse_write(ESP_EFUSE_1, 3); // ESP_EFUSE_1 == 1, here we write a new
↪value = 3. The changes will be burn by the commit() function.
esp_efuse_read...(ESP_EFUSE_1); // this function returns ESP_EFUSE_1 == 1
↪because uncommitted changes are not readable, it will be available only
↪after commit.
...

// esp_efuse_batch_write APIs can be called recursively.
esp_efuse_batch_write_begin();
esp_efuse_set_write_protect(EFUSE_BLKx);
esp_efuse_batch_write_commit(); // the burn will be skipped here, it will be
↪done in the last commit().

...

// Write all of these fields to the efuse registers
esp_efuse_batch_write_commit();
esp_efuse_read...(ESP_EFUSE_1); // this function returns ESP_EFUSE_1 == 3.
```

Note: Please note that reading in the batch mode does not show uncommitted changes.

Returns

- ESP_OK: Successful.

esp_err_t **esp_efuse_batch_write_cancel** (void)

Reset the batch mode of writing fields.

It will reset the batch writing mode and any written changes.

Returns

- ESP_OK: Successful.
- ESP_ERR_INVALID_STATE: The batch mode was not set.

esp_err_t **esp_efuse_batch_write_commit** (void)

Writes all prepared data for the batch mode.

Must be called to ensure changes are written to the efuse registers. After this the batch writing mode will be reset.

Returns

- ESP_OK: Successful.
- ESP_ERR_INVALID_STATE: The deferred writing mode was not set.

bool **esp_efuse_block_is_empty** (*esp_efuse_block_t* block)

Checks that the given block is empty.

Returns

- True: The block is empty.
- False: The block is not empty or was an error.

bool **esp_efuse_get_key_dis_read** (*esp_efuse_block_t* block)

Returns a read protection for the key block.

Parameters **block** -- [in] A key block in the range EFUSE_BLK_KEY0..EFUSE_BLK_KEY_MAX

Returns True: The key block is read protected False: The key block is readable.

esp_err_t **esp_efuse_set_key_dis_read** (*esp_efuse_block_t* block)

Sets a read protection for the key block.

Parameters **block** -- [in] A key block in the range EFUSE_BLK_KEY0..EFUSE_BLK_KEY_MAX

Returns

- ESP_OK: Successful.
- ESP_ERR_INVALID_ARG: Error in the passed arguments.
- ESP_ERR_EFUSE_REPEATED_PROG: Error repeated programming of programmed bits is strictly forbidden.
- ESP_ERR_CODING: Error range of data does not match the coding scheme.

bool **esp_efuse_get_key_dis_write** (*esp_efuse_block_t* block)

Returns a write protection for the key block.

Parameters **block** -- [in] A key block in the range EFUSE_BLK_KEY0..EFUSE_BLK_KEY_MAX

Returns True: The key block is write protected False: The key block is writeable.

esp_err_t **esp_efuse_set_key_dis_write** (*esp_efuse_block_t* block)

Sets a write protection for the key block.

Parameters **block** -- [in] A key block in the range EFUSE_BLK_KEY0..EFUSE_BLK_KEY_MAX

Returns

- ESP_OK: Successful.
- ESP_ERR_INVALID_ARG: Error in the passed arguments.

- `ESP_ERR_EFUSE_REPEATED_PROG`: Error repeated programming of programmed bits is strictly forbidden.
- `ESP_ERR_CODING`: Error range of data does not match the coding scheme.

bool **esp_efuse_key_block_unused** (*esp_efuse_block_t* block)

Returns true if the key block is unused, false otherwise.

An unused key block is all zero content, not read or write protected, and has purpose 0 (`ESP_EFUSE_KEY_PURPOSE_USER`)

Parameters **block** -- key block to check.

Returns

- True if key block is unused,
- False if key block is used or the specified block index is not a key block.

bool **esp_efuse_find_purpose** (*esp_efuse_purpose_t* purpose, *esp_efuse_block_t* *block)

Find a key block with the particular purpose set.

Parameters

- **purpose** -- [in] Purpose to search for.
- **block** -- [out] Pointer in the range `EFUSE_BLK_KEY0..EFUSE_BLK_KEY_MAX` which will be set to the key block if found. Can be NULL, if only need to test the key block exists.

Returns

- True: If found,
- False: If not found (value at block pointer is unchanged).

bool **esp_efuse_get_keypurpose_dis_write** (*esp_efuse_block_t* block)

Returns a write protection of the key purpose field for an efuse key block.

Note: For ESP32: no keypurpose, it returns always True.

Parameters **block** -- [in] A key block in the range `EFUSE_BLK_KEY0..EFUSE_BLK_KEY_MAX`

Returns True: The key purpose is write protected. False: The key purpose is writeable.

esp_efuse_purpose_t **esp_efuse_get_key_purpose** (*esp_efuse_block_t* block)

Returns the current purpose set for an efuse key block.

Parameters **block** -- [in] A key block in the range `EFUSE_BLK_KEY0..EFUSE_BLK_KEY_MAX`

Returns

- Value: If Successful, it returns the value of the purpose related to the given key block.
- `ESP_EFUSE_KEY_PURPOSE_MAX`: Otherwise.

const *esp_efuse_desc_t* ****esp_efuse_get_purpose_field** (*esp_efuse_block_t* block)

Returns a pointer to a key purpose for an efuse key block.

To get the value of this field use `esp_efuse_read_field_blob()` or `esp_efuse_get_key_purpose()`.

Parameters **block** -- [in] A key block in the range `EFUSE_BLK_KEY0..EFUSE_BLK_KEY_MAX`

Returns Pointer: If Successful returns a pointer to the corresponding efuse field otherwise NULL.

const *esp_efuse_desc_t* ****esp_efuse_get_key** (*esp_efuse_block_t* block)

Returns a pointer to a key block.

Parameters **block** -- [in] A key block in the range `EFUSE_BLK_KEY0..EFUSE_BLK_KEY_MAX`

Returns Pointer: If Successful returns a pointer to the corresponding efuse field otherwise NULL.

esp_err_t **esp_efuse_set_key_purpose** (*esp_efuse_block_t* block, *esp_efuse_purpose_t* purpose)

Sets a key purpose for an efuse key block.

Parameters

- **block** -- [in] A key block in the range EFUSE_BLK_KEY0..EFUSE_BLK_KEY_MAX
- **purpose** -- [in] Key purpose.

Returns

- ESP_OK: Successful.
- ESP_ERR_INVALID_ARG: Error in the passed arguments.
- ESP_ERR_EFUSE_REPEATED_PROG: Error repeated programming of programmed bits is strictly forbidden.
- ESP_ERR_CODING: Error range of data does not match the coding scheme.

esp_err_t **esp_efuse_set_keypurpose_dis_write** (*esp_efuse_block_t* block)

Sets a write protection of the key purpose field for an efuse key block.

Parameters **block** -- [in] A key block in the range EFUSE_BLK_KEY0..EFUSE_BLK_KEY_MAX

Returns

- ESP_OK: Successful.
- ESP_ERR_INVALID_ARG: Error in the passed arguments.
- ESP_ERR_EFUSE_REPEATED_PROG: Error repeated programming of programmed bits is strictly forbidden.
- ESP_ERR_CODING: Error range of data does not match the coding scheme.

esp_efuse_block_t **esp_efuse_find_unused_key_block** (void)

Search for an unused key block and return the first one found.

See esp_efuse_key_block_unused for a description of an unused key block.

Returns First unused key block, or EFUSE_BLK_KEY_MAX if no unused key block is found.

unsigned **esp_efuse_count_unused_key_blocks** (void)

Return the number of unused efuse key blocks in the range EFUSE_BLK_KEY0..EFUSE_BLK_KEY_MAX.

bool **esp_efuse_get_digest_revoke** (unsigned num_digest)

Returns the status of the Secure Boot public key digest revocation bit.

Parameters **num_digest** -- [in] The number of digest in range 0..2

Returns

- True: If key digest is revoked,
- False; If key digest is not revoked.

esp_err_t **esp_efuse_set_digest_revoke** (unsigned num_digest)

Sets the Secure Boot public key digest revocation bit.

Parameters **num_digest** -- [in] The number of digest in range 0..2

Returns

- ESP_OK: Successful.
- ESP_ERR_INVALID_ARG: Error in the passed arguments.
- ESP_ERR_EFUSE_REPEATED_PROG: Error repeated programming of programmed bits is strictly forbidden.
- ESP_ERR_CODING: Error range of data does not match the coding scheme.

bool **esp_efuse_get_write_protect_of_digest_revoke** (unsigned num_digest)

Returns a write protection of the Secure Boot public key digest revocation bit.

Parameters **num_digest** -- [in] The number of digest in range 0..2

Returns True: The revocation bit is write protected. False: The revocation bit is writeable.

esp_err_t **esp_efuse_set_write_protect_of_digest_revoke** (unsigned num_digest)

Sets a write protection of the Secure Boot public key digest revocation bit.

Parameters **num_digest** -- [in] The number of digest in range 0..2

Returns

- ESP_OK: Successful.
- ESP_ERR_INVALID_ARG: Error in the passed arguments.
- ESP_ERR_EFUSE_REPEATED_PROG: Error repeated programming of programmed bits is strictly forbidden.

- **ESP_ERR_CODING**: Error range of data does not match the coding scheme.

esp_err_t **esp_efuse_write_key** (*esp_efuse_block_t* block, *esp_efuse_purpose_t* purpose, const void *key, size_t key_size_bytes)

Program a block of key data to an efuse block.

The burn of a key, protection bits, and a purpose happens in batch mode.

Note: This API also enables the read protection efuse bit for certain key blocks like XTS-AES, HMAC, ECDSA etc. This ensures that the key is only accessible to hardware peripheral.

Note: For SoC's with capability `SOC_EFUSE_ECDSA_USE_HARDWARE_K` (e.g., ESP32-H2), this API writes an additional efuse bit for ECDSA key purpose to enforce hardware TRNG generated k mode in the peripheral.

Parameters

- **block** -- **[in]** Block to read purpose for. Must be in range `EFUSE_BLK_KEY0` to `EFUSE_BLK_KEY_MAX`. Key block must be unused (`esp_efuse_key_block_unused`).
- **purpose** -- **[in]** Purpose to set for this key. Purpose must be already unset.
- **key** -- **[in]** Pointer to data to write.
- **key_size_bytes** -- **[in]** Bytes length of data to write.

Returns

- **ESP_OK**: Successful.
- **ESP_ERR_INVALID_ARG**: Error in the passed arguments.
- **ESP_ERR_INVALID_STATE**: Error in efuses state, unused block not found.
- **ESP_ERR_EFUSE_REPEATED_PROG**: Error repeated programming of programmed bits is strictly forbidden.
- **ESP_ERR_CODING**: Error range of data does not match the coding scheme.

esp_err_t **esp_efuse_write_keys** (const *esp_efuse_purpose_t* purposes[], uint8_t keys[][32], unsigned number_of_keys)

Program keys to unused efuse blocks.

The burn of keys, protection bits, and purposes happens in batch mode.

Note: This API also enables the read protection efuse bit for certain key blocks like XTS-AES, HMAC, ECDSA etc. This ensures that the key is only accessible to hardware peripheral.

Note: For SoC's with capability `SOC_EFUSE_ECDSA_USE_HARDWARE_K` (e.g., ESP32-H2), this API writes an additional efuse bit for ECDSA key purpose to enforce hardware TRNG generated k mode in the peripheral.

Parameters

- **purposes** -- **[in]** Array of purposes (`purpose[number_of_keys]`).
- **keys** -- **[in]** Array of keys (`uint8_t keys[number_of_keys][32]`). Each key is 32 bytes long.
- **number_of_keys** -- **[in]** The number of keys to write (up to 6 keys).

Returns

- **ESP_OK**: Successful.
- **ESP_ERR_INVALID_ARG**: Error in the passed arguments.
- **ESP_ERR_INVALID_STATE**: Error in efuses state, unused block not found.
- **ESP_ERR_NOT_ENOUGH_UNUSED_KEY_BLOCKS**: Error not enough unused key blocks available

- `ESP_ERR_EFUSE_REPEATED_PROG`: Error repeated programming of programmed bits is strictly forbidden.
- `ESP_ERR_CODING`: Error range of data does not match the coding scheme.

esp_err_t **esp_secure_boot_read_key_digests** (*esp_secure_boot_key_digests_t* *trusted_key_digests)

Read key digests from efuse. Any revoked/missing digests will be marked as NULL.

Parameters `trusted_key_digests` -- [out] Trusted keys digests, stored in this parameter after successfully completing this function. The number of digests depends on the SOC's capabilities.

Returns

- `ESP_OK`: Successful.
- `ESP_FAIL`: If `trusted_keys` is NULL or there is no valid digest.

esp_err_t **esp_efuse_check_errors** (void)

Checks eFuse errors in BLOCK0.

It does a BLOCK0 check if eFuse `EFUSE_ERR_RST_ENABLE` is set. If BLOCK0 has an error, it prints the error and returns `ESP_FAIL`, which should be treated as `esp_restart`.

Note: Refers to ESP32-C3 only.

Returns

- `ESP_OK`: No errors in BLOCK0.
- `ESP_FAIL`: Error in BLOCK0 requiring reboot.

bool **esp_efuse_is_ecdsa_p192_curve_supported** (void)

Checks if 192-bit ECDSA curve operations are supported.

This function checks if the current eFuse configuration supports 192-bit ECDSA curve operations.

bool **esp_efuse_is_ecdsa_p256_curve_supported** (void)

Checks if 256-bit ECDSA curve operations are supported.

This function checks if the current eFuse configuration supports 256-bit ECDSA curve operations.

Structures

struct **esp_efuse_desc_t**

Type definition for an eFuse field.

Public Members

esp_efuse_block_t **efuse_block**

Block of eFuse

uint8_t **bit_start**

Start bit [0..255]

uint16_t **bit_count**

Length of bit field [1..-]

struct **esp_secure_boot_key_digests_t**

Pointers to the trusted key digests.

The number of digests depends on the SOC's capabilities.

Public Members

const void ***key_digests**[3]

Pointers to the key digests

Macros

ESP_ERR_EFUSE

Base error code for efuse api.

ESP_OK_EFUSE_CNT

OK the required number of bits is set.

ESP_ERR_EFUSE_CNT_IS_FULL

Error field is full.

ESP_ERR_EFUSE_REPEATED_PROG

Error repeated programming of programmed bits is strictly forbidden.

ESP_ERR_CODING

Error while a encoding operation.

ESP_ERR_NOT_ENOUGH_UNUSED_KEY_BLOCKS

Error not enough unused key blocks available

ESP_ERR_DAMAGED_READING

Error. Burn or reset was done during a reading operation leads to damage read data. This error is internal to the efuse component and not returned by any public API.

Enumerations

enum **esp_efuse_rom_log_scheme_t**

Type definition for ROM log scheme.

Values:

enumerator **ESP_EFUSE_ROM_LOG_ALWAYS_ON**

Always enable ROM logging

enumerator **ESP_EFUSE_ROM_LOG_ON_GPIO_LOW**

ROM logging is enabled when specific GPIO level is low during start up

enumerator **ESP_EFUSE_ROM_LOG_ON_GPIO_HIGH**

ROM logging is enabled when specific GPIO level is high during start up

enumerator **ESP_EFUSE_ROM_LOG_ALWAYS_OFF**

Disable ROM logging permanently

2.9.8 Error Code and Helper Functions

This section lists definitions of common ESP-IDF error codes and several helper functions related to error handling.

For general information about error codes in ESP-IDF, see [Error Handling](#).

For the full list of error codes defined in ESP-IDF, see [Error Codes Reference](#).

API Reference

Header File

- [components/esp_common/include/esp_check.h](#)
- This header file can be included with:

```
#include "esp_check.h"
```

Macros

ESP_RETURN_ON_ERROR (x, log_tag, format, ...)

Macro which can be used to check the error code. If the code is not ESP_OK, it prints the message and returns. In the future, we want to switch to C++20. We also want to become compatible with clang. Hence, we provide two versions of the following macros. The first one is using the GNU extension `##__VA_ARGS__`. The second one is using the C++20 feature `VA_OPT(,)`. This allows users to compile their code with standard C++20 enabled instead of the GNU extension. Below C++20, we haven't found any good alternative to using `##__VA_ARGS__`. Macro which can be used to check the error code. If the code is not ESP_OK, it prints the message and returns.

ESP_RETURN_ON_ERROR_ISR (x, log_tag, format, ...)

A version of ESP_RETURN_ON_ERROR() macro that can be called from ISR.

ESP_GOTO_ON_ERROR (x, goto_tag, log_tag, format, ...)

Macro which can be used to check the error code. If the code is not ESP_OK, it prints the message, sets the local variable 'ret' to the code, and then exits by jumping to 'goto_tag'.

ESP_GOTO_ON_ERROR_ISR (x, goto_tag, log_tag, format, ...)

A version of ESP_GOTO_ON_ERROR() macro that can be called from ISR.

ESP_RETURN_ON_FALSE (a, err_code, log_tag, format, ...)

Macro which can be used to check the condition. If the condition is not 'true', it prints the message and returns with the supplied 'err_code'.

ESP_RETURN_ON_FALSE_ISR (a, err_code, log_tag, format, ...)

A version of ESP_RETURN_ON_FALSE() macro that can be called from ISR.

ESP_GOTO_ON_FALSE (a, err_code, goto_tag, log_tag, format, ...)

Macro which can be used to check the condition. If the condition is not 'true', it prints the message, sets the local variable 'ret' to the supplied 'err_code', and then exits by jumping to 'goto_tag'.

ESP_GOTO_ON_FALSE_ISR (a, err_code, goto_tag, log_tag, format, ...)

A version of ESP_GOTO_ON_FALSE() macro that can be called from ISR.

Header File

- [components/esp_common/include/esp_err.h](#)
- This header file can be included with:

```
#include "esp_err.h"
```

Functions

const char ***esp_err_to_name** (*esp_err_t* code)

Returns string for esp_err_t error codes.

This function finds the error code in a pre-generated lookup-table and returns its string representation.

The function is generated by the Python script tools/gen_esp_err_to_name.py which should be run each time an esp_err_t error is modified, created or removed from the IDF project.

Parameters **code** -- esp_err_t error code

Returns string error message

const char ***esp_err_to_name_r** (*esp_err_t* code, char *buf, size_t buflen)

Returns string for esp_err_t and system error codes.

This function finds the error code in a pre-generated lookup-table of esp_err_t errors and returns its string representation. If the error code is not found then it is attempted to be found among system errors.

The function is generated by the Python script tools/gen_esp_err_to_name.py which should be run each time an esp_err_t error is modified, created or removed from the IDF project.

Parameters

- **code** -- esp_err_t error code
- **buf** -- [out] buffer where the error message should be written
- **buflen** -- Size of buffer buf. At most buflen bytes are written into the buf buffer (including the terminating null byte).

Returns buf containing the string error message

Macros

ESP_OK

esp_err_t value indicating success (no error)

ESP_FAIL

Generic esp_err_t code indicating failure

ESP_ERR_NO_MEM

Out of memory

ESP_ERR_INVALID_ARG

Invalid argument

ESP_ERR_INVALID_STATE

Invalid state

ESP_ERR_INVALID_SIZE

Invalid size

ESP_ERR_NOT_FOUND

Requested resource not found

ESP_ERR_NOT_SUPPORTED

Operation or feature not supported

ESP_ERR_TIMEOUT

Operation timed out

ESP_ERR_INVALID_RESPONSE

Received response was invalid

ESP_ERR_INVALID_CRC

CRC or checksum was invalid

ESP_ERR_INVALID_VERSION

Version was invalid

ESP_ERR_INVALID_MAC

MAC address was invalid

ESP_ERR_NOT_FINISHED

Operation has not fully completed

ESP_ERR_NOT_ALLOWED

Operation is not allowed

ESP_ERR_WIFI_BASE

Starting number of WiFi error codes

ESP_ERR_MESH_BASE

Starting number of MESH error codes

ESP_ERR_FLASH_BASE

Starting number of flash error codes

ESP_ERR_HW_CRYPTO_BASE

Starting number of HW cryptography module error codes

ESP_ERR_MEMPROT_BASE

Starting number of Memory Protection API error codes

ESP_ERROR_CHECK (x)

Macro which can be used to check the error code, and terminate the program in case the code is not ESP_OK. Prints the error code, error location, and the failed statement to serial output.

Disabled if assertions are disabled.

ESP_ERROR_CHECK_WITHOUT_ABORT (x)

Macro which can be used to check the error code. Prints the error code, error location, and the failed statement to serial output. In comparison with ESP_ERROR_CHECK(), this prints the same error message but isn't terminating the program.

Type Definitions

typedef int **esp_err_t**

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.9.9 ESP HTTPS OTA

Overview

`esp_https_ota` provides simplified APIs to perform firmware upgrades over HTTPS. It is an abstraction layer over the existing OTA APIs.

Application Example

```
esp_err_t do_firmware_upgrade()
{
    esp_http_client_config_t config = {
        .url = CONFIG_FIRMWARE_UPGRADE_URL,
        .cert_pem = (char *)server_cert_pem_start,
    };
    esp_https_ota_config_t ota_config = {
        .http_config = &config,
    };
    esp_err_t ret = esp_https_ota(&ota_config);
    if (ret == ESP_OK) {
        esp_restart();
    } else {
        return ESP_FAIL;
    }
    return ESP_OK;
}
```

Server Verification

Please refer to [ESP-TLS: TLS Server Verification](#) for more information on server verification. The root certificate in PEM format needs to be provided to the `esp_http_client_config_t::cert_pem` member.

Note: The server-endpoint **root** certificate should be used for verification instead of any intermediate ones from the certificate chain. The reason is that the root certificate has the maximum validity and usually remains the same for a long period of time. Users can also use the `esp_http_client_config_t::crt_bundle_attach` member for verification by the ESP x509 Certificate Bundle feature, which covers most of the trusted root certificates.

Partial Image Download over HTTPS

To use the partial image download feature, enable `partial_http_download` configuration in `esp_https_ota_config_t`. When this configuration is enabled, firmware image will be downloaded in

multiple HTTP requests of specified sizes. Maximum content length of each request can be specified by setting `max_http_request_size` to the required value.

This option is useful while fetching image from a service like AWS S3, where mbedTLS Rx buffer size ([CONFIG_MBEDTLS_SSL_IN_CONTENT_LEN](#)) can be set to a lower value which is not possible without enabling this configuration.

Default value of mbedTLS Rx buffer size is set to 16 KB. By using `partial_http_download` with `max_http_request_size` of 4 KB, size of mbedTLS Rx buffer can be reduced to 4 KB. With this configuration, memory saving of around 12 KB is expected.

Signature Verification

For additional security, signature of OTA firmware images can be verified. For more information, please refer to [Secure OTA Updates Without Secure Boot](#).

Advanced APIs

`esp_https_ota` also provides advanced APIs which can be used if more information and control is needed during the OTA process.

Example that uses advanced ESP_HTTPS_OTA APIs: [system/ota/advanced_https_ota](#).

OTA Upgrades with Pre-Encrypted Firmware

To perform OTA upgrades with pre-encrypted firmware, please enable [CONFIG_ESP_HTTPS_OTA_DECRYPT_CB](#) in component menuconfig.

Example that performs OTA upgrade with pre-encrypted firmware: [system/ota/pre_encrypted_ota](#).

OTA System Events

ESP HTTPS OTA has various events for which a handler can be triggered by the [Event Loop Library](#) when the particular event occurs. The handler has to be registered using [esp_event_handler_register\(\)](#). This helps the event handling for ESP HTTPS OTA.

[esp_https_ota_event_t](#) has all the events which can happen when performing OTA upgrade using ESP HTTPS OTA.

Event Handler Example

```
/* Event handler for catching system events */
static void event_handler(void* arg, esp_event_base_t event_base,
                          int32_t event_id, void* event_data)
{
    if (event_base == ESP_HTTPS_OTA_EVENT) {
        switch (event_id) {
            case ESP_HTTPS_OTA_START:
                ESP_LOGI(TAG, "OTA started");
                break;
            case ESP_HTTPS_OTA_CONNECTED:
                ESP_LOGI(TAG, "Connected to server");
                break;
            case ESP_HTTPS_OTA_GET_IMG_DESC:
                ESP_LOGI(TAG, "Reading Image Description");
                break;
            case ESP_HTTPS_OTA_VERIFY_CHIP_ID:
                ESP_LOGI(TAG, "Verifying chip id of new image: %d", *(esp_
↳chip_id_t *)event_data);
                break;
        }
    }
}
```

(continues on next page)

(continued from previous page)

```

        break;
    case ESP_HTTPS_OTA_DECRYPT_CB:
        ESP_LOGI(TAG, "Callback to decrypt function");
        break;
    case ESP_HTTPS_OTA_WRITE_FLASH:
        ESP_LOGD(TAG, "Writing to flash: %d written", *(int_
↪*)event_data);
        break;
    case ESP_HTTPS_OTA_UPDATE_BOOT_PARTITION:
        ESP_LOGI(TAG, "Boot partition updated. Next Partition: %d
↪", *(esp_partition_subtype_t *)event_data);
        break;
    case ESP_HTTPS_OTA_FINISH:
        ESP_LOGI(TAG, "OTA finish");
        break;
    case ESP_HTTPS_OTA_ABORT:
        ESP_LOGI(TAG, "OTA abort");
        break;
    }
}
}

```

Expected data type for different ESP HTTPS OTA events in the system event loop:

- ESP_HTTPS_OTA_START : NULL
- ESP_HTTPS_OTA_CONNECTED : NULL
- ESP_HTTPS_OTA_GET_IMG_DESC : NULL
- ESP_HTTPS_OTA_VERIFY_CHIP_ID : esp_chip_id_t
- ESP_HTTPS_OTA_DECRYPT_CB : NULL
- ESP_HTTPS_OTA_WRITE_FLASH : int
- ESP_HTTPS_OTA_UPDATE_BOOT_PARTITION : esp_partition_subtype_t
- ESP_HTTPS_OTA_FINISH : NULL
- ESP_HTTPS_OTA_ABORT : NULL

API Reference

Header File

- [components/esp_https_ota/include/esp_https_ota.h](#)
- This header file can be included with:

```
#include "esp_https_ota.h"
```

- This header file is a part of the API provided by the esp_https_ota component. To declare that your component depends on esp_https_ota, add the following to your CMakeLists.txt:

```
REQUIRES esp_https_ota
```

or

```
PRIV_REQUIRES esp_https_ota
```

Functions

esp_err_t esp_https_ota (const [esp_https_ota_config_t](#) *ota_config)

HTTPS OTA Firmware upgrade.

This function allocates HTTPS OTA Firmware upgrade context, establishes HTTPS connection, reads image data from HTTP stream and writes it to OTA partition and finishes HTTPS OTA Firmware upgrade operation. This API supports URL redirection, but if CA cert of URLs differ then it should be appended to cert_pem member of ota_config->http_config.

Note: This API handles the entire OTA operation, so if this API is being used then no other APIs from `esp_https_ota` component should be called. If more information and control is needed during the HTTPS OTA process, then one can use `esp_https_ota_begin` and subsequent APIs. If this API returns successfully, `esp_restart()` must be called to boot from the new firmware image.

Parameters `ota_config` -- [in] pointer to `esp_https_ota_config_t` structure.

Returns

- `ESP_OK`: OTA data updated, next reboot will use specified partition.
- `ESP_FAIL`: For generic failure.
- `ESP_ERR_INVALID_ARG`: Invalid argument
- `ESP_ERR_OTA_VALIDATE_FAILED`: Invalid app image
- `ESP_ERR_NO_MEM`: Cannot allocate memory for OTA operation.
- `ESP_ERR_FLASH_OP_TIMEOUT` or `ESP_ERR_FLASH_OP_FAIL`: Flash write failed.
- For other return codes, refer OTA documentation in esp-idf's app_update component.

`esp_err_t esp_https_ota_begin` (const `esp_https_ota_config_t` *ota_config, `esp_https_ota_handle_t` *handle)

Start HTTPS OTA Firmware upgrade.

This function initializes ESP HTTPS OTA context and establishes HTTPS connection. This function must be invoked first. If this function returns successfully, then `esp_https_ota_perform` should be called to continue with the OTA process and there should be a call to `esp_https_ota_finish` on completion of OTA operation or on failure in subsequent operations. This API supports URL redirection, but if CA cert of URLs differ then it should be appended to `cert_pem` member of `http_config`, which is a part of `ota_config`. In case of error, this API explicitly sets `handle` to NULL.

Note: This API is blocking, so setting `is_async` member of `http_config` structure will result in an error.

Parameters

- `ota_config` -- [in] pointer to `esp_https_ota_config_t` structure
- `handle` -- [out] pointer to an allocated data of type `esp_https_ota_handle_t` which will be initialised in this function

Returns

- `ESP_OK`: HTTPS OTA Firmware upgrade context initialised and HTTPS connection established
- `ESP_FAIL`: For generic failure.
- `ESP_ERR_INVALID_ARG`: Invalid argument (missing/incorrect config, certificate, etc.)
- For other return codes, refer documentation in app_update component and `esp_http_client` component in esp-idf.

`esp_err_t esp_https_ota_perform` (`esp_https_ota_handle_t` https_ota_handle)

Read image data from HTTP stream and write it to OTA partition.

This function reads image data from HTTP stream and writes it to OTA partition. This function must be called only if `esp_https_ota_begin()` returns successfully. This function must be called in a loop since it returns after every HTTP read operation thus giving you the flexibility to stop OTA operation midway.

Parameters `https_ota_handle` -- [in] pointer to `esp_https_ota_handle_t` structure

Returns

- `ESP_ERR_HTTPS_OTA_IN_PROGRESS`: OTA update is in progress, call this API again to continue.
- `ESP_OK`: OTA update was successful
- `ESP_FAIL`: OTA update failed
- `ESP_ERR_INVALID_ARG`: Invalid argument
- `ESP_ERR_INVALID_VERSION`: Invalid chip revision in image header

- `ESP_ERR_OTA_VALIDATE_FAILED`: Invalid app image
- `ESP_ERR_NO_MEM`: Cannot allocate memory for OTA operation.
- `ESP_ERR_FLASH_OP_TIMEOUT` or `ESP_ERR_FLASH_OP_FAIL`: Flash write failed.
- For other return codes, refer OTA documentation in esp-idf's app_update component.

bool **esp_https_ota_is_complete_data_received** ([*esp_https_ota_handle_t*](#) https_ota_handle)

Checks if complete data was received or not.

Note: This API can be called just before `esp_https_ota_finish()` to validate if the complete image was indeed received.

Parameters `https_ota_handle` -- [in] pointer to `esp_https_ota_handle_t` structure

Returns

- false
- true

[*esp_err_t*](#) **esp_https_ota_finish** ([*esp_https_ota_handle_t*](#) https_ota_handle)

Clean-up HTTPS OTA Firmware upgrade and close HTTPS connection.

This function closes the HTTP connection and frees the ESP HTTPS OTA context. This function switches the boot partition to the OTA partition containing the new firmware image.

Note: If this API returns successfully, `esp_restart()` must be called to boot from the new firmware image. `esp_https_ota_finish` should not be called after calling `esp_https_ota_abort`

Parameters `https_ota_handle` -- [in] pointer to `esp_https_ota_handle_t` structure

Returns

- `ESP_OK`: Clean-up successful
- `ESP_ERR_INVALID_STATE`
- `ESP_ERR_INVALID_ARG`: Invalid argument
- `ESP_ERR_OTA_VALIDATE_FAILED`: Invalid app image

[*esp_err_t*](#) **esp_https_ota_abort** ([*esp_https_ota_handle_t*](#) https_ota_handle)

Clean-up HTTPS OTA Firmware upgrade and close HTTPS connection.

This function closes the HTTP connection and frees the ESP HTTPS OTA context.

Note: `esp_https_ota_abort` should not be called after calling `esp_https_ota_finish`

Parameters `https_ota_handle` -- [in] pointer to `esp_https_ota_handle_t` structure

Returns

- `ESP_OK`: Clean-up successful
- `ESP_ERR_INVALID_STATE`: Invalid ESP HTTPS OTA state
- `ESP_FAIL`: OTA not started
- `ESP_ERR_NOT_FOUND`: OTA handle not found
- `ESP_ERR_INVALID_ARG`: Invalid argument

[*esp_err_t*](#) **esp_https_ota_get_img_desc** ([*esp_https_ota_handle_t*](#) https_ota_handle, [*esp_app_desc_t*](#) *new_app_info)

Reads app description from image header. The app description provides information like the "Firmware version" of the image.

Note: This API can be called only after `esp_https_ota_begin()` and before `esp_https_ota_perform()`. Calling this API is not mandatory.

Parameters

- **https_ota_handle** -- [in] pointer to `esp_https_ota_handle_t` structure
- **new_app_info** -- [out] pointer to an allocated `esp_app_desc_t` structure

Returns

- `ESP_ERR_INVALID_ARG`: Invalid arguments
- `ESP_ERR_INVALID_STATE`: Invalid state to call this API. `esp_https_ota_begin()` not called yet.
- `ESP_FAIL`: Failed to read image descriptor
- `ESP_OK`: Successfully read image descriptor

int **esp_https_ota_get_image_len_read** (`esp_https_ota_handle_t` https_ota_handle)

This function returns OTA image data read so far.

Note: This API should be called only if `esp_https_ota_perform()` has been called atleast once or if `esp_https_ota_get_img_desc` has been called before.

Parameters **https_ota_handle** -- [in] pointer to `esp_https_ota_handle_t` structure

Returns

- -1 On failure
- total bytes read so far

int **esp_https_ota_get_status_code** (`esp_https_ota_handle_t` https_ota_handle)

This function returns the HTTP status code of the last HTTP response.

Note: This API should be called only after `esp_https_ota_begin()` has been called. This can be used to check the HTTP status code of the OTA download process.

Parameters **https_ota_handle** -- [in] pointer to `esp_https_ota_handle_t` structure

Returns

- -1 On failure
- HTTP status code

int **esp_https_ota_get_image_size** (`esp_https_ota_handle_t` https_ota_handle)

This function returns OTA image total size.

Note: This API should be called after `esp_https_ota_begin()` has been already called. This can be used to create some sort of progress indication (in combination with `esp_https_ota_get_image_len_read()`)

Parameters **https_ota_handle** -- [in] pointer to `esp_https_ota_handle_t` structure

Returns

- -1 On failure or chunked encoding
- total bytes of image

Structures

struct **esp_https_ota_config_t**

ESP HTTPS OTA configuration.

Public Members

const *esp_http_client_config_t* ***http_config**

ESP HTTP client configuration

http_client_init_cb_t **http_client_init_cb**

Callback after ESP HTTP client is initialised

bool **bulk_flash_erase**

Erase entire flash partition during initialization. By default flash partition is erased during write operation and in chunk of 4K sector size

bool **partial_http_download**

Enable Firmware image to be downloaded over multiple HTTP requests

int **max_http_request_size**

Maximum request size for partial HTTP download

Macros

ESP_ERR_HTTPS_OTA_BASE

ESP_ERR_HTTPS_OTA_IN_PROGRESS

Type Definitions

typedef void ***esp_https_ota_handle_t**

typedef *esp_err_t* (***http_client_init_cb_t**)(*esp_http_client_handle_t*)

Enumerations

enum **esp_https_ota_event_t**

Events generated by OTA process.

Values:

enumerator **ESP_HTTPS_OTA_START**

OTA started

enumerator **ESP_HTTPS_OTA_CONNECTED**

Connected to server

enumerator **ESP_HTTPS_OTA_GET_IMG_DESC**

Read app description from image header

enumerator **ESP_HTTPS_OTA_VERIFY_CHIP_ID**

Verify chip id of new image

enumerator **ESP_HTTPS_OTA_DECRYPT_CB**

Callback to decrypt function

enumerator **ESP_HTTPS_OTA_WRITE_FLASH**

Flash write operation

enumerator **ESP_HTTPS_OTA_UPDATE_BOOT_PARTITION**

Boot partition update after successful ota update

enumerator **ESP_HTTPS_OTA_FINISH**

OTA finished

enumerator **ESP_HTTPS_OTA_ABORT**

OTA aborted

2.9.10 Event Loop Library

Overview

The event loop library allows components to declare events so that other components can register handlers -- codes that executes when those events occur. This allows loosely-coupled components to attach desired behavior to state changes of other components without application involvement. This also simplifies event processing by serializing and deferring code execution to another context.

Using `esp_event` APIs

There are two objects of concern for users of this library: events and event loops.

An event indicates an important occurrence, such as a successful Wi-Fi connection to an access point. A two-part identifier should be used when referencing events, see [declaring and defining events](#) for details. The event loop is the bridge between events and event handlers. The event source publishes events to the event loop using the APIs provided by the event loop library, and event handlers registered to the event loop respond to specific types of events.

Using this library roughly entails the following flow:

1. The user defines a function that should run when an event is posted to a loop. This function is referred to as the event handler, and should have the same signature as `esp_event_handler_t`.
2. An event loop is created using `esp_event_loop_create()`, which outputs a handle to the loop of type `esp_event_loop_handle_t`. Event loops created using this API are referred to as user event loops. There is, however, a special type of event loop called the default event loop which is discussed in [default event loop](#).
3. Components register event handlers to the loop using `esp_event_handler_register_with()`. Handlers can be registered with multiple loops, see [notes on handler registration](#).
4. Event sources post an event to the loop using `esp_event_post_to()`.
5. Components wanting to remove their handlers from being called can do so by unregistering from the loop using `esp_event_handler_unregister_with()`.
6. Event loops that are no longer needed can be deleted using `esp_event_loop_delete()`.

In code, the flow above may look like as follows:

```

// 1. Define the event handler
void run_on_event(void* handler_arg, esp_event_base_t base, int32_t id, void* _
↳event_data)
{
    // Event handler logic
}

void app_main()
{
    // 2. A configuration structure of type esp_event_loop_args_t is needed to_
↳specify the properties of the loop to be created. A handle of type esp_event_
↳loop_handle_t is obtained, which is needed by the other APIs to reference the_
↳loop to perform their operations.
    esp_event_loop_args_t loop_args = {
        .queue_size = ...,
        .task_name = ...
        .task_priority = ...,
        .task_stack_size = ...,
        .task_core_id = ...
    };

    esp_event_loop_handle_t loop_handle;

    esp_event_loop_create(&loop_args, &loop_handle);

    // 3. Register event handler defined in (1). MY_EVENT_BASE and MY_EVENT_ID_
↳specify a hypothetical event that handler run_on_event should execute when it_
↳gets posted to the loop.
    esp_event_handler_register_with(loop_handle, MY_EVENT_BASE, MY_EVENT_ID, run_
↳on_event, ...);

    ...

    // 4. Post events to the loop. This queues the event on the event loop. At_
↳some point, the event loop executes the event handler registered to the posted_
↳event, in this case, run_on_event. To simplify the process, this example calls_
↳esp_event_post_to from app_main, but posting can be done from any other task_
↳(which is the more interesting use case).
    esp_event_post_to(loop_handle, MY_EVENT_BASE, MY_EVENT_ID, ...);

    ...

    // 5. Unregistering an unneeded handler
    esp_event_handler_unregister_with(loop_handle, MY_EVENT_BASE, MY_EVENT_ID, run_
↳on_event);

    ...

    // 6. Deleting an unneeded event loop
    esp_event_loop_delete(loop_handle);
}

```

Declaring and Defining Events

As mentioned previously, events consist of two-part identifiers: the event base and the event ID. The event base identifies an independent group of events; the event ID identifies the event within that group. Think of the event base and event ID as a person's last name and first name, respectively. A last name identifies a family, and the first name identifies a person within that family.

The event loop library provides macros to declare and define the event base easily.

Event base declaration:

```
ESP_EVENT_DECLARE_BASE (EVENT_BASE) ;
```

Event base definition:

```
ESP_EVENT_DEFINE_BASE (EVENT_BASE) ;
```

Note: In ESP-IDF, the base identifiers for system events are uppercase and are postfixed with `_EVENT`. For example, the base for Wi-Fi events is declared and defined as `WIFI_EVENT`, the Ethernet event base `ETHERNET_EVENT`, and so on. The purpose is to have event bases look like constants (although they are global variables considering the definitions of macros `ESP_EVENT_DECLARE_BASE` and `ESP_EVENT_DEFINE_BASE`).

For event IDs, declaring them as enumerations is recommended. Once again, for visibility, these are typically placed in public header files.

Event ID:

```
enum {
    EVENT_ID_1,
    EVENT_ID_2,
    EVENT_ID_3,
    ...
}
```

Default Event Loop

The default event loop is a special type of loop used for system events (Wi-Fi events, for example). The handle for this loop is hidden from the user, and the creation, deletion, handler registration/deregistration, and posting of events are done through a variant of the APIs for user event loops. The table below enumerates those variants, and the user event loops equivalent.

User Event Loops	Default Event Loops
<code>esp_event_loop_create()</code>	<code>esp_event_loop_create_default()</code>
<code>esp_event_loop_delete()</code>	<code>esp_event_loop_delete_default()</code>
<code>esp_event_handler_register_with()</code>	<code>esp_event_handler_register()</code>
<code>esp_event_handler_unregister_with()</code>	<code>esp_event_handler_unregister()</code>
<code>esp_event_post_to()</code>	<code>esp_event_post()</code>

If you compare the signatures for both, they are mostly similar except for the lack of loop handle specification for the default event loop APIs.

Other than the API difference and the special designation to which system events are posted, there is no difference in how default event loops and user event loops behave. It is even possible for users to post their own events to the default event loop, should the user opt to not create their own loops to save memory.

Notes on Handler Registration

It is possible to register a single handler to multiple events individually by using multiple calls to [`esp\_event\_handler\_register\_with\(\)`](#). For those multiple calls, the specific event base and event ID can be specified with which the handler should execute.

However, in some cases, it is desirable for a handler to execute on the following situations:

- (1) all events that get posted to a loop
- (2) all events of a particular base identifier

This is possible using the special event base identifier `ESP_EVENT_ANY_BASE` and special event ID `ESP_EVENT_ANY_ID`. These special identifiers may be passed as the event base and event ID arguments for `esp_event_handler_register_with()`.

Therefore, the valid arguments to `esp_event_handler_register_with()` are:

1. <event base>, <event ID> - handler executes when the event with base <event base> and event ID <event ID> gets posted to the loop
2. <event base>, `ESP_EVENT_ANY_ID` - handler executes when any event with base <event base> gets posted to the loop
3. `ESP_EVENT_ANY_BASE`, `ESP_EVENT_ANY_ID` - handler executes when any event gets posted to the loop

As an example, suppose the following handler registrations were performed:

```
esp_event_handler_register_with(loop_handle, MY_EVENT_BASE, MY_EVENT_ID, run_on_
↳event_1, ...);
esp_event_handler_register_with(loop_handle, MY_EVENT_BASE, ESP_EVENT_ANY_ID, run_
↳on_event_2, ...);
esp_event_handler_register_with(loop_handle, ESP_EVENT_ANY_BASE, ESP_EVENT_ANY_ID,
↳run_on_event_3, ...);
```

If the hypothetical event `MY_EVENT_BASE`, `MY_EVENT_ID` is posted, all three handlers `run_on_event_1`, `run_on_event_2`, and `run_on_event_3` would execute.

If the hypothetical event `MY_EVENT_BASE`, `MY_OTHER_EVENT_ID` is posted, only `run_on_event_2` and `run_on_event_3` would execute.

If the hypothetical event `MY_OTHER_EVENT_BASE`, `MY_OTHER_EVENT_ID` is posted, only `run_on_event_3` would execute.

Handler Un-Registering Itself In general, an event handler run by an event loop is **not allowed to do any registering/unregistering activity on that event loop**. There is one exception, though: un-registering itself is allowed for the handler. E.g., it is possible to do the following:

```
void run_on_event(void* handler_arg, esp_event_base_t base, int32_t id, void*
↳event_data)
{
    esp_event_loop_handle_t *loop_handle = (esp_event_loop_handle_t*) handler_arg;
    esp_event_handler_unregister_with(*loop_handle, MY_EVENT_BASE, MY_EVENT_ID,
↳run_on_event);
}

void app_main(void)
{
    esp_event_loop_handle_t loop_handle;
    esp_event_loop_create(&loop_args, &loop_handle);
    esp_event_handler_register_with(loop_handle, MY_EVENT_BASE, MY_EVENT_ID, run_
↳on_event, &loop_handle);
    // ... post-event MY_EVENT_BASE, MY_EVENT_ID and run loop at some point
}
```

Handler Registration and Handler Dispatch Order The general rule is that, for handlers that match a certain posted event during dispatch, those which are registered first also get executed first. The user can then control which handlers get executed first by registering them before other handlers, provided that all registrations are performed using a single task. If the user plans to take advantage of this behavior, caution must be exercised if there are multiple tasks registering handlers. While the 'first registered, first executed' behavior still holds true, the task which gets executed first also gets its handlers registered first. Handlers registered one after the other by a single task are still dispatched in the order relative to each other, but if that task gets pre-empted in between registration by another task that also registers handlers; then during dispatch those handlers also get executed in between.

Event Loop Profiling

A configuration option `CONFIG_ESP_EVENT_LOOP_PROFILING` can be enabled in order to activate statistics collection for all event loops created. The function `esp_event_dump()` can be used to output the collected statistics to a file stream. More details on the information included in the dump can be found in the `esp_event_dump()` API Reference.

Application Example

Examples of using the `esp_event` library can be found in [system/esp_event](#). The examples cover event declaration, loop creation, handler registration and deregistration, and event posting.

Other examples which also adopt `esp_event` library:

- [NMEA Parser](#) , which decodes the statements received from GPS.

API Reference

Header File

- [components/esp_event/include/esp_event.h](#)
- This header file can be included with:

```
#include "esp_event.h"
```

- This header file is a part of the API provided by the `esp_event` component. To declare that your component depends on `esp_event`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_event
```

or

```
PRIV_REQUIRES esp_event
```

Functions

`esp_err_t esp_event_loop_create` (const `esp_event_loop_args_t` *event_loop_args, `esp_event_loop_handle_t` *event_loop)

Create a new event loop.

Parameters

- **event_loop_args** -- [in] configuration structure for the event loop to create
- **event_loop** -- [out] handle to the created event loop

Returns

- `ESP_OK`: Success
- `ESP_ERR_INVALID_ARG`: event_loop_args or event_loop was NULL
- `ESP_ERR_NO_MEM`: Cannot allocate memory for event loops list
- `ESP_FAIL`: Failed to create task loop
- Others: Fail

`esp_err_t esp_event_loop_delete` (`esp_event_loop_handle_t` event_loop)

Delete an existing event loop.

Parameters **event_loop** -- [in] event loop to delete, must not be NULL

Returns

- `ESP_OK`: Success
- Others: Fail

`esp_err_t esp_event_loop_create_default` (void)

Create default event loop.

Returns

- ESP_OK: Success
- ESP_ERR_NO_MEM: Cannot allocate memory for event loops list
- ESP_ERR_INVALID_STATE: Default event loop has already been created
- ESP_FAIL: Failed to create task loop
- Others: Fail

esp_err_t **esp_event_loop_delete_default** (void)

Delete the default event loop.

Returns

- ESP_OK: Success
- Others: Fail

esp_err_t **esp_event_loop_run** (*esp_event_loop_handle_t* event_loop, TickType_t ticks_to_run)

Dispatch events posted to an event loop.

This function is used to dispatch events posted to a loop with no dedicated task, i.e. task name was set to NULL in event_loop_args argument during loop creation. This function includes an argument to limit the amount of time it runs, returning control to the caller when that time expires (or some time afterwards). There is no guarantee that a call to this function will exit at exactly the time of expiry. There is also no guarantee that events have been dispatched during the call, as the function might have spent all the allotted time waiting on the event queue. Once an event has been dequeued, however, it is guaranteed to be dispatched. This guarantee contributes to not being able to exit exactly at time of expiry as (1) blocking on internal mutexes is necessary for dispatching the dequeued event, and (2) during dispatch of the dequeued event there is no way to control the time occupied by handler code execution. The guaranteed time of exit is therefore the allotted time + amount of time required to dispatch the last dequeued event.

In cases where waiting on the queue times out, ESP_OK is returned and not ESP_ERR_TIMEOUT, since it is normal behavior.

Note: encountering an unknown event that has been posted to the loop will only generate a warning, not an error.

Parameters

- **event_loop** -- [in] event loop to dispatch posted events from, must not be NULL
- **ticks_to_run** -- [in] number of ticks to run the loop

Returns

- ESP_OK: Success
- Others: Fail

esp_err_t **esp_event_handler_register** (*esp_event_base_t* event_base, int32_t event_id, *esp_event_handler_t* event_handler, void *event_handler_arg)

Register an event handler to the system event loop (legacy).

This function can be used to register a handler for either: (1) specific events, (2) all events of a certain event base, or (3) all events known by the system event loop.

- specific events: specify exact event_base and event_id
- all events of a certain base: specify exact event_base and use ESP_EVENT_ANY_ID as the event_id
- all events known by the loop: use ESP_EVENT_ANY_BASE for event_base and ESP_EVENT_ANY_ID as the event_id

Registering multiple handlers to events is possible. Registering a single handler to multiple events is also possible. However, registering the same handler to the same event multiple times would cause the overwriting of the event_handler_arg but the handler will be kept at the same position in the list associated with the event that triggers it. It means that the call order of registered handlers for that event will remain the same.

Note: the event loop library does not maintain a copy of `event_handler_arg`, therefore the user should ensure that `event_handler_arg` still points to a valid location by the time the handler gets called

Parameters

- **event_base** -- **[in]** the base ID of the event to register the handler for
- **event_id** -- **[in]** the ID of the event to register the handler for
- **event_handler** -- **[in]** the handler function which gets called when the event is dispatched
- **event_handler_arg** -- **[in]** data, aside from event data, that is passed to the handler when it is called

Returns

- ESP_OK: Success
- ESP_ERR_NO_MEM: Cannot allocate memory for the handler
- ESP_ERR_INVALID_ARG: Invalid combination of event base and event ID
- Others: Fail

```
esp_err_t esp_event_handler_register_with(esp_event_loop_handle_t event_loop, esp_event_base_t event_base, int32_t event_id, esp_event_handler_t event_handler, void *event_handler_arg)
```

Register an event handler to a specific loop (legacy).

This function behaves in the same manner as `esp_event_handler_register`, except the additional specification of the event loop to register the handler to.

Note: the event loop library does not maintain a copy of `event_handler_arg`, therefore the user should ensure that `event_handler_arg` still points to a valid location by the time the handler gets called

Parameters

- **event_loop** -- **[in]** the event loop to register this handler function to, must not be NULL
- **event_base** -- **[in]** the base ID of the event to register the handler for
- **event_id** -- **[in]** the ID of the event to register the handler for
- **event_handler** -- **[in]** the handler function which gets called when the event is dispatched
- **event_handler_arg** -- **[in]** data, aside from event data, that is passed to the handler when it is called

Returns

- ESP_OK: Success
- ESP_ERR_NO_MEM: Cannot allocate memory for the handler
- ESP_ERR_INVALID_ARG: Invalid combination of event base and event ID
- Others: Fail

```
esp_err_t esp_event_handler_instance_register_with(esp_event_loop_handle_t event_loop, esp_event_base_t event_base, int32_t event_id, esp_event_handler_t event_handler, void *event_handler_arg, esp_event_handler_instance_t *instance)
```

Register an instance of event handler to a specific loop.

This function can be used to register a handler for either: (1) specific events, (2) all events of a certain event base, or (3) all events known by the system event loop.

- specific events: specify exact `event_base` and `event_id`
- all events of a certain base: specify exact `event_base` and use `ESP_EVENT_ANY_ID` as the `event_id`

- all events known by the loop: use `ESP_EVENT_ANY_BASE` for `event_base` and `ESP_EVENT_ANY_ID` as the `event_id`

Besides the error, the function returns an instance object as output parameter to identify each registration. This is necessary to remove (unregister) the registration before the event loop is deleted.

Registering multiple handlers to events, registering a single handler to multiple events as well as registering the same handler to the same event multiple times is possible. Each registration yields a distinct instance object which identifies it over the registration lifetime.

Note: the event loop library does not maintain a copy of `event_handler_arg`, therefore the user should ensure that `event_handler_arg` still points to a valid location by the time the handler gets called

Parameters

- **event_loop** -- **[in]** the event loop to register this handler function to, must not be NULL
- **event_base** -- **[in]** the base ID of the event to register the handler for
- **event_id** -- **[in]** the ID of the event to register the handler for
- **event_handler** -- **[in]** the handler function which gets called when the event is dispatched
- **event_handler_arg** -- **[in]** data, aside from event data, that is passed to the handler when it is called
- **instance** -- **[out]** An event handler instance object related to the registered event handler and data, can be NULL. This needs to be kept if the specific callback instance should be unregistered before deleting the whole event loop. Registering the same event handler multiple times is possible and yields distinct instance objects. The data can be the same for all registrations. If no unregistration is needed, but the handler should be deleted when the event loop is deleted, instance can be NULL.

Returns

- `ESP_OK`: Success
- `ESP_ERR_NO_MEM`: Cannot allocate memory for the handler
- `ESP_ERR_INVALID_ARG`: Invalid combination of event base and event ID or instance is NULL
- Others: Fail

```
esp_err_t esp_event_handler_instance_register(esp_event_base_t event_base, int32_t event_id,  
                                             esp_event_handler_t event_handler, void  
                                             *event_handler_arg,  
                                             esp_event_handler_instance_t *instance)
```

Register an instance of event handler to the default loop.

This function does the same as `esp_event_handler_instance_register_with`, except that it registers the handler to the default event loop.

Note: the event loop library does not maintain a copy of `event_handler_arg`, therefore the user should ensure that `event_handler_arg` still points to a valid location by the time the handler gets called

Parameters

- **event_base** -- **[in]** the base ID of the event to register the handler for
- **event_id** -- **[in]** the ID of the event to register the handler for
- **event_handler** -- **[in]** the handler function which gets called when the event is dispatched
- **event_handler_arg** -- **[in]** data, aside from event data, that is passed to the handler when it is called
- **instance** -- **[out]** An event handler instance object related to the registered event handler and data, can be NULL. This needs to be kept if the specific callback instance should be unregistered before deleting the whole event loop. Registering the same event handler

multiple times is possible and yields distinct instance objects. The data can be the same for all registrations. If no unregistration is needed, but the handler should be deleted when the event loop is deleted, instance can be NULL.

Returns

- `ESP_OK`: Success
- `ESP_ERR_NO_MEM`: Cannot allocate memory for the handler
- `ESP_ERR_INVALID_ARG`: Invalid combination of event base and event ID or instance is NULL
- Others: Fail

esp_err_t **esp_event_handler_unregister** (*esp_event_base_t* event_base, *int32_t* event_id, *esp_event_handler_t* event_handler)

Unregister a handler with the system event loop (legacy).

Unregisters a handler, so it will no longer be called during dispatch. Handlers can be unregistered for any combination of event_base and event_id which were previously registered. To unregister a handler, the event_base and event_id arguments must match exactly the arguments passed to esp_event_handler_register() when that handler was registered. Passing ESP_EVENT_ANY_BASE and/or ESP_EVENT_ANY_ID will only unregister handlers that were registered with the same wildcard arguments.

Note: When using ESP_EVENT_ANY_ID, handlers registered to specific event IDs using the same base will not be unregistered. When using ESP_EVENT_ANY_BASE, events registered to specific bases will also not be unregistered. This avoids accidental unregistration of handlers registered by other users or components.

Parameters

- **event_base** -- **[in]** the base of the event with which to unregister the handler
- **event_id** -- **[in]** the ID of the event with which to unregister the handler
- **event_handler** -- **[in]** the handler to unregister

Returns `ESP_OK` success

Returns `ESP_ERR_INVALID_ARG` invalid combination of event base and event ID

Returns others fail

esp_err_t **esp_event_handler_unregister_with** (*esp_event_loop_handle_t* event_loop, *esp_event_base_t* event_base, *int32_t* event_id, *esp_event_handler_t* event_handler)

Unregister a handler from a specific event loop (legacy).

This function behaves in the same manner as esp_event_handler_unregister, except the additional specification of the event loop to unregister the handler with.

Parameters

- **event_loop** -- **[in]** the event loop with which to unregister this handler function, must not be NULL
- **event_base** -- **[in]** the base of the event with which to unregister the handler
- **event_id** -- **[in]** the ID of the event with which to unregister the handler
- **event_handler** -- **[in]** the handler to unregister

Returns

- `ESP_OK`: Success
- `ESP_ERR_INVALID_ARG`: Invalid combination of event base and event ID
- Others: Fail

esp_err_t **esp_event_handler_instance_unregister_with** (*esp_event_loop_handle_t* event_loop, *esp_event_base_t* event_base, *int32_t* event_id, *esp_event_handler_instance_t* instance)

Unregister a handler instance from a specific event loop.

Unregisters a handler instance, so it will no longer be called during dispatch. Handler instances can be unregistered for any combination of `event_base` and `event_id` which were previously registered. To unregister a handler instance, the `event_base` and `event_id` arguments must match exactly the arguments passed to `esp_event_handler_instance_register()` when that handler instance was registered. Passing `ESP_EVENT_ANY_BASE` and/or `ESP_EVENT_ANY_ID` will only unregister handler instances that were registered with the same wildcard arguments.

Note: When using `ESP_EVENT_ANY_ID`, handlers registered to specific event IDs using the same base will not be unregistered. When using `ESP_EVENT_ANY_BASE`, events registered to specific bases will also not be unregistered. This avoids accidental unregistration of handlers registered by other users or components.

Parameters

- **event_loop** -- **[in]** the event loop with which to unregister this handler function, must not be NULL
- **event_base** -- **[in]** the base of the event with which to unregister the handler
- **event_id** -- **[in]** the ID of the event with which to unregister the handler
- **instance** -- **[in]** the instance object of the registration to be unregistered

Returns

- `ESP_OK`: Success
- `ESP_ERR_INVALID_ARG`: Invalid combination of event base and event ID
- Others: Fail

esp_err_t **esp_event_handler_instance_unregister** (*esp_event_base_t* event_base, *int32_t* event_id, *esp_event_handler_instance_t* instance)

Unregister a handler from the system event loop.

This function does the same as `esp_event_handler_instance_unregister_with`, except that it unregisters the handler instance from the default event loop.

Parameters

- **event_base** -- **[in]** the base of the event with which to unregister the handler
- **event_id** -- **[in]** the ID of the event with which to unregister the handler
- **instance** -- **[in]** the instance object of the registration to be unregistered

Returns

- `ESP_OK`: Success
- `ESP_ERR_INVALID_ARG`: Invalid combination of event base and event ID
- Others: Fail

esp_err_t **esp_event_post** (*esp_event_base_t* event_base, *int32_t* event_id, *const void **event_data, *size_t* event_data_size, *TickType_t* ticks_to_wait)

Posts an event to the system default event loop. The event loop library keeps a copy of `event_data` and manages the copy's lifetime automatically (allocation + deletion); this ensures that the data the handler receives is always valid.

Parameters

- **event_base** -- **[in]** the event base that identifies the event
- **event_id** -- **[in]** the event ID that identifies the event
- **event_data** -- **[in]** the data, specific to the event occurrence, that gets passed to the handler
- **event_data_size** -- **[in]** the size of the event data
- **ticks_to_wait** -- **[in]** number of ticks to block on a full event queue

Returns

- `ESP_OK`: Success
- `ESP_ERR_TIMEOUT`: Time to wait for event queue to unblock expired, queue full when posting from ISR
- `ESP_ERR_INVALID_ARG`: Invalid combination of event base and event ID
- Others: Fail

esp_err_t **esp_event_post_to** (*esp_event_loop_handle_t* event_loop, esp_event_base_t event_base, int32_t event_id, const void *event_data, size_t event_data_size, TickType_t ticks_to_wait)

Posts an event to the specified event loop. The event loop library keeps a copy of event_data and manages the copy's lifetime automatically (allocation + deletion); this ensures that the data the handler receives is always valid.

This function behaves in the same manner as esp_event_post, except the additional specification of the event loop to post the event to.

Parameters

- **event_loop** -- [in] the event loop to post to, must not be NULL
- **event_base** -- [in] the event base that identifies the event
- **event_id** -- [in] the event ID that identifies the event
- **event_data** -- [in] the data, specific to the event occurrence, that gets passed to the handler
- **event_data_size** -- [in] the size of the event data
- **ticks_to_wait** -- [in] number of ticks to block on a full event queue

Returns

- ESP_OK: Success
- ESP_ERR_TIMEOUT: Time to wait for event queue to unblock expired, queue full when posting from ISR
- ESP_ERR_INVALID_ARG: Invalid combination of event base and event ID
- Others: Fail

esp_err_t **esp_event_isr_post** (esp_event_base_t event_base, int32_t event_id, const void *event_data, size_t event_data_size, BaseType_t *task_unblocked)

Special variant of esp_event_post for posting events from interrupt handlers.

Note: this function is only available when CONFIG_ESP_EVENT_POST_FROM_ISR is enabled

Note: when this function is called from an interrupt handler placed in IRAM, this function should be placed in IRAM as well by enabling CONFIG_ESP_EVENT_POST_FROM_IRAM_ISR

Parameters

- **event_base** -- [in] the event base that identifies the event
- **event_id** -- [in] the event ID that identifies the event
- **event_data** -- [in] the data, specific to the event occurrence, that gets passed to the handler
- **event_data_size** -- [in] the size of the event data; max is 4 bytes
- **task_unblocked** -- [out] an optional parameter (can be NULL) which indicates that an event task with higher priority than currently running task has been unblocked by the posted event; a context switch should be requested before the interrupt is existed.

Returns

- ESP_OK: Success
- ESP_FAIL: Event queue for the default event loop full
- ESP_ERR_INVALID_ARG: Invalid combination of event base and event ID, data size of more than 4 bytes
- Others: Fail

esp_err_t **esp_event_isr_post_to** (*esp_event_loop_handle_t* event_loop, esp_event_base_t event_base, int32_t event_id, const void *event_data, size_t event_data_size, BaseType_t *task_unblocked)

Special variant of esp_event_post_to for posting events from interrupt handlers.

Note: this function is only available when CONFIG_ESP_EVENT_POST_FROM_ISR is enabled

Note: when this function is called from an interrupt handler placed in IRAM, this function should be placed in IRAM as well by enabling CONFIG_ESP_EVENT_POST_FROM_IRAM_ISR

Parameters

- **event_loop** -- **[in]** the event loop to post to, must not be NULL
- **event_base** -- **[in]** the event base that identifies the event
- **event_id** -- **[in]** the event ID that identifies the event
- **event_data** -- **[in]** the data, specific to the event occurrence, that gets passed to the handler
- **event_data_size** -- **[in]** the size of the event data
- **task_unblocked** -- **[out]** an optional parameter (can be NULL) which indicates that an event task with higher priority than currently running task has been unblocked by the posted event; a context switch should be requested before the interrupt is existed.

Returns

- ESP_OK: Success
- ESP_FAIL: Event queue for the loop full
- ESP_ERR_INVALID_ARG: Invalid combination of event base and event ID, data size of more than 4 bytes
- Others: Fail

esp_err_t **esp_event_dump** (FILE *file)

Dumps statistics of all event loops.

Dumps event loop info in the format:

```

event loop
  handler
  handler
  ...
event loop
  handler
  handler
  ...

where:

event loop
  format: address,name rx:total_received dr:total_dropped
  where:
    address - memory address of the event loop
    name - name of the event loop, 'none' if no dedicated task
    total_received - number of successfully posted events
    total_dropped - number of events unsuccessfully posted due to queue.
↳being full

handler
  format: address ev:base,id inv:total_invoked run:total_runtime
  where:
    address - address of the handler function
    base,id - the event specified by event base and ID this handler.
↳executes
    total_invoked - number of times this handler has been invoked
    total_runtime - total amount of time used for invoking this handler

```

Note: this function is a noop when CONFIG_ESP_EVENT_LOOP_PROFILING is disabled

Parameters **file** -- [in] the file stream to output to

Returns

- ESP_OK: Success
- ESP_ERR_NO_MEM: Cannot allocate memory for event loops list
- Others: Fail

Structures

struct **esp_event_loop_args_t**

Configuration for creating event loops.

Public Members

int32_t **queue_size**

size of the event loop queue

const char ***task_name**

name of the event loop task; if NULL, a dedicated task is not created for event loop

UBaseType_t **task_priority**

priority of the event loop task, ignored if task name is NULL

uint32_t **task_stack_size**

stack size of the event loop task, ignored if task name is NULL

BaseType_t **task_core_id**

core to which the event loop task is pinned to, ignored if task name is NULL

Header File

- [components/esp_event/include/esp_event_base.h](#)
- This header file can be included with:

```
#include "esp_event_base.h"
```

- This header file is a part of the API provided by the `esp_event` component. To declare that your component depends on `esp_event`, add the following to your CMakeLists.txt:

```
REQUIRES esp_event
```

or

```
PRIV_REQUIRES esp_event
```

Macros

ESP_EVENT_DECLARE_BASE (id)

ESP_EVENT_DEFINE_BASE (id)

ESP_EVENT_ANY_BASE

register handler for any event base

ESP_EVENT_ANY_ID

register handler for any event id

Type Definitions

typedef void ***esp_event_loop_handle_t**

a number that identifies an event with respect to a base

typedef void (***esp_event_handler_t**)(void *event_handler_arg, esp_event_base_t event_base, int32_t event_id, void *event_data)

function called when an event is posted to the queue

typedef void ***esp_event_handler_instance_t**

context identifying an instance of a registered event handler

Related Documents**2.9.11 FreeRTOS Overview****Overview**

FreeRTOS is an open source RTOS (real-time operating system) kernel that is integrated into ESP-IDF as a component. Thus, all ESP-IDF applications and many ESP-IDF components are written based on FreeRTOS. The FreeRTOS kernel is ported to all architectures (i.e., Xtensa and RISC-V) available of ESP chips.

Furthermore, ESP-IDF provides different implementations of FreeRTOS in order to support SMP (Symmetric Multiprocessing) on multi-core ESP chips. This document provides an overview of the FreeRTOS component, the different FreeRTOS implementations offered by ESP-IDF, and the common aspects across all implementations.

Implementations

The [official FreeRTOS](#) (henceforth referred to as Vanilla FreeRTOS) is a single-core RTOS. In order to support the various multi-core ESP targets, ESP-IDF supports different FreeRTOS implementations as listed below:

ESP-IDF FreeRTOS ESP-IDF FreeRTOS is a FreeRTOS implementation based on Vanilla FreeRTOS v10.5.1, but contains significant modifications to support SMP. ESP-IDF FreeRTOS only supports two cores at most (i.e., dual core SMP), but is more optimized for this scenario by design. For more details regarding ESP-IDF FreeRTOS and its modifications, please refer to the [FreeRTOS \(IDF\)](#) document.

Note: ESP-IDF FreeRTOS is currently the default FreeRTOS implementation for ESP-IDF.

Configuration

Kernel Configuration Vanilla FreeRTOS requires that ports and applications configure the kernel by adding various `#define config...` macro definitions to the `FreeRTOSConfig.h` header file. Vanilla FreeRTOS supports a list of kernel configuration options which allow various kernel behaviors and features to be enabled or disabled.

However, for all FreeRTOS ports in ESP-IDF, the `FreeRTOSConfig.h` header file is considered private and must not be modified by users. A large number of kernel configuration options in `FreeRTOSConfig.h` are hard-coded as they are either required/not supported by ESP-IDF. All kernel configuration options that are configurable by the user are exposed via menuconfig under `Component Config/FreeRTOS/Kernel`.

For the full list of user configurable kernel options, see [Project Configuration](#). The list below highlights some commonly used kernel configuration options:

- `CONFIG_FREERTOS_UNICORE` runs FreeRTOS only on Core 0. Note that this is **not equivalent to running Vanilla FreeRTOS**. Furthermore, this option may affect behavior of components other than `freertos`. For more details regarding the effects of running FreeRTOS on a single core, refer to [Single-Core Mode](#) (if using ESP-IDF FreeRTOS) or the official Amazon SMP FreeRTOS documentation. Alternatively, users can also search for occurrences of `CONFIG_FREERTOS_UNICORE` in the ESP-IDF components.
- `CONFIG_FREERTOS_ENABLE_BACKWARD_COMPATIBILITY` enables backward compatibility with some FreeRTOS macros/types/functions that were deprecated from v8.0 onwards.

Port Configuration All other FreeRTOS related configuration options that are not part of the kernel configuration are exposed via menuconfig under `Component Config/FreeRTOS/Port`. These options configure aspects such as:

- The FreeRTOS ports themselves (e.g., tick timer selection, ISR stack size)
- Additional features added to the FreeRTOS implementation or ports

Using FreeRTOS

Application Entry Point Unlike Vanilla FreeRTOS, users of FreeRTOS in ESP-IDF **must never call** `vTaskStartScheduler()` and `vTaskEndScheduler()`. Instead, ESP-IDF starts FreeRTOS automatically. Users must define a `void app_main(void)` function which acts as the entry point for user's application and is automatically invoked on ESP-IDF startup.

- Typically, users would spawn the rest of their application's task from `app_main`.
- The `app_main` function is allowed to return at any point (i.e., before the application terminates).
- The `app_main` function is called from the `main` task.

Background Tasks During startup, ESP-IDF and the FreeRTOS kernel automatically create multiple tasks that run in the background (listed in the table below).

Table 5: List of Tasks Created During Startup

Task Name	Description	Stack Size	Affinity	Priority
Idle Tasks (IDLE _x)	An idle task (IDLE _x) is created for (and pinned to) each core, where _x is the core's number. _x is dropped when single-core configuration is enabled.	CONFIG_FREERTOS_IDLE_TASK_STACK_SIZE	Core 0	0
FreeRTOS Timer Task (Tmr Svc)	FreeRTOS will create the Timer Service/Daemon Task if any FreeRTOS Timer APIs are called by the application	CONFIG_FREERTOS_TIMER_TASK_STACK_SIZE	Core 0	CONFIG_FREERTOS_TIMER_TASK_PRIORITY
Main Task (main)	Task that simply calls app_main. This task will self delete when app_main returns	CONFIG_ESP_MAIN_TASK_STACK_SIZE	Core 0	1
IPC Tasks (ipcx)	When CONFIG_FREERTOS_UNICORE is false, an IPC task (ipcx) is created for (and pinned to) each core. IPC tasks are used to implement the Inter-processor Call (IPC) feature.	CONFIG_ESP_IPC_TASK_STACK_SIZE	Core 24	24
ESP Timer Task (esp_timer)	ESP-IDF creates the ESP Timer Task used to process ESP Timer callbacks	CONFIG_ESP_TIMER_TASK_STACK_SIZE	Core 22	22

Note: Note that if an application uses other ESP-IDF features (e.g., Wi-Fi or Bluetooth), those features may create their own background tasks in addition to the tasks listed in the table above.

FreeRTOS Additions

ESP-IDF provides some supplemental features to FreeRTOS such as Ring Buffers, ESP-IDF style Tick and Idle Hooks, and TLSP deletion callbacks. See [FreeRTOS \(Supplemental Features\)](#) for more details.

FreeRTOS Heap

Vanilla FreeRTOS provides its own [selection of heap implementations](#). However, ESP-IDF already implements its own heap (see [Heap Memory Allocation](#)), thus ESP-IDF does not make use of the heap implementations provided by Vanilla FreeRTOS. All FreeRTOS ports in ESP-IDF map FreeRTOS memory allocation or free calls (e.g., `pvPortMalloc()` and `pvPortFree()`) to ESP-IDF heap API (i.e., [heap_caps_malloc\(\)](#) and [heap_caps_free\(\)](#)). However, the FreeRTOS ports ensure that all dynamic memory allocated by FreeRTOS is placed in internal memory.

Note: If users wish to place FreeRTOS tasks/objects in external memory, users can use the following methods:

- Allocate the task or object using one of the `...CreateWithCaps()` API, such as [xTaskCreateWithCaps\(\)](#) and [xQueueCreateWithCaps\(\)](#) (see [IDF Additional API](#) for more details).
- Manually allocate external memory for those objects using [heap_caps_malloc\(\)](#), then create the objects from the allocated memory using one of the `...CreateStatic()` FreeRTOS functions.

2.9.12 FreeRTOS (IDF)

This document provides information regarding the dual-core SMP implementation of FreeRTOS inside ESP-IDF. This document is split into the following sections:

Sections

- *FreeRTOS (IDF)*
 - *Overview*
 - *Symmetric Multiprocessing*
 - *Tasks*
 - *SMP Scheduler*
 - *Critical Sections*
 - *Misc*
 - *Single-Core Mode*
 - *API Reference*

Overview

The original FreeRTOS (hereinafter referred to as **Vanilla FreeRTOS**) is a compact and efficient real-time operating system supported on numerous single-core MCUs and SoCs. However, to support dual-core ESP targets, such as ESP32, ESP32-S3, and ESP32-P4, ESP-IDF provides a unique implementation of FreeRTOS with dual-core symmetric multiprocessing (SMP) capabilities (hereinafter referred to as **IDF FreeRTOS**).

IDF FreeRTOS source code is based on Vanilla FreeRTOS v10.5.1 but contains significant modifications to both kernel behavior and API in order to support dual-core SMP. However, IDF FreeRTOS can also be configured for single-core by enabling the `CONFIG_FREERTOS_UNICORE` option (see *Single-Core Mode* for more details).

Note: This document assumes that the reader has a requisite understanding of Vanilla FreeRTOS, i.e., its features, behavior, and API usage. Refer to the [Vanilla FreeRTOS documentation](#) for more details.

Symmetric Multiprocessing

Basic Concepts Symmetric multiprocessing is a computing architecture where two or more identical CPU cores are connected to a single shared main memory and controlled by a single operating system. In general, an SMP system:

- has multiple cores running independently. Each core has its own register file, interrupts, and interrupt handling.
- presents an identical view of memory to each core. Thus, a piece of code that accesses a particular memory address has the same effect regardless of which core it runs on.

The main advantages of an SMP system compared to single-core or asymmetric multiprocessing systems are that:

- the presence of multiple cores allows for multiple hardware threads, thus increasing overall processing throughput.
- having symmetric memory means that threads can switch cores during execution. This, in general, can lead to better CPU utilization.

Although an SMP system allows threads to switch cores, there are scenarios where a thread must/should only run on a particular core. Therefore, threads in an SMP system also have a core affinity that specifies which particular core the thread is allowed to run on.

- A thread that is pinned to a particular core is only able to run on that core.
- A thread that is unpinned will be allowed to switch between cores during execution instead of being pinned to a particular core.

SMP on an ESP Target ESP targets such as ESP32, ESP32-S3, and ESP32-P4 are dual-core SMP SoCs. These targets have the following hardware features that make them SMP-capable:

- Two identical cores are known as Core 0 and Core 1. This means that the execution of a piece of code is identical regardless of which core it runs on.
- Symmetric memory (with some small exceptions).
 - If multiple cores access the same memory address simultaneously, their access will be serialized by the memory bus.
 - True atomic access to the same memory address is achieved via an atomic compare-and-swap instruction provided by the ISA.
- Cross-core interrupts that allow one core to trigger an interrupt on the other core. This allows cores to signal events to each other (such as requesting a context switch on the other core).

Note: Within ESP-IDF, Core 0 and Core 1 are sometimes referred to as `PRO_CPU` and `APP_CPU` respectively. The aliases exist in ESP-IDF as they reflect how typical ESP-IDF applications utilize the two cores. Typically, the tasks responsible for handling protocol related processing such as Wi-Fi or Bluetooth are pinned to Core 0 (thus the name `PRO_CPU`), where as the tasks handling the remainder of the application are pinned to Core 1, (thus the name `APP_CPU`).

Tasks

Creation Vanilla FreeRTOS provides the following functions to create a task:

- `xTaskCreate()` creates a task. The task's memory is dynamically allocated.
- `xTaskCreateStatic()` creates a task. The task's memory is statically allocated, i.e., provided by the user.

However, in an SMP system, tasks need to be assigned a particular affinity. Therefore, ESP-IDF provides a `...PinnedToCore()` version of Vanilla FreeRTOS's task creation functions:

- `xTaskCreatePinnedToCore()` creates a task with a particular core affinity. The task's memory is dynamically allocated.
- `xTaskCreateStaticPinnedToCore()` creates a task with a particular core affinity. The task's memory is statically allocated, i.e., provided by the user.

The `...PinnedToCore()` versions of the task creation function API differ from their vanilla counterparts by having an extra `xCoreID` parameter that is used to specify the created task's core affinity. The valid values for core affinity are:

- 0, which pins the created task to Core 0
- 1, which pins the created task to Core 1
- `tskNO_AFFINITY`, which allows the task to be run on both cores

Note that IDF FreeRTOS still supports the vanilla versions of the task creation functions. However, these standard functions have been modified to essentially invoke their respective `...PinnedToCore()` counterparts while setting the core affinity to `tskNO_AFFINITY`.

Note: IDF FreeRTOS also changes the units of `ulStackDepth` in the task creation functions. Task stack sizes in Vanilla FreeRTOS are specified in a number of words, whereas in IDF FreeRTOS, the task stack sizes are specified in bytes.

Execution The anatomy of a task in IDF FreeRTOS is the same as in Vanilla FreeRTOS. More specifically, IDF FreeRTOS tasks:

- Can only be in one of the following states: Running, Ready, Blocked, or Suspended.
- Task functions are typically implemented as an infinite loop.
- Task functions should never return.

Deletion Task deletion in Vanilla FreeRTOS is called via `vTaskDelete()`. The function allows deletion of another task or the currently running task if the provided task handle is `NULL`. The actual freeing of the task's memory is sometimes delegated to the idle task if the task being deleted is the currently running task.

IDF FreeRTOS provides the same `vTaskDelete()` function. However, due to the dual-core nature, there are some behavioral differences when calling `vTaskDelete()` in IDF FreeRTOS:

- When deleting a task that is currently running on the other core, a yield is triggered on the other core, and the task's memory is freed by one of the idle tasks.
- A deleted task's memory is freed immediately if it is not running on either core.

Please avoid deleting a task that is running on another core as it is difficult to determine what the task is performing, which may lead to unpredictable behavior such as:

- Deleting a task that is holding a mutex.
- Deleting a task that has yet to free memory it previously allocated.

Where possible, please design your own application so that when calling `vTaskDelete()`, the deleted task is in a known state. For example:

- Tasks self-deleting via `vTaskDelete(NULL)` when their execution is complete and have also cleaned up all resources used within the task.
- Tasks placing themselves in the suspend state via `vTaskSuspend()` before being deleted by another task.

SMP Scheduler

The Vanilla FreeRTOS scheduler is best described as a **fixed priority preemptive scheduler with time slicing** meaning that:

- Each task is given a constant priority upon creation. The scheduler executes the highest priority ready-state task.
- The scheduler can switch execution to another task without the cooperation of the currently running task.
- The scheduler periodically switches execution between ready-state tasks of the same priority in a round-robin fashion. Time slicing is governed by a tick interrupt.

The IDF FreeRTOS scheduler supports the same scheduling features, i.e., Fixed Priority, Preemption, and Time Slicing, albeit with some small behavioral differences.

Fixed Priority In Vanilla FreeRTOS, when the scheduler selects a new task to run, it always selects the current highest priority ready-state task. In IDF FreeRTOS, each core independently schedules tasks to run. When a particular core selects a task, the core will select the highest priority ready-state task that can be run by the core. A task can be run by the core if:

- The task has a compatible affinity, i.e., is either pinned to that core or is unpinned.
- The task is not currently being run by another core.

However, please do not assume that the two highest priority ready-state tasks are always run by the scheduler, as a task's core affinity must also be accounted for. For example, given the following tasks:

- Task A of priority 10 pinned to Core 0
- Task B of priority 9 pinned to Core 0
- Task C of priority 8 pinned to Core 1

The resulting schedule will have Task A running on Core 0 and Task C running on Core 1. Task B is not run even though it is the second-highest priority task.

Preemption In Vanilla FreeRTOS, the scheduler can preempt the currently running task if a higher priority task becomes ready to execute. Likewise in IDF FreeRTOS, each core can be individually preempted by the scheduler if the scheduler determines that a higher-priority task can run on that core.

However, there are some instances where a higher-priority task that becomes ready can be run on multiple cores. In this case, the scheduler only preempts one core. The scheduler always gives preference to the current core when

multiple cores can be preempted. In other words, if the higher priority ready task is unpinned and has a higher priority than the current priority of both cores, the scheduler will always choose to preempt the current core. For example, given the following tasks:

- Task A of priority 8 currently running on Core 0
- Task B of priority 9 currently running on Core 1
- Task C of priority 10 that is unpinned and was unblocked by Task B

The resulting schedule will have Task A running on Core 0 and Task C preempting Task B given that the scheduler always gives preference to the current core.

Time Slicing The Vanilla FreeRTOS scheduler implements time slicing, which means that if the current highest ready priority contains multiple ready tasks, the scheduler will switch between those tasks periodically in a round-robin fashion.

However, in IDF FreeRTOS, it is not possible to implement perfect Round Robin time slicing due to the fact that a particular task may not be able to run on a particular core due to the following reasons:

- The task is pinned to another core.
- For unpinned tasks, the task is already being run by another core.

Therefore, when a core searches the ready-state task list for a task to run, the core may need to skip over a few tasks in the same priority list or drop to a lower priority in order to find a ready-state task that the core can run.

The IDF FreeRTOS scheduler implements a Best Effort Round Robin time slicing for ready-state tasks of the same priority by ensuring that tasks that have been selected to run are placed at the back of the list, thus giving unselected tasks a higher priority on the next scheduling iteration (i.e., the next tick interrupt or yield).

The following example demonstrates the Best Effort Round Robin time slicing in action. Assume that:

- There are four ready-state tasks of the same priority AX, B0, C1, and D1 where:
 - The priority is the current highest priority with ready-state .
 - The first character represents the task's name, i.e., A, B, C, D.
 - The second character represents the task's core pinning, and X means unpinned.
- The task list is always searched from the head.

1. Starting state. None of the ready-state tasks have been selected to run.

```
Head [ AX , B0 , C1 , D0 ] Tail
```

2. Core 0 has a tick interrupt and searches for a task to run. Task A is selected and moved to the back of the list.

```
Core 0 ┌
      │
      ▼
Head [ AX , B0 , C1 , D0 ] Tail

                [0]
Head [ B0 , C1 , D0 , AX ] Tail
```

3. Core 1 has a tick interrupt and searches for a task to run. Task B cannot be run due to incompatible affinity, so Core 1 skips to Task C. Task C is selected and moved to the back of the list.

```
Core 1 ┌───┐
      │   │
      ▼   [0]
Head [ B0 , C1 , D0 , AX ] Tail

                [0] [1]
Head [ B0 , D0 , AX , C1 ] Tail
```

4. Core 0 has another tick interrupt and searches for a task to run. Task B is selected and moved to the back of the list.

```
Core 0 ┌
      │
      ▼
Head [ B0 , D0 , AX , C1 ] Tail

                [1]
```

(continues on next page)

(continued from previous page)

```

          [1]  [0]
Head [ D0 , AX , C1 , B0 ] Tail

```

5. Core 1 has another tick and searches for a task to run. Task D cannot be run due to incompatible affinity, so Core 1 skips to Task A. Task A is selected and moved to the back of the list.

```

Core 1 ————
          ▼
          [0]
Head [ D0 , AX , C1 , B0 ] Tail

          [0]  [1]
Head [ D0 , C1 , B0 , AX ] Tail

```

The implications to users regarding the Best Effort Round Robin time slicing:

- Users cannot expect multiple ready-state tasks of the same priority to run sequentially as is the case in Vanilla FreeRTOS. As demonstrated in the example above, a core may need to skip over tasks.
- However, given enough ticks, a task will eventually be given some processing time.
- If a core cannot find a task runnable task at the highest ready-state priority, it will drop to a lower priority to search for tasks.
- To achieve ideal round-robin time slicing, users should ensure that all tasks of a particular priority are pinned to the same core.

Tick Interrupts Vanilla FreeRTOS requires that a periodic tick interrupt occurs. The tick interrupt is responsible for:

- Incrementing the scheduler's tick count
- Unblocking any blocked tasks that have timed out
- Checking if time slicing is required, i.e., triggering a context switch
- Executing the application tick hook

In IDF FreeRTOS, each core receives a periodic interrupt and independently runs the tick interrupt. The tick interrupts on each core are of the same period but can be out of phase. However, the tick responsibilities listed above are not run by all cores:

- Core 0 executes all of the tick interrupt responsibilities listed above
- Core 1 only checks for time slicing and executes the application tick hook

Note: Core 0 is solely responsible for keeping time in IDF FreeRTOS. Therefore, anything that prevents Core 0 from incrementing the tick count, such as suspending the scheduler on Core 0, will cause the entire scheduler's timekeeping to lag behind.

Idle Tasks Vanilla FreeRTOS will implicitly create an idle task of priority 0 when the scheduler is started. The idle task runs when no other task is ready to run, and it has the following responsibilities:

- Freeing the memory of deleted tasks
- Executing the application idle hook

In IDF FreeRTOS, a separate pinned idle task is created for each core. The idle tasks on each core have the same responsibilities as their vanilla counterparts.

Scheduler Suspension Vanilla FreeRTOS allows the scheduler to be suspended/resumed by calling `vTaskSuspendAll()` and `xTaskResumeAll()` respectively. While the scheduler is suspended:

- Task switching is disabled but interrupts are left enabled.
- Calling any blocking/yielding function is forbidden, and time slicing is disabled.
- The tick count is frozen, but the tick interrupt still occurs to execute the application tick hook.

On scheduler resumption, `xTaskResumeAll()` catches up all of the lost ticks and unblock any timed-out tasks.

In IDF FreeRTOS, suspending the scheduler across multiple cores is not possible. Therefore when `vTaskSuspendAll()` is called on a particular core (e.g., core A):

- Task switching is disabled only on core A but interrupts for core A are left enabled.
- Calling any blocking/yielding function on core A is forbidden. Time slicing is disabled on core A.
- If an interrupt on core A unblocks any tasks, tasks with affinity to core A will go into core A's own pending ready task list. Unpinned tasks or tasks with affinity to other cores can be scheduled on cores with the scheduler running.
- If the scheduler is suspended on all cores, tasks unblocked by an interrupt will be directed to the pending ready task lists of their pinned cores. For unpinned tasks, they will be placed in the pending ready list of the core where the interrupt occurred.
- If core A is on Core 0, the tick count is frozen, and a pended tick count is incremented instead. However, the tick interrupt will still occur in order to execute the application tick hook.

When `xTaskResumeAll()` is called on a particular core (e.g., core A):

- Any tasks added to core A's pending ready task list will be resumed.
- If core A is Core 0, the pended tick count is unwound to catch up with the lost ticks.

Warning: Given that scheduler suspension on IDF FreeRTOS only suspends scheduling on a particular core, scheduler suspension is **NOT** a valid method of ensuring mutual exclusion between tasks when accessing shared data. Users should use proper locking primitives such as mutexes or spinlocks if they require mutual exclusion.

Critical Sections

Disabling Interrupts Vanilla FreeRTOS allows interrupts to be disabled and enabled by calling `taskDISABLE_INTERRUPTS` and `taskENABLE_INTERRUPTS` respectively. IDF FreeRTOS provides the same API. However, interrupts are only disabled or enabled on the current core.

Disabling interrupts is a valid method of achieving mutual exclusion in Vanilla FreeRTOS (and single-core systems in general). **However, in an SMP system, disabling interrupts is not a valid method of ensuring mutual exclusion.** Critical sections that utilize a spinlock should be used instead.

API Changes Vanilla FreeRTOS implements critical sections by disabling interrupts, which prevents preemptive context switches and the servicing of ISRs during a critical section. Thus a task/ISR that enters a critical section is guaranteed to be the sole entity to access a shared resource. Critical sections in Vanilla FreeRTOS have the following API:

- `taskENTER_CRITICAL()` enters a critical section by disabling interrupts
- `taskEXIT_CRITICAL()` exits a critical section by reenabling interrupts
- `taskENTER_CRITICAL_FROM_ISR()` enters a critical section from an ISR by disabling interrupt nesting
- `taskEXIT_CRITICAL_FROM_ISR()` exits a critical section from an ISR by reenabling interrupt nesting

However, in an SMP system, merely disabling interrupts does not constitute a critical section as the presence of other cores means that a shared resource can still be concurrently accessed. Therefore, critical sections in IDF FreeRTOS are implemented using spinlocks. To accommodate the spinlocks, the IDF FreeRTOS critical section APIs contain an additional spinlock parameter as shown below:

- Spinlocks are of `portMUX_TYPE` (**not to be confused to FreeRTOS mutexes**)
- `taskENTER_CRITICAL(&spinlock)` enters a critical from a task context
- `taskEXIT_CRITICAL(&spinlock)` exits a critical section from a task context
- `taskENTER_CRITICAL_ISR(&spinlock)` enters a critical section from an interrupt context
- `taskEXIT_CRITICAL_ISR(&spinlock)` exits a critical section from an interrupt context

Note: The critical section API can be called recursively, i.e., nested critical sections. Entering a critical section multiple times recursively is valid so long as the critical section is exited the same number of times it was entered.

However, given that critical sections can target different spinlocks, users should take care to avoid deadlocking when entering critical sections recursively.

Spinlocks can be allocated statically or dynamically. As such, macros are provided for both static and dynamic initialization of spinlocks, as demonstrated by the following code snippets.

- Allocating a static spinlock and initializing it using `portMUX_INITIALIZER_UNLOCKED`:

```
// Statically allocate and initialize the spinlock
static portMUX_TYPE my_spinlock = portMUX_INITIALIZER_UNLOCKED;

void some_function(void)
{
    taskENTER_CRITICAL(&my_spinlock);
    // We are now in a critical section
    taskEXIT_CRITICAL(&my_spinlock);
}
```

- Allocating a dynamic spinlock and initializing it using `portMUX_INITIALIZE()`:

```
// Allocate the spinlock dynamically
portMUX_TYPE *my_spinlock = malloc(sizeof(portMUX_TYPE));
// Initialize the spinlock dynamically
portMUX_INITIALIZE(my_spinlock);

...

taskENTER_CRITICAL(my_spinlock);
// Access the resource
taskEXIT_CRITICAL(my_spinlock);
```

Implementation In IDF FreeRTOS, the process of a particular core entering and exiting a critical section is as follows:

- For `taskENTER_CRITICAL(&spinlock)` or `taskENTER_CRITICAL_ISR(&spinlock)`
 1. The core disables its interrupts or interrupt nesting up to `configMAX_SYSCALL_INTERRUPT_PRIORITY`.
 2. The core then spins on the spinlock using an atomic compare-and-set instruction until it acquires the lock. A lock is acquired when the core is able to set the lock's owner value to the core's ID.
 3. Once the spinlock is acquired, the function returns. The remainder of the critical section runs with interrupts or interrupt nesting disabled.
- For `taskEXIT_CRITICAL(&spinlock)` or `taskEXIT_CRITICAL_ISR(&spinlock)`
 1. The core releases the spinlock by clearing the spinlock's owner value.
 2. The core re-enables interrupts or interrupt nesting.

Restrictions and Considerations Given that interrupts (or interrupt nesting) are disabled during a critical section, there are multiple restrictions regarding what can be done within critical sections. During a critical section, users should keep the following restrictions and considerations in mind:

- Critical sections should be kept as short as possible
 - The longer the critical section lasts, the longer a pending interrupt can be delayed.
 - A typical critical section should only access a few data structures and/or hardware registers.
 - If possible, defer as much processing and/or event handling to the outside of critical sections.
- FreeRTOS API should not be called from within a critical section
- Users should never call any blocking or yielding functions within a critical section

Misc

Floating Point Usage Usually, when a context switch occurs:

- the current state of a core's registers are saved to the stack of the task being switched out
- the previously saved state of the core's registers is loaded from the stack of the task being switched in

However, IDF FreeRTOS implements Lazy Context Switching for the Floating Point Unit (FPU) registers of a core. In other words, when a context switch occurs on a particular core (e.g., Core 0), the state of the core's FPU registers is not immediately saved to the stack of the task getting switched out (e.g., Task A). The FPU registers are left untouched until:

- A different task (e.g., Task B) runs on the same core and uses FPU. This will trigger an exception that saves the FPU registers to Task A's stack.
- Task A gets scheduled to the same core and continues execution. Saving and restoring the FPU registers is not necessary in this case.

However, given that tasks can be unpinned and thus can be scheduled on different cores (e.g., Task A switches to Core 1), it is unfeasible to copy and restore the FPU registers across cores. Therefore, when a task utilizes FPU by using a `float` type in its call flow, IDF FreeRTOS will automatically pin the task to the current core it is running on. This ensures that all tasks that use FPU are always pinned to a particular core.

Furthermore, IDF FreeRTOS by default does not support the usage of FPU within an interrupt context given that the FPU register state is tied to a particular task.

Note: ESP targets that contain an FPU do not support hardware acceleration for double precision floating point arithmetic (`double`). Instead, `double` is implemented via software, hence the behavioral restrictions regarding the `float` type do not apply to `double`. Note that due to the lack of hardware acceleration, `double` operations may consume significantly more CPU time in comparison to `float`.

Single-Core Mode

Although IDF FreeRTOS is modified for dual-core SMP, IDF FreeRTOS can also be built for single-core by enabling the `CONFIG_FREERTOS_UNICORE` option.

For single-core targets (such as ESP32-S2 and ESP32-C3), the `CONFIG_FREERTOS_UNICORE` option is always enabled. For multi-core targets (such as ESP32 and ESP32-S3), `CONFIG_FREERTOS_UNICORE` can also be set, but will result in the application only running Core 0.

When building in single-core mode, IDF FreeRTOS is designed to be identical to Vanilla FreeRTOS, thus all aforementioned SMP changes to kernel behavior are removed. As a result, building IDF FreeRTOS in single-core mode has the following characteristics:

- All operations performed by the kernel inside critical sections are now deterministic (i.e., no walking of linked lists inside critical sections).
- Vanilla FreeRTOS scheduling algorithm is restored (including perfect Round Robin time slicing).
- All SMP specific data is removed from single-core builds.

SMP APIs can still be called in single-core mode. These APIs remain exposed to allow source code to be built for single-core and multi-core, without needing to call a different set of APIs. However, SMP APIs will not exhibit any SMP behavior in single-core mode, thus becoming equivalent to their single-core counterparts. For example:

- any `...ForCore(..., BaseType_t xCoreID)` SMP API will only accept 0 as a valid value for `xCoreID`.
- `...PinnedToCore()` task creation APIs will simply ignore the `xCoreID` core affinity argument.
- Critical section APIs will still require a spinlock argument, but no spinlock will be taken and critical sections revert to simply disabling/enabling interrupts.

API Reference

This section introduces FreeRTOS types, functions, and macros. It is automatically generated from FreeRTOS header files.

Task API

Header File

- `components/freertos/FreeRTOS-Kernel/include/freertos/task.h`
- This header file can be included with:

```
#include "freertos/task.h"
```

Functions

static inline BaseType_t **xTaskCreate** (TaskFunction_t pxTaskCode, const char *const pcName, const configSTACK_DEPTH_TYPE usStackDepth, void *const pvParameters, UBaseType_t uxPriority, *TaskHandle_t* *const pxCreatedTask)

Create a new task and add it to the list of tasks that are ready to run.

Internally, within the FreeRTOS implementation, tasks use two blocks of memory. The first block is used to hold the task's data structures. The second block is used by the task as its stack. If a task is created using `xTaskCreate()` then both blocks of memory are automatically dynamically allocated inside the `xTaskCreate()` function. (see <https://www.FreeRTOS.org/a00111.html>). If a task is created using `xTaskCreateStatic()` then the application writer must provide the required memory. `xTaskCreateStatic()` therefore allows a task to be created without using any dynamic memory allocation.

See `xTaskCreateStatic()` for a version that does not use any dynamic memory allocation.

`xTaskCreate()` can only be used to create a task that has unrestricted access to the entire microcontroller memory map. Systems that include MPU support can alternatively create an MPU constrained task using `xTaskCreateRestricted()`.

Example usage:

```
// Task to be created.
void vTaskCode( void * pvParameters )
{
    for( ;; )
    {
        // Task code goes here.
    }
}

// Function that creates a task.
void vOtherFunction( void )
{
    static uint8_t ucParameterToPass;
    TaskHandle_t xHandle = NULL;

    // Create the task, storing the handle. Note that the passed parameter
    ↪ucParameterToPass
    // must exist for the lifetime of the task, so in this case is declared
    ↪static. If it was just an
    // an automatic stack variable it might no longer exist, or at least have
    ↪been corrupted, by the time
    // the new task attempts to access it.
    xTaskCreate( vTaskCode, "NAME", STACK_SIZE, &ucParameterToPass, tskIDLE_
    ↪PRIORITY, &xHandle );
    configASSERT( xHandle );

    // Use the handle to delete the task.
    if( xHandle != NULL )
```

(continues on next page)

(continued from previous page)

```

{
    vTaskDelete( xHandle );
}
}

```

Note: If `configNUMBER_OF_CORES > 1`, this function will create an unpinned task (see `tskNO_AFFINITY` for more details).

Note: If program uses thread local variables (ones specified with "`__thread`" keyword) then storage for them will be allocated on the task's stack.

Parameters

- **pxTaskCode** -- Pointer to the task entry function. Tasks must be implemented to never return (i.e. continuous loop).
- **pcName** -- A descriptive name for the task. This is mainly used to facilitate debugging. Max length defined by `configMAX_TASK_NAME_LEN` - default is 16.
- **usStackDepth** -- The size of the task stack specified as the NUMBER OF BYTES. Note that this differs from vanilla FreeRTOS.
- **pvParameters** -- Pointer that will be used as the parameter for the task being created.
- **uxPriority** -- The priority at which the task should run. Systems that include MPU support can optionally create tasks in a privileged (system) mode by setting bit `portPRIVILEGE_BIT` of the priority parameter. For example, to create a privileged task at priority 2 the `uxPriority` parameter should be set to `( 2 | portPRIVILEGE_BIT )`.
- **pxCreatedTask** -- Used to pass back a handle by which the created task can be referenced.

Returns `pdPASS` if the task was successfully created and added to a ready list, otherwise an error code defined in the file `projdefs.h`

```

static inline TaskHandle_t xTaskCreateStatic (TaskFunction_t pxTaskCode, const char *const pcName,
                                             const uint32_t ulStackDepth, void *const pvParameters,
                                             UBaseType_t uxPriority, StackType_t *const
                                             puxStackBuffer, StaticTask_t *const pxTaskBuffer)

```

Create a new task and add it to the list of tasks that are ready to run.

Internally, within the FreeRTOS implementation, tasks use two blocks of memory. The first block is used to hold the task's data structures. The second block is used by the task as its stack. If a task is created using `xTaskCreate()` then both blocks of memory are automatically dynamically allocated inside the `xTaskCreate()` function. (see <https://www.FreeRTOS.org/a00111.html>). If a task is created using `xTaskCreateStatic()` then the application writer must provide the required memory. `xTaskCreateStatic()` therefore allows a task to be created without using any dynamic memory allocation.

Example usage:

```

// Dimensions the buffer that the task being created will use as its stack.
// NOTE: This is the number of bytes the stack will hold, not the number of
// words as found in vanilla FreeRTOS.
#define STACK_SIZE 200

// Structure that will hold the TCB of the task being created.
StaticTask_t xTaskBuffer;

// Buffer that the task being created will use as its stack. Note this is
// an array of StackType_t variables. The size of StackType_t is dependent_
→ on

```

(continues on next page)

(continued from previous page)

```

// the RTOS port.
StackType_t xStack[ STACK_SIZE ];

// Function that implements the task being created.
void vTaskCode( void * pvParameters )
{
    // The parameter value is expected to be 1 as 1 is passed in the
    // pvParameters value in the call to xTaskCreateStatic().
    configASSERT( ( uint32_t ) pvParameters == 1UL );

    for( ;; )
    {
        // Task code goes here.
    }
}

// Function that creates a task.
void vOtherFunction( void )
{
    TaskHandle_t xHandle = NULL;

    // Create the task without using any dynamic memory allocation.
    xHandle = xTaskCreateStatic(
        vTaskCode,          // Function that implements the task.
        "NAME",             // Text name for the task.
        STACK_SIZE,        // Stack size in bytes.
        ( void * ) 1,       // Parameter passed into the task.
        tskIDLE_PRIORITY,  // Priority at which the task is created.
        xStack,             // Array to use as the task's stack.
        &xTaskBuffer );     // Variable to hold the task's data
    ↪ structure.

    // puxStackBuffer and pxTaskBuffer were not NULL, so the task will have
    // been created, and xHandle will be the task's handle. Use the handle
    // to suspend the task.
    vTaskSuspend( xHandle );
}

```

Note: If `configNUMBER_OF_CORES > 1`, this function will create an unpinned task (see `tskNO_AFFINITY` for more details).

Note: If program uses thread local variables (ones specified with `__thread` keyword) then storage for them will be allocated on the task's stack.

Parameters

- **pxTaskCode** -- Pointer to the task entry function. Tasks must be implemented to never return (i.e. continuous loop).
- **pcName** -- A descriptive name for the task. This is mainly used to facilitate debugging. The maximum length of the string is defined by `configMAX_TASK_NAME_LEN` in `FreeRTOSConfig.h`.
- **ulStackDepth** -- The size of the task stack specified as the NUMBER OF BYTES. Note that this differs from vanilla FreeRTOS.
- **pvParameters** -- Pointer that will be used as the parameter for the task being created.
- **uxPriority** -- The priority at which the task will run.
- **puvStackBuffer** -- Must point to a `StackType_t` array that has at least `ulStackDepth` indexes - the array will then be used as the task's stack, removing the need for the stack to

be allocated dynamically.

- **pxTaskBuffer** -- Must point to a variable of type StaticTask_t, which will then be used to hold the task's data structures, removing the need for the memory to be allocated dynamically.

Returns If neither puxStackBuffer nor pxTaskBuffer are NULL, then the task will be created and a handle to the created task is returned. If either puxStackBuffer or pxTaskBuffer are NULL then the task will not be created and NULL is returned.

void **vTaskAllocateMPURegions** (*TaskHandle_t* xTask, const MemoryRegion_t *const pxRegions)

Memory regions are assigned to a restricted task when the task is created by a call to xTaskCreateRestricted(). These regions can be redefined using vTaskAllocateMPURegions().

Example usage:

```
// Define an array of MemoryRegion_t structures that configures an MPU region
// allowing read/write access for 1024 bytes starting at the beginning of the
// ucOneKByte array. The other two of the maximum 3 definable regions are
// unused so set to zero.
static const MemoryRegion_t xAltRegions[ portNUM_CONFIGURABLE_REGIONS ] =
{
    // Base address      Length      Parameters
    { ucOneKByte,        1024,        portMPU_REGION_READ_WRITE },
    { 0,                  0,          0 },
    { 0,                  0,          0 }
};

void vATask( void *pvParameters )
{
    // This task was created such that it has access to certain regions of
    // memory as defined by the MPU configuration. At some point it is
    // desired that these MPU regions are replaced with that defined in the
    // xAltRegions const struct above. Use a call to vTaskAllocateMPURegions()
    // for this purpose. NULL is used as the task handle to indicate that this
    // function should modify the MPU regions of the calling task.
    vTaskAllocateMPURegions( NULL, xAltRegions );

    // Now the task can continue its function, but from this point on can only
    // access its stack and the ucOneKByte array (unless any other statically
    // defined or shared regions have been declared elsewhere).
}
```

Parameters

- **xTask** -- The handle of the task being updated.
- **pxRegions** -- A pointer to a MemoryRegion_t structure that contains the new memory region definitions.

void **vTaskDelete** (*TaskHandle_t* xTaskToDelete)

INCLUDE_vTaskDelete must be defined as 1 for this function to be available. See the configuration section for more information.

Remove a task from the RTOS real time kernel's management. The task being deleted will be removed from all ready, blocked, suspended and event lists.

NOTE: The idle task is responsible for freeing the kernel allocated memory from tasks that have been deleted. It is therefore important that the idle task is not starved of microcontroller processing time if your application makes any calls to vTaskDelete (). Memory allocated by the task code is not automatically freed, and should be freed before the task is deleted.

See the demo application file death.c for sample code that utilises vTaskDelete ().

Example usage:

```
void vOtherFunction( void )
{
    TaskHandle_t xHandle;

    // Create the task, storing the handle.
    xTaskCreate( vTaskCode, "NAME", STACK_SIZE, NULL, tskIDLE_PRIORITY, &xHandle );

    // Use the handle to delete the task.
    vTaskDelete( xHandle );
}
```

Parameters **xTaskToDelete** -- The handle of the task to be deleted. Passing NULL will cause the calling task to be deleted.

void **vTaskDelay** (const TickType_t xTicksToDelay)

Delay a task for a given number of ticks. The actual time that the task remains blocked depends on the tick rate. The constant `portTICK_PERIOD_MS` can be used to calculate real time from the tick rate - with the resolution of one tick period.

`INCLUDE_vTaskDelay` must be defined as 1 for this function to be available. See the configuration section for more information.

`vTaskDelay()` specifies a time at which the task wishes to unblock relative to the time at which `vTaskDelay()` is called. For example, specifying a block period of 100 ticks will cause the task to unblock 100 ticks after `vTaskDelay()` is called. `vTaskDelay()` does not therefore provide a good method of controlling the frequency of a periodic task as the path taken through the code, as well as other task and interrupt activity, will affect the frequency at which `vTaskDelay()` gets called and therefore the time at which the task next executes. See `xTaskDelayUntil()` for an alternative API function designed to facilitate fixed frequency execution. It does this by specifying an absolute time (rather than a relative time) at which the calling task should unblock.

Example usage:

```
void vTaskFunction( void * pvParameters )
{
    // Block for 500ms.
    const TickType_t xDelay = 500 / portTICK_PERIOD_MS;

    for( ;; )
    {
        // Simply toggle the LED every 500ms, blocking between each toggle.
        vToggleLED();
        vTaskDelay( xDelay );
    }
}
```

Parameters **xTicksToDelay** -- The amount of time, in tick periods, that the calling task should block.

BaseType_t **xTaskDelayUntil** (TickType_t *const pxPreviousWakeTime, const TickType_t xTimeIncrement)

`INCLUDE_xTaskDelayUntil` must be defined as 1 for this function to be available. See the configuration section for more information.

Delay a task until a specified time. This function can be used by periodic tasks to ensure a constant execution frequency.

This function differs from `vTaskDelay()` in one important aspect: `vTaskDelay()` will cause a task to block for the specified number of ticks from the time `vTaskDelay()` is called. It is therefore difficult to use `vTaskDelay()` by itself to generate a fixed execution frequency as the time between a task starting to execute and that task calling `vTaskDelay()` may not be fixed [the task may take a different path though the code between calls, or may get interrupted or preempted a different number of times each time it executes].

Whereas `vTaskDelay()` specifies a wake time relative to the time at which the function is called, `xTaskDelayUntil()` specifies the absolute (exact) time at which it wishes to unblock.

The macro `pdMS_TO_TICKS()` can be used to calculate the number of ticks from a time specified in milliseconds with a resolution of one tick period.

Example usage:

```
// Perform an action every 10 ticks.
void vTaskFunction( void * pvParameters )
{
    TickType_t xLastWakeTime;
    const TickType_t xFrequency = 10;
    BaseType_t xWasDelayed;

    // Initialise the xLastWakeTime variable with the current time.
    xLastWakeTime = xTaskGetTickCount ();
    for ( ;; )
    {
        // Wait for the next cycle.
        xWasDelayed = xTaskDelayUntil( &xLastWakeTime, xFrequency );

        // Perform action here. xWasDelayed value can be used to determine
        // whether a deadline was missed if the code here took too long.
    }
}
```

Parameters

- **pxPreviousWakeTime** -- Pointer to a variable that holds the time at which the task was last unblocked. The variable must be initialised with the current time prior to its first use (see the example below). Following this the variable is automatically updated within `xTaskDelayUntil()`.
- **xTimeIncrement** -- The cycle time period. The task will be unblocked at time `*pxPreviousWakeTime + xTimeIncrement`. Calling `xTaskDelayUntil` with the same `xTimeIncrement` parameter value will cause the task to execute with a fixed interface period.

Returns Value which can be used to check whether the task was actually delayed. Will be `pdTRUE` if the task was delayed and `pdFALSE` otherwise. A task will not be delayed if the next expected wake time is in the past.

`BaseType_t xTaskAbortDelay (TaskHandle_t xTask)`

`INCLUDE_xTaskAbortDelay` must be defined as 1 in `FreeRTOSConfig.h` for this function to be available.

A task will enter the Blocked state when it is waiting for an event. The event it is waiting for can be a temporal event (waiting for a time), such as when `vTaskDelay()` is called, or an event on an object, such as when `xQueueReceive()` or `ulTaskNotifyTake()` is called. If the handle of a task that is in the Blocked state is used in a call to `xTaskAbortDelay()` then the task will leave the Blocked state, and return from whichever function call placed the task into the Blocked state.

There is no 'FromISR' version of this function as an interrupt would need to know which object a task was blocked on in order to know which actions to take. For example, if the task was blocked on a queue the interrupt handler would then need to know if the queue was locked.

Parameters **xTask** -- The handle of the task to remove from the Blocked state.

Returns If the task referenced by `xTask` was not in the Blocked state then `pdFAIL` is returned. Otherwise `pdPASS` is returned.

UBaseType_t **uxTaskPriorityGet** (const *TaskHandle_t* xTask)

INCLUDE_uxTaskPriorityGet must be defined as 1 for this function to be available. See the configuration section for more information.

Obtain the priority of any task.

Example usage:

```
void vAFunction( void )
{
    TaskHandle_t xHandle;

    // Create a task, storing the handle.
    xTaskCreate( vTaskCode, "NAME", STACK_SIZE, NULL, tskIDLE_PRIORITY, &xHandle_
↪ );

    // ...

    // Use the handle to obtain the priority of the created task.
    // It was created with tskIDLE_PRIORITY, but may have changed
    // it itself.
    if( uxTaskPriorityGet( xHandle ) != tskIDLE_PRIORITY )
    {
        // The task has changed it's priority.
    }

    // ...

    // Is our priority higher than the created task?
    if( uxTaskPriorityGet( xHandle ) < uxTaskPriorityGet( NULL ) )
    {
        // Our priority (obtained using NULL handle) is higher.
    }
}
```

Parameters **xTask** -- Handle of the task to be queried. Passing a NULL handle results in the priority of the calling task being returned.

Returns The priority of xTask.

UBaseType_t **uxTaskPriorityGetFromISR** (const *TaskHandle_t* xTask)

A version of uxTaskPriorityGet() that can be used from an ISR.

eTaskState **eTaskGetState** (*TaskHandle_t* xTask)

INCLUDE_eTaskGetState must be defined as 1 for this function to be available. See the configuration section for more information.

Obtain the state of any task. States are encoded by the eTaskState enumerated type.

Parameters **xTask** -- Handle of the task to be queried.

Returns The state of xTask at the time the function was called. Note the state of the task might change between the function being called, and the functions return value being tested by the calling task.

void **vTaskGetInfo** (*TaskHandle_t* xTask, *TaskStatus_t* *pxTaskStatus, BaseType_t xGetFreeStackSize, *eTaskState* eState)

configUSE_TRACE_FACILITY must be defined as 1 for this function to be available. See the configuration section for more information.

Populates a TaskStatus_t structure with information about a task.

Example usage:

```
void vAFunction( void )
{
    TaskHandle_t xHandle;
    TaskStatus_t xTaskDetails;

    // Obtain the handle of a task from its name.
    xHandle = xTaskGetHandle( "Task_Name" );

    // Check the handle is not NULL.
    configASSERT( xHandle );

    // Use the handle to obtain further information about the task.
    vTaskGetInfo( xHandle,
                  &xTaskDetails,
                  pdTRUE, // Include the high water mark in xTaskDetails.
                  eInvalid ); // Include the task state in xTaskDetails.
}
```

Parameters

- **xTask** -- Handle of the task being queried. If xTask is NULL then information will be returned about the calling task.
- **pxTaskStatus** -- A pointer to the TaskStatus_t structure that will be filled with information about the task referenced by the handle passed using the xTask parameter.
- **xGetFreeStackSpace** -- The TaskStatus_t structure contains a member to report the stack high water mark of the task being queried. Calculating the stack high water mark takes a relatively long time, and can make the system temporarily unresponsive - so the xGetFreeStackSpace parameter is provided to allow the high water mark checking to be skipped. The high watermark value will only be written to the TaskStatus_t structure if xGetFreeStackSpace is not set to pdFALSE;
- **eState** -- The TaskStatus_t structure contains a member to report the state of the task being queried. Obtaining the task state is not as fast as a simple assignment - so the eState parameter is provided to allow the state information to be omitted from the TaskStatus_t structure. To obtain state information then set eState to eInvalid - otherwise the value passed in eState will be reported as the task state in the TaskStatus_t structure.

void **vTaskPrioritySet** (*TaskHandle_t* xTask, *UBaseType_t* uxNewPriority)

INCLUDE_vTaskPrioritySet must be defined as 1 for this function to be available. See the configuration section for more information.

Set the priority of any task.

A context switch will occur before the function returns if the priority being set is higher than the currently executing task.

Example usage:

```
void vAFunction( void )
{
    TaskHandle_t xHandle;

    // Create a task, storing the handle.
    xTaskCreate( vTaskCode, "NAME", STACK_SIZE, NULL, tskIDLE_PRIORITY, &xHandle_
    ↪ );

    // ...

    // Use the handle to raise the priority of the created task.
}
```

(continues on next page)

(continued from previous page)

```

vTaskPrioritySet( xHandle, tskIDLE_PRIORITY + 1 );

// ...

// Use a NULL handle to raise our priority to the same value.
vTaskPrioritySet( NULL, tskIDLE_PRIORITY + 1 );
}

```

Parameters

- **xTask** -- Handle to the task for which the priority is being set. Passing a NULL handle results in the priority of the calling task being set.
- **uxNewPriority** -- The priority to which the task will be set.

void **vTaskSuspend** (*TaskHandle_t* xTaskToSuspend)

INCLUDE_vTaskSuspend must be defined as 1 for this function to be available. See the configuration section for more information.

Suspend any task. When suspended a task will never get any microcontroller processing time, no matter what its priority.

Calls to vTaskSuspend are not accumulative - i.e. calling vTaskSuspend () twice on the same task still only requires one call to vTaskResume () to ready the suspended task.

Example usage:

```

void vAFunction( void )
{
    TaskHandle_t xHandle;

    // Create a task, storing the handle.
    xTaskCreate( vTaskCode, "NAME", STACK_SIZE, NULL, tskIDLE_PRIORITY, &xHandle_
    → );

    // ...

    // Use the handle to suspend the created task.
    vTaskSuspend( xHandle );

    // ...

    // The created task will not run during this period, unless
    // another task calls vTaskResume( xHandle ).

    //...

    // Suspend ourselves.
    vTaskSuspend( NULL );

    // We cannot get here unless another task calls vTaskResume
    // with our handle as the parameter.
}

```

Parameters **xTaskToSuspend** -- Handle to the task being suspended. Passing a NULL handle will cause the calling task to be suspended.

void **vTaskResume** (*TaskHandle_t* xTaskToResume)

INCLUDE_vTaskSuspend must be defined as 1 for this function to be available. See the configuration section for more information.

Resumes a suspended task.

A task that has been suspended by one or more calls to `vTaskSuspend ()` will be made available for running again by a single call to `vTaskResume ()`.

Example usage:

```
void vAFunction( void )
{
    TaskHandle_t xHandle;

    // Create a task, storing the handle.
    xTaskCreate( vTaskCode, "NAME", STACK_SIZE, NULL, tskIDLE_PRIORITY, &xHandle );

    // ...

    // Use the handle to suspend the created task.
    vTaskSuspend( xHandle );

    // ...

    // The created task will not run during this period, unless
    // another task calls vTaskResume( xHandle ).

    //...

    // Resume the suspended task ourselves.
    vTaskResume( xHandle );

    // The created task will once again get microcontroller processing
    // time in accordance with its priority within the system.
}
```

Parameters `xTaskToResume` -- Handle to the task being readied.

`BaseType_t xTaskResumeFromISR ( TaskHandle_t xTaskToResume )`

`INCLUDE_xTaskResumeFromISR` must be defined as 1 for this function to be available. See the configuration section for more information.

An implementation of `vTaskResume()` that can be called from within an ISR.

A task that has been suspended by one or more calls to `vTaskSuspend ()` will be made available for running again by a single call to `xTaskResumeFromISR ()`.

`xTaskResumeFromISR()` should not be used to synchronise a task with an interrupt if there is a chance that the interrupt could arrive prior to the task being suspended - as this can lead to interrupts being missed. Use of a semaphore as a synchronisation mechanism would avoid this eventuality.

Parameters `xTaskToResume` -- Handle to the task being readied.

Returns `pdTRUE` if resuming the task should result in a context switch, otherwise `pdFALSE`. This is used by the ISR to determine if a context switch may be required following the ISR.

`void vTaskSuspendAll (void)`

Suspends the scheduler without disabling interrupts. Context switches will not occur while the scheduler is suspended.

After calling `vTaskSuspendAll ()` the calling task will continue to execute without risk of being swapped out until a call to `xTaskResumeAll ()` has been made.

API functions that have the potential to cause a context switch (for example, `xTaskDelayUntil()`, `xQueueSend()`, etc.) must not be called while the scheduler is suspended.

Example usage:

```
void vTask1( void * pvParameters )
{
    for( ;; )
    {
        // Task code goes here.

        // ...

        // At some point the task wants to perform a long operation during
        // which it does not want to get swapped out. It cannot use
        // taskENTER_CRITICAL ()/taskEXIT_CRITICAL () as the length of the
        // operation may cause interrupts to be missed - including the
        // ticks.

        // Prevent the real time kernel swapping out the task.
        vTaskSuspendAll ();

        // Perform the operation here. There is no need to use critical
        // sections as we have all the microcontroller processing time.
        // During this time interrupts will still operate and the kernel
        // tick count will be maintained.

        // ...

        // The operation is complete. Restart the kernel.
        xTaskResumeAll ();
    }
}
```

BaseType_t **xTaskResumeAll** (void)

Resumes scheduler activity after it was suspended by a call to vTaskSuspendAll().

xTaskResumeAll() only resumes the scheduler. It does not unsuspend tasks that were previously suspended by a call to vTaskSuspend().

Example usage:

```
void vTask1( void * pvParameters )
{
    for( ;; )
    {
        // Task code goes here.

        // ...

        // At some point the task wants to perform a long operation during
        // which it does not want to get swapped out. It cannot use
        // taskENTER_CRITICAL ()/taskEXIT_CRITICAL () as the length of the
        // operation may cause interrupts to be missed - including the
        // ticks.

        // Prevent the real time kernel swapping out the task.
        vTaskSuspendAll ();

        // Perform the operation here. There is no need to use critical
        // sections as we have all the microcontroller processing time.
        // During this time interrupts will still operate and the real
        // time kernel tick count will be maintained.
```

(continues on next page)

(continued from previous page)

```

// ...

// The operation is complete. Restart the kernel. We want to force
// a context switch - but there is no point if resuming the scheduler
// caused a context switch already.
if ( !xTaskResumeAll () )
{
    taskYIELD ();
}
}
}

```

Returns If resuming the scheduler caused a context switch then pdTRUE is returned, otherwise pdFALSE is returned.

TickType_t **xTaskGetTickCount** (void)

Returns The count of ticks since vTaskStartScheduler was called.

TickType_t **xTaskGetTickCountFromISR** (void)

This is a version of xTaskGetTickCount() that is safe to be called from an ISR - provided that TickType_t is the natural word size of the microcontroller being used or interrupt nesting is either not supported or not being used.

Returns The count of ticks since vTaskStartScheduler was called.

UBaseType_t **uxTaskGetNumberOfTasks** (void)

Returns The number of tasks that the real time kernel is currently managing. This includes all ready, blocked and suspended tasks. A task that has been deleted but not yet freed by the idle task will also be included in the count.

char ***pcTaskGetName** ([TaskHandle_t](#) xTaskToQuery)

Returns The text (human readable) name of the task referenced by the handle xTaskToQuery. A task can query its own name by either passing in its own handle, or by setting xTaskToQuery to NULL.

[TaskHandle_t](#) **xTaskGetHandle** (const char *pcNameToQuery)

NOTE: This function takes a relatively long time to complete and should be used sparingly.

Returns The handle of the task that has the human readable name pcNameToQuery. NULL is returned if no matching name is found. INCLUDE_xTaskGetHandle must be set to 1 in FreeRTOSConfig.h for pcTaskGetHandle() to be available.

BaseType_t **xTaskGetStaticBuffers** ([TaskHandle_t](#) xTask, StackType_t **ppuxStackBuffer, StaticTask_t **ppxTaskBuffer)

Retrieve pointers to a statically created task's data structure buffer and stack buffer. These are the same buffers that are supplied at the time of creation.

Parameters

- **xTask** -- The task for which to retrieve the buffers.
- **ppuxStackBuffer** -- Used to return a pointer to the task's stack buffer.
- **ppxTaskBuffer** -- Used to return a pointer to the task's data structure buffer.

Returns pdTRUE if buffers were retrieved, pdFALSE otherwise.

UBaseType_t **uxTaskGetStackHighWaterMark** ([TaskHandle_t](#) xTask)

INCLUDE_uxTaskGetStackHighWaterMark must be set to 1 in FreeRTOSConfig.h for this function to be available.

Returns the high water mark of the stack associated with xTask. That is, the minimum free stack space there has been (in words, so on a 32 bit machine a value of 1 means 4 bytes) since the task started. The smaller the returned number the closer the task has come to overflowing its stack.

uxTaskGetStackHighWaterMark() and uxTaskGetStackHighWaterMark2() are the same except for their return type. Using configSTACK_DEPTH_TYPE allows the user to determine the return type. It gets around the problem of the value overflowing on 8-bit types without breaking backward compatibility for applications that expect an 8-bit return type.

Parameters **xTask** -- Handle of the task associated with the stack to be checked. Set xTask to NULL to check the stack of the calling task.

Returns The smallest amount of free stack space there has been (in words, so actual spaces on the stack rather than bytes) since the task referenced by xTask was created.

configSTACK_DEPTH_TYPE **uxTaskGetStackHighWaterMark2** (*TaskHandle_t* xTask)

INCLUDE_uxTaskGetStackHighWaterMark2 must be set to 1 in FreeRTOSConfig.h for this function to be available.

Returns the high water mark of the stack associated with xTask. That is, the minimum free stack space there has been (in words, so on a 32 bit machine a value of 1 means 4 bytes) since the task started. The smaller the returned number the closer the task has come to overflowing its stack.

uxTaskGetStackHighWaterMark() and uxTaskGetStackHighWaterMark2() are the same except for their return type. Using configSTACK_DEPTH_TYPE allows the user to determine the return type. It gets around the problem of the value overflowing on 8-bit types without breaking backward compatibility for applications that expect an 8-bit return type.

Parameters **xTask** -- Handle of the task associated with the stack to be checked. Set xTask to NULL to check the stack of the calling task.

Returns The smallest amount of free stack space there has been (in words, so actual spaces on the stack rather than bytes) since the task referenced by xTask was created.

void **vTaskSetApplicationTaskTag** (*TaskHandle_t* xTask, *TaskHookFunction_t* pxHookFunction)

Sets pxHookFunction to be the task hook function used by the task xTask. Passing xTask as NULL has the effect of setting the calling tasks hook function.

TaskHookFunction_t **xTaskGetApplicationTaskTag** (*TaskHandle_t* xTask)

Returns the pxHookFunction value assigned to the task xTask. Do not call from an interrupt service routine - call xTaskGetApplicationTaskTagFromISR() instead.

TaskHookFunction_t **xTaskGetApplicationTaskTagFromISR** (*TaskHandle_t* xTask)

Returns the pxHookFunction value assigned to the task xTask. Can be called from an interrupt service routine.

void **vTaskSetThreadLocalStoragePointer** (*TaskHandle_t* xTaskToSet, BaseType_t xIndex, void *pvValue)

Each task contains an array of pointers that is dimensioned by the configNUM_THREAD_LOCAL_STORAGE_POINTERS setting in FreeRTOSConfig.h. The kernel does not use the pointers itself, so the application writer can use the pointers for any purpose they wish. The following two functions are used to set and query a pointer respectively.

void ***pvTaskGetThreadLocalStoragePointer** (*TaskHandle_t* xTaskToQuery, BaseType_t xIndex)

void **vApplicationGetIdleTaskMemory** (StaticTask_t **ppxIdleTaskTCBBuffer, StackType_t **ppxIdleTaskStackBuffer, uint32_t *pulIdleTaskStackSize)

This function is used to provide a statically allocated block of memory to FreeRTOS to hold the Idle Task TCB. This function is required when configSUPPORT_STATIC_ALLOCATION is set. For more information see this URI: https://www.FreeRTOS.org/a00110.html#configSUPPORT_STATIC_ALLOCATION

Parameters

- **ppxIdleTaskTCBBuffer** -- A handle to a statically allocated TCB buffer
- **ppxIdleTaskStackBuffer** -- A handle to a statically allocated Stack buffer for the idle task

- **pulIdleTaskStackSize** -- A pointer to the number of elements that will fit in the allocated stack buffer

BaseType_t **xTaskCallApplicationTaskHook** (*TaskHandle_t* xTask, void *pvParameter)

Calls the hook function associated with xTask. Passing xTask as NULL has the effect of calling the Running tasks (the calling task) hook function.

pvParameter is passed to the hook function for the task to interpret as it wants. The return value is the value returned by the task hook function registered by the user.

TaskHandle_t **xTaskGetIdleTaskHandle** (void)

xTaskGetIdleTaskHandle() is only available if INCLUDE_xTaskGetIdleTaskHandle is set to 1 in FreeRTOSConfig.h.

Simply returns the handle of the idle task of the current core. It is not valid to call xTaskGetIdleTaskHandle() before the scheduler has been started.

UBaseType_t **uxTaskGetSystemState** (*TaskStatus_t* *const pxTaskStatusArray, const UBaseType_t uxArraySize, configRUN_TIME_COUNTER_TYPE *const pulTotalRunTime)

configUSE_TRACE_FACILITY must be defined as 1 in FreeRTOSConfig.h for uxTaskGetSystemState() to be available.

uxTaskGetSystemState() populates an TaskStatus_t structure for each task in the system. TaskStatus_t structures contain, among other things, members for the task handle, task name, task priority, task state, and total amount of run time consumed by the task. See the TaskStatus_t structure definition in this file for the full member list.

NOTE: This function is intended for debugging use only as its use results in the scheduler remaining suspended for an extended period.

Example usage:

```
// This example demonstrates how a human readable table of run time stats
// information is generated from raw data provided by uxTaskGetSystemState().
// The human readable table is written to pcWriteBuffer
void vTaskGetRunTimeStats( char *pcWriteBuffer )
{
    TaskStatus_t *pxTaskStatusArray;
    volatile UBaseType_t uxArraySize, x;
    configRUN_TIME_COUNTER_TYPE ulTotalRunTime, ulStatsAsPercentage;

    // Make sure the write buffer does not contain a string.
    pcWriteBuffer = 0x00;

    // Take a snapshot of the number of tasks in case it changes while this
    // function is executing.
    uxArraySize = uxTaskGetNumberOfTasks();

    // Allocate a TaskStatus_t structure for each task. An array could be
    // allocated statically at compile time.
    pxTaskStatusArray = pvPortMalloc( uxArraySize * sizeof( TaskStatus_t ) );

    if( pxTaskStatusArray != NULL )
    {
        // Generate raw status information about each task.
        uxArraySize = uxTaskGetSystemState( pxTaskStatusArray, uxArraySize, &
        ↪ulTotalRunTime );

        // For percentage calculations.
        ulTotalRunTime /= 100UL;
    }
}
```

(continues on next page)

(continued from previous page)

```

// Avoid divide by zero errors.
if( ulTotalRunTime > 0 )
{
    // For each populated position in the pxTaskStatusArray array,
    // format the raw data as human readable ASCII data
    for( x = 0; x < uxArraySize; x++ )
    {
        // What percentage of the total run time has the task used?
        // This will always be rounded down to the nearest integer.
        // ulTotalRunTimeDiv100 has already been divided by 100.
        ulStatsAsPercentage = pxTaskStatusArray[ x ].ulRunTimeCounter_
→/ ulTotalRunTime;

        if( ulStatsAsPercentage > 0UL )
        {
            sprintf( pcWriteBuffer, "%s\t\t%lu\t\t%lu%%\r\n",_
→pxTaskStatusArray[ x ].pcTaskName, pxTaskStatusArray[ x ].ulRunTimeCounter,_
→ulStatsAsPercentage );
        }
        else
        {
            // If the percentage is zero here then the task has
            // consumed less than 1% of the total run time.
            sprintf( pcWriteBuffer, "%s\t\t%lu\t\t<1%%\r\n",_
→pxTaskStatusArray[ x ].pcTaskName, pxTaskStatusArray[ x ].ulRunTimeCounter );
        }

        pcWriteBuffer += strlen( ( char * ) pcWriteBuffer );
    }

    // The array is no longer needed, free the memory it consumes.
    vPortFree( pxTaskStatusArray );
}
}

```

Parameters

- **pxTaskStatusArray** -- A pointer to an array of TaskStatus_t structures. The array must contain at least one TaskStatus_t structure for each task that is under the control of the RTOS. The number of tasks under the control of the RTOS can be determined using the uxTaskGetNumberOfTasks() API function.
- **uxArraySize** -- The size of the array pointed to by the pxTaskStatusArray parameter. The size is specified as the number of indexes in the array, or the number of TaskStatus_t structures contained in the array, not by the number of bytes in the array.
- **pulTotalRunTime** -- If configGENERATE_RUN_TIME_STATS is set to 1 in FreeRTOSConfig.h then *pulTotalRunTime is set by uxTaskGetSystemState() to the total run time (as defined by the run time stats clock, see <https://www.FreeRTOS.org/rtos-run-time-stats.html>) since the target booted. pulTotalRunTime can be set to NULL to omit the total run time information.

Returns The number of TaskStatus_t structures that were populated by uxTaskGetSystemState(). This should equal the number returned by the uxTaskGetNumberOfTasks() API function, but will be zero if the value passed in the uxArraySize parameter was too small.

void **vTaskList** (char *pcWriteBuffer)

configUSE_TRACE_FACILITY and configUSE_STATS_FORMATTING_FUNCTIONS must both be defined as 1 for this function to be available. See the configuration section of the FreeRTOS.org website for more information.

NOTE 1: This function will disable interrupts for its duration. It is not intended for normal application runtime use but as a debug aid.

Lists all the current tasks, along with their current state and stack usage high water mark.

Tasks are reported as blocked ('B'), ready ('R'), deleted ('D') or suspended ('S').

PLEASE NOTE:

This function is provided for convenience only, and is used by many of the demo applications. Do not consider it to be part of the scheduler.

`vTaskList()` calls `uxTaskGetSystemState()`, then formats part of the `uxTaskGetSystemState()` output into a human readable table that displays task: names, states, priority, stack usage and task number. Stack usage specified as the number of unused `StackType_t` words stack can hold on top of stack - not the number of bytes.

`vTaskList()` has a dependency on the `sprintf()` C library function that might bloat the code size, use a lot of stack, and provide different results on different platforms. An alternative, tiny, third party, and limited functionality implementation of `sprintf()` is provided in many of the FreeRTOS/Demo sub-directories in a file called `printf-stdarg.c` (note `printf-stdarg.c` does not provide a full `snprintf()` implementation!).

It is recommended that production systems call `uxTaskGetSystemState()` directly to get access to raw stats data, rather than indirectly through a call to `vTaskList()`.

Parameters `pcWriteBuffer` -- A buffer into which the above mentioned details will be written, in ASCII form. This buffer is assumed to be large enough to contain the generated report. Approximately 40 bytes per task should be sufficient.

void **`vTaskGetRunTimeStats`** (char *pcWriteBuffer)

`configGENERATE_RUN_TIME_STATS` and `configUSE_STATS_FORMATTING_FUNCTIONS` must both be defined as 1 for this function to be available. The application must also then provide definitions for `portCONFIGURE_TIMER_FOR_RUN_TIME_STATS()` and `portGET_RUN_TIME_COUNTER_VALUE()` to configure a peripheral timer/counter and return the timers current count value respectively. The counter should be at least 10 times the frequency of the tick count.

NOTE 1: This function will disable interrupts for its duration. It is not intended for normal application runtime use but as a debug aid.

Setting `configGENERATE_RUN_TIME_STATS` to 1 will result in a total accumulated execution time being stored for each task. The resolution of the accumulated time value depends on the frequency of the timer configured by the `portCONFIGURE_TIMER_FOR_RUN_TIME_STATS()` macro. Calling `vTaskGetRunTimeStats()` writes the total execution time of each task into a buffer, both as an absolute count value and as a percentage of the total system execution time.

NOTE 2:

This function is provided for convenience only, and is used by many of the demo applications. Do not consider it to be part of the scheduler.

`vTaskGetRunTimeStats()` calls `uxTaskGetSystemState()`, then formats part of the `uxTaskGetSystemState()` output into a human readable table that displays the amount of time each task has spent in the Running state in both absolute and percentage terms.

`vTaskGetRunTimeStats()` has a dependency on the `sprintf()` C library function that might bloat the code size, use a lot of stack, and provide different results on different platforms. An alternative, tiny, third party, and limited functionality implementation of `sprintf()` is provided in many of the FreeRTOS/Demo sub-directories in a file called `printf-stdarg.c` (note `printf-stdarg.c` does not provide a full `snprintf()` implementation!).

It is recommended that production systems call `uxTaskGetSystemState()` directly to get access to raw stats data, rather than indirectly through a call to `vTaskGetRunTimeStats()`.

Parameters `pcWriteBuffer` -- A buffer into which the execution times will be written, in ASCII form. This buffer is assumed to be large enough to contain the generated report. Approximately 40 bytes per task should be sufficient.

`configRUN_TIME_COUNTER_TYPE` **`ulTaskGetIdleRunTimeCounter`** (void)

`configGENERATE_RUN_TIME_STATS`, `configUSE_STATS_FORMATTING_FUNCTIONS` and `INCLUDE_xTaskGetIdleTaskHandle` must all be defined as 1 for these functions to be available. The application must also then provide definitions for `portCONFIGURE_TIMER_FOR_RUN_TIME_STATS()`

and `portGET_RUN_TIME_COUNTER_VALUE()` to configure a peripheral timer/counter and return the timers current count value respectively. The counter should be at least 10 times the frequency of the tick count.

Setting `configGENERATE_RUN_TIME_STATS` to 1 will result in a total accumulated execution time being stored for each task. The resolution of the accumulated time value depends on the frequency of the timer configured by the `portCONFIGURE_TIMER_FOR_RUN_TIME_STATS()` macro. While `uxTaskGetSystemState()` and `vTaskGetRunTimeStats()` writes the total execution time of each task into a buffer, `ulTaskGetIdleRunTimeCounter()` returns the total execution time of just the idle task and `ulTaskGetIdleRunTimePercent()` returns the percentage of the CPU time used by just the idle task.

Note the amount of idle time is only a good measure of the slack time in a system if there are no other tasks executing at the idle priority, tickless idle is not used, and `configIDLE_SHOULD_YIELD` is set to 0.

Note: If `configNUMBER_OF_CORES > 1`, calling this function will query the idle task of the current core.

Returns The total run time of the idle task or the percentage of the total run time consumed by the idle task. This is the amount of time the idle task has actually been executing. The unit of time is dependent on the frequency configured using the `portCONFIGURE_TIMER_FOR_RUN_TIME_STATS()` and `portGET_RUN_TIME_COUNTER_VALUE()` macros.

`configRUN_TIME_COUNTER_TYPE` **ulTaskGetIdleRunTimePercent** (void)

`BaseType_t` **xTaskGenericNotifyWait** (`UBaseType_t` `uxIndexToWaitOn`, `uint32_t` `ulBitsToClearOnEntry`, `uint32_t` `ulBitsToClearOnExit`, `uint32_t *``pulNotificationValue`, `TickType_t` `xTicksToWait`)

Waits for a direct to task notification to be pending at a given index within an array of direct to task notifications.

See <https://www.FreeRTOS.org/RTOS-task-notifications.html> for details.

`configUSE_TASK_NOTIFICATIONS` must be undefined or defined as 1 for this function to be available.

Each task has a private array of "notification values" (or 'notifications'), each of which is a 32-bit unsigned integer (`uint32_t`). The constant `configTASK_NOTIFICATION_ARRAY_ENTRIES` sets the number of indexes in the array, and (for backward compatibility) defaults to 1 if left undefined. Prior to FreeRTOS V10.4.0 there was only one notification value per task.

Events can be sent to a task using an intermediary object. Examples of such objects are queues, semaphores, mutexes and event groups. Task notifications are a method of sending an event directly to a task without the need for such an intermediary object.

A notification sent to a task can optionally perform an action, such as update, overwrite or increment one of the task's notification values. In that way task notifications can be used to send data to a task, or be used as light weight and fast binary or counting semaphores.

A notification sent to a task will remain pending until it is cleared by the task calling `xTaskNotifyWaitIndexed()` or `ulTaskNotifyTakeIndexed()` (or their un-indexed equivalents). If the task was already in the Blocked state to wait for a notification when the notification arrives then the task will automatically be removed from the Blocked state (unblocked) and the notification cleared.

A task can use `xTaskNotifyWaitIndexed()` to [optionally] block to wait for a notification to be pending, or `ulTaskNotifyTakeIndexed()` to [optionally] block to wait for a notification value to have a non-zero value. The task does not consume any CPU time while it is in the Blocked state.

NOTE Each notification within the array operates independently - a task can only block on one notification within the array at a time and will not be unblocked by a notification sent to any other array index.

Backward compatibility information: Prior to FreeRTOS V10.4.0 each task had a single "notification value", and all task notification API functions operated on that value. Replacing the single notification value with an array of notification values necessitated a new set of API functions that could address specific notifications within the array. `xTaskNotifyWait()` is the original API function, and remains backward compatible by always

operating on the notification value at index 0 in the array. Calling `xTaskNotifyWait()` is equivalent to calling `xTaskNotifyWaitIndexed()` with the `uxIndexToWaitOn` parameter set to 0.

Parameters

- **`uxIndexToWaitOn`** -- The index within the calling task's array of notification values on which the calling task will wait for a notification to be received. `uxIndexToWaitOn` must be less than `configTASK_NOTIFICATION_ARRAY_ENTRIES`. `xTaskNotifyWait()` does not have this parameter and always waits for notifications on index 0.
- **`ulBitsToClearOnEntry`** -- Bits that are set in `ulBitsToClearOnEntry` value will be cleared in the calling task's notification value before the task checks to see if any notifications are pending, and optionally blocks if no notifications are pending. Setting `ulBitsToClearOnEntry` to `ULONG_MAX` (if `limits.h` is included) or `0xffffffffUL` (if `limits.h` is not included) will have the effect of resetting the task's notification value to 0. Setting `ulBitsToClearOnEntry` to 0 will leave the task's notification value unchanged.
- **`ulBitsToClearOnExit`** -- If a notification is pending or received before the calling task exits the `xTaskNotifyWait()` function then the task's notification value (see the `xTaskNotify()` API function) is passed out using the `pulNotificationValue` parameter. Then any bits that are set in `ulBitsToClearOnExit` will be cleared in the task's notification value (note `*pulNotificationValue` is set before any bits are cleared). Setting `ulBitsToClearOnExit` to `ULONG_MAX` (if `limits.h` is included) or `0xffffffffUL` (if `limits.h` is not included) will have the effect of resetting the task's notification value to 0 before the function exits. Setting `ulBitsToClearOnExit` to 0 will leave the task's notification value unchanged when the function exits (in which case the value passed out in `pulNotificationValue` will match the task's notification value).
- **`pulNotificationValue`** -- Used to pass the task's notification value out of the function. Note the value passed out will not be effected by the clearing of any bits caused by `ulBitsToClearOnExit` being non-zero.
- **`xTicksToWait`** -- The maximum amount of time that the task should wait in the Blocked state for a notification to be received, should a notification not already be pending when `xTaskNotifyWait()` was called. The task will not consume any processing time while it is in the Blocked state. This is specified in kernel ticks, the macro `pdMS_TO_TICKS( value_in_ms )` can be used to convert a time specified in milliseconds to a time specified in ticks.

Returns If a notification was received (including notifications that were already pending when `xTaskNotifyWait` was called) then `pdPASS` is returned. Otherwise `pdFAIL` is returned.

void **`vTaskGenericNotifyGiveFromISR`** (*TaskHandle_t* xTaskToNotify, *UBaseType_t* uxIndexToNotify, *BaseType_t* *pxHigherPriorityTaskWoken)

A version of `xTaskNotifyGiveIndexed()` that can be called from an interrupt service routine (ISR).

See <https://www.FreeRTOS.org/RTOS-task-notifications.html> for more details.

`configUSE_TASK_NOTIFICATIONS` must be undefined or defined as 1 for this macro to be available.

Each task has a private array of "notification values" (or 'notifications'), each of which is a 32-bit unsigned integer (`uint32_t`). The constant `configTASK_NOTIFICATION_ARRAY_ENTRIES` sets the number of indexes in the array, and (for backward compatibility) defaults to 1 if left undefined. Prior to FreeRTOS V10.4.0 there was only one notification value per task.

Events can be sent to a task using an intermediary object. Examples of such objects are queues, semaphores, mutexes and event groups. Task notifications are a method of sending an event directly to a task without the need for such an intermediary object.

A notification sent to a task can optionally perform an action, such as update, overwrite or increment one of the task's notification values. In that way task notifications can be used to send data to a task, or be used as light weight and fast binary or counting semaphores.

`vTaskNotifyGiveIndexedFromISR()` is intended for use when task notifications are used as light weight and faster binary or counting semaphore equivalents. Actual FreeRTOS semaphores are given from an ISR using the `xSemaphoreGiveFromISR()` API function, the equivalent action that instead uses a task notification is `vTaskNotifyGiveIndexedFromISR()`.

When task notifications are being used as a binary or counting semaphore equivalent then the task being notified should wait for the notification using the `ulTaskNotifyTakeIndexed()` API function rather than the `xTaskNotifyWaitIndexed()` API function.

NOTE Each notification within the array operates independently - a task can only block on one notification within the array at a time and will not be unblocked by a notification sent to any other array index.

Backward compatibility information: Prior to FreeRTOS V10.4.0 each task had a single "notification value", and all task notification API functions operated on that value. Replacing the single notification value with an array of notification values necessitated a new set of API functions that could address specific notifications within the array. `xTaskNotifyFromISR()` is the original API function, and remains backward compatible by always operating on the notification value at index 0 within the array. Calling `xTaskNotifyGiveFromISR()` is equivalent to calling `xTaskNotifyGiveIndexedFromISR()` with the `uxIndexToNotify` parameter set to 0.

Parameters

- **`xTaskToNotify`** -- The handle of the task being notified. The handle to a task can be returned from the `xTaskCreate()` API function used to create the task, and the handle of the currently running task can be obtained by calling `xTaskGetCurrentTaskHandle()`.
- **`uxIndexToNotify`** -- The index within the target task's array of notification values to which the notification is to be sent. `uxIndexToNotify` must be less than `configTASK_NOTIFICATION_ARRAY_ENTRIES`. `xTaskNotifyGiveFromISR()` does not have this parameter and always sends notifications to index 0.
- **`pxHigherPriorityTaskWoken`** -- `vTaskNotifyGiveFromISR()` will set `*pxHigherPriorityTaskWoken` to `pdTRUE` if sending the notification caused the task to which the notification was sent to leave the Blocked state, and the unblocked task has a priority higher than the currently running task. If `vTaskNotifyGiveFromISR()` sets this value to `pdTRUE` then a context switch should be requested before the interrupt is exited. How a context switch is requested from an ISR is dependent on the port - see the documentation page for the port in use.

BaseType_t **`xTaskGenericNotifyStateClear`** (*TaskHandle_t* xTask, UBaseType_t uxIndexToClear)

See <https://www.FreeRTOS.org/RTOS-task-notifications.html> for details.

`configUSE_TASK_NOTIFICATIONS` must be undefined or defined as 1 for these functions to be available.

Each task has a private array of "notification values" (or 'notifications'), each of which is a 32-bit unsigned integer (`uint32_t`). The constant `configTASK_NOTIFICATION_ARRAY_ENTRIES` sets the number of indexes in the array, and (for backward compatibility) defaults to 1 if left undefined. Prior to FreeRTOS V10.4.0 there was only one notification value per task.

If a notification is sent to an index within the array of notifications then the notification at that index is said to be 'pending' until it is read or explicitly cleared by the receiving task. `xTaskNotifyStateClearIndexed()` is the function that clears a pending notification without reading the notification value. The notification value at the same array index is not altered. Set `xTask` to `NULL` to clear the notification state of the calling task.

Backward compatibility information: Prior to FreeRTOS V10.4.0 each task had a single "notification value", and all task notification API functions operated on that value. Replacing the single notification value with an array of notification values necessitated a new set of API functions that could address specific notifications within the array. `xTaskNotifyStateClear()` is the original API function, and remains backward compatible by always operating on the notification value at index 0 within the array. Calling `xTaskNotifyStateClear()` is equivalent to calling `xTaskNotifyStateClearIndexed()` with the `uxIndexToNotify` parameter set to 0.

Parameters

- **`xTask`** -- The handle of the RTOS task that will have a notification state cleared. Set `xTask` to `NULL` to clear a notification state in the calling task. To obtain a task's handle create the task using `xTaskCreate()` and make use of the `pxCreatedTask` parameter, or create the task using `xTaskCreateStatic()` and store the returned value, or use the task's name in a call to `xTaskGetHandle()`.
- **`uxIndexToClear`** -- The index within the target task's array of notification values to act upon. For example, setting `uxIndexToClear` to 1 will clear the state of the notification at index 1 within the array. `uxIndexToClear` must be less than `configTASK_NOTIFICATION_ARRAY_ENTRIES`. `ulTaskNotifyStateClear()` does not have

this parameter and always acts on the notification at index 0.

Returns pdTRUE if the task's notification state was set to eNotWaitingNotification, otherwise pdFALSE.

uint32_t **ulTaskGenericNotifyValueClear** (*TaskHandle_t* xTask, UBaseType_t uxIndexToClear, uint32_t ulBitsToClear)

See <https://www.FreeRTOS.org/RTOS-task-notifications.html> for details.

configUSE_TASK_NOTIFICATIONS must be undefined or defined as 1 for these functions to be available.

Each task has a private array of "notification values" (or 'notifications'), each of which is a 32-bit unsigned integer (uint32_t). The constant configTASK_NOTIFICATION_ARRAY_ENTRIES sets the number of indexes in the array, and (for backward compatibility) defaults to 1 if left undefined. Prior to FreeRTOS V10.4.0 there was only one notification value per task.

ulTaskNotifyValueClearIndexed() clears the bits specified by the ulBitsToClear bit mask in the notification value at array index uxIndexToClear of the task referenced by xTask.

Backward compatibility information: Prior to FreeRTOS V10.4.0 each task had a single "notification value", and all task notification API functions operated on that value. Replacing the single notification value with an array of notification values necessitated a new set of API functions that could address specific notifications within the array. ulTaskNotifyValueClear() is the original API function, and remains backward compatible by always operating on the notification value at index 0 within the array. Calling ulTaskNotifyValueClear() is equivalent to calling ulTaskNotifyValueClearIndexed() with the uxIndexToClear parameter set to 0.

Parameters

- **xTask** -- The handle of the RTOS task that will have bits in one of its notification values cleared. Set xTask to NULL to clear bits in a notification value of the calling task. To obtain a task's handle create the task using xTaskCreate() and make use of the pxCreated-Task parameter, or create the task using xTaskCreateStatic() and store the returned value, or use the task's name in a call to xTaskGetHandle().
- **uxIndexToClear** -- The index within the target task's array of notification values in which to clear the bits. uxIndexToClear must be less than configTASK_NOTIFICATION_ARRAY_ENTRIES. ulTaskNotifyValueClear() does not have this parameter and always clears bits in the notification value at index 0.
- **ulBitsToClear** -- Bit mask of the bits to clear in the notification value of xTask. Set a bit to 1 to clear the corresponding bits in the task's notification value. Set ulBitsToClear to 0xffffffff (UINT_MAX on 32-bit architectures) to clear the notification value to 0. Set ulBitsToClear to 0 to query the task's notification value without clearing any bits.

Returns The value of the target task's notification value before the bits specified by ulBitsToClear were cleared.

void **vTaskSetTimeoutState** (Timeout_t *const pxTimeout)

Capture the current time for future use with xTaskCheckForTimeout().

Parameters **pxTimeout** -- Pointer to a timeout object into which the current time is to be captured. The captured time includes the tick count and the number of times the tick count has overflowed since the system first booted.

BaseType_t **xTaskCheckForTimeout** (Timeout_t *const pxTimeout, TickType_t *const pxTicksToWait)

Determines if pxTicksToWait ticks has passed since a time was captured using a call to vTaskSetTimeoutState(). The captured time includes the tick count and the number of times the tick count has overflowed.

Example Usage:

```
// Driver library function used to receive uxWantedBytes from an Rx buffer
// that is filled by a UART interrupt. If there are not enough bytes in the
// Rx buffer then the task enters the Blocked state until it is notified that
// more data has been placed into the buffer. If there is still not enough
// data then the task re-enters the Blocked state, and xTaskCheckForTimeout()
// is used to re-calculate the Block time to ensure the total amount of time
```

(continues on next page)

(continued from previous page)

```

// spent in the Blocked state does not exceed MAX_TIME_TO_WAIT. This
// continues until either the buffer contains at least uxWantedBytes bytes,
// or the total amount of time spent in the Blocked state reaches
// MAX_TIME_TO_WAIT - at which point the task reads however many bytes are
// available up to a maximum of uxWantedBytes.

size_t xUART_Receive( uint8_t *pucBuffer, size_t uxWantedBytes )
{
    size_t uxReceived = 0;
    TickType_t xTicksToWait = MAX_TIME_TO_WAIT;
    TimeOut_t xTimeOut;

    // Initialize xTimeOut. This records the time at which this function
    // was entered.
    vTaskSetTimeOutState( &xTimeOut );

    // Loop until the buffer contains the wanted number of bytes, or a
    // timeout occurs.
    while( UART_bytes_in_rx_buffer( pxUARTInstance ) < uxWantedBytes )
    {
        // The buffer didn't contain enough data so this task is going to
        // enter the Blocked state. Adjusting xTicksToWait to account for
        // any time that has been spent in the Blocked state within this
        // function so far to ensure the total amount of time spent in the
        // Blocked state does not exceed MAX_TIME_TO_WAIT.
        if( xTaskCheckForTimeOut( &xTimeOut, &xTicksToWait ) != pdFALSE )
        {
            //Timed out before the wanted number of bytes were available,
            // exit the loop.
            break;
        }

        // Wait for a maximum of xTicksToWait ticks to be notified that the
        // receive interrupt has placed more data into the buffer.
        ulTaskNotifyTake( pdTRUE, xTicksToWait );
    }

    // Attempt to read uxWantedBytes from the receive buffer into pucBuffer.
    // The actual number of bytes read (which might be less than
    // uxWantedBytes) is returned.
    uxReceived = UART_read_from_receive_buffer( pxUARTInstance,
                                                pucBuffer,
                                                uxWantedBytes );

    return uxReceived;
}

```

See also:
<https://www.FreeRTOS.org/xTaskCheckForTimeOut.html>
Parameters

- **pxTimeOut** -- The time status as captured previously using vTaskSetTimeOutState. If the timeout has not yet occurred, it is updated to reflect the current time status.
- **pxTicksToWait** -- The number of ticks to check for timeout i.e. if pxTicksToWait ticks have passed since pxTimeOut was last updated (either by vTaskSetTimeOutState() or xTaskCheckForTimeOut()), the timeout has occurred. If the timeout has not occurred, pxTicksToWait is updated to reflect the number of remaining ticks.

Returns If timeout has occurred, pdTRUE is returned. Otherwise pdFALSE is returned and pxTicksToWait is updated to reflect the number of remaining ticks.

BaseType_t **xTaskCatchUpTicks** (TickType_t xTicksToCatchUp)

This function corrects the tick count value after the application code has held interrupts disabled for an extended period resulting in tick interrupts having been missed.

This function is similar to vTaskStepTick(), however, unlike vTaskStepTick(), xTaskCatchUpTicks() may move the tick count forward past a time at which a task should be removed from the blocked state. That means tasks may have to be removed from the blocked state as the tick count is moved.

Parameters **xTicksToCatchUp** -- The number of tick interrupts that have been missed due to interrupts being disabled. Its value is not computed automatically, so must be computed by the application writer.

Returns pdTRUE if moving the tick count forward resulted in a task leaving the blocked state and a context switch being performed. Otherwise pdFALSE.

Structures

struct **xTASK_STATUS**

Used with the uxTaskGetSystemState() function to return the state of each task in the system.

Public Members

TaskHandle_t **xHandle**

The handle of the task to which the rest of the information in the structure relates.

const char ***pcTaskName**

A pointer to the task's name. This value will be invalid if the task was deleted since the structure was populated!

UBaseType_t **xTaskNumber**

A number unique to the task.

eTaskState **eCurrentState**

The state in which the task existed when the structure was populated.

UBaseType_t **uxCurrentPriority**

The priority at which the task was running (may be inherited) when the structure was populated.

UBaseType_t **uxBasePriority**

The priority to which the task will return if the task's current priority has been inherited to avoid unbounded priority inversion when obtaining a mutex. Only valid if configUSE_MUTEXES is defined as 1 in FreeRTOSConfig.h.

configRUN_TIME_COUNTER_TYPE **ulRunTimeCounter**

The total run time allocated to the task so far, as defined by the run time stats clock. See <https://www.FreeRTOS.org/rtos-run-time-stats.html>. Only valid when configGENERATE_RUN_TIME_STATS is defined as 1 in FreeRTOSConfig.h.

StackType_t ***pxStackBase**

Points to the lowest address of the task's stack area.

configSTACK_DEPTH_TYPE **usStackHighWaterMark**

The minimum amount of stack space that has remained for the task since the task was created. The closer this value is to zero the closer the task has come to overflowing its stack.

BaseType_t **xCoreID**

Core this task is pinned to (0, 1, or tskNO_AFFINITY). If configNUMBER_OF_CORES == 1, this will always be 0.

Macros

tskIDLE_PRIORITY

Defines the priority used by the idle task. This must not be modified.

tskNO_AFFINITY

Macro representing and unpinned (i.e., "no affinity") task in xCoreID parameters

taskVALID_CORE_ID (xCoreID)

Macro to check if an xCoreID value is valid

Returns pdTRUE if valid, pdFALSE otherwise.

taskYIELD ()

Macro for forcing a context switch.

taskENTER_CRITICAL (x)

Macro to mark the start of a critical code region. Preemptive context switches cannot occur when in a critical region.

NOTE: This may alter the stack (depending on the portable implementation) so must be used with care!

taskENTER_CRITICAL_FROM_ISR ()

taskENTER_CRITICAL_ISR (x)

taskEXIT_CRITICAL (x)

Macro to mark the end of a critical code region. Preemptive context switches cannot occur when in a critical region.

NOTE: This may alter the stack (depending on the portable implementation) so must be used with care!

taskEXIT_CRITICAL_FROM_ISR (x)

taskEXIT_CRITICAL_ISR (x)

taskDISABLE_INTERRUPTS ()

Macro to disable all maskable interrupts.

taskENABLE_INTERRUPTS ()

Macro to enable microcontroller interrupts.

taskSCHEDULER_SUSPENDED

Definitions returned by xTaskGetSchedulerState(). taskSCHEDULER_SUSPENDED is 0 to generate more optimal code when configASSERT() is defined as the constant is used in assert() statements.

taskSCHEDULER_NOT_STARTED

taskSCHEDULER_RUNNING

xTaskNotifyIndexed (xTaskToNotify, uxIndexToNotify, ulValue, eAction)

See <https://www.FreeRTOS.org/RTOS-task-notifications.html> for details.

configUSE_TASK_NOTIFICATIONS must be undefined or defined as 1 for these functions to be available.

Sends a direct to task notification to a task, with an optional value and action.

Each task has a private array of "notification values" (or 'notifications'), each of which is a 32-bit unsigned integer (uint32_t). The constant configTASK_NOTIFICATION_ARRAY_ENTRIES sets the number of indexes in the array, and (for backward compatibility) defaults to 1 if left undefined. Prior to FreeRTOS V10.4.0 there was only one notification value per task.

Events can be sent to a task using an intermediary object. Examples of such objects are queues, semaphores, mutexes and event groups. Task notifications are a method of sending an event directly to a task without the need for such an intermediary object.

A notification sent to a task can optionally perform an action, such as update, overwrite or increment one of the task's notification values. In that way task notifications can be used to send data to a task, or be used as light weight and fast binary or counting semaphores.

A task can use xTaskNotifyWaitIndexed() or ulTaskNotifyTakeIndexed() to [optionally] block to wait for a notification to be pending. The task does not consume any CPU time while it is in the Blocked state.

A notification sent to a task will remain pending until it is cleared by the task calling xTaskNotifyWaitIndexed() or ulTaskNotifyTakeIndexed() (or their un-indexed equivalents). If the task was already in the Blocked state to wait for a notification when the notification arrives then the task will automatically be removed from the Blocked state (unblocked) and the notification cleared.

NOTE Each notification within the array operates independently - a task can only block on one notification within the array at a time and will not be unblocked by a notification sent to any other array index.

Backward compatibility information: Prior to FreeRTOS V10.4.0 each task had a single "notification value", and all task notification API functions operated on that value. Replacing the single notification value with an array of notification values necessitated a new set of API functions that could address specific notifications within the array. xTaskNotify() is the original API function, and remains backward compatible by always operating on the notification value at index 0 in the array. Calling xTaskNotify() is equivalent to calling xTaskNotifyIndexed() with the uxIndexToNotify parameter set to 0.

eSetBits - The target notification value is bitwise ORed with ulValue. xTaskNotifyIndexed() always returns pdPASS in this case.

eIncrement - The target notification value is incremented. ulValue is not used and xTaskNotifyIndexed() always returns pdPASS in this case.

eSetValueWithOverwrite - The target notification value is set to the value of ulValue, even if the task being notified had not yet processed the previous notification at the same array index (the task already had a notification pending at that index). xTaskNotifyIndexed() always returns pdPASS in this case.

eSetValueWithoutOverwrite - If the task being notified did not already have a notification pending at the same array index then the target notification value is set to ulValue and xTaskNotifyIndexed() will return pdPASS. If the task being notified already had a notification pending at the same array index then no action is performed and pdFAIL is returned.

eNoAction - The task receives a notification at the specified array index without the notification value at that index being updated. ulValue is not used and xTaskNotifyIndexed() always returns pdPASS in this case.

pulPreviousNotificationValue - Can be used to pass out the subject task's notification value before any bits are modified by the notify function.

Parameters

- **xTaskToNotify** -- The handle of the task being notified. The handle to a task can be returned from the xTaskCreate() API function used to create the task, and the handle of the currently running task can be obtained by calling xTaskGetCurrentTaskHandle().
- **uxIndexToNotify** -- The index within the target task's array of notification values to which the notification is to be sent. uxIndexToNotify must be less than configTASK_NOTIFICATION_ARRAY_ENTRIES. xTaskNotify() does not have this parameter and always sends notifications to index 0.
- **ulValue** -- Data that can be sent with the notification. How the data is used depends on the value of the eAction parameter.

- **eAction** -- Specifies how the notification updates the task's notification value, if at all.
Valid values for eAction are as follows:

Returns Dependent on the value of eAction. See the description of the eAction parameter.

xTaskNotifyAndQueryIndexed (xTaskToNotify, uxIndexToNotify, ulValue, eAction, pulPreviousNotifyValue)

See <https://www.FreeRTOS.org/RTOS-task-notifications.html> for details.

xTaskNotifyAndQueryIndexed() performs the same operation as xTaskNotifyIndexed() with the addition that it also returns the subject task's prior notification value (the notification value at the time the function is called rather than when the function returns) in the additional pulPreviousNotifyValue parameter.

xTaskNotifyAndQuery() performs the same operation as xTaskNotify() with the addition that it also returns the subject task's prior notification value (the notification value as it was at the time the function is called, rather than when the function returns) in the additional pulPreviousNotifyValue parameter.

xTaskNotifyIndexedFromISR (xTaskToNotify, uxIndexToNotify, ulValue, eAction, pxHigherPriorityTaskWoken)

See <https://www.FreeRTOS.org/RTOS-task-notifications.html> for details.

configUSE_TASK_NOTIFICATIONS must be undefined or defined as 1 for these functions to be available.

A version of xTaskNotifyIndexed() that can be used from an interrupt service routine (ISR).

Each task has a private array of "notification values" (or 'notifications'), each of which is a 32-bit unsigned integer (uint32_t). The constant configTASK_NOTIFICATION_ARRAY_ENTRIES sets the number of indexes in the array, and (for backward compatibility) defaults to 1 if left undefined. Prior to FreeRTOS V10.4.0 there was only one notification value per task.

Events can be sent to a task using an intermediary object. Examples of such objects are queues, semaphores, mutexes and event groups. Task notifications are a method of sending an event directly to a task without the need for such an intermediary object.

A notification sent to a task can optionally perform an action, such as update, overwrite or increment one of the task's notification values. In that way task notifications can be used to send data to a task, or be used as light weight and fast binary or counting semaphores.

A task can use xTaskNotifyWaitIndexed() to [optionally] block to wait for a notification to be pending, or ulTaskNotifyTakeIndexed() to [optionally] block to wait for a notification value to have a non-zero value. The task does not consume any CPU time while it is in the Blocked state.

A notification sent to a task will remain pending until it is cleared by the task calling xTaskNotifyWaitIndexed() or ulTaskNotifyTakeIndexed() (or their un-indexed equivalents). If the task was already in the Blocked state to wait for a notification when the notification arrives then the task will automatically be removed from the Blocked state (unblocked) and the notification cleared.

NOTE Each notification within the array operates independently - a task can only block on one notification within the array at a time and will not be unblocked by a notification sent to any other array index.

Backward compatibility information: Prior to FreeRTOS V10.4.0 each task had a single "notification value", and all task notification API functions operated on that value. Replacing the single notification value with an array of notification values necessitated a new set of API functions that could address specific notifications within the array. xTaskNotifyFromISR() is the original API function, and remains backward compatible by always operating on the notification value at index 0 within the array. Calling xTaskNotifyFromISR() is equivalent to calling xTaskNotifyIndexedFromISR() with the uxIndexToNotify parameter set to 0.

eSetBits - The task's notification value is bitwise ORed with ulValue. xTaskNotify() always returns pdPASS in this case.

eIncrement - The task's notification value is incremented. ulValue is not used and xTaskNotify() always returns pdPASS in this case.

eSetValueWithOverwrite - The task's notification value is set to the value of `ulValue`, even if the task being notified had not yet processed the previous notification (the task already had a notification pending). `xTaskNotify()` always returns `pdPASS` in this case.

eSetValueWithoutOverwrite - If the task being notified did not already have a notification pending then the task's notification value is set to `ulValue` and `xTaskNotify()` will return `pdPASS`. If the task being notified already had a notification pending then no action is performed and `pdFAIL` is returned.

eNoAction - The task receives a notification without its notification value being updated. `ulValue` is not used and `xTaskNotify()` always returns `pdPASS` in this case.

Parameters

- **uxIndexToNotify** -- The index within the target task's array of notification values to which the notification is to be sent. `uxIndexToNotify` must be less than `configTASK_NOTIFICATION_ARRAY_ENTRIES`. `xTaskNotifyFromISR()` does not have this parameter and always sends notifications to index 0.
- **xTaskToNotify** -- The handle of the task being notified. The handle to a task can be returned from the `xTaskCreate()` API function used to create the task, and the handle of the currently running task can be obtained by calling `xTaskGetCurrentTaskHandle()`.
- **ulValue** -- Data that can be sent with the notification. How the data is used depends on the value of the `eAction` parameter.
- **eAction** -- Specifies how the notification updates the task's notification value, if at all. Valid values for `eAction` are as follows:
- **pxHigherPriorityTaskWoken** -- `xTaskNotifyFromISR()` will set `*pxHigherPriorityTaskWoken` to `pdTRUE` if sending the notification caused the task to which the notification was sent to leave the Blocked state, and the unblocked task has a priority higher than the currently running task. If `xTaskNotifyFromISR()` sets this value to `pdTRUE` then a context switch should be requested before the interrupt is exited. How a context switch is requested from an ISR is dependent on the port - see the documentation page for the port in use.

Returns Dependent on the value of `eAction`. See the description of the `eAction` parameter.

xTaskNotifyAndQueryIndexedFromISR (`xTaskToNotify`, `uxIndexToNotify`, `ulValue`, `eAction`, `pulPreviousNotificationValue`, `pxHigherPriorityTaskWoken`)

See <https://www.FreeRTOS.org/RTOS-task-notifications.html> for details.

`xTaskNotifyAndQueryIndexedFromISR()` performs the same operation as `xTaskNotifyIndexedFromISR()` with the addition that it also returns the subject task's prior notification value (the notification value at the time the function is called rather than at the time the function returns) in the additional `pulPreviousNotificationValue` parameter.

`xTaskNotifyAndQueryFromISR()` performs the same operation as `xTaskNotifyFromISR()` with the addition that it also returns the subject task's prior notification value (the notification value at the time the function is called rather than at the time the function returns) in the additional `pulPreviousNotificationValue` parameter.

xTaskNotifyWait (`ulBitsToClearOnEntry`, `ulBitsToClearOnExit`, `pulNotificationValue`, `xTicksToWait`)

xTaskNotifyWaitIndexed (`uxIndexToWaitOn`, `ulBitsToClearOnEntry`, `ulBitsToClearOnExit`, `pulNotificationValue`, `xTicksToWait`)

xTaskNotifyGiveIndexed (`xTaskToNotify`, `uxIndexToNotify`)

Sends a direct to task notification to a particular index in the target task's notification array in a manner similar to giving a counting semaphore.

See <https://www.FreeRTOS.org/RTOS-task-notifications.html> for more details.

`configUSE_TASK_NOTIFICATIONS` must be undefined or defined as 1 for these macros to be available.

Each task has a private array of "notification values" (or 'notifications'), each of which is a 32-bit unsigned integer (`uint32_t`). The constant `configTASK_NOTIFICATION_ARRAY_ENTRIES` sets the number of indexes in the array, and (for backward compatibility) defaults to 1 if left undefined. Prior to FreeRTOS V10.4.0 there was only one notification value per task.

Events can be sent to a task using an intermediary object. Examples of such objects are queues, semaphores, mutexes and event groups. Task notifications are a method of sending an event directly to a task without the need for such an intermediary object.

A notification sent to a task can optionally perform an action, such as update, overwrite or increment one of the task's notification values. In that way task notifications can be used to send data to a task, or be used as light weight and fast binary or counting semaphores.

`xTaskNotifyGiveIndexed()` is a helper macro intended for use when task notifications are used as light weight and faster binary or counting semaphore equivalents. Actual FreeRTOS semaphores are given using the `xSemaphoreGive()` API function, the equivalent action that instead uses a task notification is `xTaskNotifyGiveIndexed()`.

When task notifications are being used as a binary or counting semaphore equivalent then the task being notified should wait for the notification using the `ulTaskNotifyTakeIndexed()` API function rather than the `xTaskNotifyWaitIndexed()` API function.

NOTE Each notification within the array operates independently - a task can only block on one notification within the array at a time and will not be unblocked by a notification sent to any other array index.

Backward compatibility information: Prior to FreeRTOS V10.4.0 each task had a single "notification value", and all task notification API functions operated on that value. Replacing the single notification value with an array of notification values necessitated a new set of API functions that could address specific notifications within the array. `xTaskNotifyGive()` is the original API function, and remains backward compatible by always operating on the notification value at index 0 in the array. Calling `xTaskNotifyGive()` is equivalent to calling `xTaskNotifyGiveIndexed()` with the `uxIndexToNotify` parameter set to 0.

Parameters

- **`xTaskToNotify`** -- The handle of the task being notified. The handle to a task can be returned from the `xTaskCreate()` API function used to create the task, and the handle of the currently running task can be obtained by calling `xTaskGetCurrentTaskHandle()`.
- **`uxIndexToNotify`** -- The index within the target task's array of notification values to which the notification is to be sent. `uxIndexToNotify` must be less than `configTASK_NOTIFICATION_ARRAY_ENTRIES`. `xTaskNotifyGive()` does not have this parameter and always sends notifications to index 0.

Returns `xTaskNotifyGive()` is a macro that calls `xTaskNotify()` with the `eAction` parameter set to `eIncrement` - so `pdPASS` is always returned.

`vTaskNotifyGiveFromISR` (`xTaskToNotify`, `pxHigherPriorityTaskWoken`)

`vTaskNotifyGiveIndexedFromISR` (`xTaskToNotify`, `uxIndexToNotify`, `pxHigherPriorityTaskWoken`)

`ulTaskNotifyTakeIndexed` (`uxIndexToWaitOn`, `xClearCountOnExit`, `xTicksToWait`)

Waits for a direct to task notification on a particular index in the calling task's notification array in a manner similar to taking a counting semaphore.

See <https://www.FreeRTOS.org/RTOS-task-notifications.html> for details.

`configUSE_TASK_NOTIFICATIONS` must be undefined or defined as 1 for this function to be available.

Each task has a private array of "notification values" (or 'notifications'), each of which is a 32-bit unsigned integer (`uint32_t`). The constant `configTASK_NOTIFICATION_ARRAY_ENTRIES` sets the number of indexes in the array, and (for backward compatibility) defaults to 1 if left undefined. Prior to FreeRTOS V10.4.0 there was only one notification value per task.

Events can be sent to a task using an intermediary object. Examples of such objects are queues, semaphores, mutexes and event groups. Task notifications are a method of sending an event directly to a task without the need for such an intermediary object.

A notification sent to a task can optionally perform an action, such as update, overwrite or increment one of the task's notification values. In that way task notifications can be used to send data to a task, or be used as light weight and fast binary or counting semaphores.

`ulTaskNotifyTakeIndexed()` is intended for use when a task notification is used as a faster and lighter weight binary or counting semaphore alternative. Actual FreeRTOS semaphores are taken using the `xSemaphoreTake()` API function, the equivalent action that instead uses a task notification is `ulTaskNotifyTakeIndexed()`.

When a task is using its notification value as a binary or counting semaphore other tasks should send notifications to it using the `xTaskNotifyGiveIndexed()` macro, or `xTaskNotifyIndex()` function with the `eAction` parameter set to `eIncrement`.

`ulTaskNotifyTakeIndexed()` can either clear the task's notification value at the array index specified by the `uxIndexToWaitOn` parameter to zero on exit, in which case the notification value acts like a binary semaphore, or decrement the notification value on exit, in which case the notification value acts like a counting semaphore.

A task can use `ulTaskNotifyTakeIndexed()` to [optionally] block to wait for a notification. The task does not consume any CPU time while it is in the Blocked state.

Where as `xTaskNotifyWaitIndexed()` will return when a notification is pending, `ulTaskNotifyTakeIndexed()` will return when the task's notification value is not zero.

NOTE Each notification within the array operates independently - a task can only block on one notification within the array at a time and will not be unblocked by a notification sent to any other array index.

Backward compatibility information: Prior to FreeRTOS V10.4.0 each task had a single "notification value", and all task notification API functions operated on that value. Replacing the single notification value with an array of notification values necessitated a new set of API functions that could address specific notifications within the array. `ulTaskNotifyTake()` is the original API function, and remains backward compatible by always operating on the notification value at index 0 in the array. Calling `ulTaskNotifyTake()` is equivalent to calling `ulTaskNotifyTakeIndexed()` with the `uxIndexToWaitOn` parameter set to 0.

Parameters

- **`uxIndexToWaitOn`** -- The index within the calling task's array of notification values on which the calling task will wait for a notification to be non-zero. `uxIndexToWaitOn` must be less than `configTASK_NOTIFICATION_ARRAY_ENTRIES`. `xTaskNotifyTake()` does not have this parameter and always waits for notifications on index 0.
- **`xClearCountOnExit`** -- if `xClearCountOnExit` is `pdFALSE` then the task's notification value is decremented when the function exits. In this way the notification value acts like a counting semaphore. If `xClearCountOnExit` is not `pdFALSE` then the task's notification value is cleared to zero when the function exits. In this way the notification value acts like a binary semaphore.
- **`xTicksToWait`** -- The maximum amount of time that the task should wait in the Blocked state for the task's notification value to be greater than zero, should the count not already be greater than zero when `ulTaskNotifyTake()` was called. The task will not consume any processing time while it is in the Blocked state. This is specified in kernel ticks, the macro `pdMS_TO_TICKS( value_in_ms )` can be used to convert a time specified in milliseconds to a time specified in ticks.

Returns The task's notification count before it is either cleared to zero or decremented (see the `xClearCountOnExit` parameter).

`xTaskNotifyStateClear` (`xTask`)

`xTaskNotifyStateClearIndexed` (`xTask`, `uxIndexToClear`)

`ulTaskNotifyValueClear` (`xTask`, `ulBitsToClear`)

`ulTaskNotifyValueClearIndexed` (`xTask`, `uxIndexToClear`, `ulBitsToClear`)

Type Definitions

```
typedef struct tskTaskControlBlock *TaskHandle_t
```

```
typedef BaseType_t (*TaskHookFunction_t)(void*)
```

Defines the prototype to which the application task hook function must conform.

typedef struct *xTASK_STATUS* **TaskStatus_t**

Used with the uxTaskGetSystemState() function to return the state of each task in the system.

Enumerations

enum **eTaskState**

Task states returned by eTaskGetState.

Values:

enumerator **eRunning**

A task is querying the state of itself, so must be running.

enumerator **eReady**

The task being queried is in a ready or pending ready list.

enumerator **eBlocked**

The task being queried is in the Blocked state.

enumerator **eSuspended**

The task being queried is in the Suspended state, or is in the Blocked state with an infinite time out.

enumerator **eDeleted**

The task being queried has been deleted, but its TCB has not yet been freed.

enumerator **eInvalid**

Used as an 'invalid state' value.

enum **eNotifyAction**

Actions that can be performed when vTaskNotify() is called.

Values:

enumerator **eNoAction**

Notify the task without updating its notify value.

enumerator **eSetBits**

Set bits in the task's notification value.

enumerator **eIncrement**

Increment the task's notification value.

enumerator **eSetValueWithOverwrite**

Set the task's notification value to a specific value even if the previous value has not yet been read by the task.

enumerator **eSetValueWithoutOverwrite**

Set the task's notification value if the previous value has been read by the task.

enum eSleepModeStatus

Possible return values for eTaskConfirmSleepModeStatus().

Values:

enumerator eAbortSleep

A task has been made ready or a context switch pended since portSUPPRESS_TICKS_AND_SLEEP() was called - abort entering a sleep mode.

enumerator eStandardSleep

Enter a sleep mode that will not last any longer than the expected idle time.

Queue API

Header File

- [components/freertos/FreeRTOS-Kernel/include/freertos/queue.h](#)
- This header file can be included with:

```
#include "freertos/queue.h"
```

Functions

BaseType_t **xQueueGenericSend** (*QueueHandle_t* xQueue, const void *const pvItemToQueue, TickType_t xTicksToWait, const BaseType_t xCopyPosition)

It is preferred that the macros xQueueSend(), xQueueSendToFront() and xQueueSendToBack() are used in place of calling this function directly.

Post an item on a queue. The item is queued by copy, not by reference. This function must not be called from an interrupt service routine. See xQueueSendFromISR () for an alternative which may be used in an ISR.

Example usage:

```
struct AMessage
{
    char ucMessageID;
    char ucData[ 20 ];
} xMessage;

uint32_t ulVar = 10UL;

void vATask( void *pvParameters )
{
    QueueHandle_t xQueue1, xQueue2;
    struct AMessage *pxMessage;

    // Create a queue capable of containing 10 uint32_t values.
    xQueue1 = xQueueCreate( 10, sizeof( uint32_t ) );

    // Create a queue capable of containing 10 pointers to AMessage structures.
    // These should be passed by pointer as they contain a lot of data.
    xQueue2 = xQueueCreate( 10, sizeof( struct AMessage * ) );

    // ...

    if( xQueue1 != 0 )
```

(continues on next page)

(continued from previous page)

```

{
    // Send an uint32_t. Wait for 10 ticks for space to become
    // available if necessary.
    if( xQueueGenericSend( xQueue1, ( void * ) &ulVar, ( TickType_t ) 10,
→queueSEND_TO_BACK ) != pdPASS )
    {
        // Failed to post the message, even after 10 ticks.
    }
}

if( xQueue2 != 0 )
{
    // Send a pointer to a struct AMessage object. Don't block if the
    // queue is already full.
    pxMessage = &xMessage;
    xQueueGenericSend( xQueue2, ( void * ) &pxMessage, ( TickType_t ) 0,
→queueSEND_TO_BACK );
}

// ... Rest of task code.
}

```

Parameters

- **xQueue** -- The handle to the queue on which the item is to be posted.
- **pvItemToQueue** -- A pointer to the item that is to be placed on the queue. The size of the items the queue will hold was defined when the queue was created, so this many bytes will be copied from pvItemToQueue into the queue storage area.
- **xTicksToWait** -- The maximum amount of time the task should block waiting for space to become available on the queue, should it already be full. The call will return immediately if this is set to 0 and the queue is full. The time is defined in tick periods so the constant portTICK_PERIOD_MS should be used to convert to real time if this is required.
- **xCopyPosition** -- Can take the value queueSEND_TO_BACK to place the item at the back of the queue, or queueSEND_TO_FRONT to place the item at the front of the queue (for high priority messages).

Returns pdTRUE if the item was successfully posted, otherwise errQUEUE_FULL.

BaseType_t **xQueuePeek** (*QueueHandle_t* xQueue, void *const pvBuffer, TickType_t xTicksToWait)

Receive an item from a queue without removing the item from the queue. The item is received by copy so a buffer of adequate size must be provided. The number of bytes copied into the buffer was defined when the queue was created.

Successfully received items remain on the queue so will be returned again by the next call, or a call to xQueueReceive().

This macro must not be used in an interrupt service routine. See xQueuePeekFromISR() for an alternative that can be called from an interrupt service routine.

Example usage:

```

struct AMessage
{
    char ucMessageID;
    char ucData[ 20 ];
} xMessage;

QueueHandle_t xQueue;

```

(continues on next page)

(continued from previous page)

```

// Task to create a queue and post a value.
void vATask( void *pvParameters )
{
    struct AMessage *pxMessage;

    // Create a queue capable of containing 10 pointers to AMessage structures.
    // These should be passed by pointer as they contain a lot of data.
    xQueue = xQueueCreate( 10, sizeof( struct AMessage * ) );
    if( xQueue == 0 )
    {
        // Failed to create the queue.
    }

    // ...

    // Send a pointer to a struct AMessage object. Don't block if the
    // queue is already full.
    pxMessage = & xMessage;
    xQueueSend( xQueue, ( void * ) &pxMessage, ( TickType_t ) 0 );

    // ... Rest of task code.
}

// Task to peek the data from the queue.
void vADifferentTask( void *pvParameters )
{
    struct AMessage *pRxedMessage;

    if( xQueue != 0 )
    {
        // Peek a message on the created queue. Block for 10 ticks if a
        // message is not immediately available.
        if( xQueuePeek( xQueue, &( pRxedMessage ), ( TickType_t ) 10 ) )
        {
            // pRxedMessage now points to the struct AMessage variable posted
            // by vATask, but the item still remains on the queue.
        }
    }

    // ... Rest of task code.
}

```

Parameters

- **xQueue** -- The handle to the queue from which the item is to be received.
- **pvBuffer** -- Pointer to the buffer into which the received item will be copied.
- **xTicksToWait** -- The maximum amount of time the task should block waiting for an item to receive should the queue be empty at the time of the call. The time is defined in tick periods so the constant `portTICK_PERIOD_MS` should be used to convert to real time if this is required. `xQueuePeek()` will return immediately if `xTicksToWait` is 0 and the queue is empty.

Returns `pdTRUE` if an item was successfully received from the queue, otherwise `pdFALSE`.

BaseType_t **xQueuePeekFromISR** (*QueueHandle_t* xQueue, void *const pvBuffer)

A version of `xQueuePeek()` that can be called from an interrupt service routine (ISR).

Receive an item from a queue without removing the item from the queue. The item is received by copy so a buffer of adequate size must be provided. The number of bytes copied into the buffer was defined when the queue was created.

Successfully received items remain on the queue so will be returned again by the next call, or a call to `xQueueReceive()`.

Parameters

- **xQueue** -- The handle to the queue from which the item is to be received.
- **pvBuffer** -- Pointer to the buffer into which the received item will be copied.

Returns pdTRUE if an item was successfully received from the queue, otherwise pdFALSE.

BaseType_t **xQueueReceive** (*QueueHandle_t* xQueue, void *const pvBuffer, TickType_t xTicksToWait)

Receive an item from a queue. The item is received by copy so a buffer of adequate size must be provided. The number of bytes copied into the buffer was defined when the queue was created.

Successfully received items are removed from the queue.

This function must not be used in an interrupt service routine. See xQueueReceiveFromISR for an alternative that can.

Example usage:

```

struct AMessage
{
    char ucMessageID;
    char ucData[ 20 ];
} xMessage;

QueueHandle_t xQueue;

// Task to create a queue and post a value.
void vATask( void *pvParameters )
{
    struct AMessage *pxMessage;

    // Create a queue capable of containing 10 pointers to AMessage structures.
    // These should be passed by pointer as they contain a lot of data.
    xQueue = xQueueCreate( 10, sizeof( struct AMessage * ) );
    if( xQueue == 0 )
    {
        // Failed to create the queue.
    }

    // ...

    // Send a pointer to a struct AMessage object. Don't block if the
    // queue is already full.
    pxMessage = & xMessage;
    xQueueSend( xQueue, ( void * ) &pxMessage, ( TickType_t ) 0 );

    // ... Rest of task code.
}

// Task to receive from the queue.
void vADifferentTask( void *pvParameters )
{
    struct AMessage *pRxedMessage;

    if( xQueue != 0 )
    {
        // Receive a message on the created queue. Block for 10 ticks if a
        // message is not immediately available.
        if( xQueueReceive( xQueue, &( pRxedMessage ), ( TickType_t ) 10 ) )
        {
            // pRxedMessage now points to the struct AMessage variable posted
            // by vATask.
        }
    }
}

```

(continues on next page)

(continued from previous page)

```
// ... Rest of task code.
}
```

Parameters

- **xQueue** -- The handle to the queue from which the item is to be received.
- **pvBuffer** -- Pointer to the buffer into which the received item will be copied.
- **xTicksToWait** -- The maximum amount of time the task should block waiting for an item to receive should the queue be empty at the time of the call. `xQueueReceive()` will return immediately if `xTicksToWait` is zero and the queue is empty. The time is defined in tick periods so the constant `portTICK_PERIOD_MS` should be used to convert to real time if this is required.

Returns `pdTRUE` if an item was successfully received from the queue, otherwise `pdFALSE`.

`UBaseType_t uxQueueMessagesWaiting` (const [QueueHandle_t](#) xQueue)

Return the number of messages stored in a queue.

Parameters **xQueue** -- A handle to the queue being queried.

Returns The number of messages available in the queue.

`UBaseType_t uxQueueSpacesAvailable` (const [QueueHandle_t](#) xQueue)

Return the number of free spaces available in a queue. This is equal to the number of items that can be sent to the queue before the queue becomes full if no items are removed.

Parameters **xQueue** -- A handle to the queue being queried.

Returns The number of spaces available in the queue.

`void vQueueDelete` ([QueueHandle_t](#) xQueue)

Delete a queue - freeing all the memory allocated for storing of items placed on the queue.

Parameters **xQueue** -- A handle to the queue to be deleted.

`BaseType_t xQueueGenericSendFromISR` ([QueueHandle_t](#) xQueue, const void *const pvItemToQueue, `BaseType_t` *const pxHigherPriorityTaskWoken, const `BaseType_t` xCopyPosition)

It is preferred that the macros `xQueueSendFromISR()`, `xQueueSendToFrontFromISR()` and `xQueueSendToBackFromISR()` be used in place of calling this function directly. `xQueueGiveFromISR()` is an equivalent for use by semaphores that don't actually copy any data.

Post an item on a queue. It is safe to use this function from within an interrupt service routine.

Items are queued by copy not reference so it is preferable to only queue small items, especially when called from an ISR. In most cases it would be preferable to store a pointer to the item being queued.

Example usage for buffered IO (where the ISR can obtain more than one value per call):

```
void vBufferISR( void )
{
    char cIn;
    BaseType_t xHigherPriorityTaskWokenByPost;

    // We have not woken a task at the start of the ISR.
    xHigherPriorityTaskWokenByPost = pdFALSE;

    // Loop until the buffer is empty.
    do
    {
        // Obtain a byte from the buffer.
        cIn = portINPUT_BYTE( RX_REGISTER_ADDRESS );
```

(continues on next page)

(continued from previous page)

```

    // Post each byte.
    xQueueGenericSendFromISR( xRxQueue, &cIn, &xHigherPriorityTaskWokenByPost,
    ↪ queueSEND_TO_BACK );

} while( portINPUT_BYTE( BUFFER_COUNT ) );

// Now the buffer is empty we can switch context if necessary. Note that the
// name of the yield function required is port specific.
if( xHigherPriorityTaskWokenByPost )
{
    portYIELD_FROM_ISR();
}
}

```

Parameters

- **xQueue** -- The handle to the queue on which the item is to be posted.
- **pvItemToQueue** -- A pointer to the item that is to be placed on the queue. The size of the items the queue will hold was defined when the queue was created, so this many bytes will be copied from pvItemToQueue into the queue storage area.
- **pxHigherPriorityTaskWoken** -- xQueueGenericSendFromISR() will set *pxHigherPriorityTaskWoken to pdTRUE if sending to the queue caused a task to unblock, and the unblocked task has a priority higher than the currently running task. If xQueueGenericSendFromISR() sets this value to pdTRUE then a context switch should be requested before the interrupt is exited.
- **xCopyPosition** -- Can take the value queueSEND_TO_BACK to place the item at the back of the queue, or queueSEND_TO_FRONT to place the item at the front of the queue (for high priority messages).

Returns pdTRUE if the data was successfully sent to the queue, otherwise errQUEUE_FULL.

BaseType_t **xQueueGiveFromISR** ([QueueHandle_t](#) xQueue, BaseType_t *const pxHigherPriorityTaskWoken)

BaseType_t **xQueueReceiveFromISR** ([QueueHandle_t](#) xQueue, void *const pvBuffer, BaseType_t *const pxHigherPriorityTaskWoken)

Receive an item from a queue. It is safe to use this function from within an interrupt service routine.

Example usage:

```

QueueHandle_t xQueue;

// Function to create a queue and post some values.
void vAFunction( void *pvParameters )
{
    char cValueToPost;
    const TickType_t xTicksToWait = ( TickType_t )0xff;

    // Create a queue capable of containing 10 characters.
    xQueue = xQueueCreate( 10, sizeof( char ) );
    if( xQueue == 0 )
    {
        // Failed to create the queue.
    }

    // ...

    // Post some characters that will be used within an ISR. If the queue
    // is full then this task will block for xTicksToWait ticks.
    cValueToPost = 'a';

```

(continues on next page)

(continued from previous page)

```

xQueueSend( xQueue, ( void * ) &cValueToPost, xTicksToWait );
cValueToPost = 'b';
xQueueSend( xQueue, ( void * ) &cValueToPost, xTicksToWait );

// ... keep posting characters ... this task may block when the queue
// becomes full.

cValueToPost = 'c';
xQueueSend( xQueue, ( void * ) &cValueToPost, xTicksToWait );
}

// ISR that outputs all the characters received on the queue.
void vISR_Routine( void )
{
    BaseType_t xTaskWokenByReceive = pdFALSE;
    char cRxedChar;

    while( xQueueReceiveFromISR( xQueue, ( void * ) &cRxedChar, &
    ↪xTaskWokenByReceive) )
    {
        // A character was received. Output the character now.
        vOutputCharacter( cRxedChar );

        // If removing the character from the queue woke the task that was
        // posting onto the queue xTaskWokenByReceive will have been set to
        // pdTRUE. No matter how many times this loop iterates only one
        // task will be woken.
    }

    if( xTaskWokenByReceive != ( char ) pdFALSE;
    {
        taskYIELD ();
    }
}

```

Parameters

- **xQueue** -- The handle to the queue from which the item is to be received.
- **pvBuffer** -- Pointer to the buffer into which the received item will be copied.
- **pxHigherPriorityTaskWoken** -- A task may be blocked waiting for space to become available on the queue. If xQueueReceiveFromISR causes such a task to unblock *pxTaskWoken will get set to pdTRUE, otherwise *pxTaskWoken will remain unchanged.

Returns pdTRUE if an item was successfully received from the queue, otherwise pdFALSE.

BaseType_t **xQueueIsQueueEmptyFromISR** (const [QueueHandle_t](#) xQueue)

Queries a queue to determine if the queue is empty. This function should only be used in an ISR.

Parameters **xQueue** -- The handle of the queue being queried

Returns pdFALSE if the queue is not empty, or pdTRUE if the queue is empty.

BaseType_t **xQueueIsQueueFullFromISR** (const [QueueHandle_t](#) xQueue)

Queries a queue to determine if the queue is full. This function should only be used in an ISR.

Parameters **xQueue** -- The handle of the queue being queried

Returns pdFALSE if the queue is not full, or pdTRUE if the queue is full.

UBaseType_t **uxQueueMessagesWaitingFromISR** (const [QueueHandle_t](#) xQueue)

A version of uxQueueMessagesWaiting() that can be called from an ISR. Return the number of messages stored in a queue.

Parameters **xQueue** -- A handle to the queue being queried.

Returns The number of messages available in the queue.

void **vQueueAddToRegistry** (*QueueHandle_t* xQueue, const char *pcQueueName)

The registry is provided as a means for kernel aware debuggers to locate queues, semaphores and mutexes. Call vQueueAddToRegistry() add a queue, semaphore or mutex handle to the registry if you want the handle to be available to a kernel aware debugger. If you are not using a kernel aware debugger then this function can be ignored.

configQUEUE_REGISTRY_SIZE defines the maximum number of handles the registry can hold. configQUEUE_REGISTRY_SIZE must be greater than 0 within FreeRTOSConfig.h for the registry to be available. Its value does not affect the number of queues, semaphores and mutexes that can be created - just the number that the registry can hold.

If vQueueAddToRegistry is called more than once with the same xQueue parameter, the registry will store the pcQueueName parameter from the most recent call to vQueueAddToRegistry.

Parameters

- **xQueue** -- The handle of the queue being added to the registry. This is the handle returned by a call to xQueueCreate(). Semaphore and mutex handles can also be passed in here.
- **pcQueueName** -- The name to be associated with the handle. This is the name that the kernel aware debugger will display. The queue registry only stores a pointer to the string - so the string must be persistent (global or preferably in ROM/Flash), not on the stack.

void **vQueueUnregisterQueue** (*QueueHandle_t* xQueue)

The registry is provided as a means for kernel aware debuggers to locate queues, semaphores and mutexes. Call vQueueAddToRegistry() add a queue, semaphore or mutex handle to the registry if you want the handle to be available to a kernel aware debugger, and vQueueUnregisterQueue() to remove the queue, semaphore or mutex from the register. If you are not using a kernel aware debugger then this function can be ignored.

Parameters **xQueue** -- The handle of the queue being removed from the registry.

const char ***pcQueueGetName** (*QueueHandle_t* xQueue)

The queue registry is provided as a means for kernel aware debuggers to locate queues, semaphores and mutexes. Call pcQueueGetName() to look up and return the name of a queue in the queue registry from the queue's handle.

Parameters **xQueue** -- The handle of the queue the name of which will be returned.

Returns If the queue is in the registry then a pointer to the name of the queue is returned. If the queue is not in the registry then NULL is returned.

QueueSetHandle_t **xQueueCreateSet** (const UBaseType_t uxEventQueueLength)

Queue sets provide a mechanism to allow a task to block (pend) on a read operation from multiple queues or semaphores simultaneously.

See FreeRTOS/Source/Demo/Common/Minimal/QueueSet.c for an example using this function.

A queue set must be explicitly created using a call to xQueueCreateSet() before it can be used. Once created, standard FreeRTOS queues and semaphores can be added to the set using calls to xQueueAddToSet(). xQueueSelectFromSet() is then used to determine which, if any, of the queues or semaphores contained in the set is in a state where a queue read or semaphore take operation would be successful.

Note 1: See the documentation on <https://www.freertos.org/Documentation/02-Kernel/04-API-references/07-Queue-sets/00-RTOS-queue-sets> for reasons why queue sets are very rarely needed in practice as there are simpler methods of blocking on multiple objects.

Note 2: Blocking on a queue set that contains a mutex will not cause the mutex holder to inherit the priority of the blocked task.

Note 3: An additional 4 bytes of RAM is required for each space in a every queue added to a queue set. Therefore counting semaphores that have a high maximum count value should not be added to a queue set.

Note 4: A receive (in the case of a queue) or take (in the case of a semaphore) operation must not be performed on a member of a queue set unless a call to xQueueSelectFromSet() has first returned a handle to that set member.

Parameters **uxEventQueueLength** -- Queue sets store events that occur on the queues and semaphores contained in the set. uxEventQueueLength specifies the maximum number of

events that can be queued at once. To be absolutely certain that events are not lost `uxEventQueueLength` should be set to the total sum of the length of the queues added to the set, where binary semaphores and mutexes have a length of 1, and counting semaphores have a length set by their maximum count value. Examples:

- If a queue set is to hold a queue of length 5, another queue of length 12, and a binary semaphore, then `uxEventQueueLength` should be set to $(5 + 12 + 1)$, or 18.
- If a queue set is to hold three binary semaphores then `uxEventQueueLength` should be set to $(1 + 1 + 1)$, or 3.
- If a queue set is to hold a counting semaphore that has a maximum count of 5, and a counting semaphore that has a maximum count of 3, then `uxEventQueueLength` should be set to $(5 + 3)$, or 8.

Returns If the queue set is created successfully then a handle to the created queue set is returned. Otherwise NULL is returned.

BaseType_t **xQueueAddToSet** (*QueueSetMemberHandle_t* xQueueOrSemaphore, *QueueSetHandle_t* xQueueSet)

Adds a queue or semaphore to a queue set that was previously created by a call to `xQueueCreateSet()`.

See FreeRTOS/Source/Demo/Common/Minimal/QueueSet.c for an example using this function.

Note 1: A receive (in the case of a queue) or take (in the case of a semaphore) operation must not be performed on a member of a queue set unless a call to `xQueueSelectFromSet()` has first returned a handle to that set member.

Parameters

- **xQueueOrSemaphore** -- The handle of the queue or semaphore being added to the queue set (cast to an `QueueSetMemberHandle_t` type).
- **xQueueSet** -- The handle of the queue set to which the queue or semaphore is being added.

Returns If the queue or semaphore was successfully added to the queue set then `pdPASS` is returned. If the queue could not be successfully added to the queue set because it is already a member of a different queue set then `pdFAIL` is returned.

BaseType_t **xQueueRemoveFromSet** (*QueueSetMemberHandle_t* xQueueOrSemaphore, *QueueSetHandle_t* xQueueSet)

Removes a queue or semaphore from a queue set. A queue or semaphore can only be removed from a set if the queue or semaphore is empty.

See FreeRTOS/Source/Demo/Common/Minimal/QueueSet.c for an example using this function.

Parameters

- **xQueueOrSemaphore** -- The handle of the queue or semaphore being removed from the queue set (cast to an `QueueSetMemberHandle_t` type).
- **xQueueSet** -- The handle of the queue set in which the queue or semaphore is included.

Returns If the queue or semaphore was successfully removed from the queue set then `pdPASS` is returned. If the queue was not in the queue set, or the queue (or semaphore) was not empty, then `pdFAIL` is returned.

QueueSetMemberHandle_t **xQueueSelectFromSet** (*QueueSetHandle_t* xQueueSet, const TickType_t xTicksToWait)

`xQueueSelectFromSet()` selects from the members of a queue set a queue or semaphore that either contains data (in the case of a queue) or is available to take (in the case of a semaphore). `xQueueSelectFromSet()` effectively allows a task to block (pend) on a read operation on all the queues and semaphores in a queue set simultaneously.

See FreeRTOS/Source/Demo/Common/Minimal/QueueSet.c for an example using this function.

Note 1: See the documentation on <https://www.freertos.org/Documentation/02-Kernel/04-API-references/07-Queue-sets/00-RTOS-queue-sets> for reasons why queue sets are very rarely needed in practice as there are simpler methods of blocking on multiple objects.

Note 2: Blocking on a queue set that contains a mutex will not cause the mutex holder to inherit the priority of the blocked task.

Note 3: A receive (in the case of a queue) or take (in the case of a semaphore) operation must not be performed on a member of a queue set unless a call to `xQueueSelectFromSet()` has first returned a handle to that set member.

Parameters

- **xQueueSet** -- The queue set on which the task will (potentially) block.
- **xTicksToWait** -- The maximum time, in ticks, that the calling task will remain in the Blocked state (with other tasks executing) to wait for a member of the queue set to be ready for a successful queue read or semaphore take operation.

Returns `xQueueSelectFromSet()` will return the handle of a queue (cast to a `QueueSetMemberHandle_t` type) contained in the queue set that contains data, or the handle of a semaphore (cast to a `QueueSetMemberHandle_t` type) contained in the queue set that is available, or NULL if no such queue or semaphore exists before the specified block time expires.

QueueSetMemberHandle_t **xQueueSelectFromSetFromISR** (*QueueSetHandle_t* xQueueSet)

A version of `xQueueSelectFromSet()` that can be used from an ISR.

Macros

xQueueCreate (uxQueueLength, uxItemSize)

Creates a new queue instance, and returns a handle by which the new queue can be referenced.

Internally, within the FreeRTOS implementation, queues use two blocks of memory. The first block is used to hold the queue's data structures. The second block is used to hold items placed into the queue. If a queue is created using `xQueueCreate()` then both blocks of memory are automatically dynamically allocated inside the `xQueueCreate()` function. (see <https://www.FreeRTOS.org/a00111.html>). If a queue is created using `xQueueCreateStatic()` then the application writer must provide the memory that will get used by the queue. `xQueueCreateStatic()` therefore allows a queue to be created without using any dynamic memory allocation.

<https://www.FreeRTOS.org/Embedded-RTOS-Queues.html>

Example usage:

```
struct AMessage
{
    char ucMessageID;
    char ucData[ 20 ];
};

void vATask( void *pvParameters )
{
    QueueHandle_t xQueue1, xQueue2;

    // Create a queue capable of containing 10 uint32_t values.
    xQueue1 = xQueueCreate( 10, sizeof( uint32_t ) );
    if( xQueue1 == 0 )
    {
        // Queue was not created and must not be used.
    }

    // Create a queue capable of containing 10 pointers to AMessage structures.
    // These should be passed by pointer as they contain a lot of data.
    xQueue2 = xQueueCreate( 10, sizeof( struct AMessage * ) );
    if( xQueue2 == 0 )
    {
        // Queue was not created and must not be used.
    }

    // ... Rest of task code.
}
```

Parameters

- **uxQueueLength** -- The maximum number of items that the queue can contain.
- **uxItemSize** -- The number of bytes each item in the queue will require. Items are queued by copy, not by reference, so this is the number of bytes that will be copied for each posted item. Each item on the queue must be the same size.

Returns If the queue is successfully create then a handle to the newly created queue is returned. If the queue cannot be created then 0 is returned.

xQueueCreateStatic (uxQueueLength, uxItemSize, pucQueueStorage, pxQueueBuffer)

Creates a new queue instance, and returns a handle by which the new queue can be referenced.

Internally, within the FreeRTOS implementation, queues use two blocks of memory. The first block is used to hold the queue's data structures. The second block is used to hold items placed into the queue. If a queue is created using xQueueCreate() then both blocks of memory are automatically dynamically allocated inside the xQueueCreate() function. (see <https://www.FreeRTOS.org/a00111.html>). If a queue is created using xQueueCreateStatic() then the application writer must provide the memory that will get used by the queue. xQueueCreateStatic() therefore allows a queue to be created without using any dynamic memory allocation.

<https://www.FreeRTOS.org/Embedded-RTOS-Queues.html>

Example usage:

```
struct AMessage
{
    char ucMessageID;
    char ucData[ 20 ];
};

#define QUEUE_LENGTH 10
#define ITEM_SIZE sizeof( uint32_t )

// xQueueBuffer will hold the queue structure.
StaticQueue_t xQueueBuffer;

// ucQueueStorage will hold the items posted to the queue. Must be at least
// [(queue length) * ( queue item size)] bytes long.
uint8_t ucQueueStorage[ QUEUE_LENGTH * ITEM_SIZE ];

void vATask( void *pvParameters )
{
    QueueHandle_t xQueue1;

    // Create a queue capable of containing 10 uint32_t values.
    xQueue1 = xQueueCreate( QUEUE_LENGTH, // The number of items the queue can
        ↪hold.
                               ITEM_SIZE // The size of each item in the queue
        &( ucQueueStorage[ 0 ] ), // The buffer that will
        ↪hold the items in the queue.
        &xQueueBuffer ); // The buffer that will hold the
        ↪queue structure.

    // The queue is guaranteed to be created successfully as no dynamic memory
    // allocation is used. Therefore xQueue1 is now a handle to a valid queue.

    // ... Rest of task code.
}
```

Parameters

- **uxQueueLength** -- The maximum number of items that the queue can contain.
- **uxItemSize** -- The number of bytes each item in the queue will require. Items are queued by copy, not by reference, so this is the number of bytes that will be copied for

each posted item. Each item on the queue must be the same size.

- **pucQueueStorage** -- If uxItemSize is not zero then pucQueueStorage must point to a uint8_t array that is at least large enough to hold the maximum number of items that can be in the queue at any one time - which is (uxQueueLength * uxItemsSize) bytes. If uxItemSize is zero then pucQueueStorage can be NULL.
- **pxQueueBuffer** -- Must point to a variable of type StaticQueue_t, which will be used to hold the queue's data structure.

Returns If the queue is created then a handle to the created queue is returned. If pxQueueBuffer is NULL then NULL is returned.

xQueueGetStaticBuffers (xQueue, ppucQueueStorage, ppxStaticQueue)

Retrieve pointers to a statically created queue's data structure buffer and storage area buffer. These are the same buffers that are supplied at the time of creation.

Parameters

- **xQueue** -- The queue for which to retrieve the buffers.
- **ppucQueueStorage** -- Used to return a pointer to the queue's storage area buffer.
- **ppxStaticQueue** -- Used to return a pointer to the queue's data structure buffer.

Returns pdTRUE if buffers were retrieved, pdFALSE otherwise.

xQueueSendToFront (xQueue, pvItemToQueue, xTicksToWait)

Post an item to the front of a queue. The item is queued by copy, not by reference. This function must not be called from an interrupt service routine. See xQueueSendFromISR () for an alternative which may be used in an ISR.

Example usage:

```

struct AMessage
{
    char ucMessageID;
    char ucData[ 20 ];
} xMessage;

uint32_t ulVar = 10UL;

void vATask( void *pvParameters )
{
    QueueHandle_t xQueue1, xQueue2;
    struct AMessage *pxMessage;

    // Create a queue capable of containing 10 uint32_t values.
    xQueue1 = xQueueCreate( 10, sizeof( uint32_t ) );

    // Create a queue capable of containing 10 pointers to AMessage structures.
    // These should be passed by pointer as they contain a lot of data.
    xQueue2 = xQueueCreate( 10, sizeof( struct AMessage * ) );

    // ...

    if( xQueue1 != 0 )
    {
        // Send an uint32_t. Wait for 10 ticks for space to become
        // available if necessary.
        if( xQueueSendToFront( xQueue1, ( void * ) &ulVar, ( TickType_t ) 10 ) !=
        ↪pdPASS )
        {
            // Failed to post the message, even after 10 ticks.
        }
    }
}

```

(continues on next page)

(continued from previous page)

```

if( xQueue2 != 0 )
{
    // Send a pointer to a struct AMessage object. Don't block if the
    // queue is already full.
    pxMessage = & xMessage;
    xQueueSendToFront( xQueue2, ( void * ) &pxMessage, ( TickType_t ) 0 );
}

// ... Rest of task code.
}

```

Parameters

- **xQueue** -- The handle to the queue on which the item is to be posted.
- **pvItemToQueue** -- A pointer to the item that is to be placed on the queue. The size of the items the queue will hold was defined when the queue was created, so this many bytes will be copied from pvItemToQueue into the queue storage area.
- **xTicksToWait** -- The maximum amount of time the task should block waiting for space to become available on the queue, should it already be full. The call will return immediately if this is set to 0 and the queue is full. The time is defined in tick periods so the constant portTICK_PERIOD_MS should be used to convert to real time if this is required.

Returns pdTRUE if the item was successfully posted, otherwise errQUEUE_FULL.

xQueueSendToBack (xQueue, pvItemToQueue, xTicksToWait)

This is a macro that calls xQueueGenericSend().

Post an item to the back of a queue. The item is queued by copy, not by reference. This function must not be called from an interrupt service routine. See xQueueSendFromISR () for an alternative which may be used in an ISR.

Example usage:

```

struct AMessage
{
    char ucMessageID;
    char ucData[ 20 ];
} xMessage;

uint32_t ulVar = 10UL;

void vATask( void *pvParameters )
{
    QueueHandle_t xQueue1, xQueue2;
    struct AMessage *pxMessage;

    // Create a queue capable of containing 10 uint32_t values.
    xQueue1 = xQueueCreate( 10, sizeof( uint32_t ) );

    // Create a queue capable of containing 10 pointers to AMessage structures.
    // These should be passed by pointer as they contain a lot of data.
    xQueue2 = xQueueCreate( 10, sizeof( struct AMessage * ) );

    // ...

    if( xQueue1 != 0 )
    {
        // Send an uint32_t. Wait for 10 ticks for space to become
        // available if necessary.
    }
}

```

(continues on next page)

(continued from previous page)

```

    if( xQueueSendToBack( xQueue1, ( void * ) &ulVar, ( TickType_t ) 10 ) !=
    ↪pdPASS )
    {
        // Failed to post the message, even after 10 ticks.
    }
}

if( xQueue2 != 0 )
{
    // Send a pointer to a struct AMessage object. Don't block if the
    // queue is already full.
    pxMessage = &xMessage;
    xQueueSendToBack( xQueue2, ( void * ) &pxMessage, ( TickType_t ) 0 );
}

// ... Rest of task code.
}

```

Parameters

- **xQueue** -- The handle to the queue on which the item is to be posted.
- **pvItemToQueue** -- A pointer to the item that is to be placed on the queue. The size of the items the queue will hold was defined when the queue was created, so this many bytes will be copied from pvItemToQueue into the queue storage area.
- **xTicksToWait** -- The maximum amount of time the task should block waiting for space to become available on the queue, should it already be full. The call will return immediately if this is set to 0 and the queue is full. The time is defined in tick periods so the constant portTICK_PERIOD_MS should be used to convert to real time if this is required.

Returns pdTRUE if the item was successfully posted, otherwise errQUEUE_FULL.

xQueueSend (xQueue, pvItemToQueue, xTicksToWait)

This is a macro that calls xQueueGenericSend(). It is included for backward compatibility with versions of FreeRTOS.org that did not include the xQueueSendToFront() and xQueueSendToBack() macros. It is equivalent to xQueueSendToBack().

Post an item on a queue. The item is queued by copy, not by reference. This function must not be called from an interrupt service routine. See xQueueSendFromISR () for an alternative which may be used in an ISR.

Example usage:

```

struct AMessage
{
    char ucMessageID;
    char ucData[ 20 ];
} xMessage;

uint32_t ulVar = 10UL;

void vATask( void *pvParameters )
{
    QueueHandle_t xQueue1, xQueue2;
    struct AMessage *pxMessage;

    // Create a queue capable of containing 10 uint32_t values.
    xQueue1 = xQueueCreate( 10, sizeof( uint32_t ) );

    // Create a queue capable of containing 10 pointers to AMessage structures.
    // These should be passed by pointer as they contain a lot of data.

```

(continues on next page)

(continued from previous page)

```

xQueue2 = xQueueCreate( 10, sizeof( struct AMessage * ) );

// ...

if( xQueue1 != 0 )
{
    // Send an uint32_t. Wait for 10 ticks for space to become
    // available if necessary.
    if( xQueueSend( xQueue1, ( void * ) &ulVar, ( TickType_t ) 10 ) != pdPASS_
→)
    {
        // Failed to post the message, even after 10 ticks.
    }
}

if( xQueue2 != 0 )
{
    // Send a pointer to a struct AMessage object. Don't block if the
    // queue is already full.
    pxMessage = & xMessage;
    xQueueSend( xQueue2, ( void * ) &pxMessage, ( TickType_t ) 0 );
}

// ... Rest of task code.
}

```

Parameters

- **xQueue** -- The handle to the queue on which the item is to be posted.
- **pvItemToQueue** -- A pointer to the item that is to be placed on the queue. The size of the items the queue will hold was defined when the queue was created, so this many bytes will be copied from pvItemToQueue into the queue storage area.
- **xTicksToWait** -- The maximum amount of time the task should block waiting for space to become available on the queue, should it already be full. The call will return immediately if this is set to 0 and the queue is full. The time is defined in tick periods so the constant portTICK_PERIOD_MS should be used to convert to real time if this is required.

Returns pdTRUE if the item was successfully posted, otherwise errQUEUE_FULL.

xQueueOverwrite (xQueue, pvItemToQueue)

Only for use with queues that have a length of one - so the queue is either empty or full.

Post an item on a queue. If the queue is already full then overwrite the value held in the queue. The item is queued by copy, not by reference.

This function must not be called from an interrupt service routine. See xQueueOverwriteFromISR () for an alternative which may be used in an ISR.

Example usage:

```

void vFunction( void *pvParameters )
{
    QueueHandle_t xQueue;
    uint32_t ulVarToSend, ulValReceived;

    // Create a queue to hold one uint32_t value. It is strongly
    // recommended *not* to use xQueueOverwrite() on queues that can
    // contain more than one value, and doing so will trigger an assertion
    // if configASSERT() is defined.

```

(continues on next page)

(continued from previous page)

```

xQueue = xQueueCreate( 1, sizeof( uint32_t ) );

// Write the value 10 to the queue using xQueueOverwrite().
ulVarToSend = 10;
xQueueOverwrite( xQueue, &ulVarToSend );

// Peeking the queue should now return 10, but leave the value 10 in
// the queue. A block time of zero is used as it is known that the
// queue holds a value.
ulValReceived = 0;
xQueuePeek( xQueue, &ulValReceived, 0 );

if( ulValReceived != 10 )
{
    // Error unless the item was removed by a different task.
}

// The queue is still full. Use xQueueOverwrite() to overwrite the
// value held in the queue with 100.
ulVarToSend = 100;
xQueueOverwrite( xQueue, &ulVarToSend );

// This time read from the queue, leaving the queue empty once more.
// A block time of 0 is used again.
xQueueReceive( xQueue, &ulValReceived, 0 );

// The value read should be the last value written, even though the
// queue was already full when the value was written.
if( ulValReceived != 100 )
{
    // Error!
}

// ...
}

```

Parameters

- **xQueue** -- The handle of the queue to which the data is being sent.
- **pvItemToQueue** -- A pointer to the item that is to be placed on the queue. The size of the items the queue will hold was defined when the queue was created, so this many bytes will be copied from pvItemToQueue into the queue storage area.

Returns xQueueOverwrite() is a macro that calls xQueueGenericSend(), and therefore has the same return values as xQueueSendToFront(). However, pdPASS is the only value that can be returned because xQueueOverwrite() will write to the queue even when the queue is already full.

xQueueSendToFrontFromISR (xQueue, pvItemToQueue, pxHigherPriorityTaskWoken)

This is a macro that calls xQueueGenericSendFromISR().

Post an item to the front of a queue. It is safe to use this macro from within an interrupt service routine.

Items are queued by copy not reference so it is preferable to only queue small items, especially when called from an ISR. In most cases it would be preferable to store a pointer to the item being queued.

Example usage for buffered IO (where the ISR can obtain more than one value per call):

```

void vBufferISR( void )
{
    char cIn;

```

(continues on next page)

(continued from previous page)

```

BaseType_t xHigherPriorityTaskWoken;

// We have not woken a task at the start of the ISR.
xHigherPriorityTaskWoken = pdFALSE;

// Loop until the buffer is empty.
do
{
    // Obtain a byte from the buffer.
    cIn = portINPUT_BYTE( RX_REGISTER_ADDRESS );

    // Post the byte.
    xQueueSendToFrontFromISR( xRxQueue, &cIn, &xHigherPriorityTaskWoken );

} while( portINPUT_BYTE( BUFFER_COUNT ) );

// Now the buffer is empty we can switch context if necessary.
if( xHigherPriorityTaskWoken )
{
    taskYIELD ();
}
}

```

Parameters

- **xQueue** -- The handle to the queue on which the item is to be posted.
- **pvItemToQueue** -- A pointer to the item that is to be placed on the queue. The size of the items the queue will hold was defined when the queue was created, so this many bytes will be copied from pvItemToQueue into the queue storage area.
- **pxHigherPriorityTaskWoken** -- xQueueSendToFrontFromISR() will set *pxHigherPriorityTaskWoken to pdTRUE if sending to the queue caused a task to unblock, and the unblocked task has a priority higher than the currently running task. If xQueueSendToFromFromISR() sets this value to pdTRUE then a context switch should be requested before the interrupt is exited.

Returns pdTRUE if the data was successfully sent to the queue, otherwise errQUEUE_FULL.

xQueueSendToBackFromISR (xQueue, pvItemToQueue, pxHigherPriorityTaskWoken)

This is a macro that calls xQueueGenericSendFromISR().

Post an item to the back of a queue. It is safe to use this macro from within an interrupt service routine.

Items are queued by copy not reference so it is preferable to only queue small items, especially when called from an ISR. In most cases it would be preferable to store a pointer to the item being queued.

Example usage for buffered IO (where the ISR can obtain more than one value per call):

```

void vBufferISR( void )
{
    char cIn;
    BaseType_t xHigherPriorityTaskWoken;

    // We have not woken a task at the start of the ISR.
    xHigherPriorityTaskWoken = pdFALSE;

    // Loop until the buffer is empty.
    do
    {
        // Obtain a byte from the buffer.
        cIn = portINPUT_BYTE( RX_REGISTER_ADDRESS );
    }
    while( portINPUT_BYTE( BUFFER_COUNT ) );

    // Now the buffer is empty we can switch context if necessary.
    if( xHigherPriorityTaskWoken )
    {
        taskYIELD ();
    }
}

```

(continues on next page)

(continued from previous page)

```

    // Post the byte.
    xQueueSendToBackFromISR( xRxQueue, &cIn, &xHigherPriorityTaskWoken );

} while( portINPUT_BYTE( BUFFER_COUNT ) );

// Now the buffer is empty we can switch context if necessary.
if( xHigherPriorityTaskWoken )
{
    taskYIELD ();
}
}

```

Parameters

- **xQueue** -- The handle to the queue on which the item is to be posted.
- **pvItemToQueue** -- A pointer to the item that is to be placed on the queue. The size of the items the queue will hold was defined when the queue was created, so this many bytes will be copied from pvItemToQueue into the queue storage area.
- **pxHigherPriorityTaskWoken** -- xQueueSendToBackFromISR() will set *pxHigherPriorityTaskWoken to pdTRUE if sending to the queue caused a task to unblock, and the unblocked task has a priority higher than the currently running task. If xQueueSendToBackFromISR() sets this value to pdTRUE then a context switch should be requested before the interrupt is exited.

Returns pdTRUE if the data was successfully sent to the queue, otherwise errQUEUE_FULL.

xQueueOverwriteFromISR (xQueue, pvItemToQueue, pxHigherPriorityTaskWoken)

A version of xQueueOverwrite() that can be used in an interrupt service routine (ISR).

Only for use with queues that can hold a single item - so the queue is either empty or full.

Post an item on a queue. If the queue is already full then overwrite the value held in the queue. The item is queued by copy, not by reference.

Example usage:

```

QueueHandle_t xQueue;

void vFunction( void *pvParameters )
{
    // Create a queue to hold one uint32_t value. It is strongly
    // recommended *not* to use xQueueOverwriteFromISR() on queues that can
    // contain more than one value, and doing so will trigger an assertion
    // if configASSERT() is defined.
    xQueue = xQueueCreate( 1, sizeof( uint32_t ) );
}

void vAnInterruptHandler( void )
{
    // xHigherPriorityTaskWoken must be set to pdFALSE before it is used.
    BaseType_t xHigherPriorityTaskWoken = pdFALSE;
    uint32_t ulVarToSend, ulValReceived;

    // Write the value 10 to the queue using xQueueOverwriteFromISR().
    ulVarToSend = 10;
    xQueueOverwriteFromISR( xQueue, &ulVarToSend, &xHigherPriorityTaskWoken );

    // The queue is full, but calling xQueueOverwriteFromISR() again will still
    // pass because the value held in the queue will be overwritten with the

```

(continues on next page)

(continued from previous page)

```

// new value.
ulVarToSend = 100;
xQueueOverwriteFromISR( xQueue, &ulVarToSend, &xHigherPriorityTaskWoken );

// Reading from the queue will now return 100.

// ...

if( xHigherPriorityTaskWoken == pdTRUE )
{
    // Writing to the queue caused a task to unblock and the unblocked task
    // has a priority higher than or equal to the priority of the currently
    // executing task (the task this interrupt interrupted). Perform a
    → context
    // switch so this interrupt returns directly to the unblocked task.
    portYIELD_FROM_ISR(); // or portEND_SWITCHING_ISR() depending on the port.
}
}

```

Parameters

- **xQueue** -- The handle to the queue on which the item is to be posted.
- **pvItemToQueue** -- A pointer to the item that is to be placed on the queue. The size of the items the queue will hold was defined when the queue was created, so this many bytes will be copied from pvItemToQueue into the queue storage area.
- **pxHigherPriorityTaskWoken** -- xQueueOverwriteFromISR() will set *pxHigherPriorityTaskWoken to pdTRUE if sending to the queue caused a task to unblock, and the unblocked task has a priority higher than the currently running task. If xQueueOverwriteFromISR() sets this value to pdTRUE then a context switch should be requested before the interrupt is exited.

Returns xQueueOverwriteFromISR() is a macro that calls xQueueGenericSendFromISR(), and therefore has the same return values as xQueueSendToFrontFromISR(). However, pdPASS is the only value that can be returned because xQueueOverwriteFromISR() will write to the queue even when the queue is already full.

xQueueSendFromISR (xQueue, pvItemToQueue, pxHigherPriorityTaskWoken)

This is a macro that calls xQueueGenericSendFromISR(). It is included for backward compatibility with versions of FreeRTOS.org that did not include the xQueueSendToBackFromISR() and xQueueSendToFrontFromISR() macros.

Post an item to the back of a queue. It is safe to use this function from within an interrupt service routine.

Items are queued by copy not reference so it is preferable to only queue small items, especially when called from an ISR. In most cases it would be preferable to store a pointer to the item being queued.

Example usage for buffered IO (where the ISR can obtain more than one value per call):

```

void vBufferISR( void )
{
    char cIn;
    BaseType_t xHigherPriorityTaskWoken;

    // We have not woken a task at the start of the ISR.
    xHigherPriorityTaskWoken = pdFALSE;

    // Loop until the buffer is empty.
    do
    {
        // Obtain a byte from the buffer.
    }
}

```

(continues on next page)

(continued from previous page)

```

cIn = portINPUT_BYTE( RX_REGISTER_ADDRESS );

// Post the byte.
xQueueSendFromISR( xRxQueue, &cIn, &xHigherPriorityTaskWoken );

} while( portINPUT_BYTE( BUFFER_COUNT ) );

// Now the buffer is empty we can switch context if necessary.
if( xHigherPriorityTaskWoken )
{
    // Actual macro used here is port specific.
    portYIELD_FROM_ISR ();
}
}

```

Parameters

- **xQueue** -- The handle to the queue on which the item is to be posted.
- **pvItemToQueue** -- A pointer to the item that is to be placed on the queue. The size of the items the queue will hold was defined when the queue was created, so this many bytes will be copied from pvItemToQueue into the queue storage area.
- **pxHigherPriorityTaskWoken** -- xQueueSendFromISR() will set *pxHigherPriorityTaskWoken to pdTRUE if sending to the queue caused a task to unblock, and the unblocked task has a priority higher than the currently running task. If xQueueSendFromISR() sets this value to pdTRUE then a context switch should be requested before the interrupt is exited.

Returns pdTRUE if the data was successfully sent to the queue, otherwise errQUEUE_FULL.

xQueueReset (xQueue)

Reset a queue back to its original empty state. The return value is now obsolete and is always set to pdPASS.

Type Definitions

```
typedef struct QueueDefinition *QueueHandle_t
```

```
typedef struct QueueDefinition *QueueSetHandle_t
```

Type by which queue sets are referenced. For example, a call to xQueueCreateSet() returns an xQueueSet variable that can then be used as a parameter to xQueueSelectFromSet(), xQueueAddToSet(), etc.

```
typedef struct QueueDefinition *QueueSetMemberHandle_t
```

Queue sets can contain both queues and semaphores, so the QueueSetMemberHandle_t is defined as a type to be used where a parameter or return value can be either an QueueHandle_t or an SemaphoreHandle_t.

Semaphore API**Header File**

- [components/freertos/FreeRTOS-Kernel/include/freertos/semphr.h](#)
- This header file can be included with:

```
#include "freertos/semphr.h"
```

Macros

semBINARY_SEMAPHORE_QUEUE_LENGTH

semSEMAPHORE_QUEUE_ITEM_LENGTH**semGIVE_BLOCK_TIME****vSemaphoreCreateBinary** (xSemaphore)

In many usage scenarios it is faster and more memory efficient to use a direct to task notification in place of a binary semaphore! <https://www.FreeRTOS.org/RTOS-task-notifications.html>

This old vSemaphoreCreateBinary() macro is now deprecated in favour of the xSemaphoreCreateBinary() function. Note that binary semaphores created using the vSemaphoreCreateBinary() macro are created in a state such that the first call to 'take' the semaphore would pass, whereas binary semaphores created using xSemaphoreCreateBinary() are created in a state such that the semaphore must first be 'given' before it can be 'taken'.

Macro that implements a semaphore by using the existing queue mechanism. The queue length is 1 as this is a binary semaphore. The data size is 0 as we don't want to actually store any data - we just want to know if the queue is empty or full.

This type of semaphore can be used for pure synchronisation between tasks or between an interrupt and a task. The semaphore need not be given back once obtained, so one task/interrupt can continuously 'give' the semaphore while another continuously 'takes' the semaphore. For this reason this type of semaphore does not use a priority inheritance mechanism. For an alternative that does use priority inheritance see xSemaphoreCreateMutex().

Example usage:

```
SemaphoreHandle_t xSemaphore = NULL;

void vATask( void * pvParameters )
{
    // Semaphore cannot be used before a call to vSemaphoreCreateBinary ().
    // This is a macro so pass the variable in directly.
    vSemaphoreCreateBinary( xSemaphore );

    if( xSemaphore != NULL )
    {
        // The semaphore was created successfully.
        // The semaphore can now be used.
    }
}
```

Parameters

- **xSemaphore** -- Handle to the created semaphore. Should be of type SemaphoreHandle_t.

xSemaphoreCreateBinary ()

Creates a new binary semaphore instance, and returns a handle by which the new semaphore can be referenced.

In many usage scenarios it is faster and more memory efficient to use a direct to task notification in place of a binary semaphore! <https://www.FreeRTOS.org/RTOS-task-notifications.html>

Internally, within the FreeRTOS implementation, binary semaphores use a block of memory, in which the semaphore structure is stored. If a binary semaphore is created using xSemaphoreCreateBinary() then the required memory is automatically dynamically allocated inside the xSemaphoreCreateBinary() function. (see <https://www.FreeRTOS.org/a00111.html>). If a binary semaphore is created using xSemaphoreCreateBinaryStatic() then the application writer must provide the memory. xSemaphoreCreateBinaryStatic() therefore allows a binary semaphore to be created without using any dynamic memory allocation.

The old vSemaphoreCreateBinary() macro is now deprecated in favour of this xSemaphoreCreateBinary() function. Note that binary semaphores created using the vSemaphoreCreateBinary() macro are created in

a state such that the first call to 'take' the semaphore would pass, whereas binary semaphores created using `xSemaphoreCreateBinary()` are created in a state such that the semaphore must first be 'given' before it can be 'taken'.

This type of semaphore can be used for pure synchronisation between tasks or between an interrupt and a task. The semaphore need not be given back once obtained, so one task/interrupt can continuously 'give' the semaphore while another continuously 'takes' the semaphore. For this reason this type of semaphore does not use a priority inheritance mechanism. For an alternative that does use priority inheritance see `xSemaphoreCreateMutex()`.

Example usage:

```
SemaphoreHandle_t xSemaphore = NULL;

void vATask( void * pvParameters )
{
    // Semaphore cannot be used before a call to xSemaphoreCreateBinary().
    // This is a macro so pass the variable in directly.
    xSemaphore = xSemaphoreCreateBinary();

    if( xSemaphore != NULL )
    {
        // The semaphore was created successfully.
        // The semaphore can now be used.
    }
}
```

Returns Handle to the created semaphore, or NULL if the memory required to hold the semaphore's data structures could not be allocated.

xSemaphoreCreateBinaryStatic (pxStaticSemaphore)

Creates a new binary semaphore instance, and returns a handle by which the new semaphore can be referenced.

NOTE: In many usage scenarios it is faster and more memory efficient to use a direct to task notification in place of a binary semaphore! <https://www.FreeRTOS.org/RTOS-task-notifications.html>

Internally, within the FreeRTOS implementation, binary semaphores use a block of memory, in which the semaphore structure is stored. If a binary semaphore is created using `xSemaphoreCreateBinary()` then the required memory is automatically dynamically allocated inside the `xSemaphoreCreateBinary()` function. (see <https://www.FreeRTOS.org/a00111.html>). If a binary semaphore is created using `xSemaphoreCreateBinaryStatic()` then the application writer must provide the memory. `xSemaphoreCreateBinaryStatic()` therefore allows a binary semaphore to be created without using any dynamic memory allocation.

This type of semaphore can be used for pure synchronisation between tasks or between an interrupt and a task. The semaphore need not be given back once obtained, so one task/interrupt can continuously 'give' the semaphore while another continuously 'takes' the semaphore. For this reason this type of semaphore does not use a priority inheritance mechanism. For an alternative that does use priority inheritance see `xSemaphoreCreateMutex()`.

Example usage:

```
SemaphoreHandle_t xSemaphore = NULL;
StaticSemaphore_t xSemaphoreBuffer;

void vATask( void * pvParameters )
{
    // Semaphore cannot be used before a call to xSemaphoreCreateBinary().
    // The semaphore's data structures will be placed in the xSemaphoreBuffer
    // variable, the address of which is passed into the function. The
```

(continues on next page)

(continued from previous page)

```

// function's parameter is not NULL, so the function will not attempt any
// dynamic memory allocation, and therefore the function will not return
// return NULL.
xSemaphore = xSemaphoreCreateBinary( &xSemaphoreBuffer );

// Rest of task code goes here.
}

```

Parameters

- **pxStaticSemaphore** -- Must point to a variable of type StaticSemaphore_t, which will then be used to hold the semaphore's data structure, removing the need for the memory to be allocated dynamically.

Returns If the semaphore is created then a handle to the created semaphore is returned. If pxSemaphoreBuffer is NULL then NULL is returned.

xSemaphoreTake (xSemaphore, xBlockTime)

Macro to obtain a semaphore. The semaphore must have previously been created with a call to xSemaphoreCreateBinary(), xSemaphoreCreateMutex() or xSemaphoreCreateCounting().

Example usage:

```

SemaphoreHandle_t xSemaphore = NULL;

// A task that creates a semaphore.
void vATask( void * pvParameters )
{
    // Create the semaphore to guard a shared resource.
    xSemaphore = xSemaphoreCreateBinary();
}

// A task that uses the semaphore.
void vAnotherTask( void * pvParameters )
{
    // ... Do other things.

    if( xSemaphore != NULL )
    {
        // See if we can obtain the semaphore. If the semaphore is not available
        // wait 10 ticks to see if it becomes free.
        if( xSemaphoreTake( xSemaphore, ( TickType_t ) 10 ) == pdTRUE )
        {
            // We were able to obtain the semaphore and can now access the
            // shared resource.

            // ...

            // We have finished accessing the shared resource. Release the
            // semaphore.
            xSemaphoreGive( xSemaphore );
        }
        else
        {
            // We could not obtain the semaphore and can therefore not access
            // the shared resource safely.
        }
    }
}

```

Parameters

- **xSemaphore** -- A handle to the semaphore being taken - obtained when the semaphore was created.
- **xBlockTime** -- The time in ticks to wait for the semaphore to become available. The macro `portTICK_PERIOD_MS` can be used to convert this to a real time. A block time of zero can be used to poll the semaphore. A block time of `portMAX_DELAY` can be used to block indefinitely (provided `INCLUDE_vTaskSuspend` is set to 1 in `FreeRTOSConfig.h`).

Returns `pdTRUE` if the semaphore was obtained. `pdFALSE` if `xBlockTime` expired without the semaphore becoming available.

xSemaphoreTakeRecursive (xMutex, xBlockTime)

Macro to recursively obtain, or 'take', a mutex type semaphore. The mutex must have previously been created using a call to `xSemaphoreCreateRecursiveMutex()`;

`configUSE_RECURSIVE_MUTEXES` must be set to 1 in `FreeRTOSConfig.h` for this macro to be available.

This macro must not be used on mutexes created using `xSemaphoreCreateMutex()`.

A mutex used recursively can be 'taken' repeatedly by the owner. The mutex doesn't become available again until the owner has called `xSemaphoreGiveRecursive()` for each successful 'take' request. For example, if a task successfully 'takes' the same mutex 5 times then the mutex will not be available to any other task until it has also 'given' the mutex back exactly five times.

Example usage:

```
SemaphoreHandle_t xMutex = NULL;

// A task that creates a mutex.
void vATask( void * pvParameters )
{
    // Create the mutex to guard a shared resource.
    xMutex = xSemaphoreCreateRecursiveMutex();
}

// A task that uses the mutex.
void vAnotherTask( void * pvParameters )
{
    // ... Do other things.

    if( xMutex != NULL )
    {
        // See if we can obtain the mutex. If the mutex is not available
        // wait 10 ticks to see if it becomes free.
        if( xSemaphoreTakeRecursive( xSemaphore, ( TickType_t ) 10 ) == pdTRUE )
        {
            // We were able to obtain the mutex and can now access the
            // shared resource.

            // ...
            // For some reason due to the nature of the code further calls to
            // xSemaphoreTakeRecursive() are made on the same mutex. In real
            // code these would not be just sequential calls as this would make
            // no sense. Instead the calls are likely to be buried inside
            // a more complex call structure.
            xSemaphoreTakeRecursive( xMutex, ( TickType_t ) 10 );
            xSemaphoreTakeRecursive( xMutex, ( TickType_t ) 10 );

            // The mutex has now been 'taken' three times, so will not be
            // available to another task until it has also been given back
            // three times. Again it is unlikely that real code would have
            // these calls sequentially, but instead buried in a more complex
```

(continues on next page)

(continued from previous page)

```

        // call structure. This is just for illustrative purposes.
        xSemaphoreGiveRecursive( xMutex );
        xSemaphoreGiveRecursive( xMutex );
        xSemaphoreGiveRecursive( xMutex );

        // Now the mutex can be taken by other tasks.
    }
    else
    {
        // We could not obtain the mutex and can therefore not access
        // the shared resource safely.
    }
}
}

```

Parameters

- **xMutex** -- A handle to the mutex being obtained. This is the handle returned by `xSemaphoreCreateRecursiveMutex()`;
- **xBlockTime** -- The time in ticks to wait for the semaphore to become available. The macro `portTICK_PERIOD_MS` can be used to convert this to a real time. A block time of zero can be used to poll the semaphore. If the task already owns the semaphore then `xSemaphoreTakeRecursive()` will return immediately no matter what the value of `xBlockTime`.

Returns `pdTRUE` if the semaphore was obtained. `pdFALSE` if `xBlockTime` expired without the semaphore becoming available.

xSemaphoreGive (xSemaphore)

Macro to release a semaphore. The semaphore must have previously been created with a call to `xSemaphoreCreateBinary()`, `xSemaphoreCreateMutex()` or `xSemaphoreCreateCounting()`. and obtained using `xSemaphoreTake()`.

This macro must not be used from an ISR. See `xSemaphoreGiveFromISR()` for an alternative which can be used from an ISR.

This macro must also not be used on semaphores created using `xSemaphoreCreateRecursiveMutex()`.

Example usage:

```

SemaphoreHandle_t xSemaphore = NULL;

void vATask( void * pvParameters )
{
    // Create the semaphore to guard a shared resource.
    xSemaphore = vSemaphoreCreateBinary();

    if( xSemaphore != NULL )
    {
        if( xSemaphoreGive( xSemaphore ) != pdTRUE )
        {
            // We would expect this call to fail because we cannot give
            // a semaphore without first "taking" it!
        }

        // Obtain the semaphore - don't block if the semaphore is not
        // immediately available.
        if( xSemaphoreTake( xSemaphore, ( TickType_t ) 0 ) )
        {
            // We now have the semaphore and can access the shared resource.
        }
    }
}

```

(continues on next page)

(continued from previous page)

```

        // ...

        // We have finished accessing the shared resource so can free the
        // semaphore.
        if( xSemaphoreGive( xSemaphore ) != pdTRUE )
        {
            // We would not expect this call to fail because we must have
            // obtained the semaphore to get here.
        }
    }
}

```

Parameters

- **xSemaphore** -- A handle to the semaphore being released. This is the handle returned when the semaphore was created.

Returns pdTRUE if the semaphore was released. pdFALSE if an error occurred. Semaphores are implemented using queues. An error can occur if there is no space on the queue to post a message - indicating that the semaphore was not first obtained correctly.

xSemaphoreGiveRecursive (xMutex)

Macro to recursively release, or 'give', a mutex type semaphore. The mutex must have previously been created using a call to xSemaphoreCreateRecursiveMutex();

configUSE_RECURSIVE_MUTEXES must be set to 1 in FreeRTOSConfig.h for this macro to be available.

This macro must not be used on mutexes created using xSemaphoreCreateMutex().

A mutex used recursively can be 'taken' repeatedly by the owner. The mutex doesn't become available again until the owner has called xSemaphoreGiveRecursive() for each successful 'take' request. For example, if a task successfully 'takes' the same mutex 5 times then the mutex will not be available to any other task until it has also 'given' the mutex back exactly five times.

Example usage:

```

SemaphoreHandle_t xMutex = NULL;

// A task that creates a mutex.
void vATask( void * pvParameters )
{
    // Create the mutex to guard a shared resource.
    xMutex = xSemaphoreCreateRecursiveMutex();
}

// A task that uses the mutex.
void vAnotherTask( void * pvParameters )
{
    // ... Do other things.

    if( xMutex != NULL )
    {
        // See if we can obtain the mutex. If the mutex is not available
        // wait 10 ticks to see if it becomes free.
        if( xSemaphoreTakeRecursive( xMutex, ( TickType_t ) 10 ) == pdTRUE )
        {
            // We were able to obtain the mutex and can now access the
            // shared resource.
        }
    }
}

```

(continues on next page)

(continued from previous page)

```

// ...
// For some reason due to the nature of the code further calls to
// xSemaphoreTakeRecursive() are made on the same mutex. In real
// code these would not be just sequential calls as this would make
// no sense. Instead the calls are likely to be buried inside
// a more complex call structure.
xSemaphoreTakeRecursive( xMutex, ( TickType_t ) 10 );
xSemaphoreTakeRecursive( xMutex, ( TickType_t ) 10 );

// The mutex has now been 'taken' three times, so will not be
// available to another task until it has also been given back
// three times. Again it is unlikely that real code would have
// these calls sequentially, it would be more likely that the calls
// to xSemaphoreGiveRecursive() would be called as a call stack
// unwound. This is just for demonstrative purposes.
xSemaphoreGiveRecursive( xMutex );
xSemaphoreGiveRecursive( xMutex );
xSemaphoreGiveRecursive( xMutex );

// Now the mutex can be taken by other tasks.
}
else
{
    // We could not obtain the mutex and can therefore not access
    // the shared resource safely.
}
}
}

```

Parameters

- **xMutex** -- A handle to the mutex being released, or 'given'. This is the handle returned by xSemaphoreCreateMutex();

Returns pdTRUE if the semaphore was given.

xSemaphoreGiveFromISR (xSemaphore, pxHigherPriorityTaskWoken)

Macro to release a semaphore. The semaphore must have previously been created with a call to xSemaphoreCreateBinary() or xSemaphoreCreateCounting().

Mutex type semaphores (those created using a call to xSemaphoreCreateMutex()) must not be used with this macro.

This macro can be used from an ISR.

Example usage:

```

#define LONG_TIME 0xffff
#define TICKS_TO_WAIT 10
SemaphoreHandle_t xSemaphore = NULL;

// Repetitive task.
void vATask( void * pvParameters )
{
    for( ;; )
    {
        // We want this task to run every 10 ticks of a timer. The semaphore
        // was created before this task was started.

        // Block waiting for the semaphore to become available.
        if( xSemaphoreTake( xSemaphore, LONG_TIME ) == pdTRUE )

```

(continues on next page)

(continued from previous page)

```

    {
        // It is time to execute.

        // ...

        // We have finished our task. Return to the top of the loop where
        // we will block on the semaphore until it is time to execute
        // again. Note when using the semaphore for synchronisation with an
        // ISR in this manner there is no need to 'give' the semaphore back.
    }
}

// Timer ISR
void vTimerISR( void * pvParameters )
{
    static uint8_t ucLocalTickCount = 0;
    static BaseType_t xHigherPriorityTaskWoken;

    // A timer tick has occurred.

    // ... Do other time functions.

    // Is it time for vATask () to run?
    xHigherPriorityTaskWoken = pdFALSE;
    ucLocalTickCount++;
    if( ucLocalTickCount >= TICKS_TO_WAIT )
    {
        // Unblock the task by releasing the semaphore.
        xSemaphoreGiveFromISR( xSemaphore, &xHigherPriorityTaskWoken );

        // Reset the count so we release the semaphore again in 10 ticks time.
        ucLocalTickCount = 0;
    }

    if( xHigherPriorityTaskWoken != pdFALSE )
    {
        // We can force a context switch here. Context switching from an
        // ISR uses port specific syntax. Check the demo task for your port
        // to find the syntax required.
    }
}

```

Parameters

- **xSemaphore** -- A handle to the semaphore being released. This is the handle returned when the semaphore was created.
- **pxHigherPriorityTaskWoken** -- xSemaphoreGiveFromISR() will set *pxHigherPriorityTaskWoken to pdTRUE if giving the semaphore caused a task to unblock, and the unblocked task has a priority higher than the currently running task. If xSemaphoreGiveFromISR() sets this value to pdTRUE then a context switch should be requested before the interrupt is exited.

Returns pdTRUE if the semaphore was successfully given, otherwise errQUEUE_FULL.

xSemaphoreTakeFromISR (xSemaphore, pxHigherPriorityTaskWoken)

Macro to take a semaphore from an ISR. The semaphore must have previously been created with a call to xSemaphoreCreateBinary() or xSemaphoreCreateCounting().

Mutex type semaphores (those created using a call to xSemaphoreCreateMutex()) must not be used with this macro.

This macro can be used from an ISR, however taking a semaphore from an ISR is not a common operation. It

is likely to only be useful when taking a counting semaphore when an interrupt is obtaining an object from a resource pool (when the semaphore count indicates the number of resources available).

Parameters

- **xSemaphore** -- A handle to the semaphore being taken. This is the handle returned when the semaphore was created.
- **pxHigherPriorityTaskWoken** -- xSemaphoreTakeFromISR() will set *pxHigherPriorityTaskWoken to pdTRUE if taking the semaphore caused a task to unblock, and the unblocked task has a priority higher than the currently running task. If xSemaphoreTakeFromISR() sets this value to pdTRUE then a context switch should be requested before the interrupt is exited.

Returns pdTRUE if the semaphore was successfully taken, otherwise pdFALSE

xSemaphoreCreateMutex ()

Creates a new mutex type semaphore instance, and returns a handle by which the new mutex can be referenced.

Internally, within the FreeRTOS implementation, mutex semaphores use a block of memory, in which the mutex structure is stored. If a mutex is created using xSemaphoreCreateMutex() then the required memory is automatically dynamically allocated inside the xSemaphoreCreateMutex() function. (see <https://www.FreeRTOS.org/a00111.html>). If a mutex is created using xSemaphoreCreateMutexStatic() then the application writer must provide the memory. xSemaphoreCreateMutexStatic() therefore allows a mutex to be created without using any dynamic memory allocation.

Mutexes created using this function can be accessed using the xSemaphoreTake() and xSemaphoreGive() macros. The xSemaphoreTakeRecursive() and xSemaphoreGiveRecursive() macros must not be used.

This type of semaphore uses a priority inheritance mechanism so a task 'taking' a semaphore MUST ALWAYS 'give' the semaphore back once the semaphore it is no longer required.

Mutex type semaphores cannot be used from within interrupt service routines.

See xSemaphoreCreateBinary() for an alternative implementation that can be used for pure synchronisation (where one task or interrupt always 'gives' the semaphore and another always 'takes' the semaphore) and from within interrupt service routines.

Example usage:

```
SemaphoreHandle_t xSemaphore;

void vATask( void * pvParameters )
{
    // Semaphore cannot be used before a call to xSemaphoreCreateMutex().
    // This is a macro so pass the variable in directly.
    xSemaphore = xSemaphoreCreateMutex();

    if( xSemaphore != NULL )
    {
        // The semaphore was created successfully.
        // The semaphore can now be used.
    }
}
```

Returns If the mutex was successfully created then a handle to the created semaphore is returned. If there was not enough heap to allocate the mutex data structures then NULL is returned.

xSemaphoreCreateMutexStatic (pxMutexBuffer)

Creates a new mutex type semaphore instance, and returns a handle by which the new mutex can be referenced.

Internally, within the FreeRTOS implementation, mutex semaphores use a block of memory, in which the mutex structure is stored. If a mutex is created using xSemaphoreCreateMutex() then the required memory is automatically dynamically allocated inside the xSemaphoreCreateMutex() function. (see <https://www.FreeRTOS.org/a00111.html>).

[//www.FreeRTOS.org/a00111.html](http://www.FreeRTOS.org/a00111.html)). If a mutex is created using `xSemaphoreCreateMutexStatic()` then the application writer must provide the memory. `xSemaphoreCreateMutexStatic()` therefore allows a mutex to be created without using any dynamic memory allocation.

Mutexes created using this function can be accessed using the `xSemaphoreTake()` and `xSemaphoreGive()` macros. The `xSemaphoreTakeRecursive()` and `xSemaphoreGiveRecursive()` macros must not be used.

This type of semaphore uses a priority inheritance mechanism so a task 'taking' a semaphore MUST ALWAYS 'give' the semaphore back once the semaphore it is no longer required.

Mutex type semaphores cannot be used from within interrupt service routines.

See `xSemaphoreCreateBinary()` for an alternative implementation that can be used for pure synchronisation (where one task or interrupt always 'gives' the semaphore and another always 'takes' the semaphore) and from within interrupt service routines.

Example usage:

```
SemaphoreHandle_t xSemaphore;
StaticSemaphore_t xMutexBuffer;

void vATask( void * pvParameters )
{
    // A mutex cannot be used before it has been created. xMutexBuffer is
    // into xSemaphoreCreateMutexStatic() so no dynamic memory allocation is
    // attempted.
    xSemaphore = xSemaphoreCreateMutexStatic( &xMutexBuffer );

    // As no dynamic memory allocation was performed, xSemaphore cannot be NULL,
    // so there is no need to check it.
}
```

Parameters

- **pxMutexBuffer** -- Must point to a variable of type `StaticSemaphore_t`, which will be used to hold the mutex's data structure, removing the need for the memory to be allocated dynamically.

Returns If the mutex was successfully created then a handle to the created mutex is returned. If `pxMutexBuffer` was `NULL` then `NULL` is returned.

xSemaphoreCreateRecursiveMutex()

Creates a new recursive mutex type semaphore instance, and returns a handle by which the new recursive mutex can be referenced.

Internally, within the FreeRTOS implementation, recursive mutexes use a block of memory, in which the mutex structure is stored. If a recursive mutex is created using `xSemaphoreCreateRecursiveMutex()` then the required memory is automatically dynamically allocated inside the `xSemaphoreCreateRecursiveMutex()` function. (see <http://www.FreeRTOS.org/a00111.html>). If a recursive mutex is created using `xSemaphoreCreateRecursiveMutexStatic()` then the application writer must provide the memory that will get used by the mutex. `xSemaphoreCreateRecursiveMutexStatic()` therefore allows a recursive mutex to be created without using any dynamic memory allocation.

Mutexes created using this macro can be accessed using the `xSemaphoreTakeRecursive()` and `xSemaphoreGiveRecursive()` macros. The `xSemaphoreTake()` and `xSemaphoreGive()` macros must not be used.

A mutex used recursively can be 'taken' repeatedly by the owner. The mutex doesn't become available again until the owner has called `xSemaphoreGiveRecursive()` for each successful 'take' request. For example, if a task successfully 'takes' the same mutex 5 times then the mutex will not be available to any other task until it has also 'given' the mutex back exactly five times.

This type of semaphore uses a priority inheritance mechanism so a task 'taking' a semaphore MUST ALWAYS 'give' the semaphore back once the semaphore it is no longer required.

Mutex type semaphores cannot be used from within interrupt service routines.

See `xSemaphoreCreateBinary()` for an alternative implementation that can be used for pure synchronisation (where one task or interrupt always 'gives' the semaphore and another always 'takes' the semaphore) and from within interrupt service routines.

Example usage:

```
SemaphoreHandle_t xSemaphore;

void vATask( void * pvParameters )
{
    // Semaphore cannot be used before a call to xSemaphoreCreateMutex().
    // This is a macro so pass the variable in directly.
    xSemaphore = xSemaphoreCreateRecursiveMutex();

    if( xSemaphore != NULL )
    {
        // The semaphore was created successfully.
        // The semaphore can now be used.
    }
}
```

Returns xSemaphore Handle to the created mutex semaphore. Should be of type `SemaphoreHandle_t`.

xSemaphoreCreateRecursiveMutexStatic (pxStaticSemaphore)

Creates a new recursive mutex type semaphore instance, and returns a handle by which the new recursive mutex can be referenced.

Internally, within the FreeRTOS implementation, recursive mutexes use a block of memory, in which the mutex structure is stored. If a recursive mutex is created using `xSemaphoreCreateRecursiveMutex()` then the required memory is automatically dynamically allocated inside the `xSemaphoreCreateRecursiveMutex()` function. (see <https://www.FreeRTOS.org/a00111.html>). If a recursive mutex is created using `xSemaphoreCreateRecursiveMutexStatic()` then the application writer must provide the memory that will get used by the mutex. `xSemaphoreCreateRecursiveMutexStatic()` therefore allows a recursive mutex to be created without using any dynamic memory allocation.

Mutexes created using this macro can be accessed using the `xSemaphoreTakeRecursive()` and `xSemaphoreGiveRecursive()` macros. The `xSemaphoreTake()` and `xSemaphoreGive()` macros must not be used.

A mutex used recursively can be 'taken' repeatedly by the owner. The mutex doesn't become available again until the owner has called `xSemaphoreGiveRecursive()` for each successful 'take' request. For example, if a task successfully 'takes' the same mutex 5 times then the mutex will not be available to any other task until it has also 'given' the mutex back exactly five times.

This type of semaphore uses a priority inheritance mechanism so a task 'taking' a semaphore MUST ALWAYS 'give' the semaphore back once the semaphore it is no longer required.

Mutex type semaphores cannot be used from within interrupt service routines.

See `xSemaphoreCreateBinary()` for an alternative implementation that can be used for pure synchronisation (where one task or interrupt always 'gives' the semaphore and another always 'takes' the semaphore) and from within interrupt service routines.

Example usage:

```
SemaphoreHandle_t xSemaphore;
StaticSemaphore_t xMutexBuffer;
```

(continues on next page)

(continued from previous page)

```

void vATask( void * pvParameters )
{
    // A recursive semaphore cannot be used before it is created. Here a
    // recursive mutex is created using xSemaphoreCreateRecursiveMutexStatic().
    // The address of xMutexBuffer is passed into the function, and will hold
    // the mutexes data structures - so no dynamic memory allocation will be
    // attempted.
    xSemaphore = xSemaphoreCreateRecursiveMutexStatic( &xMutexBuffer );

    // As no dynamic memory allocation was performed, xSemaphore cannot be NULL,
    // so there is no need to check it.
}

```

Parameters

- **pxStaticSemaphore** -- Must point to a variable of type StaticSemaphore_t, which will then be used to hold the recursive mutex's data structure, removing the need for the memory to be allocated dynamically.

Returns If the recursive mutex was successfully created then a handle to the created recursive mutex is returned. If pxStaticSemaphore was NULL then NULL is returned.

xSemaphoreCreateCounting (uxMaxCount, uxInitialCount)

Creates a new counting semaphore instance, and returns a handle by which the new counting semaphore can be referenced.

In many usage scenarios it is faster and more memory efficient to use a direct to task notification in place of a counting semaphore! <https://www.FreeRTOS.org/RTOS-task-notifications.html>

Internally, within the FreeRTOS implementation, counting semaphores use a block of memory, in which the counting semaphore structure is stored. If a counting semaphore is created using xSemaphoreCreateCounting() then the required memory is automatically dynamically allocated inside the xSemaphoreCreateCounting() function. (see <https://www.FreeRTOS.org/a00111.html>). If a counting semaphore is created using xSemaphoreCreateCountingStatic() then the application writer can instead optionally provide the memory that will get used by the counting semaphore. xSemaphoreCreateCountingStatic() therefore allows a counting semaphore to be created without using any dynamic memory allocation.

Counting semaphores are typically used for two things:

1) Counting events.

In this usage scenario an event handler will 'give' a semaphore each time an event occurs (incrementing the semaphore count value), and a handler task will 'take' a semaphore each time it processes an event (decrementing the semaphore count value). The count value is therefore the difference between the number of events that have occurred and the number that have been processed. In this case it is desirable for the initial count value to be zero.

2) Resource management.

In this usage scenario the count value indicates the number of resources available. To obtain control of a resource a task must first obtain a semaphore - decrementing the semaphore count value. When the count value reaches zero there are no free resources. When a task finishes with the resource it 'gives' the semaphore back - incrementing the semaphore count value. In this case it is desirable for the initial count value to be equal to the maximum count value, indicating that all resources are free.

Example usage:

```

SemaphoreHandle_t xSemaphore;

void vATask( void * pvParameters )
{

```

(continues on next page)

(continued from previous page)

```

SemaphoreHandle_t xSemaphore = NULL;

// Semaphore cannot be used before a call to xSemaphoreCreateCounting().
// The max value to which the semaphore can count should be 10, and the
// initial value assigned to the count should be 0.
xSemaphore = xSemaphoreCreateCounting( 10, 0 );

if( xSemaphore != NULL )
{
    // The semaphore was created successfully.
    // The semaphore can now be used.
}
}

```

Parameters

- **uxMaxCount** -- The maximum count value that can be reached. When the semaphore reaches this value it can no longer be 'given'.
- **uxInitialCount** -- The count value assigned to the semaphore when it is created.

Returns Handle to the created semaphore. Null if the semaphore could not be created.

xSemaphoreCreateCountingStatic (uxMaxCount, uxInitialCount, pxSemaphoreBuffer)

Creates a new counting semaphore instance, and returns a handle by which the new counting semaphore can be referenced.

In many usage scenarios it is faster and more memory efficient to use a direct to task notification in place of a counting semaphore! <https://www.FreeRTOS.org/RTOS-task-notifications.html>

Internally, within the FreeRTOS implementation, counting semaphores use a block of memory, in which the counting semaphore structure is stored. If a counting semaphore is created using xSemaphoreCreateCounting() then the required memory is automatically dynamically allocated inside the xSemaphoreCreateCounting() function. (see <https://www.FreeRTOS.org/a00111.html>). If a counting semaphore is created using xSemaphoreCreateCountingStatic() then the application writer must provide the memory. xSemaphoreCreateCountingStatic() therefore allows a counting semaphore to be created without using any dynamic memory allocation.

Counting semaphores are typically used for two things:

1) Counting events.

In this usage scenario an event handler will 'give' a semaphore each time an event occurs (incrementing the semaphore count value), and a handler task will 'take' a semaphore each time it processes an event (decrementing the semaphore count value). The count value is therefore the difference between the number of events that have occurred and the number that have been processed. In this case it is desirable for the initial count value to be zero.

2) Resource management.

In this usage scenario the count value indicates the number of resources available. To obtain control of a resource a task must first obtain a semaphore - decrementing the semaphore count value. When the count value reaches zero there are no free resources. When a task finishes with the resource it 'gives' the semaphore back - incrementing the semaphore count value. In this case it is desirable for the initial count value to be equal to the maximum count value, indicating that all resources are free.

Example usage:

```

SemaphoreHandle_t xSemaphore;
StaticSemaphore_t xSemaphoreBuffer;

void vATask( void * pvParameters )

```

(continues on next page)

(continued from previous page)

```

{
SemaphoreHandle_t xSemaphore = NULL;

// Counting semaphore cannot be used before they have been created. Create
// a counting semaphore using xSemaphoreCreateCountingStatic(). The max
// value to which the semaphore can count is 10, and the initial value
// assigned to the count will be 0. The address of xSemaphoreBuffer is
// passed in and will be used to hold the semaphore structure, so no dynamic
// memory allocation will be used.
xSemaphore = xSemaphoreCreateCounting( 10, 0, &xSemaphoreBuffer );

// No memory allocation was attempted so xSemaphore cannot be NULL, so there
// is no need to check its value.
}

```

Parameters

- **uxMaxCount** -- The maximum count value that can be reached. When the semaphore reaches this value it can no longer be 'given'.
- **uxInitialCount** -- The count value assigned to the semaphore when it is created.
- **pxSemaphoreBuffer** -- Must point to a variable of type StaticSemaphore_t, which will then be used to hold the semaphore's data structure, removing the need for the memory to be allocated dynamically.

Returns If the counting semaphore was successfully created then a handle to the created counting semaphore is returned. If pxSemaphoreBuffer was NULL then NULL is returned.

vSemaphoreDelete (xSemaphore)

Delete a semaphore. This function must be used with care. For example, do not delete a mutex type semaphore if the mutex is held by a task.

Parameters

- **xSemaphore** -- A handle to the semaphore to be deleted.

xSemaphoreGetMutexHolder (xSemaphore)

If xMutex is indeed a mutex type semaphore, return the current mutex holder. If xMutex is not a mutex type semaphore, or the mutex is available (not held by a task), return NULL.

Note: This is a good way of determining if the calling task is the mutex holder, but not a good way of determining the identity of the mutex holder as the holder may change between the function exiting and the returned value being tested.

xSemaphoreGetMutexHolderFromISR (xSemaphore)

If xMutex is indeed a mutex type semaphore, return the current mutex holder. If xMutex is not a mutex type semaphore, or the mutex is available (not held by a task), return NULL.

uxSemaphoreGetCount (xSemaphore)

If the semaphore is a counting semaphore then uxSemaphoreGetCount() returns its current count value. If the semaphore is a binary semaphore then uxSemaphoreGetCount() returns 1 if the semaphore is available, and 0 if the semaphore is not available.

uxSemaphoreGetCountFromISR (xSemaphore)

semphr.h

```

UBaseType_t uxSemaphoreGetCountFromISR( SemaphoreHandle_t xSemaphore );

```

If the semaphore is a counting semaphore then uxSemaphoreGetCountFromISR() returns its current count value. If the semaphore is a binary semaphore then uxSemaphoreGetCountFromISR() returns 1 if the semaphore is available, and 0 if the semaphore is not available.

xSemaphoreGetStaticBuffer (xSemaphore, ppxSemaphoreBuffer)

Retrieve pointer to a statically created binary semaphore, counting semaphore, or mutex semaphore's data structure buffer. This is the same buffer that is supplied at the time of creation.

Parameters

- **xSemaphore** -- The semaphore for which to retrieve the buffer.
- **ppxSemaphoreBuffer** -- Used to return a pointer to the semaphore's data structure buffer.

Returns pdTRUE if buffer was retrieved, pdFALSE otherwise.

Type Definitions

typedef *QueueHandle_t* **SemaphoreHandle_t**

Timer API**Header File**

- [components/freertos/FreeRTOS-Kernel/include/freertos/timers.h](#)
- This header file can be included with:

```
#include "freertos/timers.h"
```

Functions

TimerHandle_t **xTimerCreate** (const char *const pcTimerName, const TickType_t xTimerPeriodInTicks, const BaseType_t xAutoReload, void *const pvTimerID, *TimerCallbackFunction_t* pxCallbackFunction)

Creates a new software timer instance, and returns a handle by which the created software timer can be referenced.

Internally, within the FreeRTOS implementation, software timers use a block of memory, in which the timer data structure is stored. If a software timer is created using xTimerCreate() then the required memory is automatically dynamically allocated inside the xTimerCreate() function. (see <https://www.FreeRTOS.org/a00111.html>). If a software timer is created using xTimerCreateStatic() then the application writer must provide the memory that will get used by the software timer. xTimerCreateStatic() therefore allows a software timer to be created without using any dynamic memory allocation.

Timers are created in the dormant state. The xTimerStart(), xTimerReset(), xTimerStartFromISR(), xTimerResetFromISR(), xTimerChangePeriod() and xTimerChangePeriodFromISR() API functions can all be used to transition a timer into the active state.

Example usage:

```
* #define NUM_TIMERS 5
*
* // An array to hold handles to the created timers.
* TimerHandle_t xTimers[ NUM_TIMERS ];
*
* // An array to hold a count of the number of times each timer expires.
* int32_t lExpireCounters[ NUM_TIMERS ] = { 0 };
*
* // Define a callback function that will be used by multiple timer instances.
* // The callback function does nothing but count the number of times the
* // associated timer expires, and stop the timer once the timer has expired
* // 10 times.
* void vTimerCallback( TimerHandle_t pxTimer )
```

(continues on next page)

(continued from previous page)

```

* {
* int32_t lArrayIndex;
* const int32_t xMaxExpiryCountBeforeStopping = 10;
*
* // Optionally do something if the pxTimer parameter is NULL.
* configASSERT( pxTimer );
*
* // Which timer expired?
* lArrayIndex = ( int32_t ) pvTimerGetTimerID( pxTimer );
*
* // Increment the number of times that pxTimer has expired.
* lExpireCounters[ lArrayIndex ] += 1;
*
* // If the timer has expired 10 times then stop it from running.
* if( lExpireCounters[ lArrayIndex ] == xMaxExpiryCountBeforeStopping )
* {
*     // Do not use a block time if calling a timer API function from a
*     // timer callback function, as doing so could cause a deadlock!
*     xTimerStop( pxTimer, 0 );
* }
* }
*
* void main( void )
* {
* int32_t x;
*
* // Create then start some timers. Starting the timers before the
↪scheduler
* // has been started means the timers will start running immediately that
* // the scheduler starts.
* for( x = 0; x < NUM_TIMERS; x++ )
* {
*     xTimers[ x ] = xTimerCreate( "Timer", // Just a text
↪name, not used by the kernel.
*                                ( 100 * ( x + 1 ) ), // The timer
↪period in ticks.
*                                pdTRUE, // The timers
↪will auto-reload themselves when they expire.
*                                ( void * ) x, // Assign each
↪timer a unique id equal to its array index.
*                                vTimerCallback // Each timer
↪calls the same callback when it expires.
*                                );
*
*     if( xTimers[ x ] == NULL )
*     {
*         // The timer was not created.
*     }
*     else
*     {
*         // Start the timer. No block time is specified, and even if one
↪was
*         // it would be ignored because the scheduler has not yet been
*         // started.
*         if( xTimerStart( xTimers[ x ], 0 ) != pdPASS )
*         {
*             // The timer could not be set into the Active state.
*         }
*     }
* }
* }

```

(continues on next page)

(continued from previous page)

```

*      // ...
*      // Create tasks here.
*      // ...
*
*      // Starting the scheduler will start the timers running as they have_
→already
*      // been set into the active state.
*      vTaskStartScheduler();
*
*      // Should not reach here.
*      for( ;; );
* }
*

```

Parameters

- **pcTimerName** -- A text name that is assigned to the timer. This is done purely to assist debugging. The kernel itself only ever references a timer by its handle, and never by its name.
- **xTimerPeriodInTicks** -- The timer period. The time is defined in tick periods so the constant `portTICK_PERIOD_MS` can be used to convert a time that has been specified in milliseconds. For example, if the timer must expire after 100 ticks, then `xTimerPeriodInTicks` should be set to 100. Alternatively, if the timer must expire after 500ms, then `xPeriod` can be set to $(500 / \text{portTICK\_PERIOD\_MS})$ provided `configTICK_RATE_HZ` is less than or equal to 1000. Time timer period must be greater than 0.
- **xAutoReload** -- If `xAutoReload` is set to `pdTRUE` then the timer will expire repeatedly with a frequency set by the `xTimerPeriodInTicks` parameter. If `xAutoReload` is set to `pdFALSE` then the timer will be a one-shot timer and enter the dormant state after it expires.
- **pvTimerID** -- An identifier that is assigned to the timer being created. Typically this would be used in the timer callback function to identify which timer expired when the same callback function is assigned to more than one timer.
- **pxCallbackFunction** -- The function to call when the timer expires. Callback functions must have the prototype defined by `TimerCallbackFunction_t`, which is "void vCallbackFunction(TimerHandle_t xTimer);".

Returns If the timer is successfully created then a handle to the newly created timer is returned. If the timer cannot be created because there is insufficient FreeRTOS heap remaining to allocate the timer structures then `NULL` is returned.

TimerHandle_t **xTimerCreateStatic** (const char *const pcTimerName, const TickType_t xTimerPeriodInTicks, const BaseType_t xAutoReload, void *const pvTimerID, *TimerCallbackFunction_t* pxCallbackFunction, StaticTimer_t *pxTimerBuffer)

Creates a new software timer instance, and returns a handle by which the created software timer can be referenced.

Internally, within the FreeRTOS implementation, software timers use a block of memory, in which the timer data structure is stored. If a software timer is created using `xTimerCreate()` then the required memory is automatically dynamically allocated inside the `xTimerCreate()` function. (see <https://www.FreeRTOS.org/a00111.html>). If a software timer is created using `xTimerCreateStatic()` then the application writer must provide the memory that will get used by the software timer. `xTimerCreateStatic()` therefore allows a software timer to be created without using any dynamic memory allocation.

Timers are created in the dormant state. The `xTimerStart()`, `xTimerReset()`, `xTimerStartFromISR()`, `xTimerResetFromISR()`, `xTimerChangePeriod()` and `xTimerChangePeriodFromISR()` API functions can all be used to transition a timer into the active state.

Example usage:

```

*
* // The buffer used to hold the software timer's data structure.
* static StaticTimer_t xTimerBuffer;
*
* // A variable that will be incremented by the software timer's callback
* // function.
* UBaseType_t uxVariableToIncrement = 0;
*
* // A software timer callback function that increments a variable passed to
* // it when the software timer was created. After the 5th increment the
* // callback function stops the software timer.
* static void prvTimerCallback( TimerHandle_t xExpiredTimer )
* {
*     UBaseType_t *puxVariableToIncrement;
*     BaseType_t xReturned;
*
*     // Obtain the address of the variable to increment from the timer ID.
*     puxVariableToIncrement = ( UBaseType_t * ) pvTimerGetTimerID(
*     ↪xExpiredTimer );
*
*     // Increment the variable to show the timer callback has executed.
*     ( *puxVariableToIncrement )++;
*
*     // If this callback has executed the required number of times, stop the
*     // timer.
*     if( *puxVariableToIncrement == 5 )
*     {
*         // This is called from a timer callback so must not block.
*         xTimerStop( xExpiredTimer, staticDONT_BLOCK );
*     }
* }
*
* void main( void )
* {
*     // Create the software time. xTimerCreateStatic() has an extra parameter
*     // than the normal xTimerCreate() API function. The parameter is a
*     ↪pointer
*     // to the StaticTimer_t structure that will hold the software timer
*     // structure. If the parameter is passed as NULL then the structure
*     ↪will be
*     // allocated dynamically, just as if xTimerCreate() had been called.
*     xTimer = xTimerCreateStatic( "T1", // Text name for the task.
*     ↪ Helps debugging only. Not used by FreeRTOS.
*     xTimerPeriod, // The period of the
*     ↪timer in ticks.
*     pdTRUE, // This is an auto-reload
*     ↪timer.
*     ( void * ) &uxVariableToIncrement, // A
*     ↪variable incremented by the software timer's callback function
*     prvTimerCallback, // The function to
*     ↪execute when the timer expires.
*     &xTimerBuffer ); // The buffer that will
*     ↪hold the software timer structure.
*
*     // The scheduler has not started yet so a block time is not used.
*     xReturned = xTimerStart( xTimer, 0 );
*
*     // ...
*     // Create tasks here.
*     // ...
*
*

```

(continues on next page)

(continued from previous page)

```

* // Starting the scheduler will start the timers running as they have
↪ already
* // been set into the active state.
* vTaskStartScheduler();
*
* // Should not reach here.
* for( ;; );
* }
*

```

Parameters

- **pcTimerName** -- A text name that is assigned to the timer. This is done purely to assist debugging. The kernel itself only ever references a timer by its handle, and never by its name.
- **xTimerPeriodInTicks** -- The timer period. The time is defined in tick periods so the constant portTICK_PERIOD_MS can be used to convert a time that has been specified in milliseconds. For example, if the timer must expire after 100 ticks, then xTimerPeriodInTicks should be set to 100. Alternatively, if the timer must expire after 500ms, then xPeriod can be set to (500 / portTICK_PERIOD_MS) provided configTICK_RATE_HZ is less than or equal to 1000. The timer period must be greater than 0.
- **xAutoReload** -- If xAutoReload is set to pdTRUE then the timer will expire repeatedly with a frequency set by the xTimerPeriodInTicks parameter. If xAutoReload is set to pdFALSE then the timer will be a one-shot timer and enter the dormant state after it expires.
- **pvTimerID** -- An identifier that is assigned to the timer being created. Typically this would be used in the timer callback function to identify which timer expired when the same callback function is assigned to more than one timer.
- **pxCallbackFunction** -- The function to call when the timer expires. Callback functions must have the prototype defined by TimerCallbackFunction_t, which is "void vCallbackFunction(TimerHandle_t xTimer);".
- **pxTimerBuffer** -- Must point to a variable of type StaticTimer_t, which will be then be used to hold the software timer's data structures, removing the need for the memory to be allocated dynamically.

Returns If the timer is created then a handle to the created timer is returned. If pxTimerBuffer was NULL then NULL is returned.

void ***pvTimerGetTimerID**(const *TimerHandle_t* xTimer)

Returns the ID assigned to the timer.

IDs are assigned to timers using the pvTimerID parameter of the call to xTimerCreated() that was used to create the timer, and by calling the vTimerSetTimerID() API function.

If the same callback function is assigned to multiple timers then the timer ID can be used as time specific (timer local) storage.

Example usage:

See the xTimerCreate() API function example usage scenario.

Parameters **xTimer** -- The timer being queried.

Returns The ID assigned to the timer being queried.

void **vTimerSetTimerID**(*TimerHandle_t* xTimer, void *pvNewID)

Sets the ID assigned to the timer.

IDs are assigned to timers using the pvTimerID parameter of the call to xTimerCreated() that was used to create the timer.

If the same callback function is assigned to multiple timers then the timer ID can be used as time specific (timer local) storage.

Example usage:

See the `xTimerCreate()` API function example usage scenario.

Parameters

- **xTimer** -- The timer being updated.
- **pvNewID** -- The ID to assign to the timer.

`BaseType_t xTimerIsTimerActive (TimerHandle_t xTimer)`

Queries a timer to see if it is active or dormant.

A timer will be dormant if: 1) It has been created but not started, or 2) It is an expired one-shot timer that has not been restarted.

Timers are created in the dormant state. The `xTimerStart()`, `xTimerReset()`, `xTimerStartFromISR()`, `xTimerResetFromISR()`, `xTimerChangePeriod()` and `xTimerChangePeriodFromISR()` API functions can all be used to transition a timer into the active state.

Example usage:

```
* // This function assumes xTimer has already been created.
* void vAFunction( TimerHandle_t xTimer )
* {
*     if( xTimerIsTimerActive( xTimer ) != pdFALSE ) // or more simply and
*     equivalently "if( xTimerIsTimerActive( xTimer ) )"
*     {
*         // xTimer is active, do something.
*     }
*     else
*     {
*         // xTimer is not active, do something else.
*     }
* }
*
```

Parameters **xTimer** -- The timer being queried.

Returns `pdFALSE` will be returned if the timer is dormant. A value other than `pdFALSE` will be returned if the timer is active.

`TaskHandle_t xTimerGetTimerDaemonTaskHandle (void)`

Simply returns the handle of the timer service/daemon task. It is not valid to call `xTimerGetTimerDaemonTaskHandle()` before the scheduler has been started.

`BaseType_t xTimerPendFunctionCallFromISR (PendedFunction_t xFunctionToPend, void *pvParameter1, uint32_t ulParameter2, BaseType_t *pxHigherPriorityTaskWoken)`

Used from application interrupt service routines to defer the execution of a function to the RTOS daemon task (the timer service task, hence this function is implemented in `timers.c` and is prefixed with 'Timer').

Ideally an interrupt service routine (ISR) is kept as short as possible, but sometimes an ISR either has a lot of processing to do, or needs to perform processing that is not deterministic. In these cases `xTimerPendFunctionCallFromISR()` can be used to defer processing of a function to the RTOS daemon task.

A mechanism is provided that allows the interrupt to return directly to the task that will subsequently execute the pended callback function. This allows the callback function to execute contiguously in time with the interrupt - just as if the callback had executed in the interrupt itself.

Example usage:

```
*
* // The callback function that will execute in the context of the daemon
* task.
* // Note callback functions must all use this same prototype.
* void vProcessInterface( void *pvParameter1, uint32_t ulParameter2 )
* {
*     BaseType_t xInterfaceToService;
*
*     // The interface that requires servicing is passed in the second
*     // parameter. The first parameter is not used in this case.
*     xInterfaceToService = ( BaseType_t ) ulParameter2;
*
*     // ...Perform the processing here...
* }
*
* // An ISR that receives data packets from multiple interfaces
* void vAnISR( void )
* {
*     BaseType_t xInterfaceToService, xHigherPriorityTaskWoken;
*
*     // Query the hardware to determine which interface needs processing.
*     xInterfaceToService = prvCheckInterfaces();
*
*     // The actual processing is to be deferred to a task. Request the
*     // vProcessInterface() callback function is executed, passing in the
*     // number of the interface that needs processing. The interface to
*     // service is passed in the second parameter. The first parameter is
*     // not used in this case.
*     xHigherPriorityTaskWoken = pdFALSE;
*     xTimerPendFunctionCallFromISR( vProcessInterface, NULL, ( uint32_t )
* xInterfaceToService, &xHigherPriorityTaskWoken );
*
*     // If xHigherPriorityTaskWoken is now set to pdTRUE then a context
*     // switch should be requested. The macro used is port specific and will
*     // be either portYIELD_FROM_ISR() or portEND_SWITCHING_ISR() - refer to
*     // the documentation page for the port being used.
*     portYIELD_FROM_ISR( xHigherPriorityTaskWoken );
* }
*
```

Parameters

- **xFunctionToPend** -- The function to execute from the timer service/ daemon task. The function must conform to the PendedFunction_t prototype.
- **pvParameter1** -- The value of the callback function's first parameter. The parameter has a void * type to allow it to be used to pass any type. For example, unsigned longs can be cast to a void *, or the void * can be used to point to a structure.
- **ulParameter2** -- The value of the callback function's second parameter.
- **pxHigherPriorityTaskWoken** -- As mentioned above, calling this function will result in a message being sent to the timer daemon task. If the priority of the timer daemon task (which is set using configTIMER_TASK_PRIORITY in FreeRTOSConfig.h) is higher than the priority of the currently running task (the task the interrupt interrupted) then *pxHigherPriorityTaskWoken will be set to pdTRUE within xTimerPendFunctionCallFromISR(), indicating that a context switch should be requested before the interrupt exits. For that reason *pxHigherPriorityTaskWoken must be initialised to pdFALSE. See the example code below.

Returns pdPASS is returned if the message was successfully sent to the timer daemon task, otherwise pdFALSE is returned.

BaseType_t **xTimerPendFunctionCall** (*PendedFunction_t* xFunctionToPend, void *pvParameter1, uint32_t ulParameter2, TickType_t xTicksToWait)

Used to defer the execution of a function to the RTOS daemon task (the timer service task, hence this function is implemented in timers.c and is prefixed with 'Timer').

Parameters

- **xFunctionToPend** -- The function to execute from the timer service/ daemon task. The function must conform to the PendedFunction_t prototype.
- **pvParameter1** -- The value of the callback function's first parameter. The parameter has a void * type to allow it to be used to pass any type. For example, unsigned longs can be cast to a void *, or the void * can be used to point to a structure.
- **ulParameter2** -- The value of the callback function's second parameter.
- **xTicksToWait** -- Calling this function will result in a message being sent to the timer daemon task on a queue. xTicksToWait is the amount of time the calling task should remain in the Blocked state (so not using any processing time) for space to become available on the timer queue if the queue is found to be full.

Returns pdPASS is returned if the message was successfully sent to the timer daemon task, otherwise pdFALSE is returned.

const char ***pcTimerGetName** (*TimerHandle_t* xTimer)

Returns the name that was assigned to a timer when the timer was created.

Parameters **xTimer** -- The handle of the timer being queried.

Returns The name assigned to the timer specified by the xTimer parameter.

void **vTimerSetReloadMode** (*TimerHandle_t* xTimer, const BaseType_t xAutoReload)

Updates a timer to be either an auto-reload timer, in which case the timer automatically resets itself each time it expires, or a one-shot timer, in which case the timer will only expire once unless it is manually restarted.

Parameters

- **xTimer** -- The handle of the timer being updated.
- **xAutoReload** -- If xAutoReload is set to pdTRUE then the timer will expire repeatedly with a frequency set by the timer's period (see the xTimerPeriodInTicks parameter of the xTimerCreate() API function). If xAutoReload is set to pdFALSE then the timer will be a one-shot timer and enter the dormant state after it expires.

BaseType_t **xTimerGetReloadMode** (*TimerHandle_t* xTimer)

Queries a timer to determine if it is an auto-reload timer, in which case the timer automatically resets itself each time it expires, or a one-shot timer, in which case the timer will only expire once unless it is manually restarted.

Parameters **xTimer** -- The handle of the timer being queried.

Returns If the timer is an auto-reload timer then pdTRUE is returned, otherwise pdFALSE is returned.

UBaseType_t **uxTimerGetReloadMode** (*TimerHandle_t* xTimer)

Queries a timer to determine if it is an auto-reload timer, in which case the timer automatically resets itself each time it expires, or a one-shot timer, in which case the timer will only expire once unless it is manually restarted.

Parameters **xTimer** -- The handle of the timer being queried.

Returns If the timer is an auto-reload timer then pdTRUE is returned, otherwise pdFALSE is returned.

TickType_t **xTimerGetPeriod** (*TimerHandle_t* xTimer)

Returns the period of a timer.

Parameters **xTimer** -- The handle of the timer being queried.

Returns The period of the timer in ticks.

TickType_t **xTimerGetExpiryTime** (*TimerHandle_t* xTimer)

Returns the time in ticks at which the timer will expire. If this is less than the current tick count then the expiry time has overflowed from the current time.

Parameters **xTimer** -- The handle of the timer being queried.

Returns If the timer is running then the time in ticks at which the timer will next expire is returned.
If the timer is not running then the return value is undefined.

BaseType_t **xTimerGetStaticBuffer** (*TimerHandle_t* xTimer, StaticTimer_t **ppxTimerBuffer)

Retrieve pointer to a statically created timer's data structure buffer. This is the same buffer that is supplied at the time of creation.

Parameters

- **xTimer** -- The timer for which to retrieve the buffer.
- **ppxTimerBuffer** -- Used to return a pointer to the timers's data structure buffer.

Returns pdTRUE if the buffer was retrieved, pdFALSE otherwise.

void **vApplicationGetTimerTaskMemory** (StaticTask_t **ppxTimerTaskTCBBuffer, StackType_t **ppxTimerTaskStackBuffer, uint32_t *pulTimerTaskStackSize)

This function is used to provide a statically allocated block of memory to FreeRTOS to hold the Timer Task TCB. This function is required when configSUPPORT_STATIC_ALLOCATION is set. For more information see this URI: https://www.FreeRTOS.org/a00110.html#configSUPPORT_STATIC_ALLOCATION

Parameters

- **ppxTimerTaskTCBBuffer** -- A handle to a statically allocated TCB buffer
- **ppxTimerTaskStackBuffer** -- A handle to a statically allocated Stack buffer for the idle task
- **pulTimerTaskStackSize** -- A pointer to the number of elements that will fit in the allocated stack buffer

Macros

xTimerStart (xTimer, xTicksToWait)

Timer functionality is provided by a timer service/daemon task. Many of the public FreeRTOS timer API functions send commands to the timer service task through a queue called the timer command queue. The timer command queue is private to the kernel itself and is not directly accessible to application code. The length of the timer command queue is set by the configTIMER_QUEUE_LENGTH configuration constant.

xTimerStart() starts a timer that was previously created using the xTimerCreate() API function. If the timer had already been started and was already in the active state, then xTimerStart() has equivalent functionality to the xTimerReset() API function.

Starting a timer ensures the timer is in the active state. If the timer is not stopped, deleted, or reset in the mean time, the callback function associated with the timer will get called 'n' ticks after xTimerStart() was called, where 'n' is the timers defined period.

It is valid to call xTimerStart() before the scheduler has been started, but when this is done the timer will not actually start until the scheduler is started, and the timers expiry time will be relative to when the scheduler is started, not relative to when xTimerStart() was called.

The configUSE_TIMERS configuration constant must be set to 1 for xTimerStart() to be available.

Example usage:

See the xTimerCreate() API function example usage scenario.

Parameters

- **xTimer** -- The handle of the timer being started/restarted.
- **xTicksToWait** -- Specifies the time, in ticks, that the calling task should be held in the Blocked state to wait for the start command to be successfully sent to the timer command queue, should the queue already be full when xTimerStart() was called. xTicksToWait is ignored if xTimerStart() is called before the scheduler is started.

Returns pdFAIL will be returned if the start command could not be sent to the timer command queue even after xTicksToWait ticks had passed. pdPASS will be returned if the command was successfully sent to the timer command queue. When the command is actually processed

will depend on the priority of the timer service/daemon task relative to other tasks in the system, although the timers expiry time is relative to when `xTimerStart()` is actually called. The timer service/daemon task priority is set by the `configTIMER_TASK_PRIORITY` configuration constant.

xTimerStop (xTimer, xTicksToWait)

Timer functionality is provided by a timer service/daemon task. Many of the public FreeRTOS timer API functions send commands to the timer service task through a queue called the timer command queue. The timer command queue is private to the kernel itself and is not directly accessible to application code. The length of the timer command queue is set by the `configTIMER_QUEUE_LENGTH` configuration constant.

`xTimerStop()` stops a timer that was previously started using either of the `The xTimerStart()`, `xTimerReset()`, `xTimerStartFromISR()`, `xTimerResetFromISR()`, `xTimerChangePeriod()` or `xTimerChangePeriodFromISR()` API functions.

Stopping a timer ensures the timer is not in the active state.

The `configUSE_TIMERS` configuration constant must be set to 1 for `xTimerStop()` to be available.

Example usage:

See the `xTimerCreate()` API function example usage scenario.

Parameters

- **xTimer** -- The handle of the timer being stopped.
- **xTicksToWait** -- Specifies the time, in ticks, that the calling task should be held in the Blocked state to wait for the stop command to be successfully sent to the timer command queue, should the queue already be full when `xTimerStop()` was called. `xTicksToWait` is ignored if `xTimerStop()` is called before the scheduler is started.

Returns `pdFAIL` will be returned if the stop command could not be sent to the timer command queue even after `xTicksToWait` ticks had passed. `pdPASS` will be returned if the command was successfully sent to the timer command queue. When the command is actually processed will depend on the priority of the timer service/daemon task relative to other tasks in the system. The timer service/daemon task priority is set by the `configTIMER_TASK_PRIORITY` configuration constant.

xTimerChangePeriod (xTimer, xNewPeriod, xTicksToWait)

Timer functionality is provided by a timer service/daemon task. Many of the public FreeRTOS timer API functions send commands to the timer service task through a queue called the timer command queue. The timer command queue is private to the kernel itself and is not directly accessible to application code. The length of the timer command queue is set by the `configTIMER_QUEUE_LENGTH` configuration constant.

`xTimerChangePeriod()` changes the period of a timer that was previously created using the `xTimerCreate()` API function.

`xTimerChangePeriod()` can be called to change the period of an active or dormant state timer.

The `configUSE_TIMERS` configuration constant must be set to 1 for `xTimerChangePeriod()` to be available.

Example usage:

```
* // This function assumes xTimer has already been created. If the timer
* // referenced by xTimer is already active when it is called, then the timer
* // is deleted. If the timer referenced by xTimer is not active when it is
* // called, then the period of the timer is set to 500ms and the timer is
* // started.
* void vAFunction( TimerHandle_t xTimer )
* {
*     if( xTimerIsTimerActive( xTimer ) != pdFALSE ) // or more simply and
*     →equivalently "if( xTimerIsTimerActive( xTimer ) )"
*     {
```

(continues on next page)

(continued from previous page)

```

*      // xTimer is already active - delete it.
*      xTimerDelete( xTimer );
*  }
*  else
*  {
*      // xTimer is not active, change its period to 500ms. This will also
*      // cause the timer to start. Block for a maximum of 100 ticks if the
*      // change period command cannot immediately be sent to the timer
*      // command queue.
*      if( xTimerChangePeriod( xTimer, 500 / portTICK_PERIOD_MS, 100 ) == pdPASS )
*      {
*          // The command was successfully sent.
*      }
*      else
*      {
*          // The command could not be sent, even after waiting for 100
*          // ticks
*          // to pass. Take appropriate action here.
*      }
*  }
* }
*

```

Parameters

- **xTimer** -- The handle of the timer that is having its period changed.
- **xNewPeriod** -- The new period for xTimer. Timer periods are specified in tick periods, so the constant portTICK_PERIOD_MS can be used to convert a time that has been specified in milliseconds. For example, if the timer must expire after 100 ticks, then xNewPeriod should be set to 100. Alternatively, if the timer must expire after 500ms, then xNewPeriod can be set to (500 / portTICK_PERIOD_MS) provided configTICK_RATE_HZ is less than or equal to 1000.
- **xTicksToWait** -- Specifies the time, in ticks, that the calling task should be held in the Blocked state to wait for the change period command to be successfully sent to the timer command queue, should the queue already be full when xTimerChangePeriod() was called. xTicksToWait is ignored if xTimerChangePeriod() is called before the scheduler is started.

Returns pdFAIL will be returned if the change period command could not be sent to the timer command queue even after xTicksToWait ticks had passed. pdPASS will be returned if the command was successfully sent to the timer command queue. When the command is actually processed will depend on the priority of the timer service/daemon task relative to other tasks in the system. The timer service/daemon task priority is set by the configTIMER_TASK_PRIORITY configuration constant.

xTimerDelete (xTimer, xTicksToWait)

Timer functionality is provided by a timer service/daemon task. Many of the public FreeRTOS timer API functions send commands to the timer service task through a queue called the timer command queue. The timer command queue is private to the kernel itself and is not directly accessible to application code. The length of the timer command queue is set by the configTIMER_QUEUE_LENGTH configuration constant.

xTimerDelete() deletes a timer that was previously created using the xTimerCreate() API function.

The configUSE_TIMERS configuration constant must be set to 1 for xTimerDelete() to be available.

Example usage:

See the xTimerChangePeriod() API function example usage scenario.

Parameters

- **xTimer** -- The handle of the timer being deleted.
- **xTicksToWait** -- Specifies the time, in ticks, that the calling task should be held in the Blocked state to wait for the delete command to be successfully sent to the timer command queue, should the queue already be full when xTimerDelete() was called. xTicksToWait is ignored if xTimerDelete() is called before the scheduler is started.

Returns pdFAIL will be returned if the delete command could not be sent to the timer command queue even after xTicksToWait ticks had passed. pdPASS will be returned if the command was successfully sent to the timer command queue. When the command is actually processed will depend on the priority of the timer service/daemon task relative to other tasks in the system. The timer service/daemon task priority is set by the configTIMER_TASK_PRIORITY configuration constant.

xTimerReset (xTimer, xTicksToWait)

Timer functionality is provided by a timer service/daemon task. Many of the public FreeRTOS timer API functions send commands to the timer service task through a queue called the timer command queue. The timer command queue is private to the kernel itself and is not directly accessible to application code. The length of the timer command queue is set by the configTIMER_QUEUE_LENGTH configuration constant.

xTimerReset() re-starts a timer that was previously created using the xTimerCreate() API function. If the timer had already been started and was already in the active state, then xTimerReset() will cause the timer to re-evaluate its expiry time so that it is relative to when xTimerReset() was called. If the timer was in the dormant state then xTimerReset() has equivalent functionality to the xTimerStart() API function.

Resetting a timer ensures the timer is in the active state. If the timer is not stopped, deleted, or reset in the mean time, the callback function associated with the timer will get called 'n' ticks after xTimerReset() was called, where 'n' is the timers defined period.

It is valid to call xTimerReset() before the scheduler has been started, but when this is done the timer will not actually start until the scheduler is started, and the timers expiry time will be relative to when the scheduler is started, not relative to when xTimerReset() was called.

The configUSE_TIMERS configuration constant must be set to 1 for xTimerReset() to be available.

Example usage:

```
* // When a key is pressed, an LCD back-light is switched on. If 5 seconds
→ pass
* // without a key being pressed, then the LCD back-light is switched off. In
* // this case, the timer is a one-shot timer.
*
* TimerHandle_t xBacklightTimer = NULL;
*
* // The callback function assigned to the one-shot timer. In this case the
* // parameter is not used.
* void vBacklightTimerCallback( TimerHandle_t pxTimer )
* {
*     // The timer expired, therefore 5 seconds must have passed since a key
*     // was pressed. Switch off the LCD back-light.
*     vSetBacklightState( BACKLIGHT_OFF );
* }
*
* // The key press event handler.
* void vKeyPressEventHandler( char cKey )
* {
*     // Ensure the LCD back-light is on, then reset the timer that is
*     // responsible for turning the back-light off after 5 seconds of
*     // key inactivity. Wait 10 ticks for the command to be successfully sent
*     // if it cannot be sent immediately.
*     vSetBacklightState( BACKLIGHT_ON );
*     if( xTimerReset( xBacklightTimer, 100 ) != pdPASS )
*     {
```

(continues on next page)

(continued from previous page)

```

*          // The reset command was not executed successfully. Take appropriate
*          // action here.
*      }
*
*      // Perform the rest of the key processing here.
*  }
*
*  void main( void )
*  {
*      int32_t x;
*
*      // Create then start the one-shot timer that is responsible for turning
*      // the back-light off if no keys are pressed within a 5 second period.
*      xBacklightTimer = xTimerCreate( "BacklightTimer",          // Just a
*      ↪text name, not used by the kernel.
*
*      ↪timer period in ticks.
*
*      ↪pdFALSE,
*      ↪// The timer
*      ↪is a one-shot timer.
*
*      ↪0,
*      ↪// The id is
*      ↪not used by the callback so can take any value.
*      ↪vBacklightTimerCallback // The
*      ↪callback function that switches the LCD back-light off.
*      ↪);
*
*      if( xBacklightTimer == NULL )
*      {
*          // The timer was not created.
*      }
*      else
*      {
*          // Start the timer. No block time is specified, and even if one was
*          // it would be ignored because the scheduler has not yet been
*          // started.
*          if( xTimerStart( xBacklightTimer, 0 ) != pdPASS )
*          {
*              // The timer could not be set into the Active state.
*          }
*      }
*
*      // ...
*      // Create tasks here.
*      // ...
*
*      // Starting the scheduler will start the timer running as it has already
*      // been set into the active state.
*      vTaskStartScheduler();
*
*      // Should not reach here.
*      for( ;; );
*  }
*
*

```

Parameters

- **xTimer** -- The handle of the timer being reset/started/restarted.
- **xTicksToWait** -- Specifies the time, in ticks, that the calling task should be held in the Blocked state to wait for the reset command to be successfully sent to the timer command queue, should the queue already be full when xTimerReset() was called. xTicksToWait is ignored if xTimerReset() is called before the scheduler is started.

Returns pdFAIL will be returned if the reset command could not be sent to the timer command queue even after xTicksToWait ticks had passed. pdPASS will be returned if the command

was successfully sent to the timer command queue. When the command is actually processed will depend on the priority of the timer service/daemon task relative to other tasks in the system, although the timers expiry time is relative to when xTimerStart() is actually called. The timer service/daemon task priority is set by the configTIMER_TASK_PRIORITY configuration constant.

xTimerStartFromISR (xTimer, pxHigherPriorityTaskWoken)

A version of xTimerStart() that can be called from an interrupt service routine.

Example usage:

```
* // This scenario assumes xBacklightTimer has already been created. When a
* // key is pressed, an LCD back-light is switched on. If 5 seconds pass
* // without a key being pressed, then the LCD back-light is switched off. In
* // this case, the timer is a one-shot timer, and unlike the example given for
* // the xTimerReset() function, the key press event handler is an interrupt
* // service routine.
*
* // The callback function assigned to the one-shot timer. In this case the
* // parameter is not used.
* void vBacklightTimerCallback( TimerHandle_t pxTimer )
* {
*     // The timer expired, therefore 5 seconds must have passed since a key
*     // was pressed. Switch off the LCD back-light.
*     vSetBacklightState( BACKLIGHT_OFF );
* }
*
* // The key press interrupt service routine.
* void vKeyPressEventInterruptHandler( void )
* {
*     BaseType_t xHigherPriorityTaskWoken = pdFALSE;
*
*     // Ensure the LCD back-light is on, then restart the timer that is
*     // responsible for turning the back-light off after 5 seconds of
*     // key inactivity. This is an interrupt service routine so can only
*     // call FreeRTOS API functions that end in "FromISR".
*     vSetBacklightState( BACKLIGHT_ON );
*
*     // xTimerStartFromISR() or xTimerResetFromISR() could be called here
*     // as both cause the timer to re-calculate its expiry time.
*     // xHigherPriorityTaskWoken was initialised to pdFALSE when it was
*     // declared (in this function).
*     if ( xTimerStartFromISR( xBacklightTimer, &xHigherPriorityTaskWoken ) !=
* ↪pdPASS )
*     {
*         // The start command was not executed successfully. Take appropriate
*         // action here.
*     }
*
*     // Perform the rest of the key processing here.
*
*     // If xHigherPriorityTaskWoken equals pdTRUE, then a context switch
*     // should be performed. The syntax required to perform a context switch
*     // from inside an ISR varies from port to port, and from compiler to
*     // compiler. Inspect the demos for the port you are using to find the
*     // actual syntax required.
*     if ( xHigherPriorityTaskWoken != pdFALSE )
*     {
*         // Call the interrupt safe yield function here (actual function
*         // depends on the FreeRTOS port being used).
```

(continues on next page)

(continued from previous page)

```
*
* }
*
```

Parameters

- **xTimer** -- The handle of the timer being started/restarted.
- **pxHigherPriorityTaskWoken** -- The timer service/daemon task spends most of its time in the Blocked state, waiting for messages to arrive on the timer command queue. Calling xTimerStartFromISR() writes a message to the timer command queue, so has the potential to transition the timer service/daemon task out of the Blocked state. If calling xTimerStartFromISR() causes the timer service/daemon task to leave the Blocked state, and the timer service/ daemon task has a priority equal to or greater than the currently executing task (the task that was interrupted), then *pxHigherPriorityTaskWoken will get set to pdTRUE internally within the xTimerStartFromISR() function. If xTimerStartFromISR() sets this value to pdTRUE then a context switch should be performed before the interrupt exits.

Returns pdFAIL will be returned if the start command could not be sent to the timer command queue. pdPASS will be returned if the command was successfully sent to the timer command queue. When the command is actually processed will depend on the priority of the timer service/daemon task relative to other tasks in the system, although the timers expiry time is relative to when xTimerStartFromISR() is actually called. The timer service/daemon task priority is set by the configTIMER_TASK_PRIORITY configuration constant.

xTimerStopFromISR (xTimer, pxHigherPriorityTaskWoken)

A version of xTimerStop() that can be called from an interrupt service routine.

Example usage:

```
* // This scenario assumes xTimer has already been created and started. When
* // an interrupt occurs, the timer should be simply stopped.
*
* // The interrupt service routine that stops the timer.
* void vAnExampleInterruptServiceRoutine( void )
* {
*     BaseType_t xHigherPriorityTaskWoken = pdFALSE;
*
*     // The interrupt has occurred - simply stop the timer.
*     // xHigherPriorityTaskWoken was set to pdFALSE where it was defined
*     // (within this function). As this is an interrupt service routine, only
*     // FreeRTOS API functions that end in "FromISR" can be used.
*     if( xTimerStopFromISR( xTimer, &xHigherPriorityTaskWoken ) != pdPASS )
*     {
*         // The stop command was not executed successfully. Take appropriate
*         // action here.
*     }
*
*     // If xHigherPriorityTaskWoken equals pdTRUE, then a context switch
*     // should be performed. The syntax required to perform a context switch
*     // from inside an ISR varies from port to port, and from compiler to
*     // compiler. Inspect the demos for the port you are using to find the
*     // actual syntax required.
*     if( xHigherPriorityTaskWoken != pdFALSE )
*     {
*         // Call the interrupt safe yield function here (actual function
*         // depends on the FreeRTOS port being used).
*     }
* }
*
```

Parameters

- **xTimer** -- The handle of the timer being stopped.
- **pxHigherPriorityTaskWoken** -- The timer service/daemon task spends most of its time in the Blocked state, waiting for messages to arrive on the timer command queue. Calling xTimerStopFromISR() writes a message to the timer command queue, so has the potential to transition the timer service/daemon task out of the Blocked state. If calling xTimerStopFromISR() causes the timer service/daemon task to leave the Blocked state, and the timer service/ daemon task has a priority equal to or greater than the currently executing task (the task that was interrupted), then *pxHigherPriorityTaskWoken will get set to pdTRUE internally within the xTimerStopFromISR() function. If xTimerStopFromISR() sets this value to pdTRUE then a context switch should be performed before the interrupt exits.

Returns pdFAIL will be returned if the stop command could not be sent to the timer command queue. pdPASS will be returned if the command was successfully sent to the timer command queue. When the command is actually processed will depend on the priority of the timer service/daemon task relative to other tasks in the system. The timer service/daemon task priority is set by the configTIMER_TASK_PRIORITY configuration constant.

xTimerChangePeriodFromISR (xTimer, xNewPeriod, pxHigherPriorityTaskWoken)

A version of xTimerChangePeriod() that can be called from an interrupt service routine.

Example usage:

```
* // This scenario assumes xTimer has already been created and started. When
* // an interrupt occurs, the period of xTimer should be changed to 500ms.
*
* // The interrupt service routine that changes the period of xTimer.
* void vAnExampleInterruptServiceRoutine( void )
* {
*     BaseType_t xHigherPriorityTaskWoken = pdFALSE;
*
*     // The interrupt has occurred - change the period of xTimer to 500ms.
*     // xHigherPriorityTaskWoken was set to pdFALSE where it was defined
*     // (within this function). As this is an interrupt service routine, only
*     // FreeRTOS API functions that end in "FromISR" can be used.
*     if( xTimerChangePeriodFromISR( xTimer, &xHigherPriorityTaskWoken ) !=
    ↪pdPASS )
*     {
*         // The command to change the timers period was not executed
*         // successfully. Take appropriate action here.
*     }
*
*     // If xHigherPriorityTaskWoken equals pdTRUE, then a context switch
*     // should be performed. The syntax required to perform a context switch
*     // from inside an ISR varies from port to port, and from compiler to
*     // compiler. Inspect the demos for the port you are using to find the
*     // actual syntax required.
*     if( xHigherPriorityTaskWoken != pdFALSE )
*     {
*         // Call the interrupt safe yield function here (actual function
*         // depends on the FreeRTOS port being used).
*     }
* }
*
```

Parameters

- **xTimer** -- The handle of the timer that is having its period changed.
- **xNewPeriod** -- The new period for xTimer. Timer periods are specified in tick periods, so the constant portTICK_PERIOD_MS can be used to convert a time that has been spec-

ified in milliseconds. For example, if the timer must expire after 100 ticks, then `xNewPeriod` should be set to 100. Alternatively, if the timer must expire after 500ms, then `xNewPeriod` can be set to $(500 / \text{portTICK\_PERIOD\_MS})$ provided `configTICK_RATE_HZ` is less than or equal to 1000.

- **pxHigherPriorityTaskWoken** -- The timer service/daemon task spends most of its time in the Blocked state, waiting for messages to arrive on the timer command queue. Calling `xTimerChangePeriodFromISR()` writes a message to the timer command queue, so has the potential to transition the timer service/ daemon task out of the Blocked state. If calling `xTimerChangePeriodFromISR()` causes the timer service/daemon task to leave the Blocked state, and the timer service/daemon task has a priority equal to or greater than the currently executing task (the task that was interrupted), then `*pxHigherPriorityTaskWoken` will get set to `pdTRUE` internally within the `xTimerChangePeriodFromISR()` function. If `xTimerChangePeriodFromISR()` sets this value to `pdTRUE` then a context switch should be performed before the interrupt exits.

Returns `pdFAIL` will be returned if the command to change the timers period could not be sent to the timer command queue. `pdPASS` will be returned if the command was successfully sent to the timer command queue. When the command is actually processed will depend on the priority of the timer service/daemon task relative to other tasks in the system. The timer service/daemon task priority is set by the `configTIMER_TASK_PRIORITY` configuration constant.

xTimerResetFromISR (xTimer, pxHigherPriorityTaskWoken)

A version of `xTimerReset()` that can be called from an interrupt service routine.

Example usage:

```
* // This scenario assumes xBacklightTimer has already been created. When a
* // key is pressed, an LCD back-light is switched on. If 5 seconds pass
* // without a key being pressed, then the LCD back-light is switched off. In
* // this case, the timer is a one-shot timer, and unlike the example given for
* // the xTimerReset() function, the key press event handler is an interrupt
* // service routine.
*
* // The callback function assigned to the one-shot timer. In this case the
* // parameter is not used.
* void vBacklightTimerCallback( TimerHandle_t pxTimer )
* {
*     // The timer expired, therefore 5 seconds must have passed since a key
*     // was pressed. Switch off the LCD back-light.
*     vSetBacklightState( BACKLIGHT_OFF );
* }
*
* // The key press interrupt service routine.
* void vKeyPressEventInterruptHandler( void )
* {
*     BaseType_t xHigherPriorityTaskWoken = pdFALSE;
*
*     // Ensure the LCD back-light is on, then reset the timer that is
*     // responsible for turning the back-light off after 5 seconds of
*     // key inactivity. This is an interrupt service routine so can only
*     // call FreeRTOS API functions that end in "FromISR".
*     vSetBacklightState( BACKLIGHT_ON );
*
*     // xTimerStartFromISR() or xTimerResetFromISR() could be called here
*     // as both cause the timer to re-calculate its expiry time.
*     // xHigherPriorityTaskWoken was initialised to pdFALSE when it was
*     // declared (in this function).
*     if( xTimerResetFromISR( xBacklightTimer, &xHigherPriorityTaskWoken ) !=  
→pdPASS )
```

(continues on next page)

(continued from previous page)

```

*      {
*          // The reset command was not executed successfully. Take appropriate
*          // action here.
*      }
*
*      // Perform the rest of the key processing here.
*
*      // If xHigherPriorityTaskWoken equals pdTRUE, then a context switch
*      // should be performed. The syntax required to perform a context switch
*      // from inside an ISR varies from port to port, and from compiler to
*      // compiler. Inspect the demos for the port you are using to find the
*      // actual syntax required.
*      if( xHigherPriorityTaskWoken != pdFALSE )
*      {
*          // Call the interrupt safe yield function here (actual function
*          // depends on the FreeRTOS port being used).
*      }
*  }
*

```

Parameters

- **xTimer** -- The handle of the timer that is to be started, reset, or restarted.
- **pxHigherPriorityTaskWoken** -- The timer service/daemon task spends most of its time in the Blocked state, waiting for messages to arrive on the timer command queue. Calling xTimerResetFromISR() writes a message to the timer command queue, so has the potential to transition the timer service/daemon task out of the Blocked state. If calling xTimerResetFromISR() causes the timer service/daemon task to leave the Blocked state, and the timer service/ daemon task has a priority equal to or greater than the currently executing task (the task that was interrupted), then *pxHigherPriorityTaskWoken will get set to pdTRUE internally within the xTimerResetFromISR() function. If xTimerResetFromISR() sets this value to pdTRUE then a context switch should be performed before the interrupt exits.

Returns pdFAIL will be returned if the reset command could not be sent to the timer command queue. pdPASS will be returned if the command was successfully sent to the timer command queue. When the command is actually processed will depend on the priority of the timer service/daemon task relative to other tasks in the system, although the timers expiry time is relative to when xTimerResetFromISR() is actually called. The timer service/daemon task priority is set by the configTIMER_TASK_PRIORITY configuration constant.

Type Definitions

```
typedef struct tmrTimerControl *TimerHandle_t
```

```
typedef void (*TimerCallbackFunction_t)(TimerHandle_t xTimer)
```

Defines the prototype to which timer callback functions must conform.

```
typedef void (*PendedFunction_t)(void*, uint32_t)
```

Defines the prototype to which functions used with the xTimerPendFunctionCallFromISR() function must conform.

Event Group API**Header File**

- [components/freertos/FreeRTOS-Kernel/include/freertos/event_groups.h](#)
- This header file can be included with:

```
#include "freertos/event_groups.h"
```

Functions

EventGroupHandle_t xEventGroupCreate (void)

Create a new event group.

Internally, within the FreeRTOS implementation, event groups use a [small] block of memory, in which the event group's structure is stored. If an event groups is created using xEventGroupCreate() then the required memory is automatically dynamically allocated inside the xEventGroupCreate() function. (see <https://www.FreeRTOS.org/a00111.html>). If an event group is created using xEventGroupCreateStatic() then the application writer must instead provide the memory that will get used by the event group. xEventGroupCreateStatic() therefore allows an event group to be created without using any dynamic memory allocation.

Although event groups are not related to ticks, for internal implementation reasons the number of bits available for use in an event group is dependent on the configUSE_16_BIT_TICKS setting in FreeRTOSConfig.h. If configUSE_16_BIT_TICKS is 1 then each event group contains 8 usable bits (bit 0 to bit 7). If configUSE_16_BIT_TICKS is set to 0 then each event group has 24 usable bits (bit 0 to bit 23). The EventBits_t type is used to store event bits within an event group.

Example usage:

```
// Declare a variable to hold the created event group.
EventGroupHandle_t xCreatedEventGroup;

// Attempt to create the event group.
xCreatedEventGroup = xEventGroupCreate();

// Was the event group created successfully?
if( xCreatedEventGroup == NULL )
{
    // The event group was not created because there was insufficient
    // FreeRTOS heap available.
}
else
{
    // The event group was created.
}
```

Returns If the event group was created then a handle to the event group is returned. If there was insufficient FreeRTOS heap available to create the event group then NULL is returned. See <https://www.FreeRTOS.org/a00111.html>

EventGroupHandle_t xEventGroupCreateStatic (StaticEventGroup_t *pxEventGroupBuffer)

Create a new event group.

Internally, within the FreeRTOS implementation, event groups use a [small] block of memory, in which the event group's structure is stored. If an event groups is created using xEventGroupCreate() then the required memory is automatically dynamically allocated inside the xEventGroupCreate() function. (see <https://www.FreeRTOS.org/a00111.html>). If an event group is created using xEventGroupCreateStatic() then the application writer must instead provide the memory that will get used by the event group. xEventGroupCreateStatic() therefore allows an event group to be created without using any dynamic memory allocation.

Although event groups are not related to ticks, for internal implementation reasons the number of bits available for use in an event group is dependent on the configUSE_16_BIT_TICKS setting in FreeRTOSConfig.h. If configUSE_16_BIT_TICKS is 1 then each event group contains 8 usable bits (bit 0 to bit 7). If configUSE_16_BIT_TICKS is set to 0 then each event group has 24 usable bits (bit 0 to bit 23). The EventBits_t type is used to store event bits within an event group.

Example usage:

```
// StaticEventGroup_t is a publicly accessible structure that has the same
// size and alignment requirements as the real event group structure. It is
// provided as a mechanism for applications to know the size of the event
// group (which is dependent on the architecture and configuration file
// settings) without breaking the strict data hiding policy by exposing the
// real event group internals. This StaticEventGroup_t variable is passed
// into the xSemaphoreCreateEventGroupStatic() function and is used to store
// the event group's data structures
StaticEventGroup_t xEventGroupBuffer;

// Create the event group without dynamically allocating any memory.
xEventGroup = xEventGroupCreateStatic( &xEventGroupBuffer );
```

Parameters **pxEventGroupBuffer** -- pxEventGroupBuffer must point to a variable of type StaticEventGroup_t, which will be then be used to hold the event group's data structures, removing the need for the memory to be allocated dynamically.

Returns If the event group was created then a handle to the event group is returned. If pxEventGroupBuffer was NULL then NULL is returned.

EventBits_t xEventGroupWaitBits (*EventGroupHandle_t* xEventGroup, const *EventBits_t* uxBitsToWaitFor, const BaseType_t xClearOnExit, const BaseType_t xWaitForAllBits, TickType_t xTicksToWait)

[Potentially] block to wait for one or more bits to be set within a previously created event group.

This function cannot be called from an interrupt.

Example usage:

```
#define BIT_0 ( 1 << 0 )
#define BIT_4 ( 1 << 4 )

void aFunction( EventGroupHandle_t xEventGroup )
{
    EventBits_t uxBits;
    const TickType_t xTicksToWait = 100 / portTICK_PERIOD_MS;

    // Wait a maximum of 100ms for either bit 0 or bit 4 to be set within
    // the event group. Clear the bits before exiting.
    uxBits = xEventGroupWaitBits(
        xEventGroup,    // The event group being tested.
        BIT_0 | BIT_4,  // The bits within the event group to wait_
        ↪for.            pdTRUE,          // BIT_0 and BIT_4 should be cleared before_
        ↪returning.     pdFALSE,        // Don't wait for both bits, either bit will_
        ↪do.            xTicksToWait ); // Wait a maximum of 100ms for either bit to_
        ↪be set.

    if( ( uxBits & ( BIT_0 | BIT_4 ) ) == ( BIT_0 | BIT_4 ) )
    {
        // xEventGroupWaitBits() returned because both bits were set.
    }
    else if( ( uxBits & BIT_0 ) != 0 )
    {
        // xEventGroupWaitBits() returned because just BIT_0 was set.
    }
    else if( ( uxBits & BIT_4 ) != 0 )
```

(continues on next page)

(continued from previous page)

```

{
    // xEventGroupWaitBits() returned because just BIT_4 was set.
}
else
{
    // xEventGroupWaitBits() returned because xTicksToWait ticks passed
    // without either BIT_0 or BIT_4 becoming set.
}
}

```

Parameters

- **xEventGroup** -- The event group in which the bits are being tested. The event group must have previously been created using a call to `xEventGroupCreate()`.
- **uxBitsToWaitFor** -- A bitwise value that indicates the bit or bits to test inside the event group. For example, to wait for bit 0 and/or bit 2 set `uxBitsToWaitFor` to `0x05`. To wait for bits 0 and/or bit 1 and/or bit 2 set `uxBitsToWaitFor` to `0x07`. Etc.
- **xClearOnExit** -- If `xClearOnExit` is set to `pdTRUE` then any bits within `uxBitsToWaitFor` that are set within the event group will be cleared before `xEventGroupWaitBits()` returns if the wait condition was met (if the function returns for a reason other than a timeout). If `xClearOnExit` is set to `pdFALSE` then the bits set in the event group are not altered when the call to `xEventGroupWaitBits()` returns.
- **xWaitForAllBits** -- If `xWaitForAllBits` is set to `pdTRUE` then `xEventGroupWaitBits()` will return when either all the bits in `uxBitsToWaitFor` are set or the specified block time expires. If `xWaitForAllBits` is set to `pdFALSE` then `xEventGroupWaitBits()` will return when any one of the bits set in `uxBitsToWaitFor` is set or the specified block time expires. The block time is specified by the `xTicksToWait` parameter.
- **xTicksToWait** -- The maximum amount of time (specified in 'ticks') to wait for one/all (depending on the `xWaitForAllBits` value) of the bits specified by `uxBitsToWaitFor` to become set. A value of `portMAX_DELAY` can be used to block indefinitely (provided `INCLUDE_vTaskSuspend` is set to 1 in `FreeRTOSConfig.h`).

Returns The value of the event group at the time either the bits being waited for became set, or the block time expired. Test the return value to know which bits were set. If `xEventGroupWaitBits()` returned because its timeout expired then not all the bits being waited for will be set. If `xEventGroupWaitBits()` returned because the bits it was waiting for were set then the returned value is the event group value before any bits were automatically cleared in the case that `xClearOnExit` parameter was set to `pdTRUE`.

EventBits_t xEventGroupClearBits (*EventGroupHandle_t* xEventGroup, const *EventBits_t* uxBitsToClear)

Clear bits within an event group. This function cannot be called from an interrupt.

Example usage:

```

#define BIT_0 ( 1 << 0 )
#define BIT_4 ( 1 << 4 )

void aFunction( EventGroupHandle_t xEventGroup )
{
    EventBits_t uxBits;

    // Clear bit 0 and bit 4 in xEventGroup.
    uxBits = xEventGroupClearBits(
        xEventGroup,    // The event group being updated.
        BIT_0 | BIT_4 ); // The bits being cleared.

    if( ( uxBits & ( BIT_0 | BIT_4 ) ) == ( BIT_0 | BIT_4 ) )
    {

```

(continues on next page)

(continued from previous page)

```

        // Both bit 0 and bit 4 were set before xEventGroupClearBits() was
        // called. Both will now be clear (not set).
    }
    else if( ( uxBits & BIT_0 ) != 0 )
    {
        // Bit 0 was set before xEventGroupClearBits() was called. It will
        // now be clear.
    }
    else if( ( uxBits & BIT_4 ) != 0 )
    {
        // Bit 4 was set before xEventGroupClearBits() was called. It will
        // now be clear.
    }
    else
    {
        // Neither bit 0 nor bit 4 were set in the first place.
    }
}

```

Parameters

- **xEventGroup** -- The event group in which the bits are to be cleared.
- **uxBitsToClear** -- A bitwise value that indicates the bit or bits to clear in the event group. For example, to clear bit 3 only, set `uxBitsToClear` to `0x08`. To clear bit 3 and bit 0 set `uxBitsToClear` to `0x09`.

Returns The value of the event group before the specified bits were cleared.

EventBits_t xEventGroupSetBits (*EventGroupHandle_t* xEventGroup, const *EventBits_t* uxBitsToSet)

Set bits within an event group. This function cannot be called from an interrupt. `xEventGroupSetBitsFromISR()` is a version that can be called from an interrupt.

Setting bits in an event group will automatically unblock tasks that are blocked waiting for the bits.

Example usage:

```

#define BIT_0 ( 1 << 0 )
#define BIT_4 ( 1 << 4 )

void aFunction( EventGroupHandle_t xEventGroup )
{
    EventBits_t uxBits;

    // Set bit 0 and bit 4 in xEventGroup.
    uxBits = xEventGroupSetBits(
        xEventGroup,    // The event group being updated.
        BIT_0 | BIT_4 ); // The bits being set.

    if( ( uxBits & ( BIT_0 | BIT_4 ) ) == ( BIT_0 | BIT_4 ) )
    {
        // Both bit 0 and bit 4 remained set when the function returned.
    }
    else if( ( uxBits & BIT_0 ) != 0 )
    {
        // Bit 0 remained set when the function returned, but bit 4 was
        // cleared. It might be that bit 4 was cleared automatically as a
        // task that was waiting for bit 4 was removed from the Blocked
        // state.
    }
    else if( ( uxBits & BIT_4 ) != 0 )
    {

```

(continues on next page)

(continued from previous page)

```

        // Bit 4 remained set when the function returned, but bit 0 was
        // cleared. It might be that bit 0 was cleared automatically as a
        // task that was waiting for bit 0 was removed from the Blocked
        // state.
    }
    else
    {
        // Neither bit 0 nor bit 4 remained set. It might be that a task
        // was waiting for both of the bits to be set, and the bits were
        // cleared as the task left the Blocked state.
    }
}

```

Parameters

- **xEventGroup** -- The event group in which the bits are to be set.
- **uxBitsToSet** -- A bitwise value that indicates the bit or bits to set. For example, to set bit 3 only, set `uxBitsToSet` to `0x08`. To set bit 3 and bit 0 set `uxBitsToSet` to `0x09`.

Returns The value of the event group at the time the call to `xEventGroupSetBits()` returns. There are two reasons why the returned value might have the bits specified by the `uxBitsToSet` parameter cleared. First, if setting a bit results in a task that was waiting for the bit leaving the blocked state then it is possible the bit will be cleared automatically (see the `xClearBitOnExit` parameter of `xEventGroupWaitBits()`). Second, any unblocked (or otherwise Ready state) task that has a priority above that of the task that called `xEventGroupSetBits()` will execute and may change the event group value before the call to `xEventGroupSetBits()` returns.

EventBits_t xEventGroupSync (*EventGroupHandle_t* xEventGroup, const *EventBits_t* uxBitsToSet, const *EventBits_t* uxBitsToWaitFor, *TickType_t* xTicksToWait)

Atomically set bits within an event group, then wait for a combination of bits to be set within the same event group. This functionality is typically used to synchronise multiple tasks, where each task has to wait for the other tasks to reach a synchronisation point before proceeding.

This function cannot be used from an interrupt.

The function will return before its block time expires if the bits specified by the `uxBitsToWait` parameter are set, or become set within that time. In this case all the bits specified by `uxBitsToWait` will be automatically cleared before the function returns.

Example usage:

```

// Bits used by the three tasks.
#define TASK_0_BIT    ( 1 << 0 )
#define TASK_1_BIT    ( 1 << 1 )
#define TASK_2_BIT    ( 1 << 2 )

#define ALL_SYNC_BITS ( TASK_0_BIT | TASK_1_BIT | TASK_2_BIT )

// Use an event group to synchronise three tasks. It is assumed this event
// group has already been created elsewhere.
EventGroupHandle_t xEventBits;

void vTask0( void *pvParameters )
{
    EventBits_t uxReturn;
    TickType_t xTicksToWait = 100 / portTICK_PERIOD_MS;

    for( ;; )
    {
        // Perform task functionality here.
    }
}

```

(continues on next page)

(continued from previous page)

```

// Set bit 0 in the event flag to note this task has reached the
// sync point. The other two tasks will set the other two bits defined
// by ALL_SYNC_BITS. All three tasks have reached the synchronisation
// point when all the ALL_SYNC_BITS are set. Wait a maximum of 100ms
// for this to happen.
uxReturn = xEventGroupSync( xEventBits, TASK_0_BIT, ALL_SYNC_BITS,
↪xTicksToWait );

if( ( uxReturn & ALL_SYNC_BITS ) == ALL_SYNC_BITS )
{
    // All three tasks reached the synchronisation point before the call
    // to xEventGroupSync() timed out.
}
}
}

void vTask1( void *pvParameters )
{
    for( ;; )
    {
        // Perform task functionality here.

        // Set bit 1 in the event flag to note this task has reached the
        // synchronisation point. The other two tasks will set the other two
        // bits defined by ALL_SYNC_BITS. All three tasks have reached the
        // synchronisation point when all the ALL_SYNC_BITS are set. Wait
        // indefinitely for this to happen.
        xEventGroupSync( xEventBits, TASK_1_BIT, ALL_SYNC_BITS, portMAX_DELAY );

        // xEventGroupSync() was called with an indefinite block time, so
        // this task will only reach here if the synchronisation was made by all
        // three tasks, so there is no need to test the return value.
    }
}

void vTask2( void *pvParameters )
{
    for( ;; )
    {
        // Perform task functionality here.

        // Set bit 2 in the event flag to note this task has reached the
        // synchronisation point. The other two tasks will set the other two
        // bits defined by ALL_SYNC_BITS. All three tasks have reached the
        // synchronisation point when all the ALL_SYNC_BITS are set. Wait
        // indefinitely for this to happen.
        xEventGroupSync( xEventBits, TASK_2_BIT, ALL_SYNC_BITS, portMAX_DELAY );

        // xEventGroupSync() was called with an indefinite block time, so
        // this task will only reach here if the synchronisation was made by all
        // three tasks, so there is no need to test the return value.
    }
}
}

```

Parameters

- **xEventGroup** -- The event group in which the bits are being tested. The event group must have previously been created using a call to `xEventGroupCreate()`.
- **uxBitsToSet** -- The bits to set in the event group before determining if, and possibly waiting for, all the bits specified by the `uxBitsToWait` parameter are set.
- **uxBitsToWaitFor** -- A bitwise value that indicates the bit or bits to test inside the

event group. For example, to wait for bit 0 and bit 2 set `uxBitsToWaitFor` to 0x05. To wait for bits 0 and bit 1 and bit 2 set `uxBitsToWaitFor` to 0x07. Etc.

- **xTicksToWait** -- The maximum amount of time (specified in 'ticks') to wait for all of the bits specified by `uxBitsToWaitFor` to become set.

Returns The value of the event group at the time either the bits being waited for became set, or the block time expired. Test the return value to know which bits were set. If `xEventGroupSync()` returned because its timeout expired then not all the bits being waited for will be set. If `xEventGroupSync()` returned because all the bits it was waiting for were set then the returned value is the event group value before any bits were automatically cleared.

EventBits_t **xEventGroupGetBitsFromISR** (*EventGroupHandle_t* xEventGroup)

A version of `xEventGroupGetBits()` that can be called from an ISR.

Parameters **xEventGroup** -- The event group being queried.

Returns The event group bits at the time `xEventGroupGetBitsFromISR()` was called.

void **vEventGroupDelete** (*EventGroupHandle_t* xEventGroup)

Delete an event group that was previously created by a call to `xEventGroupCreate()`. Tasks that are blocked on the event group will be unblocked and obtain 0 as the event group's value.

Parameters **xEventGroup** -- The event group being deleted.

BaseType_t **xEventGroupGetStaticBuffer** (*EventGroupHandle_t* xEventGroup, StaticEventGroup_t **ppxEventGroupBuffer)

Retrieve a pointer to a statically created event groups's data structure buffer. It is the same buffer that is supplied at the time of creation.

Parameters

- **xEventGroup** -- The event group for which to retrieve the buffer.
- **ppxEventGroupBuffer** -- Used to return a pointer to the event groups's data structure buffer.

Returns `pdTRUE` if the buffer was retrieved, `pdFALSE` otherwise.

Macros

xEventGroupClearBitsFromISR (xEventGroup, uxBitsToClear)

A version of `xEventGroupClearBits()` that can be called from an interrupt.

Setting bits in an event group is not a deterministic operation because there are an unknown number of tasks that may be waiting for the bit or bits being set. FreeRTOS does not allow nondeterministic operations to be performed while interrupts are disabled, so protects event groups that are accessed from tasks by suspending the scheduler rather than disabling interrupts. As a result event groups cannot be accessed directly from an interrupt service routine. Therefore `xEventGroupClearBitsFromISR()` sends a message to the timer task to have the clear operation performed in the context of the timer task.

Example usage:

```
#define BIT_0 ( 1 << 0 )
#define BIT_4 ( 1 << 4 )

// An event group which it is assumed has already been created by a call to
// xEventGroupCreate().
EventGroupHandle_t xEventGroup;

void anInterruptHandler( void )
{
    // Clear bit 0 and bit 4 in xEventGroup.
    xResult = xEventGroupClearBitsFromISR(
        xEventGroup,          // The event group being updated.
        BIT_0 | BIT_4 );     // The bits being set.
```

(continues on next page)

(continued from previous page)

```

if( xResult == pdPASS )
{
    // The message was posted successfully.
    portYIELD_FROM_ISR(pdTRUE);
}

```

Note: If this function returns `pdPASS` then the timer task is ready to run and a `portYIELD_FROM_ISR(pdTRUE)` should be executed to perform the needed clear on the event group. This behavior is different from `xEventGroupSetBitsFromISR` because the parameter `xHigherPriorityTaskWoken` is not present.

Parameters

- **xEventGroup** -- The event group in which the bits are to be cleared.
- **uxBitsToClear** -- A bitwise value that indicates the bit or bits to clear. For example, to clear bit 3 only, set `uxBitsToClear` to `0x08`. To clear bit 3 and bit 0 set `uxBitsToClear` to `0x09`.

Returns If the request to execute the function was posted successfully then `pdPASS` is returned, otherwise `pdFALSE` is returned. `pdFALSE` will be returned if the timer service queue was full.

xEventGroupSetBitsFromISR (xEventGroup, uxBitsToSet, pxHigherPriorityTaskWoken)

A version of `xEventGroupSetBits()` that can be called from an interrupt.

Setting bits in an event group is not a deterministic operation because there are an unknown number of tasks that may be waiting for the bit or bits being set. FreeRTOS does not allow nondeterministic operations to be performed in interrupts or from critical sections. Therefore `xEventGroupSetBitsFromISR()` sends a message to the timer task to have the set operation performed in the context of the timer task - where a scheduler lock is used in place of a critical section.

Example usage:

```

#define BIT_0 ( 1 << 0 )
#define BIT_4 ( 1 << 4 )

// An event group which it is assumed has already been created by a call to
// xEventGroupCreate().
EventGroupHandle_t xEventGroup;

void anInterruptHandler( void )
{
    BaseType_t xHigherPriorityTaskWoken, xResult;

    // xHigherPriorityTaskWoken must be initialised to pdFALSE.
    xHigherPriorityTaskWoken = pdFALSE;

    // Set bit 0 and bit 4 in xEventGroup.
    xResult = xEventGroupSetBitsFromISR(
        xEventGroup,          // The event group being updated.
        BIT_0 | BIT_4        // The bits being set.
        &xHigherPriorityTaskWoken );

    // Was the message posted successfully?
    if( xResult == pdPASS )
    {
        // If xHigherPriorityTaskWoken is now set to pdTRUE then a context

```

(continues on next page)

(continued from previous page)

```

        // switch should be requested. The macro used is port specific and
        // will be either portYIELD_FROM_ISR() or portEND_SWITCHING_ISR() -
        // refer to the documentation page for the port being used.
        portYIELD_FROM_ISR( xHigherPriorityTaskWoken );
    }
}

```

Parameters

- **xEventGroup** -- The event group in which the bits are to be set.
- **uxBitsToSet** -- A bitwise value that indicates the bit or bits to set. For example, to set bit 3 only, set `uxBitsToSet` to `0x08`. To set bit 3 and bit 0 set `uxBitsToSet` to `0x09`.
- **pxHigherPriorityTaskWoken** -- As mentioned above, calling this function will result in a message being sent to the timer daemon task. If the priority of the timer daemon task is higher than the priority of the currently running task (the task the interrupt interrupted) then `*pxHigherPriorityTaskWoken` will be set to `pdTRUE` by `xEventGroupSetBitsFromISR()`, indicating that a context switch should be requested before the interrupt exits. For that reason `*pxHigherPriorityTaskWoken` must be initialised to `pdFALSE`. See the example code below.

Returns If the request to execute the function was posted successfully then `pdPASS` is returned, otherwise `pdFALSE` is returned. `pdFALSE` will be returned if the timer service queue was full.

xEventGroupGetBits (`xEventGroup`)

Returns the current value of the bits in an event group. This function cannot be used from an interrupt.

Parameters

- **xEventGroup** -- The event group being queried.

Returns The event group bits at the time `xEventGroupGetBits()` was called.

Type Definitions

```
typedef struct EventGroupDef_t *EventGroupHandle_t
```

```
typedef TickType_t EventBits_t
```

Stream Buffer API**Header File**

- [components/freertos/FreeRTOS-Kernel/include/freertos/stream_buffer.h](#)
- This header file can be included with:

```
#include "freertos/stream_buffer.h"
```

Functions

BaseType_t xStreamBufferGetStaticBuffers ([StreamBufferHandle_t](#) `xStreamBuffer`, `uint8_t` `**ppucStreamBufferStorageArea`, `StaticStreamBuffer_t` `**ppxStaticStreamBuffer`)

Retrieve pointers to a statically created stream buffer's data structure buffer and storage area buffer. These are the same buffers that are supplied at the time of creation.

Parameters

- **xStreamBuffer** -- The stream buffer for which to retrieve the buffers.
- **ppucStreamBufferStorageArea** -- Used to return a pointer to the stream buffer's storage area buffer.

- **ppxStaticStreamBuffer** -- Used to return a pointer to the stream buffer's data structure buffer.

Returns pdTRUE if buffers were retrieved, pdFALSE otherwise.

size_t **xStreamBufferSend** (*StreamBufferHandle_t* xStreamBuffer, const void *pvTxData, size_t xDataLengthBytes, TickType_t xTicksToWait)

Sends bytes to a stream buffer. The bytes are copied into the stream buffer.

: Uniquely among FreeRTOS objects, the stream buffer implementation (so also the message buffer implementation, as message buffers are built on top of stream buffers) assumes there is only one task or interrupt that will write to the buffer (the writer), and only one task or interrupt that will read from the buffer (the reader). It is safe for the writer and reader to be different tasks or interrupts, but, unlike other FreeRTOS objects, it is not safe to have multiple different writers or multiple different readers. If there are to be multiple different writers then the application writer must place each call to a writing API function (such as xStreamBufferSend()) inside a critical section and set the send block time to 0. Likewise, if there are to be multiple different readers then the application writer must place each call to a reading API function (such as xStreamBufferReceive()) inside a critical section and set the receive block time to 0.

Use xStreamBufferSend() to write to a stream buffer from a task. Use xStreamBufferSendFromISR() to write to a stream buffer from an interrupt service routine (ISR).

Example use:

```
void vAFunction( StreamBufferHandle_t xStreamBuffer )
{
    size_t xBytesSent;
    uint8_t ucArrayToSend[] = { 0, 1, 2, 3 };
    char *pcStringToSend = "String to send";
    const TickType_t x100ms = pdMS_TO_TICKS( 100 );

    // Send an array to the stream buffer, blocking for a maximum of 100ms to
    // wait for enough space to be available in the stream buffer.
    xBytesSent = xStreamBufferSend( xStreamBuffer, ( void * ) ucArrayToSend,
    ↪ sizeof( ucArrayToSend ), x100ms );

    if( xBytesSent != sizeof( ucArrayToSend ) )
    {
        // The call to xStreamBufferSend() times out before there was enough
        // space in the buffer for the data to be written, but it did
        // successfully write xBytesSent bytes.
    }

    // Send the string to the stream buffer. Return immediately if there is not
    // enough space in the buffer.
    xBytesSent = xStreamBufferSend( xStreamBuffer, ( void * ) pcStringToSend,
    ↪ strlen( pcStringToSend ), 0 );

    if( xBytesSent != strlen( pcStringToSend ) )
    {
        // The entire string could not be added to the stream buffer because
        // there was not enough free space in the buffer, but xBytesSent bytes
        // were sent. Could try again to send the remaining bytes.
    }
}
```

Parameters

- **xStreamBuffer** -- The handle of the stream buffer to which a stream is being sent.
- **pvTxData** -- A pointer to the buffer that holds the bytes to be copied into the stream buffer.
- **xDataLengthBytes** -- The maximum number of bytes to copy from pvTxData into the stream buffer.

- **xTicksToWait** -- The maximum amount of time the task should remain in the Blocked state to wait for enough space to become available in the stream buffer, should the stream buffer contain too little space to hold the another xDataLengthBytes bytes. The block time is specified in tick periods, so the absolute time it represents is dependent on the tick frequency. The macro `pdMS_TO_TICKS()` can be used to convert a time specified in milliseconds into a time specified in ticks. Setting `xTicksToWait` to `portMAX_DELAY` will cause the task to wait indefinitely (without timing out), provided `INCLUDE_vTaskSuspend` is set to 1 in `FreeRTOSConfig.h`. If a task times out before it can write all `xDataLengthBytes` into the buffer it will still write as many bytes as possible. A task does not use any CPU time when it is in the blocked state.

Returns The number of bytes written to the stream buffer. If a task times out before it can write all `xDataLengthBytes` into the buffer it will still write as many bytes as possible.

`size_t xStreamBufferSendFromISR` (*StreamBufferHandle_t* xStreamBuffer, const void *pvTxData, `size_t` xDataLengthBytes, `BaseType_t` *const pxHigherPriorityTaskWoken)

Interrupt safe version of the API function that sends a stream of bytes to the stream buffer.

: Uniquely among FreeRTOS objects, the stream buffer implementation (so also the message buffer implementation, as message buffers are built on top of stream buffers) assumes there is only one task or interrupt that will write to the buffer (the writer), and only one task or interrupt that will read from the buffer (the reader). It is safe for the writer and reader to be different tasks or interrupts, but, unlike other FreeRTOS objects, it is not safe to have multiple different writers or multiple different readers. If there are to be multiple different writers then the application writer must place each call to a writing API function (such as `xStreamBufferSend()`) inside a critical section and set the send block time to 0. Likewise, if there are to be multiple different readers then the application writer must place each call to a reading API function (such as `xStreamBufferReceive()`) inside a critical section and set the receive block time to 0.

Use `xStreamBufferSend()` to write to a stream buffer from a task. Use `xStreamBufferSendFromISR()` to write to a stream buffer from an interrupt service routine (ISR).

Example use:

```
// A stream buffer that has already been created.
StreamBufferHandle_t xStreamBuffer;

void vAnInterruptServiceRoutine( void )
{
    size_t xBytesSent;
    char *pcStringToSend = "String to send";
    BaseType_t xHigherPriorityTaskWoken = pdFALSE; // Initialised to pdFALSE.

    // Attempt to send the string to the stream buffer.
    xBytesSent = xStreamBufferSendFromISR( xStreamBuffer,
                                           ( void * ) pcStringToSend,
                                           strlen( pcStringToSend ),
                                           &xHigherPriorityTaskWoken );

    if( xBytesSent != strlen( pcStringToSend ) )
    {
        // There was not enough free space in the stream buffer for the entire
        // string to be written, ut xBytesSent bytes were written.
    }

    // If xHigherPriorityTaskWoken was set to pdTRUE inside
    // xStreamBufferSendFromISR() then a task that has a priority above the
    // priority of the currently executing task was unblocked and a context
    // switch should be performed to ensure the ISR returns to the unblocked
    // task. In most FreeRTOS ports this is done by simply passing
    // xHigherPriorityTaskWoken into portYIELD_FROM_ISR(), which will test the
```

(continues on next page)

(continued from previous page)

```
// variables value, and perform the context switch if necessary. Check the
// documentation for the port in use for port specific instructions.
portYIELD_FROM_ISR( xHigherPriorityTaskWoken );
}
```

Parameters

- **xStreamBuffer** -- The handle of the stream buffer to which a stream is being sent.
- **pvTxData** -- A pointer to the data that is to be copied into the stream buffer.
- **xDataLengthBytes** -- The maximum number of bytes to copy from pvTxData into the stream buffer.
- **pxHigherPriorityTaskWoken** -- It is possible that a stream buffer will have a task blocked on it waiting for data. Calling xStreamBufferSendFromISR() can make data available, and so cause a task that was waiting for data to leave the Blocked state. If calling xStreamBufferSendFromISR() causes a task to leave the Blocked state, and the unblocked task has a priority higher than the currently executing task (the task that was interrupted), then, internally, xStreamBufferSendFromISR() will set *pxHigherPriorityTaskWoken to pdTRUE. If xStreamBufferSendFromISR() sets this value to pdTRUE, then normally a context switch should be performed before the interrupt is exited. This will ensure that the interrupt returns directly to the highest priority Ready state task. *pxHigherPriorityTaskWoken should be set to pdFALSE before it is passed into the function. See the example code below for an example.

Returns The number of bytes actually written to the stream buffer, which will be less than xDataLengthBytes if the stream buffer didn't have enough free space for all the bytes to be written.

size_t **xStreamBufferReceive** (*StreamBufferHandle_t* xStreamBuffer, void *pvRxData, size_t xBufferLengthBytes, TickType_t xTicksToWait)

Receives bytes from a stream buffer.

: Uniquely among FreeRTOS objects, the stream buffer implementation (so also the message buffer implementation, as message buffers are built on top of stream buffers) assumes there is only one task or interrupt that will write to the buffer (the writer), and only one task or interrupt that will read from the buffer (the reader). It is safe for the writer and reader to be different tasks or interrupts, but, unlike other FreeRTOS objects, it is not safe to have multiple different writers or multiple different readers. If there are to be multiple different writers then the application writer must place each call to a writing API function (such as xStreamBufferSend()) inside a critical section and set the send block time to 0. Likewise, if there are to be multiple different readers then the application writer must place each call to a reading API function (such as xStreamBufferReceive()) inside a critical section and set the receive block time to 0.

Use xStreamBufferReceive() to read from a stream buffer from a task. Use xStreamBufferReceiveFromISR() to read from a stream buffer from an interrupt service routine (ISR).

Example use:

```
void vAFunction( StreamBuffer_t xStreamBuffer )
{
    uint8_t ucRxData[ 20 ];
    size_t xReceivedBytes;
    const TickType_t xBlockTime = pdMS_TO_TICKS( 20 );

    // Receive up to another sizeof( ucRxData ) bytes from the stream buffer.
    // Wait in the Blocked state (so not using any CPU processing time) for a
    // maximum of 100ms for the full sizeof( ucRxData ) number of bytes to be
    // available.
    xReceivedBytes = xStreamBufferReceive( xStreamBuffer,
                                           ( void * ) ucRxData,
                                           sizeof( ucRxData ),
                                           xBlockTime );
}
```

(continues on next page)

(continued from previous page)

```

if( xReceivedBytes > 0 )
{
    // A ucRxData contains another xReceivedBytes bytes of data, which can
    // be processed here....
}

```

Parameters

- **xStreamBuffer** -- The handle of the stream buffer from which bytes are to be received.
- **pvRxData** -- A pointer to the buffer into which the received bytes will be copied.
- **xBufferLengthBytes** -- The length of the buffer pointed to by the pvRxData parameter. This sets the maximum number of bytes to receive in one call. xStreamBufferReceive will return as many bytes as possible up to a maximum set by xBufferLengthBytes.
- **xTicksToWait** -- The maximum amount of time the task should remain in the Blocked state to wait for data to become available if the stream buffer is empty. xStreamBufferReceive() will return immediately if xTicksToWait is zero. The block time is specified in tick periods, so the absolute time it represents is dependent on the tick frequency. The macro pdMS_TO_TICKS() can be used to convert a time specified in milliseconds into a time specified in ticks. Setting xTicksToWait to portMAX_DELAY will cause the task to wait indefinitely (without timing out), provided INCLUDE_vTaskSuspend is set to 1 in FreeRTOSConfig.h. A task does not use any CPU time when it is in the Blocked state.

Returns The number of bytes actually read from the stream buffer, which will be less than xBufferLengthBytes if the call to xStreamBufferReceive() timed out before xBufferLengthBytes were available.

size_t **xStreamBufferReceiveFromISR** (*StreamBufferHandle_t* xStreamBuffer, void *pvRxData, size_t xBufferLengthBytes, BaseType_t *const pxHigherPriorityTaskWoken)

An interrupt safe version of the API function that receives bytes from a stream buffer.

Use xStreamBufferReceive() to read bytes from a stream buffer from a task. Use xStreamBufferReceiveFromISR() to read bytes from a stream buffer from an interrupt service routine (ISR).

Example use:

```

// A stream buffer that has already been created.
StreamBuffer_t xStreamBuffer;

void vAnInterruptServiceRoutine( void )
{
    uint8_t ucRxData[ 20 ];
    size_t xReceivedBytes;
    BaseType_t xHigherPriorityTaskWoken = pdFALSE; // Initialised to pdFALSE.

    // Receive the next stream from the stream buffer.
    xReceivedBytes = xStreamBufferReceiveFromISR( xStreamBuffer,
                                                ( void * ) ucRxData,
                                                sizeof( ucRxData ),
                                                &xHigherPriorityTaskWoken );

    if( xReceivedBytes > 0 )
    {
        // ucRxData contains xReceivedBytes read from the stream buffer.
        // Process the stream here....
    }
}

```

(continues on next page)

(continued from previous page)

```

// If xHigherPriorityTaskWoken was set to pdTRUE inside
// xStreamBufferReceiveFromISR() then a task that has a priority above the
// priority of the currently executing task was unblocked and a context
// switch should be performed to ensure the ISR returns to the unblocked
// task. In most FreeRTOS ports this is done by simply passing
// xHigherPriorityTaskWoken into portYIELD_FROM_ISR(), which will test the
// variables value, and perform the context switch if necessary. Check the
// documentation for the port in use for port specific instructions.
portYIELD_FROM_ISR( xHigherPriorityTaskWoken );
}

```

Parameters

- **xStreamBuffer** -- The handle of the stream buffer from which a stream is being received.
- **pvRxData** -- A pointer to the buffer into which the received bytes are copied.
- **xBufferLengthBytes** -- The length of the buffer pointed to by the pvRxData parameter. This sets the maximum number of bytes to receive in one call. xStreamBufferReceive will return as many bytes as possible up to a maximum set by xBufferLengthBytes.
- **pxHigherPriorityTaskWoken** -- It is possible that a stream buffer will have a task blocked on it waiting for space to become available. Calling xStreamBufferReceiveFromISR() can make space available, and so cause a task that is waiting for space to leave the Blocked state. If calling xStreamBufferReceiveFromISR() causes a task to leave the Blocked state, and the unblocked task has a priority higher than the currently executing task (the task that was interrupted), then, internally, xStreamBufferReceiveFromISR() will set *pxHigherPriorityTaskWoken to pdTRUE. If xStreamBufferReceiveFromISR() sets this value to pdTRUE, then normally a context switch should be performed before the interrupt is exited. That will ensure the interrupt returns directly to the highest priority Ready state task. *pxHigherPriorityTaskWoken should be set to pdFALSE before it is passed into the function. See the code example below for an example.

Returns The number of bytes read from the stream buffer, if any.

void **vStreamBufferDelete** (*StreamBufferHandle_t* xStreamBuffer)

Deletes a stream buffer that was previously created using a call to xStreamBufferCreate() or xStreamBufferCreateStatic(). If the stream buffer was created using dynamic memory (that is, by xStreamBufferCreate()), then the allocated memory is freed.

A stream buffer handle must not be used after the stream buffer has been deleted.

Parameters **xStreamBuffer** -- The handle of the stream buffer to be deleted.

BaseType_t **xStreamBufferIsFull** (*StreamBufferHandle_t* xStreamBuffer)

Queries a stream buffer to see if it is full. A stream buffer is full if it does not have any free space, and therefore cannot accept any more data.

Parameters **xStreamBuffer** -- The handle of the stream buffer being queried.

Returns If the stream buffer is full then pdTRUE is returned. Otherwise pdFALSE is returned.

BaseType_t **xStreamBufferIsEmpty** (*StreamBufferHandle_t* xStreamBuffer)

Queries a stream buffer to see if it is empty. A stream buffer is empty if it does not contain any data.

Parameters **xStreamBuffer** -- The handle of the stream buffer being queried.

Returns If the stream buffer is empty then pdTRUE is returned. Otherwise pdFALSE is returned.

BaseType_t **xStreamBufferReset** (*StreamBufferHandle_t* xStreamBuffer)

Resets a stream buffer to its initial, empty, state. Any data that was in the stream buffer is discarded. A stream buffer can only be reset if there are no tasks blocked waiting to either send to or receive from the stream buffer.

Parameters **xStreamBuffer** -- The handle of the stream buffer being reset.

Returns If the stream buffer is reset then pdPASS is returned. If there was a task blocked waiting to send to or read from the stream buffer then the stream buffer is not reset and pdFAIL is returned.

size_t **xStreamBufferSpacesAvailable** (*StreamBufferHandle_t* xStreamBuffer)

Queries a stream buffer to see how much free space it contains, which is equal to the amount of data that can be sent to the stream buffer before it is full.

Parameters **xStreamBuffer** -- The handle of the stream buffer being queried.

Returns The number of bytes that can be written to the stream buffer before the stream buffer would be full.

size_t **xStreamBufferBytesAvailable** (*StreamBufferHandle_t* xStreamBuffer)

Queries a stream buffer to see how much data it contains, which is equal to the number of bytes that can be read from the stream buffer before the stream buffer would be empty.

Parameters **xStreamBuffer** -- The handle of the stream buffer being queried.

Returns The number of bytes that can be read from the stream buffer before the stream buffer would be empty.

BaseType_t **xStreamBufferSetTriggerLevel** (*StreamBufferHandle_t* xStreamBuffer, size_t xTriggerLevel)

A stream buffer's trigger level is the number of bytes that must be in the stream buffer before a task that is blocked on the stream buffer to wait for data is moved out of the blocked state. For example, if a task is blocked on a read of an empty stream buffer that has a trigger level of 1 then the task will be unblocked when a single byte is written to the buffer or the task's block time expires. As another example, if a task is blocked on a read of an empty stream buffer that has a trigger level of 10 then the task will not be unblocked until the stream buffer contains at least 10 bytes or the task's block time expires. If a reading task's block time expires before the trigger level is reached then the task will still receive however many bytes are actually available. Setting a trigger level of 0 will result in a trigger level of 1 being used. It is not valid to specify a trigger level that is greater than the buffer size.

A trigger level is set when the stream buffer is created, and can be modified using xStreamBufferSetTriggerLevel().

Parameters

- **xStreamBuffer** -- The handle of the stream buffer being updated.
- **xTriggerLevel** -- The new trigger level for the stream buffer.

Returns If xTriggerLevel was less than or equal to the stream buffer's length then the trigger level will be updated and pdTRUE is returned. Otherwise pdFALSE is returned.

BaseType_t **xStreamBufferSendCompletedFromISR** (*StreamBufferHandle_t* xStreamBuffer, BaseType_t *pxHigherPriorityTaskWoken)

For advanced users only.

The sbSEND_COMPLETED() macro is called from within the FreeRTOS APIs when data is sent to a message buffer or stream buffer. If there was a task that was blocked on the message or stream buffer waiting for data to arrive then the sbSEND_COMPLETED() macro sends a notification to the task to remove it from the Blocked state. xStreamBufferSendCompletedFromISR() does the same thing. It is provided to enable application writers to implement their own version of sbSEND_COMPLETED(), and MUST NOT BE USED AT ANY OTHER TIME.

See the example implemented in FreeRTOS/Demo/Minimal/MessageBufferAMP.c for additional information.

Parameters

- **xStreamBuffer** -- The handle of the stream buffer to which data was written.
- **pxHigherPriorityTaskWoken** -- *pxHigherPriorityTaskWoken should be initialised to pdFALSE before it is passed into xStreamBufferSendCompletedFromISR(). If calling xStreamBufferSendCompletedFromISR() removes a task from the Blocked state, and the task has a priority above the priority of the currently running task, then *pxHigherPriorityTaskWoken will get set to pdTRUE indicating that a context switch should be performed before exiting the ISR.

Returns If a task was removed from the Blocked state then pdTRUE is returned. Otherwise pdFALSE is returned.

BaseType_t **xStreamBufferReceiveCompletedFromISR** (*StreamBufferHandle_t* xStreamBuffer, BaseType_t *pxHigherPriorityTaskWoken)

For advanced users only.

The sbRECEIVE_COMPLETED() macro is called from within the FreeRTOS APIs when data is read out of a message buffer or stream buffer. If there was a task that was blocked on the message or stream buffer waiting for data to arrive then the sbRECEIVE_COMPLETED() macro sends a notification to the task to remove it from the Blocked state. xStreamBufferReceiveCompletedFromISR() does the same thing. It is provided to enable application writers to implement their own version of sbRECEIVE_COMPLETED(), and MUST NOT BE USED AT ANY OTHER TIME.

See the example implemented in FreeRTOS/Demo/Minimal/MessageBufferAMP.c for additional information.

Parameters

- **xStreamBuffer** -- The handle of the stream buffer from which data was read.
- **pxHigherPriorityTaskWoken** -- *pxHigherPriorityTaskWoken should be initialised to pdFALSE before it is passed into xStreamBufferReceiveCompletedFromISR(). If calling xStreamBufferReceiveCompletedFromISR() removes a task from the Blocked state, and the task has a priority above the priority of the currently running task, then *pxHigherPriorityTaskWoken will get set to pdTRUE indicating that a context switch should be performed before exiting the ISR.

Returns If a task was removed from the Blocked state then pdTRUE is returned. Otherwise pdFALSE is returned.

Macros

xStreamBufferCreateWithCallback (xBufferSizeBytes, xTriggerLevelBytes, pxSendCompletedCallback, pxReceiveCompletedCallback)

Creates a new stream buffer using dynamically allocated memory. See xStreamBufferCreateStatic() for a version that uses statically allocated memory (memory that is allocated at compile time).

configSUPPORT_DYNAMIC_ALLOCATION must be set to 1 or left undefined in FreeRTOSConfig.h for xStreamBufferCreate() to be available.

Example use:

```
void vAFunction( void )
{
    StreamBufferHandle_t xStreamBuffer;
    const size_t xStreamBufferSizeBytes = 100, xTriggerLevel = 10;

    // Create a stream buffer that can hold 100 bytes. The memory used to hold
    // both the stream buffer structure and the data in the stream buffer is
    // allocated dynamically.
    xStreamBuffer = xStreamBufferCreate( xStreamBufferSizeBytes, xTriggerLevel );

    if( xStreamBuffer == NULL )
    {
        // There was not enough heap memory space available to create the
        // stream buffer.
    }
    else
    {
        // The stream buffer was created successfully and can now be used.
    }
}
```

Parameters

- **xBufferSizeBytes** -- The total number of bytes the stream buffer will be able to hold at any one time.
- **xTriggerLevelBytes** -- The number of bytes that must be in the stream buffer before a task that is blocked on the stream buffer to wait for data is moved out of the blocked state. For example, if a task is blocked on a read of an empty stream buffer that has a trigger level of 1 then the task will be unblocked when a single byte is written to the buffer or the task's block time expires. As another example, if a task is blocked on a read of an empty stream buffer that has a trigger level of 10 then the task will not be unblocked until the stream buffer contains at least 10 bytes or the task's block time expires. If a reading task's block time expires before the trigger level is reached then the task will still receive however many bytes are actually available. Setting a trigger level of 0 will result in a trigger level of 1 being used. It is not valid to specify a trigger level that is greater than the buffer size.
- **pxSendCompletedCallback** -- Callback invoked when number of bytes at least equal to trigger level is sent to the stream buffer. If the parameter is NULL, it will use the default implementation provided by sbSEND_COMPLETED macro. To enable the callback, configUSE_SB_COMPLETED_CALLBACK must be set to 1 in FreeRTOSConfig.h.
- **pxReceiveCompletedCallback** -- Callback invoked when more than zero bytes are read from a stream buffer. If the parameter is NULL, it will use the default implementation provided by sbRECEIVE_COMPLETED macro. To enable the callback, configUSE_SB_COMPLETED_CALLBACK must be set to 1 in FreeRTOSConfig.h.

Returns If NULL is returned, then the stream buffer cannot be created because there is insufficient heap memory available for FreeRTOS to allocate the stream buffer data structures and storage area. A non-NULL value being returned indicates that the stream buffer has been created successfully - the returned value should be stored as the handle to the created stream buffer.

xStreamBufferCreateStaticWithCallback (xBufferSizeBytes, xTriggerLevelBytes, pucStreamBufferStorageArea, pxStaticStreamBuffer, pxSendCompletedCallback, pxReceiveCompletedCallback)

Creates a new stream buffer using statically allocated memory. See xStreamBufferCreate() for a version that uses dynamically allocated memory.

configSUPPORT_STATIC_ALLOCATION must be set to 1 in FreeRTOSConfig.h for xStreamBufferCreateStatic() to be available.

Example use:

```
// Used to dimension the array used to hold the streams. The available space
// will actually be one less than this, so 999.
#define STORAGE_SIZE_BYTES 1000

// Defines the memory that will actually hold the streams within the stream
// buffer.
static uint8_t ucStorageBuffer[ STORAGE_SIZE_BYTES ];

// The variable used to hold the stream buffer structure.
StaticStreamBuffer_t xStreamBufferStruct;

void MyFunction( void )
{
    StreamBufferHandle_t xStreamBuffer;
    const size_t xTriggerLevel = 1;

    xStreamBuffer = xStreamBufferCreateStatic( sizeof( ucStorageBuffer ),
                                                xTriggerLevel,
                                                ucStorageBuffer,
                                                &xStreamBufferStruct );
}
```

(continues on next page)

(continued from previous page)

```

// As neither the pucStreamBufferStorageArea or pxStaticStreamBuffer
// parameters were NULL, xStreamBuffer will not be NULL, and can be used to
// reference the created stream buffer in other stream buffer API calls.

// Other code that uses the stream buffer can go here.
}

```

Parameters

- **xBufferSizeBytes** -- The size, in bytes, of the buffer pointed to by the pucStreamBufferStorageArea parameter.
- **xTriggerLevelBytes** -- The number of bytes that must be in the stream buffer before a task that is blocked on the stream buffer to wait for data is moved out of the blocked state. For example, if a task is blocked on a read of an empty stream buffer that has a trigger level of 1 then the task will be unblocked when a single byte is written to the buffer or the task's block time expires. As another example, if a task is blocked on a read of an empty stream buffer that has a trigger level of 10 then the task will not be unblocked until the stream buffer contains at least 10 bytes or the task's block time expires. If a reading task's block time expires before the trigger level is reached then the task will still receive however many bytes are actually available. Setting a trigger level of 0 will result in a trigger level of 1 being used. It is not valid to specify a trigger level that is greater than the buffer size.
- **pucStreamBufferStorageArea** -- Must point to a uint8_t array that is at least xBufferSizeBytes big. This is the array to which streams are copied when they are written to the stream buffer.
- **pxStaticStreamBuffer** -- Must point to a variable of type StaticStreamBuffer_t, which will be used to hold the stream buffer's data structure.
- **pxSendCompletedCallback** -- Callback invoked when number of bytes at least equal to trigger level is sent to the stream buffer. If the parameter is NULL, it will use the default implementation provided by sbSEND_COMPLETED macro. To enable the callback, configUSE_SB_COMPLETED_CALLBACK must be set to 1 in FreeRTOSConfig.h.
- **pxReceiveCompletedCallback** -- Callback invoked when more than zero bytes are read from a stream buffer. If the parameter is NULL, it will use the default implementation provided by sbRECEIVE_COMPLETED macro. To enable the callback, configUSE_SB_COMPLETED_CALLBACK must be set to 1 in FreeRTOSConfig.h.

Returns If the stream buffer is created successfully then a handle to the created stream buffer is returned. If either pucStreamBufferStorageArea or pxStaticstreamBuffer are NULL then NULL is returned.

Type Definitions

```
typedef struct StreamBufferDef_t *StreamBufferHandle_t
```

```
typedef void (*StreamBufferCallbackFunction_t)(StreamBufferHandle_t xStreamBuffer, BaseType_t xIsInsideISR, BaseType_t *const pxHigherPriorityTaskWoken)
```

Type used as a stream buffer's optional callback.

Message Buffer API**Header File**

- [components/freertos/FreeRTOS-Kernel/include/freertos/message_buffer.h](#)
- This header file can be included with:

```
#include "freertos/message_buffer.h"
```


Macros

xMessageBufferCreateWithCallback (xBufferSizeBytes, pxSendCompletedCallback, pxReceiveCompletedCallback)

Creates a new message buffer using dynamically allocated memory. See xMessageBufferCreateStatic() for a version that uses statically allocated memory (memory that is allocated at compile time).

configSUPPORT_DYNAMIC_ALLOCATION must be set to 1 or left undefined in FreeRTOSConfig.h for xMessageBufferCreate() to be available.

Example use:

```
void vAFunction( void )
{
    MessageBufferHandle_t xMessageBuffer;
    const size_t xMessageBufferSizeBytes = 100;

    // Create a message buffer that can hold 100 bytes. The memory used to hold
    // both the message buffer structure and the messages themselves is allocated
    // dynamically. Each message added to the buffer consumes an additional 4
    // bytes which are used to hold the length of the message.
    xMessageBuffer = xMessageBufferCreate( xMessageBufferSizeBytes );

    if( xMessageBuffer == NULL )
    {
        // There was not enough heap memory space available to create the
        // message buffer.
    }
    else
    {
        // The message buffer was created successfully and can now be used.
    }
}
```

Parameters

- **xBufferSizeBytes** -- The total number of bytes (not messages) the message buffer will be able to hold at any one time. When a message is written to the message buffer an additional sizeof(size_t) bytes are also written to store the message's length. sizeof(size_t) is typically 4 bytes on a 32-bit architecture, so on most 32-bit architectures a 10 byte message will take up 14 bytes of message buffer space.
- **pxSendCompletedCallback** -- Callback invoked when a send operation to the message buffer is complete. If the parameter is NULL or xMessageBufferCreate() is called without the parameter, then it will use the default implementation provided by sbSEND_COMPLETED macro. To enable the callback, configUSE_SB_COMPLETED_CALLBACK must be set to 1 in FreeRTOSConfig.h.
- **pxReceiveCompletedCallback** -- Callback invoked when a receive operation from the message buffer is complete. If the parameter is NULL or xMessageBufferCreate() is called without the parameter, it will use the default implementation provided by sbRECEIVE_COMPLETED macro. To enable the callback, configUSE_SB_COMPLETED_CALLBACK must be set to 1 in FreeRTOSConfig.h.

Returns If NULL is returned, then the message buffer cannot be created because there is insufficient heap memory available for FreeRTOS to allocate the message buffer data structures and storage area. A non-NULL value being returned indicates that the message buffer has been created successfully - the returned value should be stored as the handle to the created message buffer.

xMessageBufferCreateStaticWithCallback (xBufferSizeBytes, pucMessageBufferStorageArea, pxStaticMessageBuffer, pxSendCompletedCallback, pxReceiveCompletedCallback)

Creates a new message buffer using statically allocated memory. See xMessageBufferCreate() for a version that uses dynamically allocated memory.

Example use:

```
// Used to dimension the array used to hold the messages. The available space
// will actually be one less than this, so 999.
#define STORAGE_SIZE_BYTES 1000

// Defines the memory that will actually hold the messages within the message
// buffer.
static uint8_t ucStorageBuffer[ STORAGE_SIZE_BYTES ];

// The variable used to hold the message buffer structure.
StaticMessageBuffer_t xMessageBufferStruct;

void MyFunction( void )
{
    MessageBufferHandle_t xMessageBuffer;

    xMessageBuffer = xMessageBufferCreateStatic( sizeof( ucStorageBuffer ),
                                                ucStorageBuffer,
                                                &xMessageBufferStruct );

    // As neither the pucMessageBufferStorageArea or pxStaticMessageBuffer
    // parameters were NULL, xMessageBuffer will not be NULL, and can be used to
    // reference the created message buffer in other message buffer API calls.

    // Other code that uses the message buffer can go here.
}
```

Parameters

- **xBufferSizeBytes** -- The size, in bytes, of the buffer pointed to by the `pucMessageBufferStorageArea` parameter. When a message is written to the message buffer an additional `sizeof( size_t )` bytes are also written to store the message's length. `sizeof( size_t )` is typically 4 bytes on a 32-bit architecture, so on most 32-bit architecture a 10 byte message will take up 14 bytes of message buffer space. The maximum number of bytes that can be stored in the message buffer is actually `(xBufferSizeBytes - 1)`.
- **pucMessageBufferStorageArea** -- Must point to a `uint8_t` array that is at least `xBufferSizeBytes` big. This is the array to which messages are copied when they are written to the message buffer.
- **pxStaticMessageBuffer** -- Must point to a variable of type `StaticMessageBuffer_t`, which will be used to hold the message buffer's data structure.
- **pxSendCompletedCallback** -- Callback invoked when a new message is sent to the message buffer. If the parameter is `NULL` or `xMessageBufferCreate()` is called without the parameter, then it will use the default implementation provided by `sbSEND_COMPLETED` macro. To enable the callback, `configUSE_SB_COMPLETED_CALLBACK` must be set to 1 in `FreeRTOSConfig.h`.
- **pxReceiveCompletedCallback** -- Callback invoked when a message is read from a message buffer. If the parameter is `NULL` or `xMessageBufferCreate()` is called without the parameter, it will use the default implementation provided by `sbRECEIVE_COMPLETED` macro. To enable the callback, `configUSE_SB_COMPLETED_CALLBACK` must be set to 1 in `FreeRTOSConfig.h`.

Returns If the message buffer is created successfully then a handle to the created message buffer is returned. If either `pucMessageBufferStorageArea` or `pxStaticmessageBuffer` are `NULL` then `NULL` is returned.

xMessageBufferGetStaticBuffers (`xMessageBuffer`, `ppucMessageBufferStorageArea`,
`ppxStaticMessageBuffer`)

Retrieve pointers to a statically created message buffer's data structure buffer and storage area buffer. These are the same buffers that are supplied at the time of creation.

Parameters

- **xMessageBuffer** -- The message buffer for which to retrieve the buffers.
- **ppucMessageBufferStorageArea** -- Used to return a pointer to the message buffer's storage area buffer.
- **ppxStaticMessageBuffer** -- Used to return a pointer to the message buffer's data structure buffer.

Returns pdTRUE if buffers were retrieved, pdFALSE otherwise.

xMessageBufferSend (xMessageBuffer, pvTxData, xDataLengthBytes, xTicksToWait)

Sends a discrete message to the message buffer. The message can be any length that fits within the buffer's free space, and is copied into the buffer.

: Uniquely among FreeRTOS objects, the stream buffer implementation (so also the message buffer implementation, as message buffers are built on top of stream buffers) assumes there is only one task or interrupt that will write to the buffer (the writer), and only one task or interrupt that will read from the buffer (the reader). It is safe for the writer and reader to be different tasks or interrupts, but, unlike other FreeRTOS objects, it is not safe to have multiple different writers or multiple different readers. If there are to be multiple different writers then the application writer must place each call to a writing API function (such as xMessageBufferSend()) inside a critical section and set the send block time to 0. Likewise, if there are to be multiple different readers then the application writer must place each call to a reading API function (such as xMessageBufferRead()) inside a critical section and set the receive block time to 0.

Use xMessageBufferSend() to write to a message buffer from a task. Use xMessageBufferSendFromISR() to write to a message buffer from an interrupt service routine (ISR).

Example use:

```
void vAFunction( MessageBufferHandle_t xMessageBuffer )
{
    size_t xBytesSent;
    uint8_t ucArrayToSend[] = { 0, 1, 2, 3 };
    char *pcStringToSend = "String to send";
    const TickType_t x100ms = pdMS_TO_TICKS( 100 );

    // Send an array to the message buffer, blocking for a maximum of 100ms to
    // wait for enough space to be available in the message buffer.
    xBytesSent = xMessageBufferSend( xMessageBuffer, ( void * ) ucArrayToSend,
    ↪ sizeof( ucArrayToSend ), x100ms );

    if( xBytesSent != sizeof( ucArrayToSend ) )
    {
        // The call to xMessageBufferSend() times out before there was enough
        // space in the buffer for the data to be written.
    }

    // Send the string to the message buffer. Return immediately if there is
    // not enough space in the buffer.
    xBytesSent = xMessageBufferSend( xMessageBuffer, ( void * ) pcStringToSend,
    ↪ strlen( pcStringToSend ), 0 );

    if( xBytesSent != strlen( pcStringToSend ) )
    {
        // The string could not be added to the message buffer because there was
        // not enough free space in the buffer.
    }
}
```

Parameters

- **xMessageBuffer** -- The handle of the message buffer to which a message is being sent.
- **pvTxData** -- A pointer to the message that is to be copied into the message buffer.

- **xDataLengthBytes** -- The length of the message. That is, the number of bytes to copy from pvTxData into the message buffer. When a message is written to the message buffer an additional sizeof(size_t) bytes are also written to store the message's length. sizeof(size_t) is typically 4 bytes on a 32-bit architecture, so on most 32-bit architecture setting xDataLengthBytes to 20 will reduce the free space in the message buffer by 24 bytes (20 bytes of message data and 4 bytes to hold the message length).
- **xTicksToWait** -- The maximum amount of time the calling task should remain in the Blocked state to wait for enough space to become available in the message buffer, should the message buffer have insufficient space when xMessageBufferSend() is called. The calling task will never block if xTicksToWait is zero. The block time is specified in tick periods, so the absolute time it represents is dependent on the tick frequency. The macro pdMS_TO_TICKS() can be used to convert a time specified in milliseconds into a time specified in ticks. Setting xTicksToWait to portMAX_DELAY will cause the task to wait indefinitely (without timing out), provided INCLUDE_vTaskSuspend is set to 1 in FreeRTOSConfig.h. Tasks do not use any CPU time when they are in the Blocked state.

Returns The number of bytes written to the message buffer. If the call to xMessageBufferSend() times out before there was enough space to write the message into the message buffer then zero is returned. If the call did not time out then xDataLengthBytes is returned.

xMessageBufferSendFromISR (xMessageBuffer, pvTxData, xDataLengthBytes, pxHigherPriorityTaskWoken)

Interrupt safe version of the API function that sends a discrete message to the message buffer. The message can be any length that fits within the buffer's free space, and is copied into the buffer.

: Uniquely among FreeRTOS objects, the stream buffer implementation (so also the message buffer implementation, as message buffers are built on top of stream buffers) assumes there is only one task or interrupt that will write to the buffer (the writer), and only one task or interrupt that will read from the buffer (the reader). It is safe for the writer and reader to be different tasks or interrupts, but, unlike other FreeRTOS objects, it is not safe to have multiple different writers or multiple different readers. If there are to be multiple different writers then the application writer must place each call to a writing API function (such as xMessageBufferSend()) inside a critical section and set the send block time to 0. Likewise, if there are to be multiple different readers then the application writer must place each call to a reading API function (such as xMessageBufferRead()) inside a critical section and set the receive block time to 0.

Use xMessageBufferSend() to write to a message buffer from a task. Use xMessageBufferSendFromISR() to write to a message buffer from an interrupt service routine (ISR).

Example use:

```
// A message buffer that has already been created.
MessageBufferHandle_t xMessageBuffer;

void vAnInterruptServiceRoutine( void )
{
    size_t xBytesSent;
    char *pcStringToSend = "String to send";
    BaseType_t xHigherPriorityTaskWoken = pdFALSE; // Initialised to pdFALSE.

    // Attempt to send the string to the message buffer.
    xBytesSent = xMessageBufferSendFromISR( xMessageBuffer,
                                            ( void * ) pcStringToSend,
                                            strlen( pcStringToSend ),
                                            &xHigherPriorityTaskWoken );

    if( xBytesSent != strlen( pcStringToSend ) )
    {
        // The string could not be added to the message buffer because there was
        // not enough free space in the buffer.
    }
}
```

(continues on next page)

(continued from previous page)

```

// If xHigherPriorityTaskWoken was set to pdTRUE inside
// xMessageBufferSendFromISR() then a task that has a priority above the
// priority of the currently executing task was unblocked and a context
// switch should be performed to ensure the ISR returns to the unblocked
// task. In most FreeRTOS ports this is done by simply passing
// xHigherPriorityTaskWoken into portYIELD_FROM_ISR(), which will test the
// variables value, and perform the context switch if necessary. Check the
// documentation for the port in use for port specific instructions.
portYIELD_FROM_ISR( xHigherPriorityTaskWoken );
}

```

Parameters

- **xMessageBuffer** -- The handle of the message buffer to which a message is being sent.
- **pvTxData** -- A pointer to the message that is to be copied into the message buffer.
- **xDataLengthBytes** -- The length of the message. That is, the number of bytes to copy from pvTxData into the message buffer. When a message is written to the message buffer an additional sizeof(size_t) bytes are also written to store the message's length. sizeof(size_t) is typically 4 bytes on a 32-bit architecture, so on most 32-bit architecture setting xDataLengthBytes to 20 will reduce the free space in the message buffer by 24 bytes (20 bytes of message data and 4 bytes to hold the message length).
- **pxHigherPriorityTaskWoken** -- It is possible that a message buffer will have a task blocked on it waiting for data. Calling xMessageBufferSendFromISR() can make data available, and so cause a task that was waiting for data to leave the Blocked state. If calling xMessageBufferSendFromISR() causes a task to leave the Blocked state, and the unblocked task has a priority higher than the currently executing task (the task that was interrupted), then, internally, xMessageBufferSendFromISR() will set *pxHigherPriorityTaskWoken to pdTRUE. If xMessageBufferSendFromISR() sets this value to pdTRUE, then normally a context switch should be performed before the interrupt is exited. This will ensure that the interrupt returns directly to the highest priority Ready state task. *pxHigherPriorityTaskWoken should be set to pdFALSE before it is passed into the function. See the code example below for an example.

Returns The number of bytes actually written to the message buffer. If the message buffer didn't have enough free space for the message to be stored then 0 is returned, otherwise xDataLengthBytes is returned.

xMessageBufferReceive (xMessageBuffer, pvRxData, xBufferLengthBytes, xTicksToWait)

Receives a discrete message from a message buffer. Messages can be of variable length and are copied out of the buffer.

: Uniquely among FreeRTOS objects, the stream buffer implementation (so also the message buffer implementation, as message buffers are built on top of stream buffers) assumes there is only one task or interrupt that will write to the buffer (the writer), and only one task or interrupt that will read from the buffer (the reader). It is safe for the writer and reader to be different tasks or interrupts, but, unlike other FreeRTOS objects, it is not safe to have multiple different writers or multiple different readers. If there are to be multiple different writers then the application writer must place each call to a writing API function (such as xMessageBufferSend()) inside a critical section and set the send block time to 0. Likewise, if there are to be multiple different readers then the application writer must place each call to a reading API function (such as xMessageBufferRead()) inside a critical section and set the receive block time to 0.

Use xMessageBufferReceive() to read from a message buffer from a task. Use xMessageBufferReceiveFromISR() to read from a message buffer from an interrupt service routine (ISR).

Example use:

```

void vAFunction( MessageBuffer_t xMessageBuffer )
{
    uint8_t ucRxData[ 20 ];
    size_t xReceivedBytes;
    const TickType_t xBlockTime = pdMS_TO_TICKS( 20 );

    // Receive the next message from the message buffer. Wait in the Blocked
    // state (so not using any CPU processing time) for a maximum of 100ms for
    // a message to become available.
    xReceivedBytes = xMessageBufferReceive( xMessageBuffer,
                                           ( void * ) ucRxData,
                                           sizeof( ucRxData ),
                                           xBlockTime );

    if( xReceivedBytes > 0 )
    {
        // A ucRxData contains a message that is xReceivedBytes long. Process
        // the message here....
    }
}

```

Parameters

- **xMessageBuffer** -- The handle of the message buffer from which a message is being received.
- **pvRxData** -- A pointer to the buffer into which the received message is to be copied.
- **xBufferLengthBytes** -- The length of the buffer pointed to by the pvRxData parameter. This sets the maximum length of the message that can be received. If xBufferLengthBytes is too small to hold the next message then the message will be left in the message buffer and 0 will be returned.
- **xTicksToWait** -- The maximum amount of time the task should remain in the Blocked state to wait for a message, should the message buffer be empty. xMessageBufferReceive() will return immediately if xTicksToWait is zero and the message buffer is empty. The block time is specified in tick periods, so the absolute time it represents is dependent on the tick frequency. The macro pdMS_TO_TICKS() can be used to convert a time specified in milliseconds into a time specified in ticks. Setting xTicksToWait to portMAX_DELAY will cause the task to wait indefinitely (without timing out), provided INCLUDE_vTaskSuspend is set to 1 in FreeRTOSConfig.h. Tasks do not use any CPU time when they are in the Blocked state.

Returns The length, in bytes, of the message read from the message buffer, if any. If xMessageBufferReceive() times out before a message became available then zero is returned. If the length of the message is greater than xBufferLengthBytes then the message will be left in the message buffer and zero is returned.

xMessageBufferReceiveFromISR (xMessageBuffer, pvRxData, xBufferLengthBytes, pxHigherPriorityTaskWoken)

An interrupt safe version of the API function that receives a discrete message from a message buffer. Messages can be of variable length and are copied out of the buffer.

: Uniquely among FreeRTOS objects, the stream buffer implementation (so also the message buffer implementation, as message buffers are built on top of stream buffers) assumes there is only one task or interrupt that will write to the buffer (the writer), and only one task or interrupt that will read from the buffer (the reader). It is safe for the writer and reader to be different tasks or interrupts, but, unlike other FreeRTOS objects, it is not safe to have multiple different writers or multiple different readers. If there are to be multiple different writers then the application writer must place each call to a writing API function (such as xMessageBufferSend()) inside a critical section and set the send block time to 0. Likewise, if there are to be multiple different readers then the application writer must place each call to a reading API function (such as xMessageBufferRead()) inside a critical section and set the receive block time to 0.

Use xMessageBufferReceive() to read from a message buffer from a task. Use xMessageBufferReceive-

FromISR() to read from a message buffer from an interrupt service routine (ISR).

Example use:

```
// A message buffer that has already been created.
MessageBuffer_t xMessageBuffer;

void vAnInterruptServiceRoutine( void )
{
    uint8_t ucRxData[ 20 ];
    size_t xReceivedBytes;
    BaseType_t xHigherPriorityTaskWoken = pdFALSE; // Initialised to pdFALSE.

    // Receive the next message from the message buffer.
    xReceivedBytes = xMessageBufferReceiveFromISR( xMessageBuffer,
                                                    ( void * ) ucRxData,
                                                    sizeof( ucRxData ),
                                                    &xHigherPriorityTaskWoken );

    if( xReceivedBytes > 0 )
    {
        // A ucRxData contains a message that is xReceivedBytes long. Process
        // the message here....
    }

    // If xHigherPriorityTaskWoken was set to pdTRUE inside
    // xMessageBufferReceiveFromISR() then a task that has a priority above the
    // priority of the currently executing task was unblocked and a context
    // switch should be performed to ensure the ISR returns to the unblocked
    // task. In most FreeRTOS ports this is done by simply passing
    // xHigherPriorityTaskWoken into portYIELD_FROM_ISR(), which will test the
    // variables value, and perform the context switch if necessary. Check the
    // documentation for the port in use for port specific instructions.
    portYIELD_FROM_ISR( xHigherPriorityTaskWoken );
}
```

Parameters

- **xMessageBuffer** -- The handle of the message buffer from which a message is being received.
- **pvRxData** -- A pointer to the buffer into which the received message is to be copied.
- **xBufferLengthBytes** -- The length of the buffer pointed to by the pvRxData parameter. This sets the maximum length of the message that can be received. If xBufferLengthBytes is too small to hold the next message then the message will be left in the message buffer and 0 will be returned.
- **pxHigherPriorityTaskWoken** -- It is possible that a message buffer will have a task blocked on it waiting for space to become available. Calling xMessageBufferReceiveFromISR() can make space available, and so cause a task that is waiting for space to leave the Blocked state. If calling xMessageBufferReceiveFromISR() causes a task to leave the Blocked state, and the unblocked task has a priority higher than the currently executing task (the task that was interrupted), then, internally, xMessageBufferReceiveFromISR() will set *pxHigherPriorityTaskWoken to pdTRUE. If xMessageBufferReceiveFromISR() sets this value to pdTRUE, then normally a context switch should be performed before the interrupt is exited. That will ensure the interrupt returns directly to the highest priority Ready state task. *pxHigherPriorityTaskWoken should be set to pdFALSE before it is passed into the function. See the code example below for an example.

Returns The length, in bytes, of the message read from the message buffer, if any.

vMessageBufferDelete (xMessageBuffer)

Deletes a message buffer that was previously created using a call to xMessageBufferCreate() or xMessage-

BufferCreateStatic(). If the message buffer was created using dynamic memory (that is, by xMessageBufferCreate()), then the allocated memory is freed.

A message buffer handle must not be used after the message buffer has been deleted.

Parameters

- **xMessageBuffer** -- The handle of the message buffer to be deleted.

xMessageBufferIsFull (xMessageBuffer)

Tests to see if a message buffer is full. A message buffer is full if it cannot accept any more messages, of any size, until space is made available by a message being removed from the message buffer.

Parameters

- **xMessageBuffer** -- The handle of the message buffer being queried.

Returns If the message buffer referenced by xMessageBuffer is full then pdTRUE is returned. Otherwise pdFALSE is returned.

xMessageBufferIsEmpty (xMessageBuffer)

Tests to see if a message buffer is empty (does not contain any messages).

Parameters

- **xMessageBuffer** -- The handle of the message buffer being queried.

Returns If the message buffer referenced by xMessageBuffer is empty then pdTRUE is returned. Otherwise pdFALSE is returned.

xMessageBufferReset (xMessageBuffer)

Resets a message buffer to its initial empty state, discarding any message it contained.

A message buffer can only be reset if there are no tasks blocked on it.

Parameters

- **xMessageBuffer** -- The handle of the message buffer being reset.

Returns If the message buffer was reset then pdPASS is returned. If the message buffer could not be reset because either there was a task blocked on the message queue to wait for space to become available, or to wait for a message to be available, then pdFAIL is returned.

xMessageBufferSpaceAvailable (xMessageBuffer)

message_buffer.h

```
size_t xMessageBufferSpaceAvailable( MessageBufferHandle_t xMessageBuffer );
```

Returns the number of bytes of free space in the message buffer.

Parameters

- **xMessageBuffer** -- The handle of the message buffer being queried.

Returns The number of bytes that can be written to the message buffer before the message buffer would be full. When a message is written to the message buffer an additional sizeof(size_t) bytes are also written to store the message's length. sizeof(size_t) is typically 4 bytes on a 32-bit architecture, so if xMessageBufferSpacesAvailable() returns 10, then the size of the largest message that can be written to the message buffer is 6 bytes.

xMessageBufferSpacesAvailable (xMessageBuffer)

xMessageBufferNextLengthBytes (xMessageBuffer)

Returns the length (in bytes) of the next message in a message buffer. Useful if xMessageBufferReceive() returned 0 because the size of the buffer passed into xMessageBufferReceive() was too small to hold the next message.

Parameters

- **xMessageBuffer** -- The handle of the message buffer being queried.

Returns The length (in bytes) of the next message in the message buffer, or 0 if the message buffer is empty.

xMessageBufferSendCompletedFromISR (xMessageBuffer, pxHigherPriorityTaskWoken)

For advanced users only.

The `sbSEND_COMPLETED()` macro is called from within the FreeRTOS APIs when data is sent to a message buffer or stream buffer. If there was a task that was blocked on the message or stream buffer waiting for data to arrive then the `sbSEND_COMPLETED()` macro sends a notification to the task to remove it from the Blocked state. `xMessageBufferSendCompletedFromISR()` does the same thing. It is provided to enable application writers to implement their own version of `sbSEND_COMPLETED()`, and **MUST NOT BE USED AT ANY OTHER TIME**.

See the example implemented in `FreeRTOS/Demo/Minimal/MessageBufferAMP.c` for additional information.

Parameters

- **xMessageBuffer** -- The handle of the stream buffer to which data was written.
- **pxHigherPriorityTaskWoken** -- *pxHigherPriorityTaskWoken should be initialised to `pdFALSE` before it is passed into `xMessageBufferSendCompletedFromISR()`. If calling `xMessageBufferSendCompletedFromISR()` removes a task from the Blocked state, and the task has a priority above the priority of the currently running task, then *pxHigherPriorityTaskWoken will get set to `pdTRUE` indicating that a context switch should be performed before exiting the ISR.

Returns If a task was removed from the Blocked state then `pdTRUE` is returned. Otherwise `pdFALSE` is returned.

xMessageBufferReceiveCompletedFromISR (xMessageBuffer, pxHigherPriorityTaskWoken)

For advanced users only.

The `sbRECEIVE_COMPLETED()` macro is called from within the FreeRTOS APIs when data is read out of a message buffer or stream buffer. If there was a task that was blocked on the message or stream buffer waiting for data to arrive then the `sbRECEIVE_COMPLETED()` macro sends a notification to the task to remove it from the Blocked state. `xMessageBufferReceiveCompletedFromISR()` does the same thing. It is provided to enable application writers to implement their own version of `sbRECEIVE_COMPLETED()`, and **MUST NOT BE USED AT ANY OTHER TIME**.

See the example implemented in `FreeRTOS/Demo/Minimal/MessageBufferAMP.c` for additional information.

Parameters

- **xMessageBuffer** -- The handle of the stream buffer from which data was read.
- **pxHigherPriorityTaskWoken** -- *pxHigherPriorityTaskWoken should be initialised to `pdFALSE` before it is passed into `xMessageBufferReceiveCompletedFromISR()`. If calling `xMessageBufferReceiveCompletedFromISR()` removes a task from the Blocked state, and the task has a priority above the priority of the currently running task, then *pxHigherPriorityTaskWoken will get set to `pdTRUE` indicating that a context switch should be performed before exiting the ISR.

Returns If a task was removed from the Blocked state then `pdTRUE` is returned. Otherwise `pdFALSE` is returned.

Type Definitions

typedef *StreamBufferHandle_t* **MessageBufferHandle_t**

Type by which message buffers are referenced. For example, a call to `xMessageBufferCreate()` returns an `MessageBufferHandle_t` variable that can then be used as a parameter to `xMessageBufferSend()`, `xMessageBufferReceive()`, etc. Message buffer is essentially built as a stream buffer hence its handle is also set to same type as a stream buffer handle.

2.9.13 FreeRTOS (Supplemental Features)

ESP-IDF provides multiple features to supplement the features offered by FreeRTOS. These supplemental features are available on all FreeRTOS implementations supported by ESP-IDF (i.e., ESP-IDF FreeRTOS and Amazon SMP FreeRTOS). This document describes these supplemental features and is split into the following sections:

Contents

- *FreeRTOS (Supplemental Features)*
 - *Overview*
 - *Ring Buffers*
 - *ESP-IDF Tick and Idle Hooks*
 - *TLSP Deletion Callbacks*
 - *IDF Additional API*
 - *Component Specific Properties*
 - *API Reference*

Overview

ESP-IDF adds various new features to supplement the capabilities of FreeRTOS as follows:

- **Ring buffers:** Ring buffers provide a FIFO buffer that can accept entries of arbitrary lengths.
- **ESP-IDF Tick and Idle Hooks:** ESP-IDF provides multiple custom tick interrupt hooks and idle task hooks that are more numerous and more flexible when compared to FreeRTOS tick and idle hooks.
- **Thread Local Storage Pointer (TLSP) Deletion Callbacks:** TLSP Deletion callbacks are run automatically when a task is deleted, thus allowing users to clean up their TLSPs automatically.
- **IDF Additional API:** ESP-IDF specific functions added to augment the features of FreeRTOS.
- **Component Specific Properties:** Currently added only one component specific property `ORIG_INCLUDE_PATH`.

Ring Buffers

FreeRTOS provides stream buffers and message buffers as the primary mechanisms to send arbitrarily sized data between tasks and ISRs. However, FreeRTOS stream buffers and message buffers have the following limitations:

- Strictly single sender and single receiver
- Data is passed by copy
- Unable to reserve buffer space for a deferred send (i.e., send acquire)

Therefore, ESP-IDF provides a separate ring buffer implementation to address the issues above.

ESP-IDF ring buffers are strictly FIFO buffers that supports arbitrarily sized items. Ring buffers are a more memory efficient alternative to FreeRTOS queues in situations where the size of items is variable. The capacity of a ring buffer is not measured by the number of items it can store, but rather by the amount of memory used for storing items.

The ring buffer provides APIs to send an item, or to allocate space for an item in the ring buffer to be filled manually by the user. For efficiency reasons, **items are always retrieved from the ring buffer by reference**. As a result, all retrieved items **must also be returned** to the ring buffer by using `vRingbufferReturnItem()` or `vRingbufferReturnItemFromISR()`, in order for them to be removed from the ring buffer completely.

The ring buffers are split into the three following types:

No-Split buffers guarantee that an item is stored in contiguous memory and does not attempt to split an item under any circumstances. Use No-Split buffers when items must occupy contiguous memory. **Only this buffer type allows reserving buffer space for deferred sending.** Refer to the documentation of the functions `xRingbufferSendAcquire()` and `xRingbufferSendComplete()` for more details.

Allow-Split buffers allow an item to be split in two parts when wrapping around the end of the buffer if there is enough space at the tail and the head of the buffer combined to store the item. Allow-Split buffers are more memory efficient than No-Split buffers but can return an item in two parts when retrieving.

Byte buffers do not store data as separate items. All data is stored as a sequence of bytes, and any number of bytes can be sent or retrieved each time. Use byte buffers when separate items do not need to be maintained, e.g., a byte stream.

Note: No-Split buffers and Allow-Split buffers always store items at 32-bit aligned addresses. Therefore, when retrieving an item, the item pointer is guaranteed to be 32-bit aligned. This is useful especially when you need to send some data to the DMA.

Note: Each item stored in No-Split or Allow-Split buffers **requires an additional 8 bytes for a header**. Item sizes are also rounded up to a 32-bit aligned size, i.e., multiple of 4 bytes. However the true item size is recorded within the header. The sizes of No-Split and Allow-Split buffers will also be rounded up when created.

Usage The following example demonstrates the usage of `xRingbufferCreate()` and `xRingbufferSend()` to create a ring buffer and then send an item to it:

```
#include "freertos/ringbuf.h"
static char tx_item[] = "test_item";

...

//Create ring buffer
RingbufHandle_t buf_handle;
buf_handle = xRingbufferCreate(1028, RINGBUF_TYPE_NOSPLIT);
if (buf_handle == NULL) {
    printf("Failed to create ring buffer\n");
}

//Send an item
UBaseType_t res = xRingbufferSend(buf_handle, tx_item, sizeof(tx_item), pdMS_
↪TO_TICKS(1000));
if (res != pdTRUE) {
    printf("Failed to send item\n");
}
```

The following example demonstrates the usage of `xRingbufferSendAcquire()` and `xRingbufferSendComplete()` instead of `xRingbufferSend()` to acquire memory on the ring buffer (of type `RINGBUF_TYPE_NOSPLIT`) and then send an item to it. This adds one more step, but allows getting the address of the memory to write to, and writing to the memory yourself.

```
#include "freertos/ringbuf.h"
#include "soc/lldesc.h"

typedef struct {
    lldesc_t dma_desc;
    uint8_t buf[1];
} dma_item_t;

#define DMA_ITEM_SIZE(N) (sizeof(lldesc_t)+((N)+3)&(~3))

...

//Retrieve space for DMA descriptor and corresponding data buffer
//This has to be done with SendAcquire, or the address may be different when
↪we copy
dma_item_t *item;
UBaseType_t res = xRingbufferSendAcquire(buf_handle,
    (void**) &item, DMA_ITEM_SIZE(buffer_size), pdMS_TO_
↪TICKS(1000));
```

(continues on next page)

(continued from previous page)

```

if (res != pdTRUE) {
    printf("Failed to acquire memory for item\n");
}
item->dma_desc = (lldesc_t) {
    .size = buffer_size,
    .length = buffer_size,
    .eof = 0,
    .owner = 1,
    .buf = item->buf,
};
//Actually send to the ring buffer for consumer to use
res = xRingbufferSendComplete(buf_handle, &item);
if (res != pdTRUE) {
    printf("Failed to send item\n");
}

```

The following example demonstrates retrieving and returning an item from a **No-Split ring buffer** using *xRingbufferReceive()* and *vRingbufferReturnItem()*

```

...

//Receive an item from no-split ring buffer
size_t item_size;
char *item = (char *)xRingbufferReceive(buf_handle, &item_size, pdMS_TO_
↪TICKS(1000));

//Check received item
if (item != NULL) {
    //Print item
    for (int i = 0; i < item_size; i++) {
        printf("%c", item[i]);
    }
    printf("\n");
    //Return Item
    vRingbufferReturnItem(buf_handle, (void *)item);
} else {
    //Failed to receive item
    printf("Failed to receive item\n");
}

```

The following example demonstrates retrieving and returning an item from an **Allow-Split ring buffer** using *xRingbufferReceiveSplit()* and *vRingbufferReturnItem()*

```

...

//Receive an item from allow-split ring buffer
size_t item_size1, item_size2;
char *item1, *item2;
BaseType_t ret = xRingbufferReceiveSplit(buf_handle, (void **)&item1, (void
↪**)&item2, &item_size1, &item_size2, pdMS_TO_TICKS(1000));

//Check received item
if (ret == pdTRUE && item1 != NULL) {
    for (int i = 0; i < item_size1; i++) {
        printf("%c", item1[i]);
    }
    vRingbufferReturnItem(buf_handle, (void *)item1);
    //Check if item was split
    if (item2 != NULL) {
        for (int i = 0; i < item_size2; i++) {
            printf("%c", item2[i]);
        }
    }
}

```

(continues on next page)

(continued from previous page)

```

    }
    vRingbufferReturnItem(buf_handle, (void *)item2);
}
printf("\n");
} else {
    //Failed to receive item
    printf("Failed to receive item\n");
}
}

```

The following example demonstrates retrieving and returning an item from a **byte buffer** using *xRingbufferReceiveUpTo()* and *vRingbufferReturnItem()*

```

...

//Receive data from byte buffer
size_t item_size;
char *item = (char *)xRingbufferReceiveUpTo(buf_handle, &item_size, pdMS_TO_
↪TICKS(1000), sizeof(tx_item));

//Check received data
if (item != NULL) {
    //Print item
    for (int i = 0; i < item_size; i++) {
        printf("%c", item[i]);
    }
    printf("\n");
    //Return Item
    vRingbufferReturnItem(buf_handle, (void *)item);
} else {
    //Failed to receive item
    printf("Failed to receive item\n");
}

```

For ISR safe versions of the functions used above, call *xRingbufferSendFromISR()*, *xRingbufferReceiveFromISR()*, *xRingbufferReceiveSplitFromISR()*, *xRingbufferReceiveUpToFromISR()*, and *vRingbufferReturnItemFromISR()*.

Note: Two calls to *RingbufferReceive[UpTo][FromISR]()* are required if the bytes wraps around the end of the ring buffer.

Sending to Ring Buffer The following diagrams illustrate the differences between No-Split and Allow-Split buffers as compared to byte buffers with regard to sending items or data. The diagrams assume that three items of sizes **18**, **3**, and **27** bytes are sent respectively to a **buffer of 128 bytes**:




Fig. 24: Sending items to No-Split or Allow-Split ring buffers

For No-Split and Allow-Split buffers, a header of 8 bytes precedes every data item. Furthermore, the space occupied by each item is **rounded up to the nearest 32-bit aligned size** in order to maintain overall 32-bit alignment.

However, the true size of the item is recorded inside the header which will be returned when the item is retrieved.

Referring to the diagram above, the 18, 3, and 27 byte items are **rounded up to 20, 4, and 28 bytes** respectively. An 8 byte header is then added in front of each item.




Fig. 25: Sending items to byte buffers

Byte buffers treat data as a sequence of bytes and does not incur any overhead (no headers). As a result, all data sent to a byte buffer is merged into a single item.

Referring to the diagram above, the 18, 3, and 27 byte items are sequentially written to the byte buffer and **merged into a single item of 48 bytes**.

Using `SendAcquire` and `SendComplete` Items in No-Split buffers are acquired (by `SendAcquire`) in strict FIFO order and must be sent to the buffer by `SendComplete` for the data to be accessible by the consumer. Multiple items can be sent or acquired without calling `SendComplete`, and the items do not necessarily need to be completed in the order they were acquired. However, the receiving of data items must occur in FIFO order, therefore not calling `SendComplete` for the earliest acquired item prevents the subsequent items from being received.

The following diagrams illustrate what will happen when `SendAcquire` and `SendComplete` do not happen in the same order. At the beginning, there is already a data item of 16 bytes sent to the ring buffer. Then `SendAcquire` is called to acquire space of 20, 8, 24 bytes on the ring buffer.




Fig. 26: `SendAcquire`/`SendComplete` items in No-Split ring buffers

After that, we fill (use) the buffers, and send them to the ring buffer by `SendComplete` in the order of 8, 24, 20. When 8 bytes and 24 bytes data are sent, the consumer still can only get the 16 bytes data item. Hence, if `SendComplete` is not called for the 20 bytes, it will not be available, nor will the data items following the 20 bytes item.

When the 20 bytes item is finally completed, all the 3 data items can be received now, in the order of 20, 8, 24 bytes, right after the 16 bytes item existing in the buffer at the beginning.

Allow-Split buffers and byte buffers do not allow using `SendAcquire` or `SendComplete` since acquired buffers are required to be complete (not wrapped).

Wrap Around The following diagrams illustrate the differences between No-Split, Allow-Split, and byte buffers when a sent item requires a wrap around. The diagrams assume a buffer of **128 bytes** with **56 bytes of free space that wraps around** and a sent item of **28 bytes**.




Fig. 27: Wrap around in No-Split buffers

No-Split buffers **only store an item in continuous free space and do not split an item under any circumstances**. When the free space at the tail of the buffer is insufficient to completely store the item and its header, the free space at the tail will be **marked as dummy data**. The buffer will then wrap around and store the item in the free space at the head of the buffer.

Referring to the diagram above, the 16 bytes of free space at the tail of the buffer is insufficient to store the 28 byte item. Therefore, the 16 bytes is marked as dummy data and the item is written to the free space at the head of the buffer instead.




Fig. 28: Wrap around in Allow-Split buffers

Allow-Split buffers will attempt to **split the item into two parts** when the free space at the tail of the buffer is insufficient to store the item data and its header. Both parts of the split item will have their own headers, therefore incurring an extra 8 bytes of overhead.

Referring to the diagram above, the 16 bytes of free space at the tail of the buffer is insufficient to store the 28 byte item. Therefore, the item is split into two parts (8 and 20 bytes) and written as two parts to the buffer.

Note: Allow-Split buffers treat both parts of the split item as two separate items, therefore call `xRingbufferReceiveSplit()` instead of `xRingbufferReceive()` to receive both parts of a split item in a thread safe manner.




Fig. 29: Wrap around in byte buffers

Byte buffers **store as much data as possible into the free space at the tail of buffer**. The remaining data will then be stored in the free space at the head of the buffer. No overhead is incurred when wrapping around in byte buffers.

Referring to the diagram above, the 16 bytes of free space at the tail of the buffer is insufficient to completely store the 28 bytes of data. Therefore, the 16 bytes of free space is filled with data, and the remaining 12 bytes are written

to the free space at the head of the buffer. The buffer now contains data in two separate continuous parts, and each continuous part is treated as a separate item by the byte buffer.

Retrieving/Returning The following diagrams illustrate the differences between No-Split and Allow-Split buffers as compared to byte buffers in retrieving and returning data:




Fig. 30: Retrieving/Returning items in No-Split and Allow-Split ring buffers

Items in No-Split buffers and Allow-Split buffers are **retrieved in strict FIFO order** and **must be returned** for the occupied space to be freed. Multiple items can be retrieved before returning, and the items do not necessarily need to be returned in the order they were retrieved. However, the freeing of space must occur in FIFO order, therefore not returning the earliest retrieved item prevents the space of subsequent items from being freed.

Referring to the diagram above, the **16, 20, and 8 byte items are retrieved in FIFO order**. However, the items are not returned in the order they were retrieved. First, the 20 byte item is returned followed by the 8 byte and the 16 byte items. The space is not freed until the first item, i.e., the 16 byte item is returned.




Fig. 31: Retrieving/Returning data in byte buffers

Byte buffers **do not allow multiple retrievals before returning** (every retrieval must be followed by a return before another retrieval is permitted). When using `xRingbufferReceive()` or `xRingbufferReceiveFromISR()`, all continuous stored data will be retrieved. `xRingbufferReceiveUpTo()` or `xRingbufferReceiveUpToFromISR()` can be used to restrict the maximum number of bytes retrieved. Since every retrieval must be followed by a return, the space is freed as soon as the data is returned.

Referring to the diagram above, the 38 bytes of continuous stored data at the tail of the buffer is retrieved, returned, and freed. The next call to `xRingbufferReceive()` or `xRingbufferReceiveFromISR()` then wraps around and does the same to the 30 bytes of continuous stored data at the head of the buffer.

Note: Retrieving items from Allow-Split buffers must be done via `xRingbufferReceiveSplit()` or `xRingbufferReceiveSplitFromISR()` instead of `xRingbufferReceive()` or `xRingbufferReceiveFromISR()`.

Ring Buffers with Queue Sets Ring buffers can be added to FreeRTOS queue sets using `xRingbufferAddToQueueSetRead()` such that every time a ring buffer receives an item or data, the queue set is notified. Once added to a queue set, every attempt to retrieve an item from a ring buffer should be preceded by a call to `xQueueSelectFromSet()`. To check whether the selected queue set member is the ring buffer, call `xRingbufferCanRead()`.

The following example demonstrates queue set usage with ring buffers:

```
#include "freertos/queue.h"
#include "freertos/ringbuf.h"

...

//Create ring buffer and queue set
RingbufHandle_t buf_handle = xRingbufferCreate(1028, RINGBUF_TYPE_NOSPLIT);
QueueSetHandle_t queue_set = xQueueCreateSet(3);

//Add ring buffer to queue set
if (xRingbufferAddToQueueSetRead(buf_handle, queue_set) != pdTRUE) {
    printf("Failed to add to queue set\n");
}

...

//Block on queue set
QueueSetMemberHandle_t member = xQueueSelectFromSet(queue_set, pdMS_TO_
↳TICKS(1000));

//Check if member is ring buffer
if (member != NULL && xRingbufferCanRead(buf_handle, member) == pdTRUE) {
    //Member is ring buffer, receive item from ring buffer
    size_t item_size;
    char *item = (char *)xRingbufferReceive(buf_handle, &item_size, 0);

    //Handle item
    ...
} else {
    ...
}
```

Ring Buffers with Static Allocation The `xRingbufferCreateStatic()` can be used to create ring buffers with specific memory requirements (such as a ring buffer being allocated in external RAM). All blocks of memory used by a ring buffer must be manually allocated beforehand, then passed to the `xRingbufferCreateStatic()` to be initialized as a ring buffer. These blocks include the following:

- The ring buffer's data structure of type `StaticRingbuffer_t`.
- The ring buffer's storage area of size `xBufferSize`. Note that `xBufferSize` must be 32-bit aligned for No-Split and Allow-Split buffers.

The manner in which these blocks are allocated depends on the users requirements (e.g., all blocks being statically declared, or dynamically allocated with specific capabilities such as external RAM).

Note: When deleting a ring buffer created via `xRingbufferCreateStatic()`, the function `vRingbufferDelete()` will not free any of the memory blocks. This must be done manually by the user after `vRingbufferDelete()` is called.

The code snippet below demonstrates a ring buffer being allocated entirely in external RAM.

```

#include "freertos/ringbuf.h"
#include "freertos/semphr.h"
#include "esp_heap_caps.h"

#define BUFFER_SIZE      400          //32-bit aligned size
#define BUFFER_TYPE      RINGBUF_TYPE_NOSPLIT
...

//Allocate ring buffer data structure and storage area into external RAM
StaticRingbuffer_t *buffer_struct = (StaticRingbuffer_t *)heap_caps_
↳malloc(sizeof(StaticRingbuffer_t), MALLOC_CAP_SPIRAM);
uint8_t *buffer_storage = (uint8_t *)heap_caps_malloc(sizeof(uint8_t)*BUFFER_SIZE,↳
↳MALLOC_CAP_SPIRAM);

//Create a ring buffer with manually allocated memory
RingbufHandle_t handle = xRingbufferCreateStatic(BUFFER_SIZE, BUFFER_TYPE, buffer_
↳storage, buffer_struct);

...

//Delete the ring buffer after used
vRingbufferDelete(handle);

//Manually free all blocks of memory
free(buffer_struct);
free(buffer_storage);

```

ESP-IDF Tick and Idle Hooks

FreeRTOS allows applications to provide a tick hook and an idle hook at compile time:

- FreeRTOS tick hook can be enabled via the [CONFIG_FREERTOS_USE_TICK_HOOK](#) option. The application must provide the void `vApplicationTickHook( void )` callback.
- FreeRTOS idle hook can be enabled via the [CONFIG_FREERTOS_USE_IDLE_HOOK](#) option. The application must provide the void `vApplicationIdleHook( void )` callback.

However, the FreeRTOS tick hook and idle hook have the following drawbacks:

- The FreeRTOS hooks are registered at compile time
- Only one of each hook can be registered
- On multi-core targets, the FreeRTOS hooks are symmetric, meaning each core's tick interrupt and idle tasks ends up calling the same hook

Therefore, ESP-IDF tick and idle hooks are provided to supplement the features of FreeRTOS tick and idle hooks. The ESP-IDF hooks have the following features:

- The hooks can be registered and deregistered at run-time
- Multiple hooks can be registered (with a maximum of 8 hooks of each type per core)
- On multi-core targets, the hooks can be asymmetric, meaning different hooks can be registered to each core

ESP-IDF hooks can be registered and deregistered using the following APIs:

- For tick hooks:
 - Register using `esp_register_freertos_tick_hook()` or `esp_register_freertos_tick_hook_for_cpu()`
 - Deregister using `esp_deregister_freertos_tick_hook()` or `esp_deregister_freertos_tick_hook_for_cpu()`
- For idle hooks:
 - Register using `esp_register_freertos_idle_hook()` or `esp_register_freertos_idle_hook_for_cpu()`
 - Deregister using `esp_deregister_freertos_idle_hook()` or `esp_deregister_freertos_idle_hook_for_cpu()`

Note: The tick interrupt stays active while the cache is disabled, therefore any tick hook (FreeRTOS or ESP-IDF) functions must be placed in internal RAM. Please refer to the [SPI flash API documentation](#) for more details.

TLSP Deletion Callbacks

Vanilla FreeRTOS provides a Thread Local Storage Pointers (TLSP) feature. These are pointers stored directly in the Task Control Block (TCB) of a particular task. TLSPs allow each task to have its own unique set of pointers to data structures. Vanilla FreeRTOS expects users to:

- set a task's TLSPs by calling `vTaskSetThreadLocalStoragePointer()` after the task has been created.
- get a task's TLSPs by calling `pvTaskGetThreadLocalStoragePointer()` during the task's lifetime.
- free the memory pointed to by the TLSPs before the task is deleted.

However, there can be instances where users may want the freeing of TLSP memory to be automatic. Therefore, ESP-IDF provides the additional feature of TLSP deletion callbacks. These user-provided deletion callbacks are called automatically when a task is deleted, thus allowing the TLSP memory to be cleaned up without needing to add the cleanup logic explicitly to the code of every task.

The TLSP deletion callbacks are set in a similar fashion to the TLSPs themselves.

- `vTaskSetThreadLocalStoragePointerAndDelCallback()` sets both a particular TLSP and its associated callback.
- Calling the Vanilla FreeRTOS function `vTaskSetThreadLocalStoragePointer()` simply sets the TLSP's associated Deletion Callback to `NULL`, meaning that no callback is called for that TLSP during task deletion.

When implementing TLSP callbacks, users should note the following:

- The callback **must never attempt to block or yield** and critical sections should be kept as short as possible.
- The callback is called shortly before a deleted task's memory is freed. Thus, the callback can either be called from `vTaskDelete()` itself, or from the idle task.

IDF Additional API

The `freertos/esp_additions/include/freertos/idf_additions.h` header contains FreeRTOS-related helper functions added by ESP-IDF. Users can include this header via `#include "freertos/idf_additions.h"`.

Component Specific Properties

Besides standard component variables that are available with basic cmake build properties, FreeRTOS component also provides arguments (only one so far) for simpler integration with other modules:

- `ORIG_INCLUDE_PATH` - contains an absolute path to freertos root include folder. Thus instead of `#include "freertos/FreeRTOS.h"` you can refer to headers directly: `#include "FreeRTOS.h"`.

API Reference

Ring Buffer API

Header File

- `components/esp_ringbuf/include/freertos/ringbuf.h`
- This header file can be included with:

```
#include "freertos/ringbuf.h"
```

- This header file is a part of the API provided by the `esp_ringbuf` component. To declare that your component depends on `esp_ringbuf`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_ringbuf
```

or

```
PRIV_REQUIRES esp_ringbuf
```

Functions

RingbufHandle_t **xRingbufferCreate** (size_t xBufferSize, *RingbufferType_t* xBufferType)

Create a ring buffer.

Note: xBufferSize of no-split/allow-split buffers will be rounded up to the nearest 32-bit aligned size.

Parameters

- **xBufferSize** -- [in] Size of the buffer in bytes. Note that items require space for a header in no-split/allow-split buffers
- **xBufferType** -- [in] Type of ring buffer, see documentation.

Returns A handle to the created ring buffer, or NULL in case of error.

RingbufHandle_t **xRingbufferCreateNoSplit** (size_t xItemSize, size_t xItemNum)

Create a ring buffer of type `RINGBUF_TYPE_NOSPLIT` for a fixed item_size.

This API is similar to `xRingbufferCreate()`, but it will internally allocate additional space for the headers.

Parameters

- **xItemSize** -- [in] Size of each item to be put into the ring buffer
- **xItemNum** -- [in] Maximum number of items the buffer needs to hold simultaneously

Returns A *RingbufHandle_t* handle to the created ring buffer, or NULL in case of error.

RingbufHandle_t **xRingbufferCreateStatic** (size_t xBufferSize, *RingbufferType_t* xBufferType, uint8_t *pucRingbufferStorage, *StaticRingbuffer_t* *pxStaticRingbuffer)

Create a ring buffer but manually provide the required memory.

Note: xBufferSize of no-split/allow-split buffers MUST be 32-bit aligned.

Parameters

- **xBufferSize** -- [in] Size of the buffer in bytes.
- **xBufferType** -- [in] Type of ring buffer, see documentation
- **pucRingbufferStorage** -- [in] Pointer to the ring buffer's storage area. Storage area must have the same size as specified by xBufferSize
- **pxStaticRingbuffer** -- [in] Pointed to a struct of type *StaticRingbuffer_t* which will be used to hold the ring buffer's data structure

Returns A handle to the created ring buffer

BaseType_t **xRingbufferSend** (*RingbufHandle_t* xRingbuffer, const void *pvItem, size_t xItemSize, TickType_t xTicksToWait)

Insert an item into the ring buffer.

Attempt to insert an item into the ring buffer. This function will block until enough free space is available or until it times out.

Note: For no-split/allow-split ring buffers, the actual size of memory that the item will occupy will be rounded up to the nearest 32-bit aligned size. This is done to ensure all items are always stored in 32-bit aligned fashion.

Note: For no-split/allow-split buffers, an `xItemSize` of 0 will result in an item with no data being set (i.e., item only contains the header). For byte buffers, an `xItemSize` of 0 will simply return `pdTRUE` without copying any data.

Parameters

- **xRingbuffer** -- [in] Ring buffer to insert the item into
- **pvItem** -- [in] Pointer to data to insert. NULL is allowed if `xItemSize` is 0.
- **xItemSize** -- [in] Size of data to insert.
- **xTicksToWait** -- [in] Ticks to wait for room in the ring buffer.

Returns

- `pdTRUE` if succeeded
- `pdFALSE` on time-out or when the data is larger than the maximum permissible size of the buffer

BaseType_t **xRingbufferSendFromISR** (*RingbufHandle_t* xRingbuffer, const void *pvItem, size_t xItemSize, BaseType_t *pxHigherPriorityTaskWoken)

Insert an item into the ring buffer in an ISR.

Attempt to insert an item into the ring buffer from an ISR. This function will return immediately if there is insufficient free space in the buffer.

Note: For no-split/allow-split ring buffers, the actual size of memory that the item will occupy will be rounded up to the nearest 32-bit aligned size. This is done to ensure all items are always stored in 32-bit aligned fashion.

Note: For no-split/allow-split buffers, an `xItemSize` of 0 will result in an item with no data being set (i.e., item only contains the header). For byte buffers, an `xItemSize` of 0 will simply return `pdTRUE` without copying any data.

Parameters

- **xRingbuffer** -- [in] Ring buffer to insert the item into
- **pvItem** -- [in] Pointer to data to insert. NULL is allowed if `xItemSize` is 0.
- **xItemSize** -- [in] Size of data to insert.
- **pxHigherPriorityTaskWoken** -- [out] Value pointed to will be set to `pdTRUE` if the function woke up a higher priority task.

Returns

- `pdTRUE` if succeeded
- `pdFALSE` when the ring buffer does not have space.

BaseType_t **xRingbufferSendAcquire** (*RingbufHandle_t* xRingbuffer, void **ppvItem, size_t xItemSize, TickType_t xTicksToWait)

Acquire memory from the ring buffer to be written to by an external source and to be sent later.

Attempt to allocate buffer for an item to be sent into the ring buffer. This function will block until enough free space is available or until it times out.

The item, as well as the following items `SendAcquire` or `Send` after it, will not be able to be read from the ring buffer until this item is actually sent into the ring buffer.

Note: Only applicable for no-split ring buffers now, the actual size of memory that the item will occupy will

be rounded up to the nearest 32-bit aligned size. This is done to ensure all items are always stored in 32-bit aligned fashion.

Note: An `xItemSize` of 0 will result in a buffer being acquired, but the buffer will have a size of 0.

Parameters

- **xRingbuffer** -- [in] Ring buffer to allocate the memory
- **ppvItem** -- [out] Double pointer to memory acquired (set to NULL if no memory were retrieved)
- **xItemSize** -- [in] Size of item to acquire.
- **xTicksToWait** -- [in] Ticks to wait for room in the ring buffer.

Returns

- `pdTRUE` if succeeded
- `pdFALSE` on time-out or when the data is larger than the maximum permissible size of the buffer

BaseType_t **xRingbufferSendComplete** (*RingbufHandle_t* xRingbuffer, void *pvItem)

Actually send an item into the ring buffer allocated before by `xRingbufferSendAcquire`.

Note: Only applicable for no-split ring buffers. Only call for items allocated by `xRingbufferSendAcquire`.

Parameters

- **xRingbuffer** -- [in] Ring buffer to insert the item into
- **pvItem** -- [in] Pointer to item in allocated memory to insert.

Returns

- `pdTRUE` if succeeded
- `pdFALSE` if fail for some reason.

void ***xRingbufferReceive** (*RingbufHandle_t* xRingbuffer, size_t *pxItemSize, TickType_t xTicksToWait)

Retrieve an item from the ring buffer.

Attempt to retrieve an item from the ring buffer. This function will block until an item is available or until it times out.

Note: A call to `vRingbufferReturnItem()` is required after this to free the item retrieved.

Note: It is possible to receive items with a `pxItemSize` of 0 on no-split buffers.

Note: To retrieve an item from an allow-split buffer, use `xRingbufferReceiveSplit()` instead.

Parameters

- **xRingbuffer** -- [in] Ring buffer to retrieve the item from
- **pxItemSize** -- [out] Pointer to a variable to which the size of the retrieved item will be written.
- **xTicksToWait** -- [in] Ticks to wait for items in the ring buffer.

Returns

- Pointer to the retrieved item on success; `*pxItemSize` filled with the length of the item.
- NULL on timeout, `*pxItemSize` is untouched in that case.

void ***xRingbufferReceiveFromISR** (*RingbufHandle_t* xRingbuffer, size_t *pxItemSize)

Retrieve an item from the ring buffer in an ISR.

Attempt to retrieve an item from the ring buffer. This function returns immediately if there are no items available for retrieval

Note: A call to `vRingbufferReturnItemFromISR()` is required after this to free the item retrieved.

Note: Byte buffers do not allow multiple retrievals before returning an item

Note: Two calls to `RingbufferReceiveFromISR()` are required if the bytes wrap around the end of the ring buffer.

Note: It is possible to receive items with a `pxItemSize` of 0 on no-split buffers.

Note: To retrieve an item from an allow-split buffer, use `xRingbufferReceiveSplitFromISR()` instead.

Parameters

- **xRingbuffer** -- [in] Ring buffer to retrieve the item from
- **pxItemSize** -- [out] Pointer to a variable to which the size of the retrieved item will be written.

Returns

- Pointer to the retrieved item on success; `*pxItemSize` filled with the length of the item.
- NULL when the ring buffer is empty, `*pxItemSize` is untouched in that case.

BaseType_t **xRingbufferReceiveSplit** (*RingbufHandle_t* xRingbuffer, void **ppvHeadItem, void **ppvTailItem, size_t *pxHeadItemSize, size_t *pxTailItemSize, TickType_t xTicksToWait)

Retrieve a split item from an allow-split ring buffer.

Attempt to retrieve a split item from an allow-split ring buffer. If the item is not split, only a single item is retrieved. If the item is split, both parts will be retrieved. This function will block until an item is available or until it times out.

Note: Call(s) to `vRingbufferReturnItem()` is required after this to free up the item(s) retrieved.

Note: This function should only be called on allow-split buffers

Note: It is possible to receive items with a `pxItemSize` of 0 on allow split buffers.

Parameters

- **xRingbuffer** -- [in] Ring buffer to retrieve the item from
- **ppvHeadItem** -- [out] Double pointer to first part (set to NULL if no items were retrieved)
- **ppvTailItem** -- [out] Double pointer to second part (set to NULL if item is not split)

- **pxHeadItemSize** -- [out] Pointer to size of first part (unmodified if no items were retrieved)
- **pxTailItemSize** -- [out] Pointer to size of second part (unmodified if item is not split)
- **xTicksToWait** -- [in] Ticks to wait for items in the ring buffer.

Returns

- pdTRUE if an item (split or unsplit) was retrieved
- pdFALSE when no item was retrieved

BaseType_t **xRingbufferReceiveSplitFromISR** (*RingbufHandle_t* xRingbuffer, void **ppvHeadItem, void **ppvTailItem, size_t *pxHeadItemSize, size_t *pxTailItemSize)

Retrieve a split item from an allow-split ring buffer in an ISR.

Attempt to retrieve a split item from an allow-split ring buffer. If the item is not split, only a single item is retrieved. If the item is split, both parts will be retrieved. This function returns immediately if there are no items available for retrieval

Note: Calls to `vRingbufferReturnItemFromISR()` is required after this to free up the item(s) retrieved.

Note: This function should only be called on allow-split buffers

Note: It is possible to receive items with a `pxItemSize` of 0 on allow split buffers.

Parameters

- **xRingbuffer** -- [in] Ring buffer to retrieve the item from
- **ppvHeadItem** -- [out] Double pointer to first part (set to NULL if no items were retrieved)
- **ppvTailItem** -- [out] Double pointer to second part (set to NULL if item is not split)
- **pxHeadItemSize** -- [out] Pointer to size of first part (unmodified if no items were retrieved)
- **pxTailItemSize** -- [out] Pointer to size of second part (unmodified if item is not split)

Returns

- pdTRUE if an item (split or unsplit) was retrieved
- pdFALSE when no item was retrieved

void ***xRingbufferReceiveUpTo** (*RingbufHandle_t* xRingbuffer, size_t *pxItemSize, TickType_t xTicksToWait, size_t xMaxSize)

Retrieve bytes from a byte buffer, specifying the maximum amount of bytes to retrieve.

Attempt to retrieve data from a byte buffer whilst specifying a maximum number of bytes to retrieve. This function will block until there is data available for retrieval or until it times out.

Note: A call to `vRingbufferReturnItem()` is required after this to free up the data retrieved.

Note: This function should only be called on byte buffers

Note: Byte buffers do not allow multiple retrievals before returning an item

Note: Two calls to `RingbufferReceiveUpTo()` are required if the bytes wrap around the end of the ring buffer.

Parameters

- **xRingbuffer** -- [in] Ring buffer to retrieve the item from
- **pxItemSize** -- [out] Pointer to a variable to which the size of the retrieved item will be written.
- **xTicksToWait** -- [in] Ticks to wait for items in the ring buffer.
- **xMaxSize** -- [in] Maximum number of bytes to return.

Returns

- Pointer to the retrieved item on success; `*pxItemSize` filled with the length of the item.
- NULL on timeout, `*pxItemSize` is untouched in that case.

void ***xRingbufferReceiveUpToFromISR** (*RingbufHandle_t* xRingbuffer, size_t *pxItemSize, size_t xMaxSize)

Retrieve bytes from a byte buffer, specifying the maximum amount of bytes to retrieve. Call this from an ISR.

Attempt to retrieve bytes from a byte buffer whilst specifying a maximum number of bytes to retrieve. This function will return immediately if there is no data available for retrieval.

Note: A call to `vRingbufferReturnItemFromISR()` is required after this to free up the data received.

Note: This function should only be called on byte buffers

Note: Byte buffers do not allow multiple retrievals before returning an item

Parameters

- **xRingbuffer** -- [in] Ring buffer to retrieve the item from
- **pxItemSize** -- [out] Pointer to a variable to which the size of the retrieved item will be written.
- **xMaxSize** -- [in] Maximum number of bytes to return. Size of 0 simply returns NULL.

Returns

- Pointer to the retrieved item on success; `*pxItemSize` filled with the length of the item.
- NULL when the ring buffer is empty, `*pxItemSize` is untouched in that case.

void **vRingbufferReturnItem** (*RingbufHandle_t* xRingbuffer, void *pvItem)

Return a previously-retrieved item to the ring buffer.

Note: If a split item is retrieved, both parts should be returned by calling this function twice

Parameters

- **xRingbuffer** -- [in] Ring buffer the item was retrieved from
- **pvItem** -- [in] Item that was received earlier

void **vRingbufferReturnItemFromISR** (*RingbufHandle_t* xRingbuffer, void *pvItem, BaseType_t *pxHigherPriorityTaskWoken)

Return a previously-retrieved item to the ring buffer from an ISR.

Note: If a split item is retrieved, both parts should be returned by calling this function twice

Parameters

- **xRingbuffer** -- [in] Ring buffer the item was retrieved from
- **pvItem** -- [in] Item that was received earlier
- **pxHigherPriorityTaskWoken** -- [out] Value pointed to will be set to pdTRUE if the function woke up a higher priority task.

void **vRingbufferDelete** (*RingbufHandle_t* xRingbuffer)

Delete a ring buffer.

Note: This function will not deallocate any memory if the ring buffer was created using `xRingbufferCreateStatic()`. Deallocation must be done manually by the user.

Parameters **xRingbuffer** -- [in] Ring buffer to delete

size_t **xRingbufferGetMaxItemSize** (*RingbufHandle_t* xRingbuffer)

Get maximum size of an item that can be placed in the ring buffer.

This function returns the maximum size an item can have if it was placed in an empty ring buffer.

Note: The max item size for a no-split buffer is limited to $((\text{buffer\_size}/2) - \text{header\_size})$. This limit is imposed so that an item of max item size can always be sent to an empty no-split buffer regardless of the internal positions of the buffer's read/write/free pointers.

Parameters **xRingbuffer** -- [in] Ring buffer to query

Returns Maximum size, in bytes, of an item that can be placed in a ring buffer.

size_t **xRingbufferGetCurFreeSize** (*RingbufHandle_t* xRingbuffer)

Get current free size available for an item/data in the buffer.

This gives the real time free space available for an item/data in the ring buffer. This represents the maximum size an item/data can have if it was currently sent to the ring buffer.

Note: An empty no-split buffer has a max current free size for an item that is limited to $((\text{buffer\_size}/2) - \text{header\_size})$. See API reference for `xRingbufferGetMaxItemSize()`.

Warning: This API is not thread safe. So, if multiple threads are accessing the same ring buffer, it is the application's responsibility to ensure atomic access to this API and the subsequent Send

Parameters **xRingbuffer** -- [in] Ring buffer to query

Returns Current free size, in bytes, available for an entry

BaseType_t **xRingbufferAddToQueueSetRead** (*RingbufHandle_t* xRingbuffer, *QueueSetHandle_t* xQueueSet)

Add the ring buffer to a queue set. Notified when data has been written to the ring buffer.

This function adds the ring buffer to a queue set, thus allowing a task to block on multiple queues/ring buffers. The queue set is notified when the new data becomes available to read on the ring buffer.

Parameters

- **xRingbuffer** -- [in] Ring buffer to add to the queue set
- **xQueueSet** -- [in] Queue set to add the ring buffer to

Returns

- pdTRUE on success, pdFALSE otherwise


```
static inline BaseType_t xRingbufferCanRead (RingbufHandle_t xRingbuffer, QueueSetMemberHandle_t xMember)
```

Check if the selected queue set member is a particular ring buffer.

This API checks if queue set member returned from `xQueueSelectFromSet()` is a particular ring buffer. If so, this indicates the ring buffer has items waiting to be retrieved.

Parameters

- **xRingbuffer** -- [in] Ring buffer to check
- **xMember** -- [in] Member returned from `xQueueSelectFromSet`

Returns

- `pdTRUE` when selected queue set member is the ring buffer
- `pdFALSE` otherwise.

```
BaseType_t xRingbufferRemoveFromQueueSetRead (RingbufHandle_t xRingbuffer, QueueSetHandle_t xQueueSet)
```

Remove the ring buffer from a queue set.

This function removes a ring buffer from a queue set. The ring buffer must have been previously added to the queue set using `xRingbufferAddToQueueSetRead()`.

Parameters

- **xRingbuffer** -- [in] Ring buffer to remove from the queue set
- **xQueueSet** -- [in] Queue set to remove the ring buffer from

Returns

- `pdTRUE` on success
- `pdFALSE` otherwise

```
void vRingbufferGetInfo (RingbufHandle_t xRingbuffer, BaseType_t *uxFree, BaseType_t *uxRead, BaseType_t *uxWrite, BaseType_t *uxAcquire, BaseType_t *uxItemsWaiting)
```

Get information about ring buffer status.

Get information of a ring buffer's current status such as free/read/write/acquire pointer positions, and number of items waiting to be retrieved. Arguments can be set to NULL if they are not required.

Parameters

- **xRingbuffer** -- [in] Ring buffer to remove from the queue set
- **uxFree** -- [out] Pointer use to store free pointer position
- **uxRead** -- [out] Pointer use to store read pointer position
- **uxWrite** -- [out] Pointer use to store write pointer position
- **uxAcquire** -- [out] Pointer use to store acquire pointer position
- **uxItemsWaiting** -- [out] Pointer use to store number of items (bytes for byte buffer) waiting to be retrieved

```
void xRingbufferPrintInfo (RingbufHandle_t xRingbuffer)
```

Debugging function to print the internal pointers in the ring buffer.

Parameters **xRingbuffer** -- Ring buffer to show

```
BaseType_t xRingbufferGetStaticBuffer (RingbufHandle_t xRingbuffer, uint8_t **ppucRingbufferStorage, StaticRingbuffer_t **ppxStaticRingbuffer)
```

Retrieve the pointers to a statically created ring buffer.

Parameters

- **xRingbuffer** -- [in] Ring buffer
- **ppucRingbufferStorage** -- [out] Used to return a pointer to the queue's storage area buffer
- **ppxStaticRingbuffer** -- [out] Used to return a pointer to the queue's data structure buffer

Returns `pdTRUE` if buffers were retrieved, `pdFALSE` otherwise.

RingbufHandle_t **xRingbufferCreateWithCaps** (size_t xBufferSize, *RingbufferType_t* xBufferType, UBaseType_t uxMemoryCaps)

Creates a ring buffer with specific memory capabilities.

This function is similar to xRingbufferCreate(), except that it allows the memory allocated for the ring buffer to have specific capabilities (e.g., MALLOC_CAP_INTERNAL).

Note: A queue created using this function must only be deleted using vRingbufferDeleteWithCaps()

Parameters

- **xBufferSize** -- [in] Size of the buffer in bytes
- **xBufferType** -- [in] Type of ring buffer, see documentation.
- **uxMemoryCaps** -- [in] Memory capabilities of the queue's memory (see esp_heap_caps.h)

Returns Handle to the created ring buffer or NULL on failure.

void **vRingbufferDeleteWithCaps** (*RingbufHandle_t* xRingbuffer)

Deletes a ring buffer previously created using xRingbufferCreateWithCaps()

Parameters **xRingbuffer** -- Ring buffer

Structures

struct **xSTATIC_RINGBUFFER**

Struct that is equivalent in size to the ring buffer's data structure.

The contents of this struct are not meant to be used directly. This structure is meant to be used when creating a statically allocated ring buffer where this struct is of the exact size required to store a ring buffer's control data structure.

Type Definitions

typedef void ***RingbufHandle_t**

Type by which ring buffers are referenced. For example, a call to xRingbufferCreate() returns a RingbufHandle_t variable that can then be used as a parameter to xRingbufferSend(), xRingbufferReceive(), etc.

typedef struct *xSTATIC_RINGBUFFER* **StaticRingbuffer_t**

Struct that is equivalent in size to the ring buffer's data structure.

The contents of this struct are not meant to be used directly. This structure is meant to be used when creating a statically allocated ring buffer where this struct is of the exact size required to store a ring buffer's control data structure.

Enumerations

enum **RingbufferType_t**

Values:

enumerator **RINGBUF_TYPE_NOSPLIT**

No-split buffers will only store an item in contiguous memory and will never split an item. Each item requires an 8 byte overhead for a header and will always internally occupy a 32-bit aligned size of space.

enumerator `RINGBUF_TYPE_ALLOWSPILT`

Allow-split buffers will split an item into two parts if necessary in order to store it. Each item requires an 8 byte overhead for a header, splitting incurs an extra header. Each item will always internally occupy a 32-bit aligned size of space.

enumerator `RINGBUF_TYPE_BYTEBUF`

Byte buffers store data as a sequence of bytes and do not maintain separate items, therefore byte buffers have no overhead. All data is stored as a sequence of byte and any number of bytes can be sent or retrieved each time.

enumerator `RINGBUF_TYPE_MAX`

Hooks API

Header File

- [components/esp_system/include/esp_freertos_hooks.h](#)
- This header file can be included with:

```
#include "esp_freertos_hooks.h"
```

Functions

esp_err_t **esp_register_freertos_idle_hook_for_cpu** (*esp_freertos_idle_cb_t* new_idle_cb, UBaseType_t cpuid)

Register a callback to be called from the specified core's idle hook. The callback should return true if it should be called by the idle hook once per interrupt (or FreeRTOS tick), and return false if it should be called repeatedly as fast as possible by the idle hook.

Warning: Idle callbacks MUST NOT, UNDER ANY CIRCUMSTANCES, CALL A FUNCTION THAT MIGHT BLOCK.

Parameters

- **new_idle_cb** -- [in] Callback to be called
- **cpuid** -- [in] id of the core

Returns

- **ESP_OK**: Callback registered to the specified core's idle hook
- **ESP_ERR_NO_MEM**: No more space on the specified core's idle hook to register callback
- **ESP_ERR_INVALID_ARG**: cpuid is invalid

esp_err_t **esp_register_freertos_idle_hook** (*esp_freertos_idle_cb_t* new_idle_cb)

Register a callback to the idle hook of the core that calls this function. The callback should return true if it should be called by the idle hook once per interrupt (or FreeRTOS tick), and return false if it should be called repeatedly as fast as possible by the idle hook.

Warning: Idle callbacks MUST NOT, UNDER ANY CIRCUMSTANCES, CALL A FUNCTION THAT MIGHT BLOCK.

Parameters **new_idle_cb** -- [in] Callback to be called

Returns

- **ESP_OK**: Callback registered to the calling core's idle hook
- **ESP_ERR_NO_MEM**: No more space on the calling core's idle hook to register callback

esp_err_t **esp_register_freertos_tick_hook_for_cpu** (*esp_freertos_tick_cb_t* new_tick_cb,
UBaseType_t cpuid)

Register a callback to be called from the specified core's tick hook.

Parameters

- **new_tick_cb** -- [in] Callback to be called
- **cpuid** -- [in] id of the core

Returns

- ESP_OK: Callback registered to specified core's tick hook
- ESP_ERR_NO_MEM: No more space on the specified core's tick hook to register the callback
- ESP_ERR_INVALID_ARG: cpuid is invalid

esp_err_t **esp_register_freertos_tick_hook** (*esp_freertos_tick_cb_t* new_tick_cb)

Register a callback to be called from the calling core's tick hook.

Parameters **new_tick_cb** -- [in] Callback to be called

Returns

- ESP_OK: Callback registered to the calling core's tick hook
- ESP_ERR_NO_MEM: No more space on the calling core's tick hook to register the callback

void **esp_deregister_freertos_idle_hook_for_cpu** (*esp_freertos_idle_cb_t* old_idle_cb,
UBaseType_t cpuid)

Unregister an idle callback from the idle hook of the specified core.

Parameters

- **old_idle_cb** -- [in] Callback to be unregistered
- **cpuid** -- [in] id of the core

void **esp_deregister_freertos_idle_hook** (*esp_freertos_idle_cb_t* old_idle_cb)

Unregister an idle callback. If the idle callback is registered to the idle hooks of both cores, the idle hook will be unregistered from both cores.

Parameters **old_idle_cb** -- [in] Callback to be unregistered

void **esp_deregister_freertos_tick_hook_for_cpu** (*esp_freertos_tick_cb_t* old_tick_cb,
UBaseType_t cpuid)

Unregister a tick callback from the tick hook of the specified core.

Parameters

- **old_tick_cb** -- [in] Callback to be unregistered
- **cpuid** -- [in] id of the core

void **esp_deregister_freertos_tick_hook** (*esp_freertos_tick_cb_t* old_tick_cb)

Unregister a tick callback. If the tick callback is registered to the tick hooks of both cores, the tick hook will be unregistered from both cores.

Parameters **old_tick_cb** -- [in] Callback to be unregistered

Type Definitions

typedef bool (***esp_freertos_idle_cb_t**)(void)

typedef void (***esp_freertos_tick_cb_t**)(void)

Additional API

Header File

- [components/freertos/esp_additions/include/freertos/idf_additions.h](#)
- This header file can be included with:

```
#include "freertos/idf_additions.h"
```

Functions

BaseType_t xTaskCreatePinnedToCore (TaskFunction_t pxTaskCode, const char *const pcName, const uint32_t ulStackDepth, void *const pvParameters, UBaseType_t uxPriority, [TaskHandle_t](#) *const pxCreatedTask, const BaseType_t xCoreID)

Create a new task that is pinned to a particular core.

This function is similar to xTaskCreate(), but allows the creation of a pinned task. The task's pinned core is specified by the xCoreID argument. If xCoreID is set to tskNO_AFFINITY, then the task is unpinned and can run on any core.

Note: If (configNUMBER_OF_CORES == 1), setting xCoreID to tskNO_AFFINITY will be treated as 0.

Parameters

- **pxTaskCode** -- Pointer to the task entry function.
- **pcName** -- A descriptive name for the task.
- **ulStackDepth** -- The size of the task stack specified as the NUMBER OF BYTES. Note that this differs from vanilla FreeRTOS.
- **pvParameters** -- Pointer that will be used as the parameter for the task being created.
- **uxPriority** -- The priority at which the task should run.
- **pxCreatedTask** -- Used to pass back a handle by which the created task can be referenced.
- **xCoreID** -- The core to which the task is pinned to, or tskNO_AFFINITY if the task has no core affinity.

Returns pdPASS if the task was successfully created and added to a ready list, otherwise an error code defined in the file projdefs.h

[TaskHandle_t](#) xTaskCreateStaticPinnedToCore (TaskFunction_t pxTaskCode, const char *const pcName, const uint32_t ulStackDepth, void *const pvParameters, UBaseType_t uxPriority, StackType_t *const puxStackBuffer, StaticTask_t *const pxTaskBuffer, const BaseType_t xCoreID)

Create a new static task that is pinned to a particular core.

This function is similar to xTaskCreateStatic(), but allows the creation of a pinned task. The task's pinned core is specified by the xCoreID argument. If xCoreID is set to tskNO_AFFINITY, then the task is unpinned and can run on any core.

Note: If (configNUMBER_OF_CORES == 1), setting xCoreID to tskNO_AFFINITY will be treated as 0.

Parameters

- **pxTaskCode** -- Pointer to the task entry function.
- **pcName** -- A descriptive name for the task.
- **ulStackDepth** -- The size of the task stack specified as the NUMBER OF BYTES. Note that this differs from vanilla FreeRTOS.
- **pvParameters** -- Pointer that will be used as the parameter for the task being created.
- **uxPriority** -- The priority at which the task should run.

- **pxStackBuffer** -- Must point to a StackType_t array that has at least ulStackDepth indexes
- **pxTaskBuffer** -- Must point to a variable of type StaticTask_t, which will then be used to hold the task's data structures,
- **xCoreID** -- The core to which the task is pinned to, or tskNO_AFFINITY if the task has no core affinity.

Returns The task handle if the task was created, NULL otherwise.

BaseType_t **xTaskGetCoreID** (*TaskHandle_t* xTask)

Get the current core ID of a particular task.

Helper function to get the core ID of a particular task. If the task is pinned to a particular core, the core ID is returned. If the task is not pinned to a particular core, tskNO_AFFINITY is returned.

If CONFIG_FREERTOS_UNICORE is enabled, this function simply returns 0.

[refactor-todo] See if this needs to be deprecated (IDF-8145)(IDF-8164)

Note: If CONFIG_FREERTOS_SMP is enabled, please call vTaskCoreAffinityGet() instead.

Note: In IDF FreerTOS when configNUMBER_OF_CORES == 1, this function will always return 0,

Parameters **xTask** -- The task to query

Returns The task's core ID or tskNO_AFFINITY

TaskHandle_t **xTaskGetIdleTaskHandleForCore** (BaseType_t xCoreID)

Get the handle of idle task for the given core.

[refactor-todo] See if this needs to be deprecated (IDF-8145)

Note: If CONFIG_FREERTOS_SMP is enabled, please call xTaskGetIdleTaskHandle() instead.

Parameters **xCoreID** -- The core to query

Returns Handle of the idle task for the queried core

TaskHandle_t **xTaskGetCurrentTaskHandleForCore** (BaseType_t xCoreID)

Get the handle of the task currently running on a certain core.

Because of the nature of SMP processing, there is no guarantee that this value will still be valid on return and should only be used for debugging purposes.

[refactor-todo] See if this needs to be deprecated (IDF-8145)

Note: If CONFIG_FREERTOS_SMP is enabled, please call xTaskGetCurrentTaskHandleCPU() instead.

Parameters **xCoreID** -- The core to query

Returns Handle of the current task running on the queried core

uint8_t ***pxTaskGetStackStart** (*TaskHandle_t* xTask)

Returns the start of the stack associated with xTask.

Returns the lowest stack memory address, regardless of whether the stack grows up or down.

[refactor-todo] Change return type to StackType_t (IDF-8158)

Parameters **xTask** -- Handle of the task associated with the stack returned. Set xTask to NULL to return the stack of the calling task.

Returns A pointer to the start of the stack.

```
void vTaskSetThreadLocalStoragePointerAndDelCallback (TaskHandle_t xTaskToSet,
                                                    BaseType_t xIndex, void *pvValue,
                                                    TlsDeleteCallbackFunction_t
                                                    pvDelCallback)
```

Set local storage pointer and deletion callback.

Each task contains an array of pointers that is dimensioned by the configNUM_THREAD_LOCAL_STORAGE_POINTERS setting in FreeRTOSConfig.h. The kernel does not use the pointers itself, so the application writer can use the pointers for any purpose they wish.

Local storage pointers set for a task can reference dynamically allocated resources. This function is similar to vTaskSetThreadLocalStoragePointer, but provides a way to release these resources when the task gets deleted. For each pointer, a callback function can be set. This function will be called when task is deleted, with the local storage pointer index and value as arguments.

Parameters

- **xTaskToSet** -- Task to set thread local storage pointer for
- **xIndex** -- The index of the pointer to set, from 0 to configNUM_THREAD_LOCAL_STORAGE_POINTERS - 1.
- **pvValue** -- Pointer value to set.
- **pvDelCallback** -- Function to call to dispose of the local storage pointer when the task is deleted.

```
BaseType_t xTaskCreatePinnedToCoreWithCaps (TaskFunction_t pvTaskCode, const char *const
                                           pcName, const configSTACK_DEPTH_TYPE
                                           usStackDepth, void *const pvParameters,
                                           UBaseType_t uxPriority, TaskHandle_t *const
                                           pvCreatedTask, const BaseType_t xCoreID,
                                           UBaseType_t uxMemoryCaps)
```

Creates a pinned task where its stack has specific memory capabilities.

This function is similar to xTaskCreatePinnedToCore(), except that it allows the memory allocated for the task's stack to have specific capabilities (e.g., MALLOC_CAP_SPIRAM).

However, the specified capabilities will NOT apply to the task's TCB as a TCB must always be in internal RAM.

Parameters

- **pvTaskCode** -- Pointer to the task entry function
- **pcName** -- A descriptive name for the task
- **usStackDepth** -- The size of the task stack specified as the number of bytes
- **pvParameters** -- Pointer that will be used as the parameter for the task being created.
- **uxPriority** -- The priority at which the task should run.
- **pvCreatedTask** -- Used to pass back a handle by which the created task can be referenced.
- **xCoreID** -- Core to which the task is pinned to, or tskNO_AFFINITY if unpinned.
- **uxMemoryCaps** -- Memory capabilities of the task stack's memory (see esp_heap_caps.h)

Returns pdPASS if the task was successfully created and added to a ready list, otherwise an error code defined in the file projdefs.h

```
static inline BaseType_t xTaskCreateWithCaps (TaskFunction_t pvTaskCode, const char *const pcName,
                                           configSTACK_DEPTH_TYPE usStackDepth, void *const
                                           pvParameters, UBaseType_t uxPriority, TaskHandle_t
                                           *pvCreatedTask, UBaseType_t uxMemoryCaps)
```

Creates a task where its stack has specific memory capabilities.

This function is similar to `xTaskCreate()`, except that it allows the memory allocated for the task's stack to have specific capabilities (e.g., `MALLOC_CAP_SPIRAM`).

However, the specified capabilities will NOT apply to the task's TCB as a TCB must always be in internal RAM.

Note: A task created using this function must only be deleted using `vTaskDeleteWithCaps()`

Parameters

- **pvTaskCode** -- Pointer to the task entry function
- **pcName** -- A descriptive name for the task
- **usStackDepth** -- The size of the task stack specified as the number of bytes
- **pvParameters** -- Pointer that will be used as the parameter for the task being created.
- **uxPriority** -- The priority at which the task should run.
- **pvCreatedTask** -- Used to pass back a handle by which the created task can be referenced.
- **uxMemoryCaps** -- Memory capabilities of the task stack's memory (see `esp_heap_caps.h`)

Returns `pdPASS` if the task was successfully created and added to a ready list, otherwise an error code defined in the file `projdefs.h`

void **vTaskDeleteWithCaps** (*TaskHandle_t* xTaskToDelete)

Deletes a task previously created using `xTaskCreateWithCaps()` or `xTaskCreatePinnedToCoreWithCaps()`

Note: It is recommended to use this API to delete tasks from another task's context, rather than self-deletion. When the task is being deleted, it is vital to ensure that it is not running on another core. This API must not be called from an interrupt context.

Parameters **xTaskToDelete** -- A handle to the task to be deleted

QueueHandle_t **xQueueCreateWithCaps** (`UBaseType_t` uxQueueLength, `UBaseType_t` uxItemSize, `UBaseType_t` uxMemoryCaps)

Creates a queue with specific memory capabilities.

This function is similar to `xQueueCreate()`, except that it allows the memory allocated for the queue to have specific capabilities (e.g., `MALLOC_CAP_INTERNAL`).

Note: A queue created using this function must only be deleted using `vQueueDeleteWithCaps()`

Parameters

- **uxQueueLength** -- The maximum number of items that the queue can contain.
- **uxItemSize** -- The number of bytes each item in the queue will require.
- **uxMemoryCaps** -- Memory capabilities of the queue's memory (see `esp_heap_caps.h`)

Returns Handle to the created queue or `NULL` on failure.

void **vQueueDeleteWithCaps** (*QueueHandle_t* xQueue)

Deletes a queue previously created using `xQueueCreateWithCaps()`

Parameters **xQueue** -- A handle to the queue to be deleted.

static inline *SemaphoreHandle_t* **xSemaphoreCreateBinaryWithCaps** (`UBaseType_t` uxMemoryCaps)

Creates a binary semaphore with specific memory capabilities.

This function is similar to `vSemaphoreCreateBinary()`, except that it allows the memory allocated for the binary semaphore to have specific capabilities (e.g., `MALLOC_CAP_INTERNAL`).

Note: A binary semaphore created using this function must only be deleted using `vSemaphoreDeleteWithCaps()`

Parameters **uxMemoryCaps** -- Memory capabilities of the binary semaphore's memory (see `esp_heap_caps.h`)

Returns Handle to the created binary semaphore or NULL on failure.

static inline [*SemaphoreHandle_t*](#) **xSemaphoreCreateCountingWithCaps** (UBaseType_t uxMaxCount,
UBaseType_t uxInitialCount,
UBaseType_t uxMemoryCaps)

Creates a counting semaphore with specific memory capabilities.

This function is similar to `xSemaphoreCreateCounting()`, except that it allows the memory allocated for the counting semaphore to have specific capabilities (e.g., `MALLOC_CAP_INTERNAL`).

Note: A counting semaphore created using this function must only be deleted using `vSemaphoreDeleteWithCaps()`

Parameters

- **uxMaxCount** -- The maximum count value that can be reached.
- **uxInitialCount** -- The count value assigned to the semaphore when it is created.
- **uxMemoryCaps** -- Memory capabilities of the counting semaphore's memory (see `esp_heap_caps.h`)

Returns Handle to the created counting semaphore or NULL on failure.

static inline [*SemaphoreHandle_t*](#) **xSemaphoreCreateMutexWithCaps** (UBaseType_t uxMemoryCaps)

Creates a mutex semaphore with specific memory capabilities.

This function is similar to `xSemaphoreCreateMutex()`, except that it allows the memory allocated for the mutex semaphore to have specific capabilities (e.g., `MALLOC_CAP_INTERNAL`).

Note: A mutex semaphore created using this function must only be deleted using `vSemaphoreDeleteWithCaps()`

Parameters **uxMemoryCaps** -- Memory capabilities of the mutex semaphore's memory (see `esp_heap_caps.h`)

Returns Handle to the created mutex semaphore or NULL on failure.

static inline [*SemaphoreHandle_t*](#) **xSemaphoreCreateRecursiveMutexWithCaps** (UBaseType_t
uxMemoryCaps)

Creates a recursive mutex with specific memory capabilities.

This function is similar to `xSemaphoreCreateRecursiveMutex()`, except that it allows the memory allocated for the recursive mutex to have specific capabilities (e.g., `MALLOC_CAP_INTERNAL`).

Note: A recursive mutex created using this function must only be deleted using `vSemaphoreDeleteWithCaps()`

Parameters **uxMemoryCaps** -- Memory capabilities of the recursive mutex's memory (see `esp_heap_caps.h`)

Returns Handle to the created recursive mutex or NULL on failure.

void **vSemaphoreDeleteWithCaps** (*SemaphoreHandle_t* xSemaphore)

Deletes a semaphore previously created using one of the xSemaphoreCreate...WithCaps() functions.

Parameters **xSemaphore** -- A handle to the semaphore to be deleted.

static inline *StreamBufferHandle_t* **xStreamBufferCreateWithCaps** (size_t xBufferSizeBytes, size_t xTriggerLevelBytes, UBaseType_t uxMemoryCaps)

Creates a stream buffer with specific memory capabilities.

This function is similar to xStreamBufferCreate(), except that it allows the memory allocated for the stream buffer to have specific capabilities (e.g., MALLOC_CAP_INTERNAL).

Note: A stream buffer created using this function must only be deleted using vStreamBufferDeleteWithCaps()

Parameters

- **xBufferSizeBytes** -- The total number of bytes the stream buffer will be able to hold at any one time.
- **xTriggerLevelBytes** -- The number of bytes that must be in the stream buffer before unblocking
- **uxMemoryCaps** -- Memory capabilities of the stream buffer's memory (see esp_heap_caps.h)

Returns Handle to the created stream buffer or NULL on failure.

static inline void **vStreamBufferDeleteWithCaps** (*StreamBufferHandle_t* xStreamBuffer)

Deletes a stream buffer previously created using xStreamBufferCreateWithCaps()

Parameters **xStreamBuffer** -- A handle to the stream buffer to be deleted.

static inline *MessageBufferHandle_t* **xMessageBufferCreateWithCaps** (size_t xBufferSizeBytes, UBaseType_t uxMemoryCaps)

Creates a message buffer with specific memory capabilities.

This function is similar to xMessageBufferCreate(), except that it allows the memory allocated for the message buffer to have specific capabilities (e.g., MALLOC_CAP_INTERNAL).

Note: A message buffer created using this function must only be deleted using vMessageBufferDeleteWithCaps()

Parameters

- **xBufferSizeBytes** -- The total number of bytes (not messages) the message buffer will be able to hold at any one time.
- **uxMemoryCaps** -- Memory capabilities of the message buffer's memory (see esp_heap_caps.h)

Returns Handle to the created message buffer or NULL on failure.

static inline void **vMessageBufferDeleteWithCaps** (*MessageBufferHandle_t* xMessageBuffer)

Deletes a stream buffer previously created using xMessageBufferCreateWithCaps()

Parameters **xMessageBuffer** -- A handle to the message buffer to be deleted.

EventGroupHandle_t **xEventGroupCreateWithCaps** (UBaseType_t uxMemoryCaps)

Creates an event group with specific memory capabilities.

This function is similar to xEventGroupCreate(), except that it allows the memory allocated for the event group to have specific capabilities (e.g., MALLOC_CAP_INTERNAL).

Note: An event group created using this function must only be deleted using `vEventGroupDeleteWithCaps()`

Parameters `uxMemoryCaps` -- Memory capabilities of the event group's memory (see `esp_heap_caps.h`)

Returns Handle to the created event group or NULL on failure.

void **vEventGroupDeleteWithCaps** (*EventGroupHandle_t* xEventGroup)

Deletes an event group previously created using `xEventGroupCreateWithCaps()`

Parameters `xEventGroup` -- A handle to the event group to be deleted.

Type Definitions

typedef void (***TlsDeleteCallbackFunction_t**)(int, void*)

Prototype of local storage pointer deletion callback.

2.9.14 Heap Memory Allocation

Stack and Heap

ESP-IDF applications use the common computer architecture patterns of **stack** (dynamic memory allocated by program control flow), **heap** (dynamic memory allocated by function calls), and **static memory** (memory allocated at compile time).

Because ESP-IDF is a multi-threaded RTOS environment, each RTOS task has its own stack. By default, each of these stacks is allocated from the heap when the task is created. See *xTaskCreateStatic()* for the alternative where stacks are statically allocated.

Because ESP32-P4 uses multiple types of RAM, it also contains multiple heaps with different capabilities. A capabilities-based memory allocator allows apps to make heap allocations for different purposes.

For most purposes, the C Standard Library's `malloc()` and `free()` functions can be used for heap allocation without any special consideration. However, in order to fully make use of all of the memory types and their characteristics, ESP-IDF also has a capabilities-based heap memory allocator. If you want to have a memory with certain properties (e.g., *DMA-Capable Memory* or executable-memory), you can create an OR-mask of the required capabilities and pass that to *heap_caps_malloc()*.

Memory Capabilities

The ESP32-P4 contains multiple types of RAM:

- DRAM (Data RAM) is memory that is connected to CPU's data bus and is used to hold data. This is the most common kind of memory accessed as a heap.
- IRAM (Instruction RAM) is memory that is connected to the CPU's instruction bus and usually holds executable data only (i.e., instructions). If accessed as generic memory, all accesses must be aligned to *32-Bit Accessible Memory*.
- D/IRAM is RAM that is connected to CPU's data bus and instruction bus, thus can be used either Instruction or Data RAM.

For more details on these internal memory types, see *Memory Types*.

It is also possible to connect external SPI RAM to the ESP32-P4. The *external RAM* is integrated into the ESP32-P4's memory map via the cache, and accessed similarly to DRAM.

All DRAM memory is single-byte accessible, thus all DRAM heaps possess the `MALLOC_CAP_8BIT` capability. Users can call `heap_caps_get_free_size(MALLOC_CAP_8BIT)` to get the free size of all DRAM heaps.

When calling `malloc()`, the ESP-IDF `malloc()` internally calls `heap_caps_malloc_default(size)`. This will allocate memory with the capability `MALLOC_CAP_DEFAULT`, which is byte-addressable.

Because `malloc()` uses the capabilities-based allocation system, memory allocated using `heap_caps_malloc()` can be freed by calling the standard `free()` function.

Available Heap

DRAM At startup, the DRAM heap contains all data memory that is not statically allocated by the app. Reducing statically-allocated buffers increases the amount of available free heap.

To find the amount of statically allocated memory, use the `idf.py size` command.

Note: At runtime, the available heap DRAM may be less than calculated at compile time, because, at startup, some memory is allocated from the heap before the FreeRTOS scheduler is started (including memory for the stacks of initial FreeRTOS tasks).

IRAM At startup, the IRAM heap contains all instruction memory that is not used by the app executable code.

The `idf.py size` command can be used to find the amount of IRAM used by the app.

D/IRAM Some memory in the ESP32-P4 is available as either DRAM or IRAM. If memory is allocated from a D/IRAM region, the free heap size for both types of memory will decrease.

Heap Sizes At startup, all ESP-IDF apps log a summary of all heap addresses (and sizes) at level Info:

```
I (252) heap_init: Initializing. RAM available for dynamic allocation:
I (259) heap_init: At 3FFAE6E0 len 00001920 (6 KiB): DRAM
I (265) heap_init: At 3FFB2EC8 len 0002D138 (180 KiB): DRAM
I (272) heap_init: At 3FFE0440 len 00003AE0 (14 KiB): D/IRAM
I (278) heap_init: At 3FFE4350 len 0001BCB0 (111 KiB): D/IRAM
I (284) heap_init: At 4008944C len 00016BB4 (90 KiB): IRAM
```

Finding Available Heap See [Heap Information](#).

Special Capabilities

DMA-Capable Memory Use the `MALLOC_CAP_DMA` flag to allocate memory which is suitable for use with hardware DMA engines (for example SPI and I2S). This capability flag excludes any external PSRAM.

32-Bit Accessible Memory If a certain memory structure is only addressed in 32-bit units, for example, an array of ints or pointers, it can be useful to allocate it with the `MALLOC_CAP_32BIT` flag. This also allows the allocator to give out IRAM memory, which is sometimes unavailable for a normal `malloc()` call. This can help to use all the available memory in the ESP32-P4.

Memory allocated with `MALLOC_CAP_32BIT` can **only** be accessed via 32-bit reads and writes, any other type of access will generate a fatal `LoadStoreError` exception.

External SPI Memory When [external RAM](#) is enabled, external SPI RAM can be allocated using standard `malloc` calls, or via `heap_caps_malloc(MALLOC_CAP_SPIRAM)`, depending on the configuration. See [Configuring External RAM](#) for more details.

Thread Safety

Heap functions are thread-safe, meaning they can be called from different tasks simultaneously without any limitations.

It is technically possible to call `malloc`, `free`, and related functions from interrupt handler (ISR) context (see [Calling Heap-Related Functions from ISR](#)). However, this is not recommended, as heap function calls may delay other interrupts. It is strongly recommended to refactor applications so that any buffers used by an ISR are pre-allocated outside of the ISR. Support for calling heap functions from ISRs may be removed in a future update.

Calling Heap-Related Functions from ISR

The following functions from the heap component can be called from the interrupt handler (ISR):

- `heap_caps_malloc()`
- `heap_caps_malloc_default()`
- `heap_caps_realloc_default()`
- `heap_caps_malloc_prefer()`
- `heap_caps_realloc_prefer()`
- `heap_caps_calloc_prefer()`
- `heap_caps_free()`
- `heap_caps_realloc()`
- `heap_caps_calloc()`
- `heap_caps_aligned_alloc()`
- `heap_caps_aligned_free()`

Note: However, this practice is strongly discouraged.

Heap Tracing & Debugging

The following features are documented on the [Heap Memory Debugging](#) page:

- [Heap Information](#) (free space, etc.)
- [Heap Allocation and Free Function Hooks](#)
- [Heap Corruption Detection](#)
- [Heap Tracing](#) (memory leak detection, monitoring, etc.)

Implementation Notes

Knowledge about the regions of memory in the chip comes from the "SoC" component, which contains memory layout information for the chip, and the different capabilities of each region. Each region's capabilities are prioritized, so that (for example) dedicated DRAM and IRAM regions are used for allocations ahead of the more versatile D/IRAM regions.

Each contiguous region of memory contains its own memory heap. The heaps are created using the [multi_heap](#) functionality. `multi_heap` allows any contiguous region of memory to be used as a heap.

The heap capabilities allocator uses knowledge of the memory regions to initialize each individual heap. Allocation functions in the heap capabilities API will find the most appropriate heap for the allocation based on desired capabilities, available space, and preferences for each region's use, and then calling `multi_heap_malloc()` for the heap situated in that particular region.

Calling `free()` involves finding the particular heap corresponding to the freed address, and then call `multi_heap_free()` on that particular `multi_heap` instance.

API Reference - Heap Allocation

Header File

- [components/heap/include/esp_heap_caps.h](#)
- This header file can be included with:

```
#include "esp_heap_caps.h"
```

Functions

***esp_err_t* heap_caps_register_failed_alloc_callback** (*esp_alloc_failed_hook_t* callback)

registers a callback function to be invoked if a memory allocation operation fails

Parameters **callback** -- caller defined callback to be invoked

Returns ESP_OK if callback was registered.

void ***heap_caps_malloc** (size_t size, uint32_t caps)

Allocate a chunk of memory which has the given capabilities.

Equivalent semantics to libc malloc(), for capability-aware memory.

Parameters

- **size** -- Size, in bytes, of the amount of memory to allocate
- **caps** -- Bitwise OR of MALLOC_CAP_* flags indicating the type of memory to be returned

Returns A pointer to the memory allocated on success, NULL on failure

void **heap_caps_free** (void *ptr)

Free memory previously allocated via heap_caps_malloc() or heap_caps_realloc().

Equivalent semantics to libc free(), for capability-aware memory.

In IDF, free(p) is equivalent to heap_caps_free(p).

Parameters **ptr** -- Pointer to memory previously returned from heap_caps_malloc() or heap_caps_realloc(). Can be NULL.

void ***heap_caps_realloc** (void *ptr, size_t size, uint32_t caps)

Reallocate memory previously allocated via heap_caps_malloc() or heap_caps_realloc().

Equivalent semantics to libc realloc(), for capability-aware memory.

In IDF, realloc(p, s) is equivalent to heap_caps_realloc(p, s, MALLOC_CAP_8BIT).

'caps' parameter can be different to the capabilities that any original 'ptr' was allocated with. In this way, realloc can be used to "move" a buffer if necessary to ensure it meets a new set of capabilities.

Parameters

- **ptr** -- Pointer to previously allocated memory, or NULL for a new allocation.
- **size** -- Size of the new buffer requested, or 0 to free the buffer.
- **caps** -- Bitwise OR of MALLOC_CAP_* flags indicating the type of memory desired for the new allocation.

Returns Pointer to a new buffer of size 'size' with capabilities 'caps', or NULL if allocation failed.

void ***heap_caps_aligned_alloc** (size_t alignment, size_t size, uint32_t caps)

Allocate an aligned chunk of memory which has the given capabilities.

Equivalent semantics to libc aligned_alloc(), for capability-aware memory.

Parameters

- **alignment** -- How the pointer received needs to be aligned must be a power of two
- **size** -- Size, in bytes, of the amount of memory to allocate
- **caps** -- Bitwise OR of MALLOC_CAP_* flags indicating the type of memory to be returned

Returns A pointer to the memory allocated on success, NULL on failure

void **heap_caps_aligned_free** (void *ptr)

Used to deallocate memory previously allocated with heap_caps_aligned_alloc.

Note: This function is deprecated, please consider using heap_caps_free() instead

Parameters **ptr** -- Pointer to the memory allocated

void ***heap_caps_aligned_malloc** (size_t alignment, size_t n, size_t size, uint32_t caps)

Allocate an aligned chunk of memory which has the given capabilities. The initialized value in the memory is set to zero.

Parameters

- **alignment** -- How the pointer received needs to be aligned must be a power of two
- **n** -- Number of continuing chunks of memory to allocate
- **size** -- Size, in bytes, of a chunk of memory to allocate
- **caps** -- Bitwise OR of MALLOC_CAP_* flags indicating the type of memory to be returned

Returns A pointer to the memory allocated on success, NULL on failure

void ***heap_caps_malloc** (size_t n, size_t size, uint32_t caps)

Allocate a chunk of memory which has the given capabilities. The initialized value in the memory is set to zero.

Equivalent semantics to libc calloc(), for capability-aware memory.

In IDF, calloc(p) is equivalent to heap_caps_malloc(p, MALLOC_CAP_8BIT).

Parameters

- **n** -- Number of continuing chunks of memory to allocate
- **size** -- Size, in bytes, of a chunk of memory to allocate
- **caps** -- Bitwise OR of MALLOC_CAP_* flags indicating the type of memory to be returned

Returns A pointer to the memory allocated on success, NULL on failure

size_t **heap_caps_get_total_size** (uint32_t caps)

Get the total size of all the regions that have the given capabilities.

This function takes all regions capable of having the given capabilities allocated in them and adds up the total space they have.

Parameters **caps** -- Bitwise OR of MALLOC_CAP_* flags indicating the type of memory

Returns total size in bytes

size_t **heap_caps_get_free_size** (uint32_t caps)

Get the total free size of all the regions that have the given capabilities.

This function takes all regions capable of having the given capabilities allocated in them and adds up the free space they have.

Note: Note that because of heap fragmentation it is probably not possible to allocate a single block of memory of this size. Use heap_caps_get_largest_free_block() for this purpose.

Parameters **caps** -- Bitwise OR of MALLOC_CAP_* flags indicating the type of memory

Returns Amount of free bytes in the regions

size_t **heap_caps_get_minimum_free_size** (uint32_t caps)

Get the total minimum free memory of all regions with the given capabilities.

This adds all the low watermarks of the regions capable of delivering the memory with the given capabilities.

Note: Note the result may be less than the global all-time minimum available heap of this kind, as "low watermarks" are tracked per-region. Individual regions' heaps may have reached their "low watermarks" at different points in time. However, this result still gives a "worst case" indication for all-time minimum free heap.

Parameters **caps** -- Bitwise OR of `MALLOC_CAP_*` flags indicating the type of memory

Returns Amount of free bytes in the regions

`size_t heap_caps_get_largest_free_block (uint32_t caps)`

Get the largest free block of memory able to be allocated with the given capabilities.

Returns the largest value of `s` for which `heap_caps_malloc(s, caps)` will succeed.

Parameters **caps** -- Bitwise OR of `MALLOC_CAP_*` flags indicating the type of memory

Returns Size of the largest free block in bytes.

`void heap_caps_get_info (multi_heap_info_t *info, uint32_t caps)`

Get heap info for all regions with the given capabilities.

Calls `multi_heap_info()` on all heaps which share the given capabilities. The information returned is an aggregate across all matching heaps. The meanings of fields are the same as defined for `multi_heap_info_t`, except that `minimum_free_bytes` has the same caveats described in `heap_caps_get_minimum_free_size()`.

Parameters

- **info** -- Pointer to a structure which will be filled with relevant heap metadata.
- **caps** -- Bitwise OR of `MALLOC_CAP_*` flags indicating the type of memory

`void heap_caps_print_heap_info (uint32_t caps)`

Print a summary of all memory with the given capabilities.

Calls `multi_heap_info` on all heaps which share the given capabilities, and prints a two-line summary for each, then a total summary.

Parameters **caps** -- Bitwise OR of `MALLOC_CAP_*` flags indicating the type of memory

`bool heap_caps_check_integrity_all (bool print_errors)`

Check integrity of all heap memory in the system.

Calls `multi_heap_check` on all heaps. Optionally print errors if heaps are corrupt.

Calling this function is equivalent to calling `heap_caps_check_integrity` with the `caps` argument set to `MALLOC_CAP_INVALID`.

Note: Please increase the value of `CONFIG_ESP_INT_WDT_TIMEOUT_MS` when using this API with PSRAM enabled.

Parameters **print_errors** -- Print specific errors if heap corruption is found.

Returns True if all heaps are valid, False if at least one heap is corrupt.

`bool heap_caps_check_integrity (uint32_t caps, bool print_errors)`

Check integrity of all heaps with the given capabilities.

Calls `multi_heap_check` on all heaps which share the given capabilities. Optionally print errors if the heaps are corrupt.

See also `heap_caps_check_integrity_all` to check all heap memory in the system and `heap_caps_check_integrity_addr` to check memory around a single address.

Note: Please increase the value of `CONFIG_ESP_INT_WDT_TIMEOUT_MS` when using this API with PSRAM capability flag.

Parameters

- **caps** -- Bitwise OR of `MALLOC_CAP_*` flags indicating the type of memory
- **print_errors** -- Print specific errors if heap corruption is found.

Returns True if all heaps are valid, False if at least one heap is corrupt.

bool **heap_caps_check_integrity_addr** (intptr_t addr, bool print_errors)

Check integrity of heap memory around a given address.

This function can be used to check the integrity of a single region of heap memory, which contains the given address.

This can be useful if debugging heap integrity for corruption at a known address, as it has a lower overhead than checking all heap regions. Note that if the corrupt address moves around between runs (due to timing or other factors) then this approach won't work, and you should call `heap_caps_check_integrity` or `heap_caps_check_integrity_all` instead.

Note: The entire heap region around the address is checked, not only the adjacent heap blocks.

Parameters

- **addr** -- Address in memory. Check for corruption in region containing this address.
- **print_errors** -- Print specific errors if heap corruption is found.

Returns True if the heap containing the specified address is valid, False if at least one heap is corrupt or the address doesn't belong to a heap region.

void **heap_caps_malloc_extmem_enable** (size_t limit)

Enable `malloc()` in external memory and set limit below which `malloc()` attempts are placed in internal memory.

When external memory is in use, the allocation strategy is to initially try to satisfy smaller allocation requests with internal memory and larger requests with external memory. This sets the limit between the two, as well as generally enabling allocation in external memory.

Parameters **limit** -- Limit, in bytes.

void ***heap_caps_malloc_prefer** (size_t size, size_t num, ...)

Allocate a chunk of memory as preference in decreasing order.

Attention The variable parameters are bitwise OR of `MALLOC_CAP_*` flags indicating the type of memory. This API prefers to allocate memory with the first parameter. If failed, allocate memory with the next parameter. It will try in this order until allocating a chunk of memory successfully or fail to allocate memories with any of the parameters.

Parameters

- **size** -- Size, in bytes, of the amount of memory to allocate
- **num** -- Number of variable parameters

Returns A pointer to the memory allocated on success, NULL on failure

void ***heap_caps_realloc_prefer** (void *ptr, size_t size, size_t num, ...)

Reallocate a chunk of memory as preference in decreasing order.

Parameters

- **ptr** -- Pointer to previously allocated memory, or NULL for a new allocation.
- **size** -- Size of the new buffer requested, or 0 to free the buffer.
- **num** -- Number of variable parameters

Returns Pointer to a new buffer of size 'size', or NULL if allocation failed.

void ***heap_caps_malloc_prefer** (size_t n, size_t size, size_t num, ...)

Allocate a chunk of memory as preference in decreasing order.

Parameters

- **n** -- Number of continuing chunks of memory to allocate
- **size** -- Size, in bytes, of a chunk of memory to allocate
- **num** -- Number of variable paramters

Returns A pointer to the memory allocated on success, NULL on failure

void **heap_caps_dump** (uint32_t caps)

Dump the full structure of all heaps with matching capabilities.

Prints a large amount of output to serial (because of locking limitations, the output bypasses stdout/stderr). For each (variable sized) block in each matching heap, the following output is printed on a single line:

- Block address (the data buffer returned by malloc is 4 bytes after this if heap debugging is set to Basic, or 8 bytes otherwise).
- Data size (the data size may be larger than the size requested by malloc, either due to heap fragmentation or because of heap debugging level).
- Address of next block in the heap.
- If the block is free, the address of the next free block is also printed.

Parameters **caps** -- Bitwise OR of MALLOC_CAP_* flags indicating the type of memory

void **heap_caps_dump_all** (void)

Dump the full structure of all heaps.

Covers all registered heaps. Prints a large amount of output to serial.

Output is the same as for heap_caps_dump.

size_t **heap_caps_get_allocated_size** (void *ptr)

Return the size that a particular pointer was allocated with.

Note: The app will crash with an assertion failure if the pointer is not valid.

Parameters **ptr** -- Pointer to currently allocated heap memory. Must be a pointer value previously returned by heap_caps_malloc, malloc, calloc, etc. and not yet freed.

Returns Size of the memory allocated at this block.

Macros

HEAP_IRAM_ATTR

MALLOC_CAP_EXEC

Flags to indicate the capabilities of the various memory systems.

Memory must be able to run executable code

MALLOC_CAP_32BIT

Memory must allow for aligned 32-bit data accesses.

MALLOC_CAP_8BIT

Memory must allow for 8/16/...-bit data accesses.

MALLOC_CAP_DMA

Memory must be able to accessed by DMA.

MALLOC_CAP_PID2

Memory must be mapped to PID2 memory space (PIDs are not currently used)

MALLOC_CAP_PID3

Memory must be mapped to PID3 memory space (PIDs are not currently used)

MALLOC_CAP_PID4

Memory must be mapped to PID4 memory space (PIDs are not currently used)

MALLOC_CAP_PID5

Memory must be mapped to PID5 memory space (PIDs are not currently used)

MALLOC_CAP_PID6

Memory must be mapped to PID6 memory space (PIDs are not currently used)

MALLOC_CAP_PID7

Memory must be mapped to PID7 memory space (PIDs are not currently used)

MALLOC_CAP_SPIRAM

Memory must be in SPI RAM.

MALLOC_CAP_INTERNAL

Memory must be internal; specifically it should not disappear when flash/spiram cache is switched off.

MALLOC_CAP_DEFAULT

Memory can be returned in a non-capability-specific memory allocation (e.g. malloc(), calloc()) call.

MALLOC_CAP_IRAM_8BIT

Memory must be in IRAM and allow unaligned access.

MALLOC_CAP_RETENTION

Memory must be able to accessed by retention DMA.

MALLOC_CAP_RTCRAM

Memory must be in RTC fast memory.

MALLOC_CAP_TCM

Memory must be in TCM memory.

MALLOC_CAP_INVALID

Memory can't be used / list end marker.

Type Definitions

```
typedef void (*esp_alloc_failed_hook_t)(size_t size, uint32_t caps, const char *function_name)
```

callback called when an allocation operation fails, if registered

Param size in bytes of failed allocation

Param caps capabilities requested of failed allocation

Param function_name function which generated the failure

API Reference - Initialisation

Header File

- [components/heap/include/esp_heap_caps_init.h](#)
- This header file can be included with:

```
#include "esp_heap_caps_init.h"
```

Functions

void **heap_caps_init** (void)

Initialize the capability-aware heap allocator.

This is called once in the IDF startup code. Do not call it at other times.

void **heap_caps_enable_nonos_stack_heaps** (void)

Enable heap(s) in memory regions where the startup stacks are located.

On startup, the pro/app CPUs have a certain memory region they use as stack, so we cannot do allocations in the regions these stack frames are. When FreeRTOS is completely started, they do not use that memory anymore and heap(s) there can be enabled.

[esp_err_t](#) **heap_caps_add_region** (intptr_t start, intptr_t end)

Add a region of memory to the collection of heaps at runtime.

Most memory regions are defined in `soc_memory_layout.c` for the SoC, and are registered via `heap_caps_init()`. Some regions can't be used immediately and are later enabled via `heap_caps_enable_nonos_stack_heaps()`.

Call this function to add a region of memory to the heap at some later time.

This function does not consider any of the "reserved" regions or other data in `soc_memory_layout`, caller needs to consider this themselves.

All memory within the region specified by start & end parameters must be otherwise unused.

The capabilities of the newly registered memory will be determined by the start address, as looked up in the regions specified in `soc_memory_layout.c`.

Use `heap_caps_add_region_with_caps()` to register a region with custom capabilities.

Note: Please refer to following example for memory regions allowed for addition to heap based on an existing region (address range for demonstration purpose only):

```
Existing region: 0x1000 <-> 0x3000
New region:      0x1000 <-> 0x3000 (Allowed)
New region:      0x1000 <-> 0x2000 (Allowed)
New region:      0x0000 <-> 0x1000 (Allowed)
New region:      0x3000 <-> 0x4000 (Allowed)
New region:      0x0000 <-> 0x2000 (NOT Allowed)
New region:      0x0000 <-> 0x4000 (NOT Allowed)
New region:      0x1000 <-> 0x4000 (NOT Allowed)
New region:      0x2000 <-> 0x4000 (NOT Allowed)
```

Parameters

- **start** -- Start address of new region.
- **end** -- End address of new region.

Returns ESP_OK on success, ESP_ERR_INVALID_ARG if a parameter is invalid, ESP_ERR_NOT_FOUND if the specified start address doesn't reside in a known region, or any error returned by heap_caps_add_region_with_caps().

esp_err_t **heap_caps_add_region_with_caps** (const uint32_t caps[], intptr_t start, intptr_t end)

Add a region of memory to the collection of heaps at runtime, with custom capabilities.

Similar to heap_caps_add_region(), only custom memory capabilities are specified by the caller.

Note: Please refer to following example for memory regions allowed for addition to heap based on an existing region (address range for demonstration purpose only):

Existing region:	0x1000 <-> 0x3000	
New region:	0x1000 <-> 0x3000	(Allowed)
New region:	0x1000 <-> 0x2000	(Allowed)
New region:	0x0000 <-> 0x1000	(Allowed)
New region:	0x3000 <-> 0x4000	(Allowed)
New region:	0x0000 <-> 0x2000	(NOT Allowed)
New region:	0x0000 <-> 0x4000	(NOT Allowed)
New region:	0x1000 <-> 0x4000	(NOT Allowed)
New region:	0x2000 <-> 0x4000	(NOT Allowed)

Parameters

- **caps** -- Ordered array of capability masks for the new region, in order of priority. Must have length SOC_MEMORY_TYPE_NO_PRIOS. Does not need to remain valid after the call returns.
- **start** -- Start address of new region.
- **end** -- End address of new region.

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if a parameter is invalid
- ESP_ERR_NO_MEM if no memory to register new heap.
- ESP_ERR_INVALID_SIZE if the memory region is too small to fit a heap
- ESP_FAIL if region overlaps the start and/or end of an existing region

API Reference - Multi-Heap API

(Note: The multi-heap API is used internally by the heap capabilities allocator. Most ESP-IDF programs never need to call this API directly.)

Header File

- [components/heap/include/multi_heap.h](#)
- This header file can be included with:

```
#include "multi_heap.h"
```

Functions

void ***multi_heap_aligned_alloc** (*multi_heap_handle_t* heap, size_t size, size_t alignment)

allocate a chunk of memory with specific alignment

Parameters

- **heap** -- Handle to a registered heap.

- **size** -- size in bytes of memory chunk
- **alignment** -- how the memory must be aligned

Returns pointer to the memory allocated, NULL on failure

void ***multi_heap_malloc** (*multi_heap_handle_t* heap, size_t size)

malloc() a buffer in a given heap

Semantics are the same as standard malloc(), only the returned buffer will be allocated in the specified heap.

Parameters

- **heap** -- Handle to a registered heap.
- **size** -- Size of desired buffer.

Returns Pointer to new memory, or NULL if allocation fails.

void **multi_heap_aligned_free** (*multi_heap_handle_t* heap, void *p)

free() a buffer aligned in a given heap.

Note: This function is deprecated, consider using multi_heap_free() instead

Parameters

- **heap** -- Handle to a registered heap.
- **p** -- NULL, or a pointer previously returned from multi_heap_aligned_alloc() for the same heap.

void **multi_heap_free** (*multi_heap_handle_t* heap, void *p)

free() a buffer in a given heap.

Semantics are the same as standard free(), only the argument 'p' must be NULL or have been allocated in the specified heap.

Parameters

- **heap** -- Handle to a registered heap.
- **p** -- NULL, or a pointer previously returned from multi_heap_malloc() or multi_heap_realloc() for the same heap.

void ***multi_heap_realloc** (*multi_heap_handle_t* heap, void *p, size_t size)

realloc() a buffer in a given heap.

Semantics are the same as standard realloc(), only the argument 'p' must be NULL or have been allocated in the specified heap.

Parameters

- **heap** -- Handle to a registered heap.
- **p** -- NULL, or a pointer previously returned from multi_heap_malloc() or multi_heap_realloc() for the same heap.
- **size** -- Desired new size for buffer.

Returns New buffer of 'size' containing contents of 'p', or NULL if reallocation failed.

size_t **multi_heap_get_allocated_size** (*multi_heap_handle_t* heap, void *p)

Return the size that a particular pointer was allocated with.

Parameters

- **heap** -- Handle to a registered heap.
- **p** -- Pointer, must have been previously returned from multi_heap_malloc() or multi_heap_realloc() for the same heap.

Returns Size of the memory allocated at this block. May be more than the original size argument, due to padding and minimum block sizes.

multi_heap_handle_t **multi_heap_register** (void *start, size_t size)

Register a new heap for use.

This function initialises a heap at the specified address, and returns a handle for future heap operations.

There is no equivalent function for deregistering a heap - if all blocks in the heap are free, you can immediately start using the memory for other purposes.

Parameters

- **start** -- Start address of the memory to use for a new heap.
- **size** -- Size (in bytes) of the new heap.

Returns Handle of a new heap ready for use, or NULL if the heap region was too small to be initialised.

void **multi_heap_set_lock** (*multi_heap_handle_t* heap, void *lock)

Associate a private lock pointer with a heap.

The lock argument is supplied to the MULTI_HEAP_LOCK() and MULTI_HEAP_UNLOCK() macros, defined in multi_heap_platform.h.

The lock in question must be recursive.

When the heap is first registered, the associated lock is NULL.

Parameters

- **heap** -- Handle to a registered heap.
- **lock** -- Optional pointer to a locking structure to associate with this heap.

void **multi_heap_dump** (*multi_heap_handle_t* heap)

Dump heap information to stdout.

For debugging purposes, this function dumps information about every block in the heap to stdout.

Parameters **heap** -- Handle to a registered heap.

bool **multi_heap_check** (*multi_heap_handle_t* heap, bool print_errors)

Check heap integrity.

Walks the heap and checks all heap data structures are valid. If any errors are detected, an error-specific message can be optionally printed to stderr. Print behaviour can be overridden at compile time by defining MULTI_CHECK_FAIL_PRINTF in multi_heap_platform.h.

Note: This function is not thread-safe as it sets a global variable with the value of print_errors.

Parameters

- **heap** -- Handle to a registered heap.
- **print_errors** -- If true, errors will be printed to stderr.

Returns true if heap is valid, false otherwise.

size_t **multi_heap_free_size** (*multi_heap_handle_t* heap)

Return free heap size.

Returns the number of bytes available in the heap.

Equivalent to the total_free_bytes member returned by multi_heap_get_heap_info().

Note that the heap may be fragmented, so the actual maximum size for a single malloc() may be lower. To know this size, see the largest_free_block member returned by multi_heap_get_heap_info().

Parameters **heap** -- Handle to a registered heap.

Returns Number of free bytes.

size_t **multi_heap_minimum_free_size** (*multi_heap_handle_t* heap)

Return the lifetime minimum free heap size.

Equivalent to the minimum_free_bytes member returned by multi_heap_get_info().

Returns the lifetime "low watermark" of possible values returned from multi_free_heap_size(), for the specified heap.

Parameters **heap** -- Handle to a registered heap.

Returns Number of free bytes.

void **multi_heap_get_info** (*multi_heap_handle_t* heap, *multi_heap_info_t* *info)

Return metadata about a given heap.

Fills a *multi_heap_info_t* structure with information about the specified heap.

Parameters

- **heap** -- Handle to a registered heap.
- **info** -- Pointer to a structure to fill with heap metadata.

void ***multi_heap_aligned_alloc_offs** (*multi_heap_handle_t* heap, size_t size, size_t alignment, size_t offset)

Perform an aligned allocation from the provided offset.

Parameters

- **heap** -- The heap in which to perform the allocation
- **size** -- The size of the allocation
- **alignment** -- How the memory must be aligned
- **offset** -- The offset at which the alignment should start

Returns void* The ptr to the allocated memory

Structures

struct **multi_heap_info_t**

Structure to access heap metadata via multi_heap_get_info.

Public Members

size_t **total_free_bytes**

Total free bytes in the heap. Equivalent to multi_free_heap_size().

size_t **total_allocated_bytes**

Total bytes allocated to data in the heap.

size_t **largest_free_block**

Size of the largest free block in the heap. This is the largest malloc-able size.

size_t **minimum_free_bytes**

Lifetime minimum free heap size. Equivalent to multi_minimum_free_heap_size().

size_t **allocated_blocks**

Number of (variable size) blocks allocated in the heap.

size_t **free_blocks**

Number of (variable size) free blocks in the heap.

size_t **total_blocks**

Total number of (variable size) blocks in the heap.

Type Definitions

typedef struct multi_heap_info ***multi_heap_handle_t**

Opaque handle to a registered heap.

2.9.15 Memory Management for MMU Supported Memory

Introduction

ESP32-P4 Memory Management Unit (MMU) is relatively simple. It can do memory address translation between physical memory addresses and virtual memory addresses. So CPU can access physical memories via virtual addresses. There are multiple types of virtual memory addresses, which have different capabilities.

ESP-IDF provides a memory mapping driver that manages the relation between these physical memory addresses and virtual memory addresses, so as to achieve some features such as reading from SPI flash via a pointer.

Memory mapping driver is actually a capabilities-based virtual memory address allocator that allows applications to make virtual memory address allocations for different purposes. In the following chapters, we call this driver `esp_mmap` driver.

ESP-IDF also provides a memory synchronization driver which can be used for potential memory desynchronization scenarios.

Physical Memory Types

Memory mapping driver currently supports mapping to following physical memory type(s):

- SPI flash
- PSRAM

Virtual Memory Capabilities


- `MMU_MEM_CAP_EXEC`: This capability indicates that the virtual memory address has the execute permission. Note this permission scope is within the MMU hardware.
- `MMU_MEM_CAP_READ`: This capability indicates that the virtual memory address has the read permission. Note this permission scope is within the MMU hardware.
- `MMU_MEM_CAP_WRITE`: This capability indicates that the virtual memory address has the write permission. Note this permission scope is within the MMU hardware.
- `MMU_MEM_CAP_32BIT`: This capability indicates that the virtual memory address allows for 32 bits or multiples of 32 bits access.
- `MMU_MEM_CAP_8BIT`: This capability indicates that the virtual memory address allows for 8 bits or multiples of 8 bits access.

You can call `esp_mmu_map_get_max_consecutive_free_block_size()` to know the largest consecutive mappable block size with certain capabilities.

Memory Management Drivers

Driver Concept


Terminology The virtual memory pool is made up with one or multiple virtual memory regions, see below figure:




- A virtual memory pool stands for the whole virtual address range that can be mapped to physical memory.
- A virtual memory region is a range of virtual address with same attributes.
- A virtual memory block is a piece of virtual address range that is dynamically mapped.
- A slot is the virtual address range between two virtual memory blocks.
- A physical memory block is a piece of physical address range that is to-be-mapped or already mapped to a virtual memory block.
- Dynamical mapping is done by calling esp_mmap driver API `esp_mmu_map()`. This API maps the given physical memory block to a virtual memory block which is allocated by the esp_mmap driver.

Relation Between Memory Blocks When mapping a physical memory block A, block A can have one of the following relations with another previously mapped physical memory block B:

- Enclosed: block A is completely enclosed within block B, see figure below:




- Identical: block A is completely the same as block B, see figure below:




Note that esp_mmap driver considers the identical scenario **the same as the enclosed scenario**.

- Overlapped: block A is overlapped with block B, see figure below:



There is a special condition, when block A entirely encloses block B, see figure below:



Note that `esp_mmap` driver considers this scenario **the same as the overlapped scenario**.

Driver Behaviour

Memory Map You can call `esp_mmu_map()` to do a dynamical mapping. This API can allocate a certain size of virtual memory block according to the virtual memory capabilities you selected, then map this virtual memory block to the physical memory block as you requested. The `esp_mmap` driver supports mapping to one or more types of physical memory, so you should specify the physical memory target when mapping.

By default, physical memory blocks and virtual memory blocks are one-to-one mapped. This means, when calling `esp_mmu_map()`:

- If it is the enclosed scenario, this API will return an `ESP_ERR_INVALID_STATE`. The `out_ptr` will be assigned to the start virtual memory address of the previously mapped one which encloses the to-be-mapped one.
- If it is the identical scenario, this API will behave exactly the same as the enclosed scenario.
- If it is the overlapped scenario, this API will by default return an `ESP_ERR_INVALID_ARG`. This means, `esp_mmap` driver by default does not allow mapping a physical memory address to multiple virtual memory addresses.

Specially, you can use `ESP_MMU_MMAP_FLAG_PADDR_SHARED`. This flag stands for one-to-multiple mapping between a physical address and multiple virtual addresses:

- If it is the overlapped scenario, this API will allocate a new virtual memory block as requested, then map to the given physical memory block.

Memory Unmap You can call `esp_mmu_unmap()` to unmap a previously mapped memory block. This API returns an `ESP_ERR_NOT_FOUND` if you are trying to unmap a virtual memory block that is not mapped to any physical memory block yet.

Memory Address Conversion The `esp_mmap` driver provides two helper APIs to do the conversion between virtual memory address and physical memory address:

- `esp_mmu_vaddr_to_paddr()` converts virtual address to physical address.
- `esp_mmu_paddr_to_vaddr()` converts physical address to virtual address.

Memory Synchronization MMU supported physical memories can be accessed by one or multiple methods.

SPI flash can be accessed by SPI1 (ESP-IDF `esp_flash` driver APIs), or by pointers. ESP-IDF `esp_flash` driver APIs have already considered the memory synchronization, so users do not need to worry about this.

PSRAM can be accessed by pointers, hardware guarantees the data consistency when PSRAM is only accessed via pointers.

Thread Safety

Following APIs in `esp_mmu_map.h` are not guaranteed to be thread-safe:

- `esp_mmu_map_dump_mapped_blocks()`

APIs in `esp_cache.h` are guaranteed to be thread-safe.

API Reference

API Reference - ESP MMAP Driver

Header File

- `components/esp_mm/include/esp_mmu_map.h`
- This header file can be included with:

```
#include "esp_mmu_map.h"
```

- This header file is a part of the API provided by the `esp_mm` component. To declare that your component depends on `esp_mm`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_mm
```

or

```
PRIV_REQUIRES esp_mm
```

Functions

`esp_err_t esp_mmu_map(esp_paddr_t paddr_start, size_t size, mmu_target_t target, mmu_mem_caps_t caps, int flags, void **out_ptr)`

Map a physical memory block to external virtual address block, with given capabilities.

Parameters

- **paddr_start** -- [in] Start address of the physical memory block
- **size** -- [in] Size to be mapped. Size will be rounded up by to the nearest multiple of MMU page size
- **target** -- [in] Physical memory target you're going to map to, see `mmu_target_t`
- **caps** -- [in] Memory capabilities, see `mmu_mem_caps_t`
- **flags** -- [in] Mmap flags
- **out_ptr** -- [out] Start address of the mapped virtual memory

Returns

- `ESP_OK`
- `ESP_ERR_INVALID_ARG`: Invalid argument, see printed logs
- `ESP_ERR_NOT_SUPPORTED`: Only on ESP32, PSRAM is not a supported physical memory target
- `ESP_ERR_NOT_FOUND`: No enough size free block to use
- `ESP_ERR_NO_MEM`: Out of memory, this API will allocate some heap memory for internal usage

- **ESP_ERR_INVALID_STATE**: Paddr is mapped already, this API will return corresponding `vaddr_start + new_block_offset` as per the previously mapped block. Only to-be-mapped paddr block is totally enclosed by a previously mapped block will lead to this error. (Identical scenario will behave similarly) `new_block_start new_block_end` |-----— New Block -----—| |-----— Block -----—|
block_start block_end

esp_err_t **esp_mmu_unmap** (void *ptr)

Unmap a previously mapped virtual memory block.

Parameters **ptr** -- [in] Start address of the virtual memory

Returns

- **ESP_OK**
- **ESP_ERR_INVALID_ARG**: Null pointer
- **ESP_ERR_NOT_FOUND**: Vaddr is not in external memory, or it's not mapped yet

esp_err_t **esp_mmu_map_get_max_consecutive_free_block_size** (mmu_mem_caps_t caps, mmu_target_t target, size_t *out_len)

Get largest consecutive free external virtual memory block size, with given capabilities and given physical target.

Parameters

- **caps** -- [in] Bitwise OR of **MMU_MEM_CAP_*** flags indicating the memory block
- **target** -- [in] Physical memory target you're going to map to, see `mmu_target_t`.
- **out_len** -- [out] Largest free block length, in bytes.

Returns

- **ESP_OK**
- **ESP_ERR_INVALID_ARG**: Invalid arguments, could be null pointer

esp_err_t **esp_mmu_map_dump_mapped_blocks** (FILE *stream)

Dump all the previously mapped blocks

Note: This API shall not be called from an ISR.

Note: This API does not guarantee thread safety

Parameters **stream** -- stream to print information to; use `stdout` or `stderr` to print to the console; use `fmemopen/open_memstream` to print to a string buffer.

Returns

- **ESP_OK**

esp_err_t **esp_mmu_vaddr_to_paddr** (void *vaddr, *esp_paddr_t* *out_paddr, mmu_target_t *out_target)

Convert virtual address to physical address.

Parameters

- **vaddr** -- [in] Virtual address
- **out_paddr** -- [out] Physical address
- **out_target** -- [out] Physical memory target, see `mmu_target_t`

Returns

- **ESP_OK**
- **ESP_ERR_INVALID_ARG**: Null pointer, or vaddr is not within external memory
- **ESP_ERR_NOT_FOUND**: Vaddr is not mapped yet

esp_err_t **esp_mmu_paddr_to_vaddr** (*esp_paddr_t* paddr, mmu_target_t target, mmu_vaddr_t type, void **out_vaddr)

Convert physical address to virtual address.

Parameters

- **paddr** -- [in] Physical address
- **target** -- [in] Physical memory target, see `mmu_target_t`
- **type** -- [in] Virtual address type, could be either instruction or data
- **out_vaddr** -- [out] Virtual address

Returns

- `ESP_OK`
- `ESP_ERR_INVALID_ARG`: Null pointer
- `ESP_ERR_NOT_FOUND`: Paddr is not mapped yet

esp_err_t **esp_mmu_paddr_find_caps** (const *esp_paddr_t* paddr, *mmu_mem_caps_t* *out_caps)

If the physical address is mapped, this API will provide the capabilities of the virtual address where the physical address is mapped to.

Note: : Only return value is `ESP_OK`(which means physically address is successfully mapped), then caps you get make sense.

Note: This API only check one page (see `CONFIG_MMU_PAGE_SIZE`), starting from the paddr

Parameters

- **paddr** -- [in] Physical address
- **out_caps** -- [out] Bitwise OR of `MMU_MEM_CAP_*` flags indicating the capabilities of a virtual address where the physical address is mapped to.

Returns

- `ESP_OK`: Physical address successfully mapped.
- `ESP_ERR_INVALID_ARG`: Null pointer
- `ESP_ERR_NOT_FOUND`: Physical address is not mapped successfully.

Macros

ESP_MMU_MMAP_FLAG_PADDR_SHARED

Share this mapping.

MMU Memory Mapping Driver APIs for MMU supported memory

Driver Backgrounds:

Type Definitions

typedef uint32_t **esp_paddr_t**

Physical memory type.

2.9.16 Memory Synchronization

Introduction

ESP32-P4 can access its connected PSRAM via these ways:

- CPU
- DMA

ESP32-P4 can access its internal memory via these ways:

- CPU
- DMA

By default, CPU accesses the above mentioned memory via cache. Whereas DMA accesses the memory directly, without going through cache.

This leads to potential cache data coherence issue:

- When a DMA transaction changes the content of a piece of memory, and the content has been cached already. Under this condition:
 - CPU may read stale data.
 - the stale data in the cache may be written back to the memory. The new data updated by the previous DMA transaction will be overwritten.
- CPU changes the content of an address. The content is in the cache, but not in the memory yet (cache will write back the content to the memory according to its own strategy). Under this condition:
 - The next DMA transactions to read this content from the memory will get stale data.

There are three common methods to address such cache data coherence issue:

1. Hardware based cache Coherent Interconnect, ESP32-P4 does not have such ability.
2. Use the DMA buffer from non-cacheable memory. Memory that CPU access it without going through cache is called non-cacheable memory.
3. Explicitly call a memory synchronization API to writeback the content in the cache back to the memory, or invalidate the content in the cache.

Memory Synchronisation Driver

The suggested way to deal with such cache data coherence issue is by using the memory synchronization API `esp_cache_msync()` provided by ESP-IDF `esp_mm` component.

Driver Concept Direction of the cache memory synchronization:

- `ESP_CACHE_MSINC_FLAG_DIR_C2M`, for synchronization from cache to memory.
- `ESP_CACHE_MSINC_FLAG_DIR_M2C`, for synchronization from memory to cache.

Type of the cache memory synchronization:

- `ESP_CACHE_MSINC_FLAG_TYPE_DATA`, for synchronization to a data address region.
- `ESP_CACHE_MSINC_FLAG_TYPE_INST`, for synchronization to an instruction address region.

Driver Behaviour Calling `esp_cache_msync()` will do a synchronization between cache and memory. The first parameter `addr` and the second parameter `size` together describe the memory region that is to be synchronized. About the third parameter `flags`:

- `ESP_CACHE_MSINC_FLAG_DIR_C2M`. With this flag, content in the specified address region is written back to the memory. This direction is usually used **after** the content of an address is updated by the CPU, e.g. a memset to the address. Operation in this direction should happen **before** a DMA operation to the same address.
- `ESP_CACHE_MSINC_FLAG_DIR_M2C`. With this flag, content in the specified address region is invalidated from the cache. This direction is usually used **after** the content of an address is updated by the DMA. Operation in this direction should happen **before** a CPU read operation to the same address.

The above two flags help select the synchronization direction. Specially, if neither of these two flags are used, `esp_cache_msync()` will by default select the `ESP_CACHE_MSINC_FLAG_DIR_C2M` direction. Users are not allowed to set both of the two flags at the same time.

- `ESP_CACHE_MSINC_FLAG_TYPE_DATA`.
- `ESP_CACHE_MSINC_FLAG_TYPE_INST`.

The above two flags help select the type of the synchronization address. Specially, if neither of these two flags are used, `esp_cache_msync()` will by default select the `ESP_CACHE_MSINC_FLAG_TYPE_DATA` direction. Users are not allowed to set both of the two flags at the same time.

- `ESP_CACHE_MSINC_FLAG_INVALIDATE`. This flag is used to trigger a cache invalidation to the specified address region, after the region is written back to the memory. This flag is mainly used for `ESP_CACHE_MSINC_FLAG_DIR_C2M` direction. For `ESP_CACHE_MSINC_FLAG_DIR_M2C` direction, behaviour is the same as if the `ESP_CACHE_MSINC_FLAG_INVALIDATE` flag is not set.
- `ESP_CACHE_MSINC_FLAG_UNALIGNED`. This flag force the `esp_cache_msync()` API to do synchronization without checking the address and size alignment. For more details, see chapter *Address Alignment Requirement* following.

Address Alignment Requirement

There is address and size alignment requirement (in bytes) for using `esp_cache_msync()`. The alignment requirement comes from cache.

- An address region whose start address and size both meet the cache memory synchronization alignment requirement is defined as an **aligned address region**.
- An address region whose start address or size does not meet the cache memory synchronization alignment requirement is defined as an **unaligned address region**.

By default, if you specify an unaligned address region, `esp_cache_msync()` will return an `ESP_ERR_INVALID_ARG` error, together with the required alignment.

Memory Allocation Helper cache memory synchronization is usually considered when DMA is involved. ESP-IDF provides an API to do memory allocation that can meet the alignment requirement from both the cache and the DMA.

- `esp_dma_malloc()`, this API allocates a chunk of memory that meets the alignment requirement from both the cache and the DMA.
- `esp_dma_calloc()`, this API allocates a chunk of memory that meets the alignment requirement from both the cache and the DMA. The initialized value in the memory is set to zero.

You can also use `ESP_DMA_MALLOC_FLAG_PSRAM` to allocate from the PSRAM.


Warning for Address Alignment Requirement You can set the `ESP_CACHE_MSINC_FLAG_UNALIGNED` flag to bypass such check. Note you should be very careful about using this flag. cache memory synchronization to an unaligned address region may silently corrupt the memory.

For example, assume:

- alignment requirement is 0x40 bytes.
- a call to `esp_cache_msync()`, with `ESP_CACHE_MSINC_FLAG_DIR_M2C` | `ESP_CACHE_MSINC_FLAG_UNALIGNED` flags, the specified address region is 0x4000_0020 ~ 0x4000_0060 (see **data C** in below graph).

Above settings will trigger a cache invalidation to the address region 0x4000_0000 ~ 0x4000_0080, see **sync item0** and **sync item1** in the below graph.

If the content in 0x4000_0000 ~ 0x4000_0020 (**data A** in the below graph) or 0x4000_0060 ~ 0x4000_0080 (**data B** in the below graph) are not written back to the memory yet, then these **data A** and **data B** will be discarded.



API Reference

API Reference - ESP Msync Driver

Header File

- [components/esp_mm/include/esp_cache.h](#)
- This header file can be included with:

```
#include "esp_cache.h"
```

- This header file is a part of the API provided by the `esp_mm` component. To declare that your component depends on `esp_mm`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_mm
```

or

```
PRIV_REQUIRES esp_mm
```

Functions

`esp_err_t esp_cache_msync` (void *addr, size_t size, int flags)

Memory sync between Cache and storage memory.

For cache-to-memory (C2M) direction:

- For cache writeback supported chips (you can refer to `SOC_CACHE_WRITEBACK_SUPPORTED` in `soc_caps.h`)
 - This API will do a writeback to synchronise between cache and storage memory

- With `ESP_CACHE_MSINC_FLAG_INVALIDATE`, this API will also invalidate the values that just written
- Note: although ESP32 is with PSRAM, but cache writeback isn't supported, so this API will do nothing on ESP32
- For other chips, this API will do nothing. The out-of-sync should be already dealt by the SDK

For memory-to-cache (M2C) direction:

- This API will by default do an invalidation

This API is cache-safe and thread-safe

Note: If you don't set direction (`ESP_CACHE_MSINC_FLAG_DIR_x` flags), this API is by default C2M direction

Note: If you don't set type (`ESP_CACHE_MSINC_FLAG_TYPE_x` flags), this API is by default doing msync for data

Note: You should not call this during any Flash operations (e.g. `esp_flash` APIs, `nvs` and some other APIs that are based on `esp_flash` APIs)

Note: If `XIP_From_PSRAM` is enabled (by enabling both `CONFIG_SPIRAM_FETCH_INSTRUCTIONS` and `CONFIG_SPIRAM_RODATA`), you can call this API during Flash operations

Parameters

- **addr** -- [in] Starting address to do the msync
- **size** -- [in] Size to do the msync
- **flags** -- [in] Flags, see `ESP_CACHE_MSINC_FLAG_x`

Returns

- `ESP_OK`:
 - Successful msync
 - For C2M direction, if this chip doesn't support cache writeback, if the input `addr` is a cache supported one, this API will return `ESP_OK`
- `ESP_ERR_INVALID_ARG`: Invalid argument, not cache supported `addr`, see printed logs

Macros

`ESP_CACHE_MSINC_FLAG_INVALIDATE`

Do an invalidation.

Cache msync flags

- For cache-to-memory (C2M) direction: setting this flag will start an invalidation after the cache writeback operation
- For memory-to-cache (M2C) direction: setting / unsetting this flag will behave similarly, trigger an invalidation

`ESP_CACHE_MSINC_FLAG_UNALIGNED`

Allow msync to a address block that are not aligned to the data cache line size.

`ESP_CACHE_MSINC_FLAG_DIR_C2M`

Cache msync direction: from Cache to memory.

Note: If you don't set direction (ESP_CACHE_MSINC_FLAG_DIR_x flags), it is by default cache-to-memory (C2M) direction

ESP_CACHE_MSINC_FLAG_DIR_M2C

Cache msync direction: from memory to Cache.

ESP_CACHE_MSINC_FLAG_TYPE_DATA

Cache msync type: data.

Note: If you don't set type (ESP_CACHE_MSINC_FLAG_TYPE_x flags), it is by default data type

ESP_CACHE_MSINC_FLAG_TYPE_INST

Cache msync type: instruction.

API Reference - ESP DMA Utils

Header File

- [components/esp_hw_support/include/esp_dma_utils.h](#)
- This header file can be included with:

```
#include "esp_dma_utils.h"
```

Functions

esp_err_t **esp_dma_malloc** (size_t size, uint32_t flags, void **out_ptr, size_t *actual_size)

Helper function for malloc a DMA capable memory buffer.

Parameters

- **size** -- **[in]** Size in bytes, the amount of memory to allocate
- **flags** -- **[in]** Flags, see ESP_DMA_MALLOC_FLAG_x
- **out_ptr** -- **[out]** A pointer to the memory allocated successfully
- **actual_size** -- **[out]** Actual size for allocation in bytes, when the size you specified doesn't meet the DMA alignment requirements, this value might be bigger than the size you specified. Set null if you don't care this value.

Returns

- ESP_OK:
- ESP_ERR_INVALID_ARG: Invalid argument
- ESP_ERR_NO_MEM: No enough memory for allocation

esp_err_t **esp_dma_calloc** (size_t n, size_t size, uint32_t flags, void **out_ptr, size_t *actual_size)

Helper function for calloc a DMA capable memory buffer.

Parameters

- **n** -- **[in]** Number of continuing chunks of memory to allocate
- **size** -- **[in]** Size of one chunk, in bytes
- **flags** -- **[in]** Flags, see ESP_DMA_MALLOC_FLAG_x
- **out_ptr** -- **[out]** A pointer to the memory allocated successfully
- **actual_size** -- **[out]** Actual size for allocation in bytes, when the size you specified doesn't meet the cache alignment requirements, this value might be bigger than the size you specified. Set null if you don't care this value.

Returns

- ESP_OK:
- ESP_ERR_INVALID_ARG: Invalid argument

- `ESP_ERR_NO_MEM`: No enough memory for allocation

bool **esp_dma_is_buffer_aligned** (const void *ptr, size_t size, [esp_dma_buf_location_t](#) location)

Helper function to check if a buffer meets DMA alignment requirements.

Parameters

- **ptr** -- [in] Pointer to the buffer
- **size** -- [in] Size of the buffer
- **location** -- [in] Location of the DMA buffer, see [esp_dma_buf_location_t](#)

Returns

- True: Buffer is aligned
- False: Buffer is not aligned, or buffer is not DMA capable

Macros

ESP_DMA_MALLOC_FLAG_PSRAM

Memory is in PSRAM.

DMA malloc flags

Enumerations

enum **esp_dma_buf_location_t**

DMA buffer location.

Values:

enumerator **ESP_DMA_BUF_LOCATION_INTERNAL**

DMA buffer is in internal memory.

enumerator **ESP_DMA_BUF_LOCATION_PSRAM**

DMA buffer is in PSRAM.

2.9.17 Heap Memory Debugging

Overview

ESP-IDF integrates tools for requesting [heap information](#), [heap corruption detection](#), and [heap tracing](#). These can help track down memory-related bugs.

For general information about the heap memory allocator, see [Heap Memory Allocation](#).

Heap Information

To obtain information about the state of the heap, call the following functions:

- [heap_caps_get_free_size\(\)](#) can be used to return the current free memory for different memory capabilities.
- [heap_caps_get_largest_free_block\(\)](#) can be used to return the largest free block in the heap, which is also the largest single allocation currently possible. Tracking this value and comparing it to the total free heap allows you to detect heap fragmentation.
- [heap_caps_get_minimum_free_size\(\)](#) can be used to track the heap "low watermark" since boot.
- [heap_caps_get_info\(\)](#) returns a [multi_heap_info_t](#) structure, which contains the information from the above functions, plus some additional heap-specific data (number of allocations, etc.).

- `heap_caps_print_heap_info()` prints a summary of the information returned by `heap_caps_get_info()` to stdout.
- `heap_caps_dump()` and `heap_caps_dump_all()` output detailed information about the structure of each block in the heap. Note that this can be a large amount of output.

Heap Allocation and Free Function Hooks

Heap allocation and free detection hooks allow you to be notified of every successful allocation and free operation:

- Providing a definition of `esp_heap_trace_alloc_hook()` allows you to be notified of every successful memory allocation operation
- Providing a definition of `esp_heap_trace_free_hook()` allows you to be notified of every successful memory-free operations

This feature can be enabled by setting the `CONFIG_HEAP_USE_HOOKS` option. `esp_heap_trace_alloc_hook()` and `esp_heap_trace_free_hook()` have weak declarations (e.g., `__attribute__((weak))`), thus it is not necessary to provide declarations for both hooks. Given that it is technically possible to allocate and free memory from an ISR (**though strongly discouraged from doing so**), the `esp_heap_trace_alloc_hook()` and `esp_heap_trace_free_hook()` can potentially be called from an ISR.

It is not recommended to perform (or call API functions to perform) blocking operations or memory allocation/free operations in the hook functions. In general, the best practice is to keep the implementation concise and leave the heavy computation outside of the hook functions.

The example below shows how to define the allocation and free function hooks:

```
#include "esp_heap_caps.h"

void esp_heap_trace_alloc_hook(void* ptr, size_t size, uint32_t caps)
{
    ...
}

void esp_heap_trace_free_hook(void* ptr)
{
    ...
}

void app_main()
{
    ...
}
```

Heap Corruption Detection

Heap corruption detection allows you to detect various types of heap memory errors:

- Out-of-bound writes & buffer overflows
- Writes to freed memory
- Reads from freed or uninitialized memory

Assertions The heap implementation (`heap/multi_heap.c`, etc.) includes numerous assertions that will fail if the heap memory is corrupted. To detect heap corruption most effectively, ensure that assertions are enabled in the project configuration via the `CONFIG_COMPILER_OPTIMIZATION_ASSERTION_LEVEL` option.

If a heap integrity assertion fails, a line will be printed like `CORRUPT HEAP: multi_heap.c:225 detected at 0x3ffbb71c`. The memory address printed is the address of the heap structure that has corrupt content.

It is also possible to manually check heap integrity by calling `heap_caps_check_integrity_all()` or related functions. This function checks all of the requested heap memory for integrity and can be used even if assertions

are disabled. If the integrity checks detects an error, it will print the error along with the address(es) of corrupt heap structures.

Memory Allocation Failed Hook Users can use `heap_caps_register_failed_alloc_callback()` to register a callback that is invoked every time an allocation operation fails.

Additionally, users can enable the `CONFIG_HEAP_ABORT_WHEN_ALLOCATION_FAILS`, which will automatically trigger a system abort if any allocation operation fails.

The example below shows how to register an allocation failure callback:

```
#include "esp_heap_caps.h"

void heap_caps_alloc_failed_hook(size_t requested_size, uint32_t caps, const char*
↪function_name)
{
    printf("%s was called but failed to allocate %d bytes with 0x%X capabilities. \n
↪", function_name, requested_size, caps);
}

void app_main()
{
    ...
    esp_err_t error = heap_caps_register_failed_alloc_callback(heap_caps_alloc_
↪failed_hook);
    ...
    void *ptr = heap_caps_malloc(allocation_size, MALLOC_CAP_DEFAULT);
    ...
}
```

Finding Heap Corruption Memory corruption can be one of the hardest classes of bugs to find and fix, as the source of the corruption could be completely unrelated to the symptoms of the corruption. Here are some tips:

- A crash with a `CORRUPT HEAP :` message usually includes a stack trace, but this stack trace is rarely useful. The crash is the symptom of memory corruption when the system realizes the heap is corrupt. But usually, the corruption happens elsewhere and earlier in time.
- Increasing the heap memory debugging [Configuration](#) level to "Light impact" or "Comprehensive" gives you a more accurate message with the first corrupt memory address.
- Adding regular calls to `heap_caps_check_integrity_all()` or `heap_caps_check_integrity_addr()` in your code helps you pin down the exact time that the corruption happened. You can move these checks around to "close in on" the section of code that corrupted the heap.
- Based on the memory address that has been corrupted, you can use [JTAG debugging](#) to set a watchpoint on this address and have the CPU halt when it is written to.
- If you do not have JTAG, but you do know roughly when the corruption happens, set a watchpoint in software just beforehand via `esp_cpu_set_watchpoint()`. A fatal exception will occur when the watchpoint triggers. The following is an example of how to use the function - `esp_cpu_set_watchpoint(0, (void *)addr, 4, ESP_WATCHPOINT_STORE)`. Note that watchpoints are per-CPU and are set on the current running CPU only. So if you do not know which CPU is corrupting memory, call this function on both CPUs.
- For buffer overflows, [heap tracing](#) in `HEAP_TRACE_ALL` mode tells which callers are allocating which addresses from the heap. See [Heap Tracing To Find Heap Corruption](#) for more details. You can try to find the function that allocates memory with an address immediately before the corrupted address, since it is probably the function that overflows the buffer.
- Calling `heap_caps_dump()` or `heap_caps_dump_all()` can give an indication of what heap blocks are surrounding the corrupted region and may have overflowed or underflowed, etc.

Configuration Temporarily increasing the heap corruption detection level can give more detailed information about heap corruption errors.

In the project configuration menu, under `Component config`, there is a menu `Heap memory debugging`. The option `CONFIG_HEAP_CORRUPTION_DETECTION` can be set to one of the following three levels:

Basic (No Poisoning) This is the default level. By default, no special heap corruption features are enabled, but the provided assertions are enabled. A heap corruption error will be printed if any of the heap's internal data structures appear overwritten or corrupted. This usually indicates a buffer overrun or out-of-bounds write.

If assertions are enabled, an assertion will also trigger if a double-free occurs (the same memory is freed twice).

Calling `heap_caps_check_integrity()` in Basic mode checks the integrity of all heap structures, and print errors if any appear to be corrupted.

Light Impact At this level, heap memory is additionally "poisoned" with head and tail "canary bytes" before and after each block that is allocated. If an application writes outside the bounds of allocated buffers, the canary bytes will be corrupted, and the integrity check will fail.

The head canary word is `0xABBA1234` (3412BAAB in byte order), and the tail canary word is `0xBAAD5678` (7856ADBA in byte order).

With basic heap corruption checks, most out-of-bound writes can be detected and the number of overrun bytes before a failure is detected depends on the properties of the heap. However, the Light Impact mode is more precise as even a single-byte overrun can be detected.

Enabling light-impact checking increases the memory usage. Each individual allocation uses 9 to 12 additional bytes of memory depending on alignment.

Each time `heap_caps_free()` is called in Light Impact mode, the head and tail canary bytes of the buffer being freed are checked against the expected values.

When `heap_caps_check_integrity()` is called, all allocated blocks of heap memory have their canary bytes checked against the expected values.

In both cases, the functions involve checking that the first 4 bytes of an allocated block (before the buffer is returned to the user) should be the word `0xABBA1234`, and the last 4 bytes of the allocated block (after the buffer is returned to the user) should be the word `0xBAAD5678`.

Different values usually indicate buffer underrun or overrun. Overrun indicates that when writing to memory, the data written exceeds the size of the allocated memory, resulting in writing to an unallocated memory area; underrun indicates that when reading memory, the data read exceeds the allocated memory and reads data from an unallocated memory area.

Comprehensive This level incorporates the "light impact" detection features plus additional checks for uninitialized-access and use-after-free bugs. In this mode, all freshly allocated memory is filled with the pattern `0xCE`, and all freed memory is filled with the pattern `0xFE`.

Enabling Comprehensive mode has a substantial impact on runtime performance, as all memory needs to be set to the allocation patterns each time a `heap_caps_malloc()` or `heap_caps_free()` completes, and the memory also needs to be checked each time. However, this mode allows easier detection of memory corruption bugs which are much more subtle to find otherwise. It is recommended to only enable this mode when debugging, not in production.

Crashes in Comprehensive Mode If an application crashes when reading or writing an address related to `0xCE-CECECE` in Comprehensive mode, it indicates that it has read uninitialized memory. The application should be changed to either use `heap_caps_calloc()` (which zeroes memory), or initialize the memory before using it. The value `0xCECECECE` may also be seen in stack-allocated automatic variables, because, in ESP-IDF, most task stacks are originally allocated from the heap, and in C, stack memory is uninitialized by default.

If an application crashes, and the exception register dump indicates that some addresses or values were `0xFEFE-FEFE`, this indicates that it is reading heap memory after it has been freed, i.e., a "use-after-free bug". The application should be changed to not access heap memory after it has been freed.

If a call to `heap_caps_malloc()` or `heap_caps_realloc()` causes a crash because it was expected to find the pattern `0xFEFEFEFE` in free memory and a different pattern was found, it indicates that the app has a use-after-free bug where it is writing to memory that has already been freed.

Manual Heap Checks in Comprehensive Mode Calls to `heap_caps_check_integrity()` may print errors relating to `0xFEFEFEFE`, `0xABBA1234`, or `0xBAAD5678`. In each case the checker is expected to find a given pattern, and will error out if not found:

- For free heap blocks, the checker expects to find all bytes set to `0xFE`. Any other values indicate a use-after-free bug where free memory has been incorrectly overwritten.
- For allocated heap blocks, the behavior is the same as for the Light Impact mode. The canary bytes `0xABBA1234` and `0xBAAD5678` are checked at the head and tail of each allocated buffer, and any variation indicates a buffer overrun or underrun.

Heap Task Tracking

Heap Task Tracking can be used to get per-task info for heap memory allocation. The application has to specify the heap capabilities for which the heap allocation is to be tracked.

Example code is provided in [system/heap_task_tracking](#).

Heap Tracing

Heap Tracing allows the tracing of code which allocates or frees memory. Two tracing modes are supported:

- Standalone. In this mode, traced data are kept on-board, so the size of the gathered information is limited by the buffer assigned for that purpose, and the analysis is done by the on-board code. There are a couple of APIs available for accessing and dumping collected info.
- Host-based. This mode does not have the limitation of the standalone mode, because traced data are sent to the host over JTAG connection using `app_trace` library. Later on, they can be analyzed using special tools.

Heap tracing can perform two functions:

- Leak checking: find memory that is allocated and never freed.
- Heap use analysis: show all functions that are allocating or freeing memory while the trace is running.

How to Diagnose Memory Leaks If you suspect a memory leak, the first step is to figure out which part of the program is leaking memory. Use the `heap_caps_get_free_size()` or related functions in [heap information](#) to track memory use over the life of the application. Try to narrow the leak down to a single function or sequence of functions where free memory always decreases and never recovers.

Standalone Mode Once you have identified the code which you think is leaking:

- Enable the `CONFIG_HEAP_TRACING_DEST` option.
- Call the function `heap_trace_init_standalone()` early in the program, to register a buffer that can be used to record the memory trace.
- Call the function `heap_trace_start()` to begin recording all mallocs or frees in the system. Call this immediately before the piece of code which you suspect is leaking memory.
- Call the function `heap_trace_stop()` to stop the trace once the suspect piece of code has finished executing.
- Call the function `heap_trace_dump()` to dump the results of the heap trace.

The following code snippet demonstrates how application code would typically initialize, start, and stop heap tracing:

```
#include "esp_heap_trace.h"

#define NUM_RECORDS 100
static heap_trace_record_t trace_record[NUM_RECORDS]; // This buffer must be in
↳ internal RAM
```

(continues on next page)

(continued from previous page)

```

...

void app_main()
{
    ...
    ESP_ERROR_CHECK( heap_trace_init_standalone(trace_record, NUM_RECORDS) );
    ...
}

void some_function()
{
    ESP_ERROR_CHECK( heap_trace_start(HEAP_TRACE_LEAKS) );

    do_something_you_suspect_is_leaking();

    ESP_ERROR_CHECK( heap_trace_stop() );
    heap_trace_dump();
    ...
}

```

The output from the heap trace has a similar format to the following example:

```

2 allocations trace (100 entry buffer)
32 bytes (@ 0x3ffaf214) allocated CPU 0 ccount 0x2e9b7384 caller
8 bytes (@ 0x3ffaf804) allocated CPU 0 ccount 0x2e9b79c0 caller
40 bytes 'leaked' in trace (2 allocations)
total allocations 2 total frees 0

```

Note: The above example output uses *IDF Monitor* to automatically decode PC addresses to their source files and line numbers.

The first line indicates how many allocation entries are in the buffer, compared to its total size.

In `HEAP_TRACE_LEAKS` mode, for each traced memory allocation that has not already been freed, a line is printed with:

- `XX bytes` is the number of bytes allocated.
- `@ 0x...` is the heap address returned from `heap_caps_malloc()` or `heap_caps_calloc()`.
- Internal or PSRAM is the general location of the allocated memory.
- `CPU x` is the CPU (0 or 1) running when the allocation was made.
- `ccount 0x...` is the CCOUNT (CPU cycle count) register value the allocation was made. The value is different for CPU 0 vs CPU 1.

Finally, the total number of the 'leaked' bytes (bytes allocated but not freed while the trace is running) is printed together with the total number of allocations it represents.

A warning will be printed if the trace buffer was not large enough to hold all the allocations happened. If you see this warning, consider either shortening the tracing period or increasing the number of records in the trace buffer.

Host-Based Mode Once you have identified the code which you think is leaking:

- In the project configuration menu, navigate to Component settings > Heap Memory Debugging > `CONFIG_HEAP_TRACING_DEST` and select Host-Based.
- In the project configuration menu, navigate to Component settings > Application Level Tracing > `CONFIG_APPTRACE_DESTINATION1` and select Trace memory.
- In the project configuration menu, navigate to Component settings > Application Level Tracing > FreeRTOS SystemView Tracing and enable `CONFIG_APPTRACE_SV_ENABLE`.

- Call the function `heap_trace_init_tohost()` early in the program, to initialize the JTAG heap tracing module.
- Call the function `heap_trace_start()` to begin recording all memory allocation and free calls in the system. Call this immediately before the piece of code which you suspect is leaking memory.
In host-based mode, the argument to this function is ignored, and the heap tracing module behaves like `HEAP_TRACE_ALL` is passed, i.e., all allocations and deallocations are sent to the host.
- Call the function `heap_trace_stop()` to stop the trace once the suspect piece of code has finished executing.

The following code snippet demonstrates how application code would typically initialize, start, and stop host-based mode heap tracing:

```
#include "esp_heap_trace.h"

...

void app_main()
{
    ...
    ESP_ERROR_CHECK( heap_trace_init_tohost() );
    ...
}

void some_function()
{
    ESP_ERROR_CHECK( heap_trace_start(HEAP_TRACE_LEAKS) );

    do_something_you_suspect_is_leaking();

    ESP_ERROR_CHECK( heap_trace_stop() );
    ...
}
```

To gather and analyze heap trace, do the following on the host:

1. Build the program and download it to the target as described in [Step 5. First Steps on ESP-IDF](#).
2. Run OpenOCD (see [JTAG Debugging](#)).

Note: In order to use this feature, you need OpenOCD version `v0.10.0-esp32-20181105` or later.

3. You can use GDB to start and/or stop tracing automatically. To do this you need to prepare a special `gdbinit` file:

```
target remote :3333

mon reset halt
maintenance flush register-cache

tb heap_trace_start
commands
mon esp sysview start file:///tmp/heap.svdat
c
end

tb heap_trace_stop
commands
mon esp sysview stop
end

c
```

Using this file GDB can connect to the target, reset it, and start tracing when the program hits breakpoint at

`heap_trace_start()`. Tracing will be stopped when the program hits breakpoint at `heap_trace_stop()`. Traced data will be saved to `/tmp/heap_log.svdat`.

4. Run GDB using `riscv32-esp-elf-gdb -x gdbinit </path/to/program/elf>`.
5. Quit GDB when the program stops at `heap_trace_stop()`. Traced data are saved in `/tmp/heap.svdat`.
6. Run processing script `$IDF_PATH/tools/esp_app_trace/sysviewtrace_proc.py -p -b </path/to/program/elf> /tmp/heap_log.svdat`.

The output from the heap trace has a similar format to the following example:

```
Parse trace from '/tmp/heap.svdat'...
Stop parsing trace. (Timeout 0.000000 sec while reading 1 bytes!)
Process events from '['/tmp/heap.svdat']'...
[0.002244575] HEAP: Allocated 1 bytes @ 0x3ffaafd8 from task "alloc" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↳sysview_heap_log.c:47
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

[0.002258425] HEAP: Allocated 2 bytes @ 0x3ffaaffe0 from task "alloc" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↳sysview_heap_log.c:48
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

[0.002563725] HEAP: Freed bytes @ 0x3ffaaffe0 from task "free" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↳sysview_heap_log.c:31 (discriminator 9)
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

[0.002782950] HEAP: Freed bytes @ 0x3ffb40b8 from task "main" on core 0 by:
/home/user/projects/esp/esp-idf/components/freertos/tasks.c:4590
/home/user/projects/esp/esp-idf/components/freertos/tasks.c:4590

[0.002798700] HEAP: Freed bytes @ 0x3ffb50bc from task "main" on core 0 by:
/home/user/projects/esp/esp-idf/components/freertos/tasks.c:4590
/home/user/projects/esp/esp-idf/components/freertos/tasks.c:4590

[0.102436025] HEAP: Allocated 2 bytes @ 0x3ffaaffe0 from task "alloc" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↳sysview_heap_log.c:47
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

[0.102449800] HEAP: Allocated 4 bytes @ 0x3ffaaffe8 from task "alloc" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↳sysview_heap_log.c:48
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

[0.102666150] HEAP: Freed bytes @ 0x3ffaaffe8 from task "free" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↳sysview_heap_log.c:31 (discriminator 9)
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

[0.202436200] HEAP: Allocated 3 bytes @ 0x3ffaaffe8 from task "alloc" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↳sysview_heap_log.c:47
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

[0.202451725] HEAP: Allocated 6 bytes @ 0x3ffaaff0 from task "alloc" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↳sysview_heap_log.c:48
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

[0.202667075] HEAP: Freed bytes @ 0x3ffaaff0 from task "free" on core 0 by:
```

(continues on next page)

(continued from previous page)

```

/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↪sysview_heap_log.c:31 (discriminator 9)
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

[0.302436000] HEAP: Allocated 4 bytes @ 0x3ffafff0 from task "alloc" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↪sysview_heap_log.c:47
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

[0.302451475] HEAP: Allocated 8 bytes @ 0x3ffb40b8 from task "alloc" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↪sysview_heap_log.c:48
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

[0.302667500] HEAP: Freed bytes @ 0x3ffb40b8 from task "free" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↪sysview_heap_log.c:31 (discriminator 9)
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

Processing completed.

Processed 1019 events

===== HEAP TRACE REPORT =====

Processed 14 heap events.

[0.002244575] HEAP: Allocated 1 bytes @ 0x3ffaafd8 from task "alloc" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↪sysview_heap_log.c:47
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

[0.102436025] HEAP: Allocated 2 bytes @ 0x3ffaaffe0 from task "alloc" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↪sysview_heap_log.c:47
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

[0.202436200] HEAP: Allocated 3 bytes @ 0x3ffaaffe8 from task "alloc" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↪sysview_heap_log.c:47
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

[0.302436000] HEAP: Allocated 4 bytes @ 0x3ffafff0 from task "alloc" on core 0 by:
/home/user/projects/esp/esp-idf/examples/system/sysview_tracing_heap_log/main/
↪sysview_heap_log.c:47
/home/user/projects/esp/esp-idf/components/freertos/port.c:355 (discriminator 1)

Found 10 leaked bytes in 4 blocks.

```

Heap Tracing To Find Heap Corruption Heap tracing can also be used to help track down heap corruption. When a region in the heap is corrupted, it may be from some other part of the program that allocated memory at a nearby address.

If you have an approximate idea of when the corruption occurred, enabling heap tracing in `HEAP_TRACE_ALL` mode allows you to record all the memory allocation functions used and the corresponding allocation addresses.

Using heap tracing in this way is very similar to memory leak detection as described above. For memories that are allocated and not freed, the output is the same. However, records will also be shown for memory that has been freed.

Performance Impact Enabling heap tracing in menuconfig increases the code size of your program, and has a very small negative impact on the performance of heap allocation or free operations even when heap tracing is not running.

When heap tracing is running, heap allocation or free operations are substantially slower than when heap tracing is stopped. Increasing the depth of stack frames recorded for each allocation (see above) also increases this performance impact.

To mitigate the performance loss when the heap tracing is enabled and active, enable `CONFIG_HEAP_TRACE_HASH_MAP`. With this configuration enabled, a hash map mechanism will be used to handle the heap trace records, thus considerably decreasing the heap allocation or free execution time. The size of the hash map can be modified by setting the value of `CONFIG_HEAP_TRACE_HASH_MAP_SIZE`.

By default, the hash map is placed into internal RAM. It can also be placed into external RAM if `CONFIG_HEAP_TRACE_HASH_MAP_IN_EXT_RAM` is enabled. In order to enable this configuration, make sure to enable `CONFIG_SPIRAM` and `CONFIG_SPIRAM_ALLOW_BSS_SEG_EXTERNAL_MEMORY`.

False-Positive Memory Leaks Not everything printed by `heap_trace_dump()` is necessarily a memory leak. The following cases may also be printed:

- Any memory that is allocated after `heap_trace_start()` but freed after `heap_trace_stop()` appears in the leaked dump.
- Allocations may be made by other tasks in the system. Depending on the timing of these tasks, it is quite possible that this memory is freed after `heap_trace_stop()` is called.
- The first time a task uses stdio - e.g., when it calls `heap_caps_printf()` - a lock, i.e., RTOS mutex semaphore, is allocated by the libc. This allocation lasts until the task is deleted.
- Certain uses of `heap_caps_printf()`, such as printing floating point numbers and allocating some memory from the heap on demand. These allocations last until the task is deleted.
- The Bluetooth, Wi-Fi, and TCP/IP libraries allocate heap memory buffers to handle incoming or outgoing data. These memory buffers are usually short-lived, but some may be shown in the heap leak trace if the data has been received or transmitted by the lower levels of the network during the heap tracing.
- TCP connections retain some memory even after they are closed due to the `TIME_WAIT` state. Once the `TIME_WAIT` period is completed, this memory will be freed.

One way to differentiate between "real" and "false positive" memory leaks is to call the suspect code multiple times while tracing is running, and look for patterns (multiple matching allocations) in the heap trace output.

API Reference - Heap Tracing

Header File

- `components/heap/include/esp_heap_trace.h`
- This header file can be included with:

```
#include "esp_heap_trace.h"
```

Functions

`esp_err_t heap_trace_init_standalone(heap_trace_record_t *record_buffer, size_t num_records)`

Initialise heap tracing in standalone mode.

This function must be called before any other heap tracing functions.

To disable heap tracing and allow the buffer to be freed, stop tracing and then call `heap_trace_init_standalone(NULL, 0)`;

Parameters

- **record_buffer** -- Provide a buffer to use for heap trace data. Note: External RAM is allowed, but it prevents recording allocations made from ISR's.
- **num_records** -- Size of the heap trace buffer, as number of record structures.

Returns

- ESP_ERR_NOT_SUPPORTED Project was compiled without heap tracing enabled in menuconfig.
- ESP_ERR_INVALID_STATE Heap tracing is currently in progress.
- ESP_OK Heap tracing initialised successfully.

esp_err_t **heap_trace_init_tohost** (void)

Initialise heap tracing in host-based mode.

This function must be called before any other heap tracing functions.

Returns

- ESP_ERR_INVALID_STATE Heap tracing is currently in progress.
- ESP_OK Heap tracing initialised successfully.

esp_err_t **heap_trace_start** (*heap_trace_mode_t* mode)

Start heap tracing. All heap allocations & frees will be traced, until heap_trace_stop() is called.

Note: heap_trace_init_standalone() must be called to provide a valid buffer, before this function is called.

Note: Calling this function while heap tracing is running will reset the heap trace state and continue tracing.

Parameters **mode** -- Mode for tracing.

- HEAP_TRACE_ALL means all heap allocations and frees are traced.
- HEAP_TRACE_LEAKS means only suspected memory leaks are traced. (When memory is freed, the record is removed from the trace buffer.)

Returns

- ESP_ERR_NOT_SUPPORTED Project was compiled without heap tracing enabled in menuconfig.
- ESP_ERR_INVALID_STATE A non-zero-length buffer has not been set via heap_trace_init_standalone().
- ESP_OK Tracing is started.

esp_err_t **heap_trace_stop** (void)

Stop heap tracing.

Returns

- ESP_ERR_NOT_SUPPORTED Project was compiled without heap tracing enabled in menuconfig.
- ESP_ERR_INVALID_STATE Heap tracing was not in progress.
- ESP_OK Heap tracing stopped..

esp_err_t **heap_trace_resume** (void)

Resume heap tracing which was previously stopped.

Unlike heap_trace_start(), this function does not clear the buffer of any pre-existing trace records.

The heap trace mode is the same as when heap_trace_start() was last called (or HEAP_TRACE_ALL if heap_trace_start() was never called).

Returns

- ESP_ERR_NOT_SUPPORTED Project was compiled without heap tracing enabled in menuconfig.
- ESP_ERR_INVALID_STATE Heap tracing was already started.
- ESP_OK Heap tracing resumed.

size_t **heap_trace_get_count** (void)

Return number of records in the heap trace buffer.

It is safe to call this function while heap tracing is running.

esp_err_t **heap_trace_get** (size_t index, *heap_trace_record_t* *record)

Return a raw record from the heap trace buffer.

Note: It is safe to call this function while heap tracing is running, however in HEAP_TRACE_LEAK mode record indexing may skip entries unless heap tracing is stopped first.

Parameters

- **index** -- Index (zero-based) of the record to return.
- **record** -- [out] Record where the heap trace record will be copied.

Returns

- ESP_ERR_NOT_SUPPORTED Project was compiled without heap tracing enabled in menuconfig.
- ESP_ERR_INVALID_STATE Heap tracing was not initialised.
- ESP_ERR_INVALID_ARG Index is out of bounds for current heap trace record count.
- ESP_OK Record returned successfully.

void **heap_trace_dump** (void)

Dump heap trace record data to stdout.

Note: It is safe to call this function while heap tracing is running, however in HEAP_TRACE_LEAK mode the dump may skip entries unless heap tracing is stopped first.

void **heap_trace_dump_caps** (const uint32_t caps)

Dump heap trace from the memory of the capabilities passed as parameter.

Parameters caps -- Capability(ies) of the memory from which to dump the trace. Set `MALLOC_CAP_INTERNAL` to dump heap trace data from internal memory. Set `MALLOC_CAP_SPIRAM` to dump heap trace data from PSRAM. Set both to dump both heap trace data.

esp_err_t **heap_trace_summary** (*heap_trace_summary_t* *summary)

Get summary information about the result of a heap trace.

Note: It is safe to call this function while heap tracing is running.

Structures

struct **heap_trace_record_t**

Trace record data type. Stores information about an allocated region of memory.

Public Members

uint32_t **ccount**

CCOUNT of the CPU when the allocation was made. LSB (bit value 1) is the CPU number (0 or 1).

void ***address**

Address which was allocated. If NULL, then this record is empty.

size_t **size**

Size of the allocation.

void ***allocated_by**[CONFIG_HEAP_TRACING_STACK_DEPTH]

Call stack of the caller which allocated the memory.

void ***freed_by**[CONFIG_HEAP_TRACING_STACK_DEPTH]

Call stack of the caller which freed the memory (all zero if not freed.)

struct **heap_trace_summary_t**

Stores information about the result of a heap trace.

Public Members

heap_trace_mode_t **mode**

The heap trace mode we just completed / are running.

size_t **total_allocations**

The total number of allocations made during tracing.

size_t **total_frees**

The total number of frees made during tracing.

size_t **count**

The number of records in the internal buffer.

size_t **capacity**

The capacity of the internal buffer.

size_t **high_water_mark**

The maximum value that 'count' got to.

size_t **has_overflowed**

True if the internal buffer overflowed at some point.

Macros

CONFIG_HEAP_TRACING_STACK_DEPTH

Type Definitions

typedef struct *heap_trace_record_t* **heap_trace_record_t**

Trace record data type. Stores information about an allocated region of memory.

Enumerations

enum **heap_trace_mode_t**

Values:

enumerator **HEAP_TRACE_ALL**

enumerator **HEAP_TRACE_LEAKS**

2.9.18 High Resolution Timer (ESP Timer)

Overview

Although FreeRTOS provides software timers, FreeRTOS software timers have a few limitations:

- Maximum resolution is equal to the RTOS tick period
- Timer callbacks are dispatched from a low-priority timer service (i.e., daemon) task. This task can be pre-empted by other tasks, leading to decreased precision and accuracy.

Although hardware timers are not subject to the limitations mentioned, they may not be as user-friendly. For instance, application components may require timer events to be triggered at specific future times, but hardware timers typically have only one "compare" value for interrupt generation. This necessitates the creation of an additional system on top of the hardware timer to keep track of pending events and ensure that callbacks are executed when the corresponding hardware interrupts occur.

`esp_timer` set of APIs provides one-shot and periodic timers, microsecond time resolution, and 52-bit range.

Internally, `esp_timer` uses a 52-bit hardware timer. The exact hardware timer implementation used depends on the target, where SYSTIMER is used for ESP32-P4.

Timer callbacks can be dispatched by two methods:

- `ESP_TIMER_TASK`.
- `ESP_TIMER_ISR`. Available only if [`CONFIG\_ESP\_TIMER\_SUPPORTS\_ISR\_DISPATCH\_METHOD`](#) is enabled (by default disabled).

`ESP_TIMER_TASK`. Timer callbacks are dispatched from a high-priority `esp_timer` task. Because all the callbacks are dispatched from the same task, it is recommended to only do the minimal possible amount of work from the callback itself, posting an event to a lower-priority task using a queue instead.

If other tasks with a priority higher than `esp_timer` are running, callback dispatching will be delayed until the `esp_timer` task has a chance to run. For example, this will happen if an SPI Flash operation is in progress.

`ESP_TIMER_ISR`. Timer callbacks are dispatched directly from the timer interrupt handler. This method is useful for some simple callbacks which aim for lower latency.

Creating and starting a timer, and dispatching the callback takes some time. Therefore, there is a lower limit to the timeout value of one-shot `esp_timer`. If [`esp\_timer\_start\_once\(\)`](#) is called with a timeout value of less than 20 us, the callback will be dispatched only after approximately 20 us.

Periodic `esp_timer` also imposes a 50 us restriction on the minimal timer period. Periodic software timers with a period of less than 50 us are not practical since they would consume most of the CPU time. Consider using dedicated hardware peripherals or DMA features if you find that a timer with a small period is required.

Using `esp_timer` APIs

A single timer is represented by [`esp\_timer\_handle\_t`](#) type. Each timer has a callback function associated with it. This callback function is called from the `esp_timer` task each time the timer elapses.

- To create a timer, call [`esp\_timer\_create\(\)`](#).
- To delete the timer when it is no longer needed, call [`esp\_timer\_delete\(\)`](#).

The timer can be started in one-shot mode or in periodic mode.

- To start the timer in one-shot mode, call [`esp\_timer\_start\_once\(\)`](#), passing the time interval after which the callback should be called. When the callback gets called, the timer is considered to be stopped.
- To start the timer in periodic mode, call [`esp\_timer\_start\_periodic\(\)`](#), passing the period with which the callback should be called. The timer keeps running until [`esp\_timer\_stop\(\)`](#) is called.

Note that the timer must not be running when `esp_timer_start_once()` or `esp_timer_start_periodic()` is called. To restart a running timer, call `esp_timer_stop()` first, then call one of the start functions.

Callback Functions

Note: Keep the callback functions as short as possible. Otherwise, it will affect all timers.

Timer callbacks that are processed by the `ESP_TIMER_ISR` method should not call the context switch call `portYIELD_FROM_ISR()`. Instead, use the `esp_timer_isr_dispatch_need_yield()` function. The context switch will be done after all ISR dispatch timers have been processed if required by the system.

esp_timer During Light-sleep

During Light-sleep, the `esp_timer` counter stops and no callback functions are called. Instead, the time is counted by the RTC counter. Upon waking up, the system gets the difference between the counters and calls a function that advances the `esp_timer` counter. Since the counter has been advanced, the system starts calling callbacks that were not called during sleep. The number of callbacks depends on the duration of the sleep and the period of the timers. It can lead to the overflow of some queues. This only applies to periodic timers, since one-shot timers are only called once.

This behavior can be changed by calling `esp_timer_stop()` before sleeping. In some cases, this can be inconvenient, and instead of the stop function, you can use the `skip_unhandled_events` option during `esp_timer_create()`. When the `skip_unhandled_events` is true, if a periodic timer expires one or more times during Light-sleep, then only one callback is called on wake.

Using the `skip_unhandled_events` option with automatic Light-sleep (see [Power Management APIs](#)) helps to reduce the power consumption of the system when it is in Light-sleep. The duration of Light-sleep is also in part determined by the next event occurs. Timers with `skip_unhandled_events` option does not wake up the system.

Handling Callbacks

`esp_timer` is designed to achieve a high-resolution and low-latency timer with the ability to handle delayed events. If the timer is late, then the callback will be called as soon as possible, and it will not be lost. In the worst case, when the timer has not been processed for more than one period (for periodic timers), the callbacks will be called one after the other without waiting for the set period. This can be bad for some applications, and the `skip_unhandled_events` option is introduced to eliminate this behavior. If `skip_unhandled_events` is set, then a periodic timer that has expired multiple times without being able to call the callback will still result in only one callback event once processing is possible.

Obtaining Current Time

`esp_timer` also provides a convenience function to obtain the time passed since start-up, with microsecond precision: `esp_timer_get_time()`. This function returns the number of microseconds since `esp_timer` was initialized, which usually happens shortly before `app_main` function is called.

Unlike `gettimeofday` function, values returned by `esp_timer_get_time()`:

- Start from zero after the chip wakes up from Deep-sleep
- Do not have timezone or DST adjustments applied

Application Example

The following example illustrates the usage of `esp_timer` APIs: [system/esp_timer](#).

API Reference

Header File

- [components/esp_timer/include/esp_timer.h](#)
- This header file can be included with:

```
#include "esp_timer.h"
```

- This header file is a part of the API provided by the `esp_timer` component. To declare that your component depends on `esp_timer`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_timer
```

or

```
PRIV_REQUIRES esp_timer
```

Functions

esp_err_t **esp_timer_early_init** (void)

Minimal initialization of `esp_timer`.

This function can be called very early in startup process, after this call only `esp_timer_get_time` function can be used.

Note: This function is called from startup code. Applications do not need to call this function before using other `esp_timer` APIs.

Returns

- `ESP_OK` on success

esp_err_t **esp_timer_init** (void)

Initialize `esp_timer` library.

This function will be called from startup code on every core if `CONFIG_ESP_TIMER_ISR_AFFINITY_NO_AFFINITY` is enabled, It allocates the timer ISR on MULTIPLE cores and creates the timer task which can be run on any core.

Note: This function is called from startup code. Applications do not need to call this function before using other `esp_timer` APIs. Before calling this function, `esp_timer_early_init` must be called by the startup code.

Returns

- `ESP_OK` on success
- `ESP_ERR_NO_MEM` if allocation has failed
- `ESP_ERR_INVALID_STATE` if already initialized
- other errors from interrupt allocator

esp_err_t **esp_timer_deinit** (void)

De-initialize `esp_timer` library.

Note: Normally this function should not be called from applications

Returns

- ESP_OK on success
- ESP_ERR_INVALID_STATE if not yet initialized

esp_err_t **esp_timer_create** (const *esp_timer_create_args_t* *create_args, *esp_timer_handle_t* *out_handle)

Create an esp_timer instance.

Note: When done using the timer, delete it with esp_timer_delete function.

Parameters

- **create_args** -- Pointer to a structure with timer creation arguments. Not saved by the library, can be allocated on the stack.
- **out_handle** -- [out] Output, pointer to esp_timer_handle_t variable which will hold the created timer handle.

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if some of the create_args are not valid
- ESP_ERR_INVALID_STATE if esp_timer library is not initialized yet
- ESP_ERR_NO_MEM if memory allocation fails

esp_err_t **esp_timer_start_once** (*esp_timer_handle_t* timer, uint64_t timeout_us)

Start one-shot timer.

Timer should not be running when this function is called.

Parameters

- **timer** -- timer handle created using esp_timer_create
- **timeout_us** -- timer timeout, in microseconds relative to the current moment

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if the handle is invalid
- ESP_ERR_INVALID_STATE if the timer is already running

esp_err_t **esp_timer_start_periodic** (*esp_timer_handle_t* timer, uint64_t period)

Start a periodic timer.

Timer should not be running when this function is called. This function will start the timer which will trigger every 'period' microseconds.

Parameters

- **timer** -- timer handle created using esp_timer_create
- **period** -- timer period, in microseconds

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if the handle is invalid
- ESP_ERR_INVALID_STATE if the timer is already running

esp_err_t **esp_timer_restart** (*esp_timer_handle_t* timer, uint64_t timeout_us)

Restart a currently running timer.

If the given timer is a one-shot timer, the timer is restarted immediately and will timeout once in timeout_us microseconds. If the given timer is a periodic timer, the timer is restarted immediately with a new period of timeout_us microseconds.

Parameters

- **timer** -- timer Handle created using esp_timer_create
- **timeout_us** -- Timeout, in microseconds relative to the current time. In case of a periodic timer, also represents the new period.

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if the handle is invalid

- `ESP_ERR_INVALID_STATE` if the timer is not running

esp_err_t **esp_timer_stop** (*esp_timer_handle_t* timer)

Stop the timer.

This function stops the timer previously started using `esp_timer_start_once` or `esp_timer_start_periodic`.

Parameters **timer** -- timer handle created using `esp_timer_create`

Returns

- `ESP_OK` on success
- `ESP_ERR_INVALID_STATE` if the timer is not running

esp_err_t **esp_timer_delete** (*esp_timer_handle_t* timer)

Delete an `esp_timer` instance.

The timer must be stopped before deleting. A one-shot timer which has expired does not need to be stopped.

Parameters **timer** -- timer handle allocated using `esp_timer_create`

Returns

- `ESP_OK` on success
- `ESP_ERR_INVALID_STATE` if the timer is running

int64_t **esp_timer_get_time** (void)

Get time in microseconds since boot.

Returns number of microseconds since underlying timer has been started

int64_t **esp_timer_get_next_alarm** (void)

Get the timestamp when the next timeout is expected to occur.

Returns Timestamp of the nearest timer event, in microseconds. The timebase is the same as for the values returned by `esp_timer_get_time`.

int64_t **esp_timer_get_next_alarm_for_wake_up** (void)

Get the timestamp when the next timeout is expected to occur skipping those which have `skip_unhandled_events` flag.

Returns Timestamp of the nearest timer event, in microseconds. The timebase is the same as for the values returned by `esp_timer_get_time`.

esp_err_t **esp_timer_get_period** (*esp_timer_handle_t* timer, *uint64_t* *period)

Get the period of a timer.

This function fetches the timeout period of a timer.

Note: The timeout period is the time interval with which a timer restarts after expiry. For one-shot timers, the period is 0 as there is no periodicity associated with such timers.

Parameters

- **timer** -- timer handle allocated using `esp_timer_create`
- **period** -- memory to store the timer period value in microseconds

Returns

- `ESP_OK` on success
- `ESP_ERR_INVALID_ARG` if the arguments are invalid

esp_err_t **esp_timer_get_expiry_time** (*esp_timer_handle_t* timer, *uint64_t* *expiry)

Get the expiry time of a one-shot timer.

This function fetches the expiry time of a one-shot timer.

Note: This API returns a valid expiry time only for a one-shot timer. It returns an error if the timer handle passed to the function is for a periodic timer.

Parameters

- **timer** -- timer handle allocated using `esp_timer_create`
- **expiry** -- memory to store the timeout value in microseconds

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if the arguments are invalid
- ESP_ERR_NOT_SUPPORTED if the timer type is periodic

esp_err_t **esp_timer_dump** (FILE *stream)

Dump the list of timers to a stream.

If CONFIG_ESP_TIMER_PROFILING option is enabled, this prints the list of all the existing timers. Otherwise, only the list active timers is printed.

The format is:

name period alarm times_armed times_triggered total_callback_run_time

where:

name —timer name (if CONFIG_ESP_TIMER_PROFILING is defined), or timer pointer period —period of timer, in microseconds, or 0 for one-shot timer alarm - time of the next alarm, in microseconds since boot, or 0 if the timer is not started

The following fields are printed if CONFIG_ESP_TIMER_PROFILING is defined:

times_armed —number of times the timer was armed via `esp_timer_start_X` times_triggered - number of times the callback was called total_callback_run_time - total time taken by callback to execute, across all calls

Parameters **stream** -- stream (such as stdout) to dump the information to

Returns

- ESP_OK on success
- ESP_ERR_NO_MEM if can not allocate temporary buffer for the output

void **esp_timer_isr_dispatch_need_yield** (void)

Requests a context switch from a timer callback function.

This only works for a timer that has an ISR dispatch method. The context switch will be called after all ISR dispatch timers have been processed.

bool **esp_timer_is_active** (*esp_timer_handle_t* timer)

Returns status of a timer, active or not.

This function is used to identify if the timer is still active or not.

Parameters **timer** -- timer handle created using `esp_timer_create`

Returns

- 1 if timer is still active
- 0 if timer is not active.

esp_err_t **esp_timer_new_etm_alarm_event** (*esp_etm_event_handle_t* *out_event)

Get the ETM event handle of `esp_timer` underlying alarm event.

Note: The created ETM event object can be deleted later by calling `esp_etm_del_event`

Note: The ETM event is generated by the underlying hardware — systimer, therefore, if the esp_timer is not clocked by systimer, then no ETM event will be generated.

Parameters **out_event** -- [out] Returned ETM event handle

Returns

- ESP_OK Success
- ESP_ERR_INVALID_ARG Parameter error

Structures

struct **esp_timer_create_args_t**

Timer configuration passed to esp_timer_create.

Public Members

esp_timer_cb_t **callback**

Function to call when timer expires.

void ***arg**

Argument to pass to the callback.

esp_timer_dispatch_t **dispatch_method**

Call the callback from task or from ISR.

const char ***name**

Timer name, used in esp_timer_dump function.

bool **skip_unhandled_events**

Skip unhandled events for periodic timers.

Type Definitions

typedef struct esp_timer ***esp_timer_handle_t**

Opaque type representing a single esp_timer.

typedef void (***esp_timer_cb_t**)(void *arg)

Timer callback function type.

Param arg pointer to opaque user-specific data

Enumerations

enum **esp_timer_dispatch_t**

Method for dispatching timer callback.

Values:

enumerator **ESP_TIMER_TASK**

Callback is called from timer task.

enumerator **ESP_TIMER_MAX**

Count of the methods for dispatching timer callback.

2.9.19 Internal and Unstable APIs

This section is listing some APIs that are internal or likely to be changed or removed in the next releases of ESP-IDF.

API Reference

Header File

- `components/esp_rom/include/esp_rom_sys.h`
- This header file can be included with:

```
#include "esp_rom_sys.h"
```

Functions

void **esp_rom_software_reset_system** (void)

Software Reset digital core include RTC.

It is not recommended to use this function in esp-idf, use `esp_restart()` instead.

void **esp_rom_software_reset_cpu** (int cpu_no)

Software Reset cpu core.

It is not recommended to use this function in esp-idf, use `esp_restart()` instead.

Parameters `cpu_no` -- : The CPU to reset, 0 for PRO CPU, 1 for APP CPU.

int **esp_rom_printf** (const char *fmt, ...)

Print formatted string to console device.

Note: float and long long data are not supported!

Parameters

- `fmt` -- Format string
- `...` -- Additional arguments, depending on the format string

Returns int: Total number of characters written on success; A negative number on failure.

void **esp_rom_delay_us** (uint32_t us)

Pauses execution for us microseconds.

Parameters `us` -- Number of microseconds to pause

void **esp_rom_install_channel_putc** (int channel, void (*putc)(char c))

`esp_rom_printf` can print message to different channels simultaneously. This function can help install the low level `putc` function for `esp_rom_printf`.

Parameters

- `channel` -- Channel number (startting from 1)
- `putc` -- Function pointer to the `putc` implementation. Set NULL can disconnect `esp_rom_printf` with `putc`.

void **esp_rom_install_uart_printf** (void)

Install UART1 as the default console channel, equivalent to `esp_rom_install_channel_putc(1, esp_rom_uart_putc)`

soc_reset_reason_t **esp_rom_get_reset_reason** (int cpu_no)

Get reset reason of CPU.

Parameters `cpu_no` -- CPU number

Returns Reset reason code (see in soc/reset_reasons.h)

void **esp_rom_route_intr_matrix** (int cpu_core, uint32_t periph_intr_id, uint32_t cpu_intr_num)

Route peripheral interrupt sources to CPU's interrupt port by matrix.

Usually there're 4 steps to use an interrupt:

- Route peripheral interrupt source to CPU. e.g. `esp_rom_route_intr_matrix(0, ETS_WIFI_MAC_INTR_SOURCE, ETS_WMAC_INUM)`
- Set interrupt handler for CPU
- Enable CPU interrupt
- Enable peripheral interrupt

Parameters

- `cpu_core` -- The CPU number, which the peripheral interrupt will inform to
- `periph_intr_id` -- The peripheral interrupt source number
- `cpu_intr_num` -- The CPU (external) interrupt number. On targets that use CLIC as their interrupt controller, this number represents the external interrupt number. For example, passing `cpu_intr_num = i` to this function would in fact bind peripheral source to CPU interrupt `CLIC_EXT_INTR_NUM_OFFSET + i`.

uint32_t **esp_rom_get_cpu_ticks_per_us** (void)

Get the real CPU ticks per us.

Returns CPU ticks per us

void **esp_rom_set_cpu_ticks_per_us** (uint32_t ticks_per_us)

Set the real CPU tick rate.

Note: Call this function when CPU frequency is changed, otherwise the `esp_rom_delay_us` can be inaccurate.

Parameters `ticks_per_us` -- CPU ticks per us

2.9.20 Inter-Processor Call (IPC)

Note: IPC stands for an "Inter-Processor Call" and NOT "Inter-Process Communication" as found on other operating systems.

Overview

Due to the dual core nature of the ESP32-P4, there are some scenarios where a certain callback must be executed from a particular core such as:

- When allocating an ISR to an interrupt source of a particular core (applies to freeing a particular core's interrupt source as well)
- On particular chips (such as the ESP32), accessing memory that is exclusive to a particular core (such as RTC Fast Memory)
- Reading the registers/state of another core

The IPC (Inter-Processor Call) feature allows a particular core (the calling core) to trigger the execution of a callback function on another core (the target core). The IPC feature allows execution of a callback function on the target core in either a task context, or an interrupt context. Depending on the context that the callback function is executed in, different restrictions apply to the implementation of the callback function.

IPC in Task Context

The IPC feature implements callback execution in a task context by creating an IPC task for each core during application startup. When the calling core needs to execute a callback on the target core, the callback will execute in the context of the target core's IPC task.

When using IPCs in a task context, users need to consider the following:

- IPC callbacks should ideally be simple and short. An IPC callback **must never block or yield**.
- The IPC tasks are created at the highest possible priority (i.e., `configMAX_PRIORITIES - 1`).
 - If `CONFIG_ESP_IPC_USES_CALLERS_PRIORITY` is enabled, the target core's IPC task will be lowered to the current priority of the target core before executing the callback.
 - If `CONFIG_ESP_IPC_USES_CALLERS_PRIORITY` is disabled, the target core will always execute the callback at the highest possible priority.
- Depending on the complexity of the callback, users may need to configure the stack size of the IPC task via `CONFIG_ESP_IPC_TASK_STACK_SIZE`.
- The IPC feature is internally protected by a mutex. Therefore, simultaneous IPC calls from two or more calling core's are serialized on a first come first serve basis.

API Usage Task Context IPC callbacks have the following restrictions:

- The callback must be of the `esp_ipc_func_t` type.
- The callback **must never block or yield** as this will result in the target core's IPC task blocking or yielding.
- The callback must avoid changing any aspect of the IPC task's state, e.g., by calling `vTaskPrioritySet(NULL, x)`.

The IPC feature offers the API listed below to execute a callback in a task context on a target core. The API allows the calling core to block until the callback's execution has completed, or return immediately once the callback's execution has started.

- `esp_ipc_call()` triggers an IPC call on the target core. This function will block until the target core's IPC task **begins** execution of the callback.
- `esp_ipc_call_blocking()` triggers an IPC on the target core. This function will block until the target core's IPC task **completes** execution of the callback.

IPC in Interrupt Context

In some cases, we need to quickly obtain the state of another core such as in a core dump, GDB stub, various unit tests, and hardware errata workarounds. The IPC ISR feature implements callback execution from a High Priority Interrupt context by reserving a High Priority Interrupt on each core for IPC usage. When a calling core needs to execute a callback on the target core, the callback will execute in the context of the High Priority Interrupt of the target core.

When using IPCs in High Priority Interrupt context, users need to consider the following:

- The priority of the reserved High Priority Interrupt is dependent on the `CONFIG_ESP_SYSTEM_CHECK_INT_LEVEL` option.

When the callback executes, users need to consider the following:

- The calling core will disable interrupts of priority level 3 and lower.

- Although the priority of the reserved interrupt depends on `CONFIG_ESP_SYSTEM_CHECK_INT_LEVEL`, during the execution of IPC ISR callback, the target core will disable all interrupts.

API Usage High Priority Interrupt IPC callbacks must be of type `esp_ipc_isr_func_t` and have the same restrictions as for regular interrupt handlers. The callback function can be written in C.

The IPC feature offers the API listed below to execute a callback in a High Priority Interrupt context:

- `esp_ipc_isr_call()` triggers an IPC call on the target core. This function will busy-wait until the target core **begins** execution of the callback.
- `esp_ipc_isr_call_blocking()` triggers an IPC call on the target core. This function will busy-wait until the target core **completes** execution of the callback.

See `examples/system/ipc/ipc_isr/riscv/main/main.c` for an example of its use.

The High Priority Interrupt IPC API also provides the following convenience functions that can stall/resume the target core. These APIs utilize the High Priority Interrupt IPC, but supply their own internal callbacks:

- `esp_ipc_isr_stall_other_cpu()` stalls the target core. The calling core disables interrupts of level 3 and lower, while the target core will busy-wait with all interrupts disabled. The target core will busy-wait until `esp_ipc_isr_release_other_cpu()` is called.
- `esp_ipc_isr_release_other_cpu()` resumes the target core.

API Reference

Header File

- `components/esp_system/include/esp_ipc.h`
- This header file can be included with:

```
#include "esp_ipc.h"
```

Functions

`esp_err_t esp_ipc_call` (uint32_t cpu_id, `esp_ipc_func_t` func, void *arg)

Execute a callback on a given CPU.

Execute a given callback on a particular CPU. The callback must be of type "esp_ipc_func_t" and will be invoked in the context of the target CPU's IPC task.

- This function will block the target CPU's IPC task has begun execution of the callback
- If another IPC call is ongoing, this function will block until the ongoing IPC call completes
- The stack size of the IPC task can be configured via the `CONFIG_ESP_IPC_TASK_STACK_SIZE` option

Note: In single-core mode, returns `ESP_ERR_INVALID_ARG` for cpu_id 1.

Parameters

- **cpu_id** -- [in] CPU where the given function should be executed (0 or 1)
- **func** -- [in] Pointer to a function of type void func(void* arg) to be executed
- **arg** -- [in] Arbitrary argument of type void* to be passed into the function

Returns

- `ESP_ERR_INVALID_ARG` if cpu_id is invalid
- `ESP_ERR_INVALID_STATE` if the FreeRTOS scheduler is not running
- `ESP_OK` otherwise

esp_err_t **esp_ipc_call_blocking** (uint32_t cpu_id, *esp_ipc_func_t* func, void *arg)

Execute a callback on a given CPU until and block until it completes.

This function is identical to `esp_ipc_call()` except that this function will block until the execution of the callback completes.

Note: In single-core mode, returns `ESP_ERR_INVALID_ARG` for `cpu_id` 1.

Parameters

- **cpu_id** -- [in] CPU where the given function should be executed (0 or 1)
- **func** -- [in] Pointer to a function of type `void func(void* arg)` to be executed
- **arg** -- [in] Arbitrary argument of type `void*` to be passed into the function

Returns

- `ESP_ERR_INVALID_ARG` if `cpu_id` is invalid
- `ESP_ERR_INVALID_STATE` if the FreeRTOS scheduler is not running
- `ESP_OK` otherwise

Type Definitions

```
typedef void (*esp_ipc_func_t)(void *arg)
```

IPC Callback.

A callback of this type should be provided as an argument when calling `esp_ipc_call()` or `esp_ipc_call_blocking()`.

Header File

- `components/esp_system/include/esp_ipc_isr.h`
- This header file can be included with:

```
#include "esp_ipc_isr.h"
```

Functions

void **esp_ipc_isr_call** (*esp_ipc_isr_func_t* func, void *arg)

Execute an ISR callback on the other CPU.

Execute a given callback on the other CPU in the context of a High Priority Interrupt.

- This function will busy-wait in a critical section until the other CPU has started execution of the callback
- The callback must be written:
 - in assembly for XTENSA chips (such as ESP32, ESP32S3). The function is invoked using a `CALLX0` instruction and can use only `a2`, `a3`, `a4` registers. See `:doc:IPC in Interrupt Context` [</api-reference/system/ipc>](#) doc for more details.
 - in C or assembly for RISC-V chips (such as ESP32P4).

Note: This function is not available in single-core mode.

Parameters

- **func** -- [in] Pointer to a function of type `void func(void* arg)` to be executed
- **arg** -- [in] Arbitrary argument of type `void*` to be passed into the function

void **esp_ipc_isr_call_blocking** (*esp_ipc_isr_func_t* func, void *arg)

Execute an ISR callback on the other CPU and busy-wait until it completes.

This function is identical to `esp_ipc_isr_call()` except that this function will busy-wait until the execution of the callback completes.

Note: This function is not available in single-core mode.

Parameters

- **func** -- [in] Pointer to a function of type `void func(void* arg)` to be executed
- **arg** -- [in] Arbitrary argument of type `void*` to be passed into the function

void **esp_ipc_isr_stall_other_cpu** (void)

Stall the other CPU.

This function will stall the other CPU. The other CPU is stalled by busy-waiting in the context of a High Priority Interrupt. The other CPU will not be resumed until `esp_ipc_isr_release_other_cpu()` is called.

- This function is internally implemented using IPC ISR
- This function is used for DPORT workaround.
- If the stall feature is paused using `esp_ipc_isr_stall_pause()`, this function will have no effect

Note: This function is not available in single-core mode.

Note: It is the caller's responsibility to avoid deadlocking on spinlocks

void **esp_ipc_isr_release_other_cpu** (void)

Release the other CPU.

This function will release the other CPU that was previously stalled from calling `esp_ipc_isr_stall_other_cpu()`

- This function is used for DPORT workaround.
- If the stall feature is paused using `esp_ipc_isr_stall_pause()`, this function will have no effect

Note: This function is not available in single-core mode.

void **esp_ipc_isr_stall_pause** (void)

Pause the CPU stall feature.

This function will pause the CPU stall feature. Once paused, calls to `esp_ipc_isr_stall_other_cpu()` and `esp_ipc_isr_release_other_cpu()` will have no effect. If a IPC ISR call is already in progress, this function will busy-wait until the call completes before pausing the CPU stall feature.

void **esp_ipc_isr_stall_abort** (void)

Abort a CPU stall.

This function will abort any stalling routine of the other CPU due to a previous call to `esp_ipc_isr_stall_other_cpu()`. This function aborts the stall in a non-recoverable manner, thus should only be called in case of a panic().

- This function is used in panic handling code

void **esp_ipc_isr_stall_resume** (void)

Resume the CPU stall feature.

This function will resume the CPU stall feature that was previously paused by calling `esp_ipc_isr_stall_pause()`. Once resumed, calls to `esp_ipc_isr_stall_other_cpu()` and `esp_ipc_isr_release_other_cpu()` will have effect again.

Macros

esp_ipc_isr_asm_call (func, arg)

Execute an ISR callback on the other CPU See `esp_ipc_isr_call()`.

esp_ipc_isr_asm_call_blocking (func, arg)

Execute an ISR callback on the other CPU and busy-wait until it completes See `esp_ipc_isr_call_blocking()`.

Type Definitions

typedef void (***esp_ipc_isr_func_t**)(void *arg)

IPC ISR Callback.

The callback must be written:

- in assembly for XTENSA chips (such as ESP32, ESP32S3).
- in C or assembly for RISC-V chips (such as ESP32P4).

A callback of this type should be provided as an argument when calling `esp_ipc_isr_call()` or `esp_ipc_isr_call_blocking()`.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.9.21 Interrupt Allocation

Overview

Because there are more interrupt sources than interrupts, sometimes it makes sense to share an interrupt in multiple drivers. The `esp_intr_alloc()` abstraction exists to hide all these implementation details.

A driver can allocate an interrupt for a certain peripheral by calling `esp_intr_alloc()` (or `esp_intr_alloc_intrstatus()`). It can use the flags passed to this function to specify the type, priority, and trigger method of the interrupt to allocate. The interrupt allocation code will then find an applicable interrupt, use the interrupt matrix to hook it up to the peripheral, and install the given interrupt handler and ISR to it.

The interrupt allocator presents two different types of interrupts, namely shared interrupts and non-shared interrupts, both of which require different handling. Non-shared interrupts will allocate a separate interrupt for every `esp_intr_alloc()` call, and this interrupt is used solely for the peripheral attached to it, with only one ISR that will get called. Shared interrupts can have multiple peripherals triggering them, with multiple ISRs being called when one of the peripherals attached signals an interrupt. Thus, ISRs that are intended for shared interrupts should check the interrupt status of the peripheral they service in order to check if any action is required.

Non-shared interrupts can be either level- or edge-triggered. Shared interrupts can only be level interrupts due to the chance of missed interrupts when edge interrupts are used.

To illustrate why shared interrupts can only be level-triggered, take the scenario where peripheral A and peripheral B share the same edge-triggered interrupt. Peripheral B triggers an interrupt and sets its interrupt signal high, causing a low-to-high edge, which in turn latches the CPU's interrupt bit and triggers the ISR. The ISR executes, checks that peripheral A did not trigger an interrupt, and proceeds to handle and clear peripheral B's interrupt signal. Before the

ISR returns, the CPU clears its interrupt bit latch. Thus, during the entire interrupt handling process, if peripheral A triggers an interrupt, it will be missed due the CPU clearing the interrupt bit latch.

IRAM-Safe Interrupt Handlers

The `ESP_INTR_FLAG_IRAM` flag registers an interrupt handler that always runs from IRAM (and reads all its data from DRAM), and therefore does not need to be disabled during flash erase and write operations.

This is useful for interrupts which need a guaranteed minimum execution latency, as flash write and erase operations can be slow (erases can take tens or hundreds of milliseconds to complete).

It can also be useful to keep an interrupt handler in IRAM if it is called very frequently, to avoid flash cache misses.

Refer to the [SPI flash API documentation](#) for more details.

Multiple Handlers Sharing A Source

Several handlers can be assigned to a same source, given that all handlers are allocated using the `ESP_INTR_FLAG_SHARED` flag. They will all be allocated to the interrupt, which the source is attached to, and called sequentially when the source is active. The handlers can be disabled and freed individually. The source is attached to the interrupt (enabled), if one or more handlers are enabled, otherwise detached. A handler will never be called when disabled, while **its source may still be triggered** if any one of its handler enabled.

Sources attached to non-shared interrupt do not support this feature.

Though the framework supports this feature, you have to use it **very carefully**. There usually exist two ways to stop an interrupt from being triggered: **disable the source** or **mask peripheral interrupt status**. ESP-IDF only handles enabling and disabling of the source itself, leaving status and mask bits to be handled by users.

Status bits shall either be masked before the handler responsible for it is disabled, or be masked and then properly handled in another enabled interrupt.

Note: Leaving some status bits unhandled without masking them, while disabling the handlers for them, will cause the interrupt(s) to be triggered indefinitely, resulting therefore in a system crash.

Troubleshooting Interrupt Allocation

On most Espressif SoCs, CPU interrupts are a limited resource. Therefore it is possible for a program to run out of CPU interrupts, for example by initializing several peripheral drivers. Typically, this will result in the driver initialization function returning `ESP_ERR_NOT_FOUND` error code.

If this happens, you can use `esp_intr_dump()` function to print the list of interrupts along with their status. The output of this function typically looks like this:

```
CPU 0 interrupt status:
Int  Level  Type   Status
0     1     Level  Reserved
1     1     Level  Reserved
2     1     Level  Used: RTC_CORE
3     1     Level  Used: TG0_LACT_LEVEL
...
```

The columns of the output have the following meaning:

- **Int:** CPU interrupt input number. This is typically not used in software directly, and is provided for reference only.
- **Level:** For interrupts which have been allocated, the priority of the interrupt. For free interrupts * is printed.

- **Type:** For interrupts which have been allocated, the type (Level or Edge) of the interrupt. For free interrupts * is printed.
- **Status: One of the possible statuses of the interrupt:**
 - **Reserved:** The interrupt is reserved either at hardware level, or by one of the parts of ESP-IDF. It can not be allocated using `esp_intr_alloc()`.
 - **Used:** <source>: The interrupt is allocated and connected to a single peripheral.
 - **Shared:** <source1> <source2> . . . : The interrupt is allocated and connected to multiple peripherals. See [Multiple Handlers Sharing A Source](#) above.
 - **Free:** The interrupt is not allocated and can be used by `esp_intr_alloc()`.

If you have confirmed that the application is indeed running out of interrupts, a combination of the following suggestions can help resolve the issue:

- On multi-core SoCs, try initializing some of the peripheral drivers from a task pinned to the second core. Interrupts are typically allocated on the same core where the peripheral driver initialization function runs. Therefore by running the initialization function on the second core, more interrupt inputs can be used.
- Determine the interrupts which can tolerate higher latency, and allocate them using `ESP_INTR_FLAG_SHARED` flag (optionally ORed with `ESP_INTR_FLAG_LOWMED`). Using this flag for two or more peripherals will let them use a single interrupt input, and therefore save interrupt inputs for other peripherals. See [Multiple Handlers Sharing A Source](#) above.
- Check if some of the peripheral drivers do not need to be used all the time, and initialize or deinitialize them on demand. This can reduce the number of simultaneously allocated interrupts.

API Reference

Header File

- [components/esp_hw_support/include/esp_intr_types.h](#)
- This header file can be included with:

```
#include "esp_intr_types.h"
```

Macros

ESP_INTR_CPU_AFFINITY_TO_CORE_ID (cpu_affinity)

Convert `esp_intr_cpu_affinity_t` to CPU core ID.

Type Definitions

`typedef void (*intr_handler_t)(void *arg)`

Function prototype for interrupt handler function

`typedef struct intr_handle_data_t *intr_handle_t`

Handle to an interrupt handler

Enumerations

enum **esp_intr_cpu_affinity_t**

Interrupt CPU core affinity.

This type specify the CPU core that the peripheral interrupt is connected to.

Values:

enumerator **ESP_INTR_CPU_AFFINITY_AUTO**

Install the peripheral interrupt to ANY CPU core, decided by on which CPU the interrupt allocator is running.

enumerator **ESP_INTR_CPU_AFFINITY_0**

Install the peripheral interrupt to CPU core 0.

enumerator **ESP_INTR_CPU_AFFINITY_1**

Install the peripheral interrupt to CPU core 1.

Header File

- [components/esp_hw_support/include/esp_intr_alloc.h](#)
- This header file can be included with:

```
#include "esp_intr_alloc.h"
```

Functions

esp_err_t **esp_intr_mark_shared** (int intno, int cpu, bool is_in_iram)

Mark an interrupt as a shared interrupt.

This will mark a certain interrupt on the specified CPU as an interrupt that can be used to hook shared interrupt handlers to.

Parameters

- **intno** -- The number of the interrupt (0-31)
- **cpu** -- CPU on which the interrupt should be marked as shared (0 or 1)
- **is_in_iram** -- Shared interrupt is for handlers that reside in IRAM and the int can be left enabled while the flash cache is disabled.

Returns ESP_ERR_INVALID_ARG if cpu or intno is invalid ESP_OK otherwise

esp_err_t **esp_intr_reserve** (int intno, int cpu)

Reserve an interrupt to be used outside of this framework.

This will mark a certain interrupt on the specified CPU as reserved, not to be allocated for any reason.

Parameters

- **intno** -- The number of the interrupt (0-31)
- **cpu** -- CPU on which the interrupt should be marked as shared (0 or 1)

Returns ESP_ERR_INVALID_ARG if cpu or intno is invalid ESP_OK otherwise

esp_err_t **esp_intr_alloc** (int source, int flags, *intr_handler_t* handler, void *arg, *intr_handle_t* *ret_handle)

Allocate an interrupt with the given parameters.

This finds an interrupt that matches the restrictions as given in the flags parameter, maps the given interrupt source to it and hooks up the given interrupt handler (with optional argument) as well. If needed, it can return a handle for the interrupt as well.

The interrupt will always be allocated on the core that runs this function.

If ESP_INTR_FLAG_IRAM flag is used, and handler address is not in IRAM or RTC_FAST_MEM, then ESP_ERR_INVALID_ARG is returned.

Parameters

- **source** -- The interrupt source. One of the ETS*_INTR_SOURCE interrupt mux sources, as defined in soc/soc.h, or one of the internal ETS_INTERNAL_*_INTR_SOURCE sources as defined in this header.
- **flags** -- An ORred mask of the ESP_INTR_FLAG_* defines. These restrict the choice of interrupts that this routine can choose from. If this value is 0, it will default to allocating a non-shared interrupt of level 1, 2 or 3. If this is ESP_INTR_FLAG_SHARED, it will

allocate a shared interrupt of level 1. Setting `ESP_INTR_FLAG_INTRDISABLED` will return from this function with the interrupt disabled.

- **handler** -- The interrupt handler. Must be NULL when an interrupt of level >3 is requested, because these types of interrupts aren't C-callable.
- **arg** -- Optional argument for passed to the interrupt handler
- **ret_handle** -- Pointer to an `intr_handle_t` to store a handle that can later be used to request details or free the interrupt. Can be NULL if no handle is required.

Returns `ESP_ERR_INVALID_ARG` if the combination of arguments is invalid.
`ESP_ERR_NOT_FOUND` No free interrupt found with the specified flags `ESP_OK` otherwise

esp_err_t **esp_intr_alloc_intrstatus** (int source, int flags, uint32_t intrstatusreg, uint32_t intrstatusmask, *intr_handler_t* handler, void *arg, *intr_handle_t* *ret_handle)

Allocate an interrupt with the given parameters.

This essentially does the same as `esp_intr_alloc`, but allows specifying a register and mask combo. For shared interrupts, the handler is only called if a read from the specified register, ANDed with the mask, returns non-zero. By passing an interrupt status register address and a fitting mask, this can be used to accelerate interrupt handling in the case a shared interrupt is triggered; by checking the interrupt statuses first, the code can decide which ISRs can be skipped

Parameters

- **source** -- The interrupt source. One of the `ETS*_INTR_SOURCE` interrupt mux sources, as defined in `soc/soc.h`, or one of the internal `ETS_INTERNAL*_INTR_SOURCE` sources as defined in this header.
- **flags** -- An ORred mask of the `ESP_INTR_FLAG_*` defines. These restrict the choice of interrupts that this routine can choose from. If this value is 0, it will default to allocating a non-shared interrupt of level 1, 2 or 3. If this is `ESP_INTR_FLAG_SHARED`, it will allocate a shared interrupt of level 1. Setting `ESP_INTR_FLAG_INTRDISABLED` will return from this function with the interrupt disabled.
- **intrstatusreg** -- The address of an interrupt status register
- **intrstatusmask** -- A mask. If a read of address `intrstatusreg` has any of the bits that are 1 in the mask set, the ISR will be called. If not, it will be skipped.
- **handler** -- The interrupt handler. Must be NULL when an interrupt of level >3 is requested, because these types of interrupts aren't C-callable.
- **arg** -- Optional argument for passed to the interrupt handler
- **ret_handle** -- Pointer to an `intr_handle_t` to store a handle that can later be used to request details or free the interrupt. Can be NULL if no handle is required.

Returns `ESP_ERR_INVALID_ARG` if the combination of arguments is invalid.
`ESP_ERR_NOT_FOUND` No free interrupt found with the specified flags `ESP_OK` otherwise

esp_err_t **esp_intr_free** (*intr_handle_t* handle)

Disable and free an interrupt.

Use an interrupt handle to disable the interrupt and release the resources associated with it. If the current core is not the core that registered this interrupt, this routine will be assigned to the core that allocated this interrupt, blocking and waiting until the resource is successfully released.

Note: When the handler shares its source with other handlers, the interrupt status bits it's responsible for should be managed properly before freeing it. see `esp_intr_disable` for more details. Please do not call this function in `esp_ipc_call_blocking`.

Parameters **handle** -- The handle, as obtained by `esp_intr_alloc` or `esp_intr_alloc_intrstatus`

Returns `ESP_ERR_INVALID_ARG` the handle is NULL `ESP_FAIL` failed to release this handle
`ESP_OK` otherwise

int **esp_intr_get_cpu** (*intr_handle_t* handle)

Get CPU number an interrupt is tied to.

Parameters **handle** -- The handle, as obtained by `esp_intr_alloc` or `esp_intr_alloc_intrstatus`

Returns The core number where the interrupt is allocated

int **esp_intr_get_intno** (*intr_handle_t* handle)

Get the allocated interrupt for a certain handle.

Parameters **handle** -- The handle, as obtained by `esp_intr_alloc` or `esp_intr_alloc_intrstatus`

Returns The interrupt number

esp_err_t **esp_intr_disable** (*intr_handle_t* handle)

Disable the interrupt associated with the handle.

Note:

- For local interrupts (ESP_INTERNAL_* sources), this function has to be called on the CPU the interrupt is allocated on. Other interrupts have no such restriction.
 - When several handlers sharing a same interrupt source, interrupt status bits, which are handled in the handler to be disabled, should be masked before the disabling, or handled in other enabled interrupts properly. Miss of interrupt status handling will cause infinite interrupt calls and finally system crash.
-

Parameters **handle** -- The handle, as obtained by `esp_intr_alloc` or `esp_intr_alloc_intrstatus`

Returns ESP_ERR_INVALID_ARG if the combination of arguments is invalid. ESP_OK otherwise

esp_err_t **esp_intr_enable** (*intr_handle_t* handle)

Enable the interrupt associated with the handle.

Note: For local interrupts (ESP_INTERNAL_* sources), this function has to be called on the CPU the interrupt is allocated on. Other interrupts have no such restriction.

Parameters **handle** -- The handle, as obtained by `esp_intr_alloc` or `esp_intr_alloc_intrstatus`

Returns ESP_ERR_INVALID_ARG if the combination of arguments is invalid. ESP_OK otherwise

esp_err_t **esp_intr_set_in_iram** (*intr_handle_t* handle, bool is_in_iram)

Set the "in IRAM" status of the handler.

Note: Does not work on shared interrupts.

Parameters

- handle** -- The handle, as obtained by `esp_intr_alloc` or `esp_intr_alloc_intrstatus`
- is_in_iram** -- Whether the handler associated with this handle resides in IRAM. Handlers residing in IRAM can be called when cache is disabled.

Returns ESP_ERR_INVALID_ARG if the combination of arguments is invalid. ESP_OK otherwise

void **esp_intr_noniram_disable** (void)

Disable interrupts that aren't specifically marked as running from IRAM.

void **esp_intr_noniram_enable** (void)

Re-enable interrupts disabled by `esp_intr_noniram_disable`.

void **esp_intr_enable_source** (int inum)

enable the interrupt source based on its number

Parameters **inum** -- interrupt number from 0 to 31

void **esp_intr_disable_source** (int inum)

disable the interrupt source based on its number

Parameters **inum** -- interrupt number from 0 to 31

static inline int **esp_intr_flags_to_level** (int flags)

Get the lowest interrupt level from the flags.

Parameters **flags** -- The same flags that pass to `esp_intr_alloc_intrstatus` API

static inline int **esp_intr_level_to_flags** (int level)

Get the interrupt flags from the supplied level (priority)

Parameters **level** -- The interrupt priority level

esp_err_t **esp_intr_dump** (FILE *stream)

Dump the status of allocated interrupts.

Parameters **stream** -- The stream to dump to, if NULL then stdout is used

Returns ESP_OK on success

Macros

ESP_INTR_FLAG_LEVEL1

Interrupt allocation flags.

These flags can be used to specify which interrupt qualities the code calling `esp_intr_alloc*` needs. Accept a Level 1 interrupt vector (lowest priority)

ESP_INTR_FLAG_LEVEL2

Accept a Level 2 interrupt vector.

ESP_INTR_FLAG_LEVEL3

Accept a Level 3 interrupt vector.

ESP_INTR_FLAG_LEVEL4

Accept a Level 4 interrupt vector.

ESP_INTR_FLAG_LEVEL5

Accept a Level 5 interrupt vector.

ESP_INTR_FLAG_LEVEL6

Accept a Level 6 interrupt vector.

ESP_INTR_FLAG_NMI

Accept a Level 7 interrupt vector (highest priority)

ESP_INTR_FLAG_SHARED

Interrupt can be shared between ISRs.

ESP_INTR_FLAG_EDGE

Edge-triggered interrupt.

ESP_INTR_FLAG_IRAM

ISR can be called if cache is disabled.

ESP_INTR_FLAG_INTRDISABLED

Return with this interrupt disabled.

ESP_INTR_FLAG_LOWMED

Low and medium prio interrupts. These can be handled in C.

ESP_INTR_FLAG_HIGH

High level interrupts. Need to be handled in assembly.

ESP_INTR_FLAG_LEVELMASK

Mask for all level flags

ETS_INTERNAL_TIMER0_INTR_SOURCE

Platform timer 0 interrupt source.

The `esp_intr_alloc*` functions can allocate an int for all `ETS_*_INTR_SOURCE` interrupt sources that are routed through the interrupt mux. Apart from these sources, each core also has some internal sources that do not pass through the interrupt mux. To allocate an interrupt for these sources, pass these pseudo-sources to the functions.

ETS_INTERNAL_TIMER1_INTR_SOURCE

Platform timer 1 interrupt source.

ETS_INTERNAL_TIMER2_INTR_SOURCE

Platform timer 2 interrupt source.

ETS_INTERNAL_SW0_INTR_SOURCE

Software int source 1.

ETS_INTERNAL_SW1_INTR_SOURCE

Software int source 2.

ETS_INTERNAL_PROFILING_INTR_SOURCE

Int source for profiling.

ETS_INTERNAL_UNUSED_INTR_SOURCE

Interrupt is not assigned to any source.

ETS_INTERNAL_INTR_SOURCE_OFF

Provides SystemView with positive IRQ IDs, otherwise scheduler events are not shown properly

ESP_INTR_ENABLE (inum)

Enable interrupt by interrupt number

ESP_INTR_DISABLE (inum)

Disable interrupt by interrupt number

2.9.22 Logging library

Overview

The logging library provides three ways for setting log verbosity:

- **At compile time:** in menuconfig, set the verbosity level using the option `CONFIG_LOG_DEFAULT_LEVEL`.
- Optionally, also in menuconfig, set the maximum verbosity level using the option `CONFIG_LOG_MAXIMUM_LEVEL`. By default, this is the same as the default level, but it can be set higher in order to compile more optional logs into the firmware.
- **At runtime:** all logs for verbosity levels lower than `CONFIG_LOG_DEFAULT_LEVEL` are enabled by default. The function `esp_log_level_set()` can be used to set a logging level on a per-module basis. Modules are identified by their tags, which are human-readable ASCII zero-terminated strings.
- **At runtime:** if `CONFIG_LOG_MASTER_LEVEL` is enabled then a Master logging level can be set using `esp_log_set_level_master()`. This option adds an additional logging level check for all compiled logs. Note that this will increase application size. This feature is useful if you want to compile in a lot of logs that are selectable at runtime, but also want to avoid the performance hit from looking up the tags and their log level when you don't want log output.

There are the following verbosity levels:

- Error (lowest)
- Warning
- Info
- Debug
- Verbose (highest)

Note: The function `esp_log_level_set()` cannot set logging levels higher than specified by `CONFIG_LOG_MAXIMUM_LEVEL`. To increase log level for a specific file above this maximum at compile time, use the macro `LOG_LOCAL_LEVEL` (see the details below).

How to use this library

In each C file that uses logging functionality, define the TAG variable as shown below:

```
static const char* TAG = "MyModule";
```

Then use one of logging macros to produce output, e.g:

```
ESP_LOGW(TAG, "Baud rate error %.1f%%. Requested: %d baud, actual: %d baud", error_
↪ * 100, baud_req, baud_real);
```

Several macros are available for different verbosity levels:

- `ESP_LOGE` - error (lowest)
- `ESP_LOGW` - warning
- `ESP_LOGI` - info
- `ESP_LOGD` - debug
- `ESP_LOGV` - verbose (highest)

Additionally, there are `ESP_EARLY_LOGx` versions for each of these macros, e.g. `ESP_EARLY_LOGE`. These versions have to be used explicitly in the early startup code only, before heap allocator and syscalls have been initialized. Normal `ESP_LOGx` macros can also be used while compiling the bootloader, but they will fall back to the same implementation as `ESP_EARLY_LOGx` macros.

There are also `ESP_DRAM_LOGx` versions for each of these macros, e.g. `ESP_DRAM_LOGE`. These versions are used in some places where logging may occur with interrupts disabled or with flash cache inaccessible. Use of this macros should be as sparing as possible, as logging in these types of code should be avoided for performance reasons.

Note: Inside critical sections interrupts are disabled so it's only possible to use `ESP_DRAM_LOGx` (preferred) or `ESP_EARLY_LOGx`. Even though it's possible to log in these situations, it's better if your program can be structured not to require it.

To override default verbosity level at file or component scope, define the `LOG_LOCAL_LEVEL` macro.

At file scope, define it before including `esp_log.h`, e.g.:

```
#define LOG_LOCAL_LEVEL ESP_LOG_VERBOSE
#include "esp_log.h"
```

At component scope, define it in the component CMakeLists:

```
target_compile_definitions(${COMPONENT_LIB} PUBLIC "-DLOG_LOCAL_LEVEL=ESP_LOG_
↳VERBOSE")
```

To configure logging output per module at runtime, add calls to the function `esp_log_level_set()` as follows:

```
esp_log_level_set("*", ESP_LOG_ERROR);           // set all components to ERROR level
esp_log_level_set("wifi", ESP_LOG_WARN);         // enable WARN logs from WiFi stack
esp_log_level_set("dhcpc", ESP_LOG_INFO);        // enable INFO logs from DHCP client
```

Note: The "DRAM" and "EARLY" log macro variants documented above do not support per module setting of log verbosity. These macros will always log at the "default" verbosity level, which can only be changed at runtime by calling `esp_log_level_set("*", level)`.

Even when logs are disabled by using a tag name they will still require a processing time of around 10.9 microseconds per entry.

Master Logging Level To enable the Master logging level feature, the `CONFIG_LOG_MASTER_LEVEL` option must be enabled. It adds an additional level check for `ESP_LOGx` macros before calling `esp_log_write()`. This allows to set a higher `CONFIG_LOG_MAXIMUM_LEVEL`, but not inflict a performance hit during normal operation (only when directed). An application may set the master logging level (`esp_log_set_level_master()`) globally to enforce a maximum log level. `ESP_LOGx` macros above this level will be skipped immediately, rather than calling `esp_log_write()` and doing a tag lookup. It is recommended to only use this in a top-level application and not in shared components as this would override the global log level for any user using the component. By default, at startup, the Master logging level is `CONFIG_LOG_DEFAULT_LEVEL`.

Note that this feature increases application size because the additional check is added into all `ESP_LOGx` macros.

The snippet below shows how it works. Setting the Master logging level to `ESP_LOG_NONE` disables all logging globally. `esp_log_level_set()` does not currently affect logging. But after the Master logging level is released, the logs will be printed as set by `esp_log_level_set()`.

```
// Master logging level is CONFIG_LOG_DEFAULT_LEVEL at start up and = ESP_LOG_INFO
ESP_LOGI("lib_name", "Message for print");           // prints a INFO message
esp_log_level_set("lib_name", ESP_LOG_WARN);         // enables WARN logs from lib_
↳name

esp_log_set_level_master(ESP_LOG_NONE);               // disables all logs globally.↳
↳esp_log_level_set has no effect at the moment.

ESP_LOGW("lib_name", "Message for print");           // no print, Master logging↳
↳level blocks it
esp_log_level_set("lib_name", ESP_LOG_INFO);         // enable INFO logs from lib_
↳name
ESP_LOGI("lib_name", "Message for print");           // no print, Master logging↳
↳level blocks it
```

(continues on next page)

(continued from previous page)

```

esp_log_set_level_master(ESP_LOG_INFO);           // enables all INFO logs_
↳ globally.

ESP_LOGI("lib_name", "Message for print");        // prints a INFO message

```

Logging to Host via JTAG By default, the logging library uses the `vprintf`-like function to write formatted output to the dedicated UART. By calling a simple API, all log output may be routed to JTAG instead, making logging several times faster. For details, please refer to Section [Logging to Host](#).

Thread Safety The log string is first written into a memory buffer and then sent to the UART for printing. Log calls are thread-safe, i.e., logs of different threads do not conflict with each other.

Application Example

The logging library is commonly used by most ESP-IDF components and examples. For demonstration of log functionality, check ESP-IDF's [examples](#) directory. The most relevant examples that deal with logging are the following:

- [system/ota](#)
- [storage/sd_card](#)
- [protocols/https_request](#)

API Reference

Header File

- [components/log/include/esp_log.h](#)
- This header file can be included with:

```
#include "esp_log.h"
```

Functions

void **esp_log_set_level_master** ([esp_log_level_t](#) level)

Master log level.

Optional master log level to check against for `ESP_LOGx` macros before calling `esp_log_write`. Allows one to set a higher `CONFIG_LOG_MAXIMUM_LEVEL` but not impose a performance hit during normal operation (only when instructed). An application may set `esp_log_set_level_master(level)` to globally enforce a maximum log level. `ESP_LOGx` macros above this level will be skipped immediately, rather than calling `esp_log_write` and doing a cache hit.

The tradeoff is increased application size.

Parameters `level` -- Master log level

[esp_log_level_t](#) **esp_log_get_level_master** (void)

Returns master log level.

Returns Master log level

void **esp_log_level_set** (const char *tag, [esp_log_level_t](#) level)

Set log level for given tag.

If logging for given component has already been enabled, changes previous setting.

Note: Note that this function can not raise log level above the level set using `CONFIG_LOG_MAXIMUM_LEVEL` setting in `menuconfig`. To raise log level above the default one for a

given file, define LOG_LOCAL_LEVEL to one of the ESP_LOG_* values, before including esp_log.h in this file.

Parameters

- **tag** -- Tag of the log entries to enable. Must be a non-NULL zero terminated string. Value "*" resets log level for all tags to the given value.
- **level** -- Selects log level to enable. Only logs at this and lower verbosity levels will be shown.

esp_log_level_t **esp_log_level_get** (const char *tag)

Get log level for a given tag, can be used to avoid expensive log statements.

Parameters **tag** -- Tag of the log to query current level. Must be a non-NULL zero terminated string.

Returns The current log level for the given tag

vprintf_like_t **esp_log_set_vprintf** (*vprintf_like_t* func)

Set function used to output log entries.

By default, log output goes to UART0. This function can be used to redirect log output to some other destination, such as file or network. Returns the original log handler, which may be necessary to return output to the previous destination.

Note: Please note that function callback here must be re-entrant as it can be invoked in parallel from multiple thread context.

Parameters **func** -- new Function used for output. Must have same signature as vprintf.

Returns func old Function used for output.

uint32_t **esp_log_timestamp** (void)

Function which returns timestamp to be used in log output.

This function is used in expansion of ESP_LOGx macros. In the 2nd stage bootloader, and at early application startup stage this function uses CPU cycle counter as time source. Later when FreeRTOS scheduler start running, it switches to FreeRTOS tick count.

For now, we ignore millisecond counter overflow.

Returns timestamp, in milliseconds

char ***esp_log_system_timestamp** (void)

Function which returns system timestamp to be used in log output.

This function is used in expansion of ESP_LOGx macros to print the system time as "HH:MM:SS.sss". The system time is initialized to 0 on startup, this can be set to the correct time with an SNTP sync, or manually with standard POSIX time functions.

Currently, this will not get used in logging from binary blobs (i.e. Wi-Fi & Bluetooth libraries), these will still print the RTOS tick time.

Returns timestamp, in "HH:MM:SS.sss"

uint32_t **esp_log_early_timestamp** (void)

Function which returns timestamp to be used in log output.

This function uses HW cycle counter and does not depend on OS, so it can be safely used after application crash.

Returns timestamp, in milliseconds

void **esp_log_write** (*esp_log_level_t* level, const char *tag, const char *format, ...)

Write message into the log.

This function is not intended to be used directly. Instead, use one of ESP_LOGE, ESP_LOGW, ESP_LOGI, ESP_LOGD, ESP_LOGV macros.

This function or these macros should not be used from an interrupt.

void **esp_log_writelv** (*esp_log_level_t* level, const char *tag, const char *format, va_list args)

Write message into the log, va_list variant.

This function is provided to ease integration toward other logging framework, so that esp_log can be used as a log sink.

See also:

esp_log_write()

Macros

ESP_LOG_BUFFER_HEX_LEVEL (tag, buffer, buff_len, level)

Log a buffer of hex bytes at specified level, separated into 16 bytes each line.

Parameters

- **tag** -- description tag
- **buffer** -- Pointer to the buffer array
- **buff_len** -- length of buffer in bytes
- **level** -- level of the log

ESP_LOG_BUFFER_CHAR_LEVEL (tag, buffer, buff_len, level)

Log a buffer of characters at specified level, separated into 16 bytes each line. Buffer should contain only printable characters.

Parameters

- **tag** -- description tag
- **buffer** -- Pointer to the buffer array
- **buff_len** -- length of buffer in bytes
- **level** -- level of the log

ESP_LOG_BUFFER_HEXDUMP (tag, buffer, buff_len, level)

Dump a buffer to the log at specified level.

The dump log shows just like the one below:

```
W (195) log_example: 0x3ffb4280  45 53 50 33 32 20 69 73  20 67 72 65 61 74_
→2c 20 |ESP32 is great, |
W (195) log_example: 0x3ffb4290  77 6f 72 6b 69 6e 67 20  61 6c 6f 6e 67 20_
→77 69 |working along wi|
W (205) log_example: 0x3ffb42a0  74 68 20 74 68 65 20 49  44 46 2e 00      _
→      |th the IDF..|
```

It is highly recommended to use terminals with over 102 text width.

Parameters

- **tag** -- description tag
- **buffer** -- Pointer to the buffer array
- **buff_len** -- length of buffer in bytes
- **level** -- level of the log

ESP_LOG_BUFFER_HEX (tag, buffer, buff_len)

Log a buffer of hex bytes at Info level.

See also:`esp_log_buffer_hex_level`**Parameters**

- **tag** -- description tag
- **buffer** -- Pointer to the buffer array
- **buff_len** -- length of buffer in bytes

ESP_LOG_BUFFER_CHAR (tag, buffer, buff_len)

Log a buffer of characters at Info level. Buffer should contain only printable characters.

See also:`esp_log_buffer_char_level`**Parameters**

- **tag** -- description tag
- **buffer** -- Pointer to the buffer array
- **buff_len** -- length of buffer in bytes

ESP_EARLY_LOGE (tag, format, ...)

macro to output logs in startup code, before heap allocator and syscalls have been initialized. Log at ESP_LOG_ERROR level.

See also:

`printf, ESP_LOGE, ESP_DRAM_LOGE` In the future, we want to become compatible with clang. Hence, we provide two versions of the following macros which are using variadic arguments. The first one is using the GNU extension `##_VA_ARGS__`. The second one is using the C++20 feature `VA_OPT(,)`. This allows users to compile their code with standard C++20 enabled instead of the GNU extension. Below C++20, we haven't found any good alternative to using `##_VA_ARGS__`.

ESP_EARLY_LOGW (tag, format, ...)

macro to output logs in startup code at ESP_LOG_WARN level.

See also:`ESP_EARLY_LOGE, ESP_LOGE, printf`**ESP_EARLY_LOGI** (tag, format, ...)

macro to output logs in startup code at ESP_LOG_INFO level.

See also:`ESP_EARLY_LOGE, ESP_LOGE, printf`**ESP_EARLY_LOGD** (tag, format, ...)

macro to output logs in startup code at ESP_LOG_DEBUG level.

See also:`ESP_EARLY_LOGE, ESP_LOGE, printf`

ESP_EARLY_LOGV (tag, format, ...)

macro to output logs in startup code at ESP_LOG_VERBOSE level.

See also:

ESP_EARLY_LOGE, ESP_LOGE, printf

_ESP_LOG_EARLY_ENABLED (log_level)

ESP_LOG_EARLY_IMPL (tag, format, log_level, log_tag_letter, ...)

ESP_LOGE (tag, format, ...)

ESP_LOGW (tag, format, ...)

ESP_LOGI (tag, format, ...)

ESP_LOGD (tag, format, ...)

ESP_LOGV (tag, format, ...)

ESP_LOG_LEVEL (level, tag, format, ...)

runtime macro to output logs at a specified level.

See also:

printf

Parameters

- **tag** -- tag of the log, which can be used to change the log level by `esp_log_level_set` at runtime.
- **level** -- level of the output log.
- **format** -- format of the output log. See `printf`
- ... -- variables to be replaced into the log. See `printf`

ESP_LOG_LEVEL_LOCAL (level, tag, format, ...)

runtime macro to output logs at a specified level. Also check the level with `LOG_LOCAL_LEVEL`. If `CONFIG_LOG_MASTER_LEVEL` set, also check first against `esp_log_get_level_master()`.

See also:

printf, ESP_LOG_LEVEL

ESP_DRAM_LOGE (tag, format, ...)

Macro to output logs when the cache is disabled. Log at ESP_LOG_ERROR level.

Similar to

Usage: `ESP_DRAM_LOGE(DRAM_STR("my_tag"), "format", or ESP_DRAM_LOGE(TAG, "format", ...)`, where TAG is a `char*` that points to a str in the DRAM.

See also:

ESP_EARLY_LOGE, the log level cannot be changed per-tag, however `esp_log_level_set("", level)` will set the default level which controls these log lines also.

See also:

esp_rom_printf, ESP_LOGE

Note: Unlike normal logging macros, it's possible to use this macro when interrupts are disabled or inside an ISR.

Note: Placing log strings in DRAM reduces available DRAM, so only use when absolutely essential.

ESP_DRAM_LOGW (tag, format, ...)

macro to output logs when the cache is disabled at ESP_LOG_WARN level.

See also:

ESP_DRAM_LOGW, ESP_LOGW, esp_rom_printf

ESP_DRAM_LOGI (tag, format, ...)

macro to output logs when the cache is disabled at ESP_LOG_INFO level.

See also:

ESP_DRAM_LOGI, ESP_LOGI, esp_rom_printf

ESP_DRAM_LOGD (tag, format, ...)

macro to output logs when the cache is disabled at ESP_LOG_DEBUG level.

See also:

ESP_DRAM_LOGD, ESP_LOGD, esp_rom_printf

ESP_DRAM_LOGV (tag, format, ...)

macro to output logs when the cache is disabled at ESP_LOG_VERBOSE level.

See also:

ESP_DRAM_LOGV, ESP_LOGV, esp_rom_printf

Type Definitions

typedef int (***vprintf_like_t**)(const char*, va_list)

Enumerations

enum **esp_log_level_t**

Log level.

Values:

enumerator **ESP_LOG_NONE**

No log output

enumerator **ESP_LOG_ERROR**

Critical errors, software module can not recover on its own

enumerator **ESP_LOG_WARN**

Error conditions from which recovery measures have been taken

enumerator **ESP_LOG_INFO**

Information messages which describe normal flow of events

enumerator **ESP_LOG_DEBUG**

Extra information which is not necessary for normal use (values, pointers, sizes, etc).

enumerator **ESP_LOG_VERBOSE**

Bigger chunks of debugging information, or frequent messages which can potentially flood the output.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.9.23 Miscellaneous System APIs

Software Reset

To perform software reset of the chip, the `esp_restart()` function is provided. When the function is called, execution of the program stops, both CPUs are reset, the application is loaded by the bootloader and starts execution again.

Additionally, the `esp_register_shutdown_handler()` function can register a routine that will be automatically called before a restart (that is triggered by `esp_restart()`) occurs. This is similar to the functionality of `atexit` POSIX function.

Reset Reason

ESP-IDF applications can be started or restarted due to a variety of reasons. To get the last reset reason, call `esp_reset_reason()` function. See description of `esp_reset_reason_t` for the list of possible reset reasons.

Heap Memory

Two heap-memory-related functions are provided:

- `esp_get_free_heap_size()` returns the current size of free heap memory.
- `esp_get_minimum_free_heap_size()` returns the minimum size of free heap memory that has ever been available (i.e., the smallest size of free heap memory in the application's lifetime).

Note that ESP-IDF supports multiple heaps with different capabilities. The functions mentioned in this section return the size of heap memory that can be allocated using the `malloc` family of functions. For further information about heap memory, see [Heap Memory Allocation](#).

MAC Address

These APIs allow querying and customizing MAC addresses for different supported network interfaces (e.g., Wi-Fi, Bluetooth, Ethernet).

To fetch the MAC address for a specific network interface (e.g., Wi-Fi, Bluetooth, Ethernet), call the function `esp_read_mac()`.

In ESP-IDF, the MAC addresses for the various network interfaces are calculated from a single **base MAC address**. By default, the Espressif base MAC address is used. This base MAC address is pre-programmed into the ESP32-P4 eFuse in the factory during production.

Interface	MAC Address (4 universally administered, default)	MAC Address (2 universally administered)
Wi-Fi Station	base_mac	base_mac
Wi-Fi SoftAP	base_mac, +1 to the last octet	<i>Local MAC</i> (derived from Wi-Fi Station MAC)
Bluetooth	base_mac, +2 to the last octet	base_mac, +1 to the last octet
Ethernet	base_mac, +3 to the last octet	<i>Local MAC</i> (derived from Bluetooth MAC)

Note: The `configuration` configures the number of universally administered MAC addresses that are provided by Espressif.

Note: Although ESP32-P4 has no integrated Ethernet MAC, it is still possible to calculate an Ethernet MAC address. However, this MAC address can only be used with an external ethernet interface such as an SPI-Ethernet device. See *Ethernet*.

Custom Interface MAC Sometimes you may need to define custom MAC addresses that are not generated from the base MAC address. To set a custom interface MAC address, use the `esp_iface_mac_addr_set()` function. This function allows you to overwrite the MAC addresses of interfaces set (or not yet set) by the base MAC address. Once a MAC address has been set for a particular interface, it will not be affected when the base MAC address is changed.

Custom Base MAC The default base MAC is pre-programmed by Espressif in eFuse BLK1. To set a custom base MAC instead, call the function `esp_iface_mac_addr_set()` with the `ESP_MAC_BASE` argument (or `esp_base_mac_addr_set()`) before initializing any network interfaces or calling the `esp_read_mac()` function. The custom MAC address can be stored in any supported storage device (e.g., flash, NVS).

The custom base MAC addresses should be allocated such that derived MAC addresses will not overlap. Based on the table above, users can configure the option `CONFIG_ESP32P4_UNIVERSAL_MAC_ADDRESSES` to set the number of valid universal MAC addresses that can be derived from the custom base MAC.

Note: It is also possible to call the function `esp_netif_set_mac()` to set the specific MAC used by a network interface after network initialization. But it is recommended to use the base MAC approach documented here to avoid the possibility of the original MAC address briefly appearing on the network before being changed.

Custom MAC Address in eFuse When reading custom MAC addresses from eFuse, ESP-IDF provides a helper function `esp_efuse_mac_get_custom()`. Users can also use `esp_read_mac()` with the `ESP_MAC_EFUSE_CUSTOM` argument. This loads the MAC address from eFuse BLK3. The `esp_efuse_mac_get_custom()` function assumes that the custom base MAC address is stored in the following format:

Field	# of bits	Range of bits
MAC address	48	200:248

Note: The eFuse BLK3 uses RS-coding during burning, which means that all eFuse fields in this block must be burnt at the same time.

Once custom eFuse MAC address has been obtained (using `esp_efuse_mac_get_custom()` or `esp_read_mac()`), you need to set it as the base MAC address. There are two ways to do it:

1. Use an old API: call `esp_base_mac_addr_set()`.
2. Use a new API: call `esp_iface_mac_addr_set()` with the `ESP_MAC_BASE` argument.

Local Versus Universal MAC Addresses ESP32-P4 comes pre-programmed with enough valid Espressif universally administered MAC addresses for all internal interfaces. The table above shows how to calculate and derive the MAC address for a specific interface according to the base MAC address.

When using a custom MAC address scheme, it is possible that not all interfaces can be assigned with a universally administered MAC address. In these cases, a locally administered MAC address is assigned. Note that these addresses are intended for use on a single local network only.

See [this article](#) for the definition of locally and universally administered MAC addresses.

Function `esp_derive_local_mac()` is called internally to derive a local MAC address from a universal MAC address. The process is as follows:

1. The U/L bit (bit value 0x2) is set in the first octet of the universal MAC address, creating a local MAC address.
2. If this bit is already set in the supplied universal MAC address (i.e., the supplied "universal" MAC address was in fact already a local MAC address), then the first octet of the local MAC address is XORed with 0x4.

Chip Version

`esp_chip_info()` function fills `esp_chip_info_t` structure with information about the chip. This includes the chip revision, number of CPU cores, and a bit mask of features enabled in the chip.

SDK Version

`esp_get_idf_version()` returns a string describing the ESP-IDF version which is used to compile the application. This is the same value as the one available through `IDF_VER` variable of the build system. The version string generally has the format of `git describe` output.

To get the version at build time, additional version macros are provided. They can be used to enable or disable parts of the program depending on the ESP-IDF version.

- `ESP_IDF_VERSION_MAJOR`, `ESP_IDF_VERSION_MINOR`, `ESP_IDF_VERSION_PATCH` are defined to integers representing major, minor, and patch version.
- `ESP_IDF_VERSION_VAL` and `ESP_IDF_VERSION` can be used when implementing version checks:

```
#include "esp_idf_version.h"

#if ESP_IDF_VERSION >= ESP_IDF_VERSION_VAL(4, 0, 0)
    // enable functionality present in ESP-IDF v4.0
#endif
```

App Version

The application version is stored in `esp_app_desc_t` structure. It is located in DROM sector and has a fixed offset from the beginning of the binary file. The structure is located after `esp_image_header_t` and `esp_image_segment_header_t` structures. The type of the field version is string and it has a maximum length of 32 chars.

To set the version in your project manually, you need to set the `PROJECT_VER` variable in the `CMakeLists.txt` of your project. In application `CMakeLists.txt`, put `set (PROJECT_VER "0.1.0.1")` before including `project.cmake`.

If the `CONFIG_APP_PROJECT_VER_FROM_CONFIG` option is set, the value of `CONFIG_APP_PROJECT_VER` will be used. Otherwise, if the `PROJECT_VER` variable is not set in the project, it will be retrieved either from the `$(PROJECT_PATH)/version.txt` file (if present) or using git command `git describe`. If neither is available, `PROJECT_VER` will be set to "1". Application can make use of this by calling `esp_app_get_description()` or `esp_ota_get_partition_description()` functions.

API Reference

Header File

- `components/esp_system/include/esp_system.h`
- This header file can be included with:

```
#include "esp_system.h"
```

Functions

`esp_err_t esp_register_shutdown_handler (shutdown_handler_t handle)`

Register shutdown handler.

This function allows you to register a handler that gets invoked before the application is restarted using `esp_restart` function.

Parameters `handle` -- function to execute on restart

Returns

- `ESP_OK` on success
- `ESP_ERR_INVALID_STATE` if the handler has already been registered
- `ESP_ERR_NO_MEM` if no more shutdown handler slots are available

`esp_err_t esp_unregister_shutdown_handler (shutdown_handler_t handle)`

Unregister shutdown handler.

This function allows you to unregister a handler which was previously registered using `esp_register_shutdown_handler` function.

- `ESP_OK` on success
- `ESP_ERR_INVALID_STATE` if the given handler hasn't been registered before

void `esp_restart` (void)

Restart PRO and APP CPUs.

This function can be called both from PRO and APP CPUs. After successful restart, CPU reset reason will be `SW_CPU_RESET`. Peripherals (except for Wi-Fi, BT, UART0, SPI1, and legacy timers) are not reset. This function does not return.

`esp_reset_reason_t esp_reset_reason` (void)

Get reason of last reset.

Returns See description of `esp_reset_reason_t` for explanation of each value.

uint32_t `esp_get_free_heap_size` (void)

Get the size of available heap.

Note: Note that the returned value may be larger than the maximum contiguous block which can be allocated.

Returns Available heap size, in bytes.

uint32_t **esp_get_free_internal_heap_size** (void)

Get the size of available internal heap.

Note: Note that the returned value may be larger than the maximum contiguous block which can be allocated.

Returns Available internal heap size, in bytes.

uint32_t **esp_get_minimum_free_heap_size** (void)

Get the minimum heap that has ever been available.

Returns Minimum free heap ever available

void **esp_system_abort** (const char *details)

Trigger a software abort.

Parameters details -- Details that will be displayed during panic handling.

Type Definitions

typedef void (***shutdown_handler_t**)(void)

Shutdown handler type

Enumerations

enum **esp_reset_reason_t**

Reset reasons.

Values:

enumerator **ESP_RST_UNKNOWN**

Reset reason can not be determined.

enumerator **ESP_RST_POWERON**

Reset due to power-on event.

enumerator **ESP_RST_EXT**

Reset by external pin (not applicable for ESP32)

enumerator **ESP_RST_SW**

Software reset via esp_restart.

enumerator **ESP_RST_PANIC**

Software reset due to exception/panic.

enumerator **ESP_RST_INT_WDT**

Reset (software or hardware) due to interrupt watchdog.

enumerator **ESP_RST_TASK_WDT**

Reset due to task watchdog.

enumerator **ESP_RST_WDT**

Reset due to other watchdogs.

enumerator **ESP_RST_DEEPSLEEP**

Reset after exiting deep sleep mode.

enumerator **ESP_RST_BROWNOUT**

Brownout reset (software or hardware)

enumerator **ESP_RST_SDIO**

Reset over SDIO.

enumerator **ESP_RST_USB**

Reset by USB peripheral.

enumerator **ESP_RST_JTAG**

Reset by JTAG.

enumerator **ESP_RST_EFUSE**

Reset due to efuse error.

enumerator **ESP_RST_PWR_GLITCH**

Reset due to power glitch detected.

enumerator **ESP_RST_CPU_LOCKUP**

Reset due to CPU lock up.

Header File

- [components/esp_common/include/esp_idf_version.h](#)
- This header file can be included with:

```
#include "esp_idf_version.h"
```

Functions

const char ***esp_get_idf_version** (void)

Return full IDF version string, same as 'git describe' output.

Note: If you are printing the ESP-IDF version in a log file or other information, this function provides more information than using the numerical version macros. For example, numerical version macros don't differentiate between development, pre-release and release versions, but the output of this function does.

Returns constant string from IDF_VER

Macros

ESP_IDF_VERSION_MAJOR

Major version number (X.x.x)

ESP_IDF_VERSION_MINOR

Minor version number (x.X.x)

ESP_IDF_VERSION_PATCH

Patch version number (x.x.X)

ESP_IDF_VERSION_VAL (major, minor, patch)

Macro to convert IDF version number into an integer

To be used in comparisons, such as ESP_IDF_VERSION >= ESP_IDF_VERSION_VAL(4, 0, 0)

ESP_IDF_VERSION

Current IDF version, as an integer

To be used in comparisons, such as ESP_IDF_VERSION >= ESP_IDF_VERSION_VAL(4, 0, 0)

Header File

- [components/esp_hw_support/include/esp_mac.h](#)
- This header file can be included with:

```
#include "esp_mac.h"
```

Functions

esp_err_t esp_base_mac_addr_set (const uint8_t *mac)

Set base MAC address with the MAC address which is stored in BLK3 of EFUSE or external storage e.g. flash and EEPROM.

Base MAC address is used to generate the MAC addresses used by network interfaces.

If using a custom base MAC address, call this API before initializing any network interfaces. Refer to the ESP-IDF Programming Guide for details about how the Base MAC is used.

Note: Base MAC must be a unicast MAC (least significant bit of first byte must be zero).

Note: If not using a valid OUI, set the "locally administered" bit (bit value 0x02 in the first byte) to avoid collisions.

Parameters **mac** -- base MAC address, length: 6 bytes. length: 6 bytes for MAC-48

Returns ESP_OK on success ESP_ERR_INVALID_ARG If mac is NULL or is not a unicast MAC

esp_err_t esp_base_mac_addr_get (uint8_t *mac)

Return base MAC address which is set using esp_base_mac_addr_set.

Note: If no custom Base MAC has been set, this returns the pre-programmed Espressif base MAC address.

Parameters **mac** -- base MAC address, length: 6 bytes. length: 6 bytes for MAC-48

Returns ESP_OK on success ESP_ERR_INVALID_ARG mac is NULL
ESP_ERR_INVALID_MAC base MAC address has not been set

esp_err_t esp_efuse_mac_get_custom (uint8_t *mac)

Return base MAC address which was previously written to BLK3 of EFUSE.

Base MAC address is used to generate the MAC addresses used by the networking interfaces. This API returns the custom base MAC address which was previously written to EFUSE BLK3 in a specified format.

Writing this EFUSE allows setting of a different (non-Espressif) base MAC address. It is also possible to store a custom base MAC address elsewhere, see `esp_base_mac_addr_set()` for details.

Note: This function is currently only supported on ESP32.

Parameters **mac** -- base MAC address, length: 6 bytes/8 bytes. length: 6 bytes for MAC-48 8 bytes for EUI-64(used for IEEE 802.15.4, if CONFIG_SOC_IEEE802154_SUPPORTED=y)

Returns ESP_OK on success ESP_ERR_INVALID_ARG mac is NULL
ESP_ERR_INVALID_MAC CUSTOM_MAC address has not been set, all zeros (for esp32-xx)
ESP_ERR_INVALID_VERSION An invalid MAC version field was read from BLK3 of EFUSE (for esp32)
ESP_ERR_INVALID_CRC An invalid MAC CRC was read from BLK3 of EFUSE (for esp32)

esp_err_t **esp_efuse_mac_get_default** (uint8_t *mac)

Return base MAC address which is factory-programmed by Espressif in EFUSE.

Parameters **mac** -- base MAC address, length: 6 bytes/8 bytes. length: 6 bytes for MAC-48 8 bytes for EUI-64(used for IEEE 802.15.4, if CONFIG_SOC_IEEE802154_SUPPORTED=y)

Returns ESP_OK on success ESP_ERR_INVALID_ARG mac is NULL

esp_err_t **esp_read_mac** (uint8_t *mac, *esp_mac_type_t* type)

Read base MAC address and set MAC address of the interface.

This function first get base MAC address using `esp_base_mac_addr_get()`. Then calculates the MAC address of the specific interface requested, refer to ESP-IDF Programming Guide for the algorithm.

The MAC address set by the `esp_iface_mac_addr_set()` function will not depend on the base MAC address.

Parameters

- **mac** -- base MAC address, length: 6 bytes/8 bytes. length: 6 bytes for MAC-48 8 bytes for EUI-64(used for IEEE 802.15.4, if CONFIG_SOC_IEEE802154_SUPPORTED=y)
- **type** -- Type of MAC address to return

Returns ESP_OK on success

esp_err_t **esp_derive_local_mac** (uint8_t *local_mac, const uint8_t *universal_mac)

Derive local MAC address from universal MAC address.

This function copies a universal MAC address and then sets the "locally

administered" bit (bit 0x2) in the first octet, creating a locally administered MAC address.

If the universal MAC address argument is already a locally administered MAC address, then the first octet is XORed with 0x4 in order to create a different locally administered MAC address.

Parameters

- **local_mac** -- base MAC address, length: 6 bytes. length: 6 bytes for MAC-48
- **universal_mac** -- Source universal MAC address, length: 6 bytes.

Returns ESP_OK on success

esp_err_t **esp_iface_mac_addr_set** (const uint8_t *mac, *esp_mac_type_t* type)

Set custom MAC address of the interface. This function allows you to overwrite the MAC addresses of the interfaces set by the base MAC address.

Parameters

- **mac** -- MAC address, length: 6 bytes/8 bytes. length: 6 bytes for MAC-48 8 bytes for EUI-64(used for ESP_MAC_IEEE802154 type, if CONFIG_SOC_IEEE802154_SUPPORTED=y)
- **type** -- Type of MAC address

Returns ESP_OK on success

size_t **esp_mac_addr_len_get** (*esp_mac_type_t* type)

Return the size of the MAC type in bytes.

If CONFIG_SOC_IEEE802154_SUPPORTED is set then for these types:

- ESP_MAC_IEEE802154 is 8 bytes.
- ESP_MAC_BASE, ESP_MAC_EFUSE_FACTORY and ESP_MAC_EFUSE_CUSTOM the MAC size is 6 bytes.
- ESP_MAC_EFUSE_EXT is 2 bytes. If CONFIG_SOC_IEEE802154_SUPPORTED is not set then for all types it returns 6 bytes.

Parameters type -- Type of MAC address

Returns 0 MAC type not found (not supported) 6 bytes for MAC-48. 8 bytes for EUI-64.

Macros

MAC2STR (a)

MACSTR

Enumerations

enum **esp_mac_type_t**

Values:

enumerator **ESP_MAC_WIFI_STA**

MAC for WiFi Station (6 bytes)

enumerator **ESP_MAC_WIFI_SOFTAP**

MAC for WiFi Soft-AP (6 bytes)

enumerator **ESP_MAC_BT**

MAC for Bluetooth (6 bytes)

enumerator **ESP_MAC_ETH**

MAC for Ethernet (6 bytes)

enumerator **ESP_MAC_IEEE802154**

if CONFIG_SOC_IEEE802154_SUPPORTED=y, MAC for IEEE802154 (8 bytes)

enumerator **ESP_MAC_BASE**

Base MAC for that used for other MAC types (6 bytes)

enumerator **ESP_MAC_EFUSE_FACTORY**

MAC_FACTORY eFuse which was burned by Espressif in production (6 bytes)

enumerator **ESP_MAC_EFUSE_CUSTOM**

MAC_CUSTOM eFuse which was can be burned by customer (6 bytes)

enumerator **ESP_MAC_EFUSE_EXT**

if CONFIG_SOC_IEEE802154_SUPPORTED=y, MAC_EXT eFuse which is used as an extender for IEEE802154 MAC (2 bytes)

Header File

- [components/esp_hw_support/include/esp_chip_info.h](#)
- This header file can be included with:

```
#include "esp_chip_info.h"
```

Functions

void **esp_chip_info** ([esp_chip_info_t](#) *out_info)

Fill an [esp_chip_info_t](#) structure with information about the chip.

Parameters **out_info** -- [out] structure to be filled

Structures

struct **esp_chip_info_t**

The structure represents information about the chip.

Public Members

[esp_chip_model_t](#) **model**

chip model, one of [esp_chip_model_t](#)

uint32_t **features**

bit mask of CHIP_FEATURE_x feature flags

uint16_t **revision**

chip revision number (in format MXX; where M - wafer major version, XX - wafer minor version)

uint8_t **cores**

number of CPU cores

Macros

CHIP_FEATURE_EMB_FLASH

Chip has embedded flash memory.

CHIP_FEATURE_WIFI_BGN

Chip has 2.4GHz WiFi.

CHIP_FEATURE_BLE

Chip has Bluetooth LE.

CHIP_FEATURE_BT

Chip has Bluetooth Classic.

CHIP_FEATURE_IEEE802154

Chip has IEEE 802.15.4.

CHIP_FEATURE_EMB_PSRAM

Chip has embedded psram.

Enumerations

enum **esp_chip_model_t**

Chip models.

Values:

enumerator **CHIP_ESP32**

ESP32.

enumerator **CHIP_ESP32S2**

ESP32-S2.

enumerator **CHIP_ESP32S3**

ESP32-S3.

enumerator **CHIP_ESP32C3**

ESP32-C3.

enumerator **CHIP_ESP32C2**

ESP32-C2.

enumerator **CHIP_ESP32C6**

ESP32-C6.

enumerator **CHIP_ESP32H2**

ESP32-H2.

enumerator **CHIP_ESP32P4**

ESP32-P4.

enumerator **CHIP_POSIX_LINUX**

The code is running on POSIX/Linux simulator.

Header File

- [components/esp_hw_support/include/esp_cpu.h](#)
- This header file can be included with:

```
#include "esp_cpu.h"
```

Functions

void **esp_cpu_stall** (int core_id)

Stall a CPU core.

Parameters **core_id** -- The core's ID

void **esp_cpu_unstall** (int core_id)

Resume a previously stalled CPU core.

Parameters **core_id** -- The core's ID

void **esp_cpu_reset** (int core_id)

Reset a CPU core.

Parameters **core_id** -- The core's ID

void **esp_cpu_wait_for_intr** (void)

Wait for Interrupt.

This function causes the current CPU core to execute its Wait For Interrupt (WFI or equivalent) instruction. After executing this function, the CPU core will stop execution until an interrupt occurs.

int **esp_cpu_get_core_id** (void)

Get the current core's ID.

This function will return the ID of the current CPU (i.e., the CPU that calls this function).

Returns The current core's ID [0..SOC_CPU_CORES_NUM - 1]

void ***esp_cpu_get_sp** (void)

Read the current stack pointer address.

Returns Stack pointer address

esp_cpu_cycle_count_t **esp_cpu_get_cycle_count** (void)

Get the current CPU core's cycle count.

Each CPU core maintains an internal counter (i.e., cycle count) that increments every CPU clock cycle.

Returns Current CPU's cycle count, 0 if not supported.

void **esp_cpu_set_cycle_count** (*esp_cpu_cycle_count_t* cycle_count)

Set the current CPU core's cycle count.

Set the given value into the internal counter that increments every CPU clock cycle.

Parameters **cycle_count** -- CPU cycle count

void ***esp_cpu_pc_to_addr** (uint32_t pc)

Convert a program counter (PC) value to address.

If the architecture does not store the true virtual address in the CPU's PC or return addresses, this function will convert the PC value to a virtual address. Otherwise, the PC is just returned

Parameters **pc** -- PC value

Returns Virtual address

void **esp_cpu_intr_get_desc** (int core_id, int intr_num, *esp_cpu_intr_desc_t* *intr_desc_ret)

Get a CPU interrupt's descriptor.

Each CPU interrupt has a descriptor describing the interrupt's capabilities and restrictions. This function gets the descriptor of a particular interrupt on a particular CPU.

Parameters

- **core_id** -- [in] The core's ID
- **intr_num** -- [in] Interrupt number
- **intr_desc_ret** -- [out] The interrupt's descriptor

void **esp_cpu_intr_set_ivt_addr** (const void *ivt_addr)

Set the base address of the current CPU's Interrupt Vector Table (IVT)

Parameters **ivt_addr** -- Interrupt Vector Table's base address

void **esp_cpu_intr_set_mtv_t_addr** (const void *mtvt_addr)

Set the base address of the current CPU's Interrupt Vector Table (MTVT)

Note: The MTVT table is only applicable when CLIC is supported

Parameters `mtvt_addr` -- Interrupt Vector Table's base address

void **esp_cpu_intr_set_type** (int intr_num, *esp_cpu_intr_type_t* intr_type)

Set the interrupt type of a particular interrupt.

Set the interrupt type (Level or Edge) of a particular interrupt on the current CPU.

Parameters

- **intr_num** -- Interrupt number (from 0 to 31)
- **intr_type** -- The interrupt's type

esp_cpu_intr_type_t **esp_cpu_intr_get_type** (int intr_num)

Get the current configured type of a particular interrupt.

Get the currently configured type (i.e., level or edge) of a particular interrupt on the current CPU.

Parameters **intr_num** -- Interrupt number (from 0 to 31)

Returns Interrupt type

void **esp_cpu_intr_set_priority** (int intr_num, int intr_priority)

Set the priority of a particular interrupt.

Set the priority of a particular interrupt on the current CPU.

Parameters

- **intr_num** -- Interrupt number (from 0 to 31)
- **intr_priority** -- The interrupt's priority

int **esp_cpu_intr_get_priority** (int intr_num)

Get the current configured priority of a particular interrupt.

Get the currently configured priority of a particular interrupt on the current CPU.

Parameters **intr_num** -- Interrupt number (from 0 to 31)

Returns Interrupt's priority

bool **esp_cpu_intr_has_handler** (int intr_num)

Check if a particular interrupt already has a handler function.

Check if a particular interrupt on the current CPU already has a handler function assigned.

Note: This function simply checks if the IVT of the current CPU already has a handler assigned.

Parameters **intr_num** -- Interrupt number (from 0 to 31)

Returns True if the interrupt has a handler function, false otherwise.

void **esp_cpu_intr_set_handler** (int intr_num, *esp_cpu_intr_handler_t* handler, void *handler_arg)

Set the handler function of a particular interrupt.

Assign a handler function (i.e., ISR) to a particular interrupt on the current CPU.

Note: This function simply sets the handler function (in the IVT) and does not actually enable the interrupt.

Parameters

- **intr_num** -- Interrupt number (from 0 to 31)
- **handler** -- Handler function
- **handler_arg** -- Argument passed to the handler function

void ***esp_cpu_intr_get_handler_arg** (int intr_num)

Get a handler function's argument of.

Get the argument of a previously assigned handler function on the current CPU.

Parameters **intr_num** -- Interrupt number (from 0 to 31)

Returns The the argument passed to the handler function

void **esp_cpu_intr_enable** (uint32_t intr_mask)

Enable particular interrupts on the current CPU.

Parameters **intr_mask** -- Bit mask of the interrupts to enable

void **esp_cpu_intr_disable** (uint32_t intr_mask)

Disable particular interrupts on the current CPU.

Parameters **intr_mask** -- Bit mask of the interrupts to disable

uint32_t **esp_cpu_intr_get_enabled_mask** (void)

Get the enabled interrupts on the current CPU.

Returns Bit mask of the enabled interrupts

void **esp_cpu_intr_edge_ack** (int intr_num)

Acknowledge an edge interrupt.

Parameters **intr_num** -- Interrupt number (from 0 to 31)

void **esp_cpu_configure_region_protection** (void)

Configure the CPU to disable access to invalid memory regions.

esp_err_t **esp_cpu_set_breakpoint** (int bp_num, const void *bp_addr)

Set and enable a hardware breakpoint on the current CPU.

Note: This function is meant to be called by the panic handler to set a breakpoint for an attached debugger during a panic.

Note: Overwrites previously set breakpoint with same breakpoint number.

Parameters

- **bp_num** -- Hardware breakpoint number [0..SOC_CPU_BREAKPOINTS_NUM - 1]
- **bp_addr** -- Address to set a breakpoint on

Returns ESP_OK if breakpoint is set. Failure otherwise

esp_err_t **esp_cpu_clear_breakpoint** (int bp_num)

Clear a hardware breakpoint on the current CPU.

Note: Clears a breakpoint regardless of whether it was previously set

Parameters **bp_num** -- Hardware breakpoint number [0..SOC_CPU_BREAKPOINTS_NUM - 1]

Returns ESP_OK if breakpoint is cleared. Failure otherwise

esp_err_t **esp_cpu_set_watchpoint** (int wp_num, const void *wp_addr, size_t size, *esp_cpu_watchpoint_trigger_t* trigger)

Set and enable a hardware watchpoint on the current CPU.

Set and enable a hardware watchpoint on the current CPU, specifying the memory range and trigger operation. Watchpoints will break/panic the CPU when the CPU accesses (according to the trigger type) on a certain memory range.

Note: Overwrites previously set watchpoint with same watchpoint number. On RISC-V chips, this API uses method0(Exact matching) and method1(NAPOT matching) according to the riscv-debug-spec-0.13 specification for address matching. If the watch region size is 1byte, it uses exact matching (method 0). If the watch region size is larger than 1byte, it uses NAPOT matching (method 1). This mode requires the watching region start address to be aligned to the watching region size.

Parameters

- **wp_num** -- Hardware watchpoint number [0..SOC_CPU_WATCHPOINTS_NUM - 1]
- **wp_addr** -- Watchpoint's base address, must be naturally aligned to the size of the region
- **size** -- Size of the region to watch. Must be one of 2^n and in the range of [1 ... SOC_CPU_WATCHPOINT_MAX_REGION_SIZE]
- **trigger** -- Trigger type

Returns ESP_ERR_INVALID_ARG on invalid arg, ESP_OK otherwise

esp_err_t **esp_cpu_clear_watchpoint** (int wp_num)

Clear a hardware watchpoint on the current CPU.

Note: Clears a watchpoint regardless of whether it was previously set

Parameters **wp_num** -- Hardware watchpoint number [0..SOC_CPU_WATCHPOINTS_NUM - 1]

Returns ESP_OK if watchpoint was cleared. Failure otherwise.

bool **esp_cpu_dbggr_is_attached** (void)

Check if the current CPU has a debugger attached.

Returns True if debugger is attached, false otherwise

void **esp_cpu_dbggr_break** (void)

Trigger a call to the current CPU's attached debugger.

intptr_t **esp_cpu_get_call_addr** (intptr_t return_address)

Given the return address, calculate the address of the preceding call instruction This is typically used to answer the question "where was the function called from?".

Parameters **return_address** -- The value of the return address register. Typically set to the value of __builtin_return_address(0).

Returns Address of the call instruction preceding the return address.

bool **esp_cpu_compare_and_set** (volatile uint32_t *addr, uint32_t compare_value, uint32_t new_value)

Atomic compare-and-set operation.

Parameters

- **addr** -- Address of atomic variable
- **compare_value** -- Value to compare the atomic variable to
- **new_value** -- New value to set the atomic variable to

Returns Whether the atomic variable was set or not

void **esp_cpu_branch_prediction_enable** (void)

Enable branch prediction.

Structures

struct **esp_cpu_intr_desc_t**

CPU interrupt descriptor.

Each particular CPU interrupt has an associated descriptor describing that particular interrupt's characteristics. Call `esp_cpu_intr_get_desc()` to get the descriptors of a particular interrupt.

Public Members

int **priority**

Priority of the interrupt if it has a fixed priority, (-1) if the priority is configurable.

esp_cpu_intr_type_t **type**

Whether the interrupt is an edge or level type interrupt, `ESP_CPU_INTR_TYPE_NA` if the type is configurable.

uint32_t **flags**

Flags indicating extra details.

Macros

ESP_CPU_INTR_DESC_FLAG_SPECIAL

Interrupt descriptor flags of *esp_cpu_intr_desc_t*.

The interrupt is a special interrupt (e.g., a CPU timer interrupt)

ESP_CPU_INTR_DESC_FLAG_RESVD

The interrupt is reserved for internal use

Type Definitions

typedef uint32_t **esp_cpu_cycle_count_t**

CPU cycle count type.

This data type represents the CPU's clock cycle count

typedef void (***esp_cpu_intr_handler_t**)(void *arg)

CPU interrupt handler type.

Enumerations

enum **esp_cpu_intr_type_t**

CPU interrupt type.

Values:

enumerator **ESP_CPU_INTR_TYPE_LEVEL**

enumerator **ESP_CPU_INTR_TYPE_EDGE**

enumerator **ESP_CPU_INTR_TYPE_NA**

enum **esp_cpu_watchpoint_trigger_t**

CPU watchpoint trigger type.

Values:

enumerator **ESP_CPU_WATCHPOINT_LOAD**

enumerator **ESP_CPU_WATCHPOINT_STORE**

enumerator **ESP_CPU_WATCHPOINT_ACCESS**

Header File

- [components/esp_app_format/include/esp_app_desc.h](#)
- This header file can be included with:

```
#include "esp_app_desc.h"
```

- This header file is a part of the API provided by the `esp_app_format` component. To declare that your component depends on `esp_app_format`, add the following to your `CMakeLists.txt`:

```
REQUIRES esp_app_format
```

or

```
PRIV_REQUIRES esp_app_format
```

Functions

const [esp_app_desc_t](#) ***esp_app_get_description** (void)

Return `esp_app_desc` structure. This structure includes app version.

Return description for running app.

Returns Pointer to `esp_app_desc` structure.

int **esp_app_get_elf_sha256** (char *dst, size_t size)

Fill the provided buffer with SHA256 of the ELF file, formatted as hexadecimal, null-terminated. If the buffer size is not sufficient to fit the entire SHA256 in hex plus a null terminator, the largest possible number of bytes will be written followed by a null.

Parameters

- **dst** -- Destination buffer
- **size** -- Size of the buffer

Returns Number of bytes written to dst (including null terminator)

char ***esp_app_get_elf_sha256_str** (void)

Return SHA256 of the ELF file which is already formatted as hexadecimal, null-terminated included. Can be used in panic handler or core dump during when cache is disabled. The length is defined by `CONFIG_APP_RETRIEVE_LEN_ELF_SHA` option.

Returns Hexadecimal SHA256 string

Structures

struct **esp_app_desc_t**

Description about application.

Public Members

uint32_t **magic_word**

Magic word ESP_APP_DESC_MAGIC_WORD

uint32_t **secure_version**

Secure version

uint32_t **reserv1**[2]

reserv1

char **version**[32]

Application version

char **project_name**[32]

Project name

char **time**[16]

Compile time

char **date**[16]

Compile date

char **idf_ver**[32]

Version IDF

uint8_t **app_elf_sha256**[32]

sha256 of elf file

uint16_t **min_efuse_blk_rev_full**

Minimal eFuse block revision supported by image, in format: major * 100 + minor

uint16_t **max_efuse_blk_rev_full**

Maximal eFuse block revision supported by image, in format: major * 100 + minor

uint32_t **reserv2**[19]

reserv2

Macros

ESP_APP_DESC_MAGIC_WORD

The magic word for the esp_app_desc structure that is in DROM.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.9.24 Over The Air Updates (OTA)

OTA Process Overview

The OTA update mechanism allows a device to update itself based on data received while the normal firmware is running (for example, over Wi-Fi or Bluetooth.)

OTA requires configuring the [Partition Tables](#) of the device with at least two OTA app slot partitions (i.e., `ota_0` and `ota_1`) and an OTA Data Partition.

The OTA operation functions write a new app firmware image to whichever OTA app slot that is currently not selected for booting. Once the image is verified, the OTA Data partition is updated to specify that this image should be used for the next boot.

OTA Data Partition

An OTA data partition (type `data`, subtype `ota`) must be included in the [Partition Tables](#) of any project which uses the OTA functions.

For factory boot settings, the OTA data partition should contain no data (all bytes erased to `0xFF`). In this case, the ESP-IDF software bootloader will boot the factory app if it is present in the partition table. If no factory app is included in the partition table, the first available OTA slot (usually `ota_0`) is booted.

After the first OTA update, the OTA data partition is updated to specify which OTA app slot partition should be booted next.

The OTA data partition is two flash sectors (`0x2000` bytes) in size, to prevent problems if there is a power failure while it is being written. Sectors are independently erased and written with matching data, and if they disagree a counter field is used to determine which sector was written more recently.

App Rollback

The main purpose of the application rollback is to keep the device working after the update. This feature allows you to roll back to the previous working application in case a new application has critical errors. When the rollback process is enabled and an OTA update provides a new version of the app, one of three things can happen:

- The application works fine, [`esp\_ota\_mark\_app\_valid\_cancel\_rollback\(\)`](#) marks the running application with the state `ESP_OTA_IMG_VALID`. There are no restrictions on booting this application.
- The application has critical errors and further work is not possible, a rollback to the previous application is required, [`esp\_ota\_mark\_app\_invalid\_rollback\_and\_reboot\(\)`](#) marks the running application with the state `ESP_OTA_IMG_INVALID` and reset. This application will not be selected by the bootloader for boot and will boot the previously working application.
- If the [`CONFIG\_BOOTLOADER\_APP\_ROLLBACK\_ENABLE`](#) option is set, and a reset occurs without calling either function then the application is rolled back.

Note: The state is not written to the binary image of the application but rather to the `otadata` partition. The partition contains a `ota_seq` counter, which is a pointer to the slot (`ota_0`, `ota_1`, ...) from which the application will be selected for boot.

App OTA State States control the process of selecting a boot app:

States	Restriction of selecting a boot app in bootloader
ESP_OTA_IMG_VALID	Under the restriction. Will be selected.
ESP_OTA_IMG_UNDETERMINED	Under the restriction. Will be selected.
ESP_OTA_IMG_INVALID	Will not be selected.
ESP_OTA_IMG_ABORTED	Will not be selected.
ESP_OTA_IMG_NEW	If CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE option is set it will be selected only once. In bootloader the state immediately changes to ESP_OTA_IMG_PENDING_VERIFY.
ESP_OTA_IMG_PENDING_VERIFY	If CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE option is set it will not be selected, and the state will change to ESP_OTA_IMG_ABORTED.

If [CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE](#) option is not enabled (by default), then the use of the following functions [esp_ota_mark_app_valid_cancel_rollback\(\)](#) and [esp_ota_mark_app_invalid_rollback_and_reboot\(\)](#) are optional, and ESP_OTA_IMG_NEW and ESP_OTA_IMG_PENDING_VERIFY states are not used.

An option in Kconfig [CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE](#) allows you to track the first boot of a new application. In this case, the application must confirm its operability by calling [esp_ota_mark_app_valid_cancel_rollback\(\)](#) function, otherwise the application will be rolled back upon reboot. It allows you to control the operability of the application during the boot phase. Thus, a new application has only one attempt to boot successfully.

Rollback Process The description of the rollback process when [CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE](#) option is enabled:

- The new application is successfully downloaded and [esp_ota_set_boot_partition\(\)](#) function makes this partition bootable and sets the state ESP_OTA_IMG_NEW. This state means that the application is new and should be monitored for its first boot.
- Reboot [esp_restart\(\)](#).
- The bootloader checks for the ESP_OTA_IMG_PENDING_VERIFY state if it is set, then it will be written to ESP_OTA_IMG_ABORTED.
- The bootloader selects a new application to boot so that the state is not set as ESP_OTA_IMG_INVALID or ESP_OTA_IMG_ABORTED.
- The bootloader checks the selected application for ESP_OTA_IMG_NEW state if it is set, then it will be written to ESP_OTA_IMG_PENDING_VERIFY. This state means that the application requires confirmation of its operability, if this does not happen and a reboot occurs, this state will be overwritten to ESP_OTA_IMG_ABORTED (see above) and this application will no longer be able to start, i.e., there will be a rollback to the previous working application.
- A new application has started and should make a self-test.
- If the self-test has completed successfully, then you must call the function [esp_ota_mark_app_valid_cancel_rollback\(\)](#) because the application is awaiting confirmation of operability (ESP_OTA_IMG_PENDING_VERIFY state).
- If the self-test fails, then call [esp_ota_mark_app_invalid_rollback_and_reboot\(\)](#) function to roll back to the previous working application, while the invalid application is set ESP_OTA_IMG_INVALID state.
- If the application has not been confirmed, the state remains ESP_OTA_IMG_PENDING_VERIFY, and the next boot it will be changed to ESP_OTA_IMG_ABORTED, which prevents re-boot of this application. There will be a rollback to the previous working application.

Unexpected Reset If a power loss or an unexpected crash occurs at the time of the first boot of a new application, it will roll back the application.

Recommendation: Perform the self-test procedure as quickly as possible, to prevent rollback due to power loss.

Only OTA partitions can be rolled back. Factory partition is not rolled back.

Booting Invalid/aborted Apps Booting an application which was previously set to `ESP_OTA_IMG_INVALID` or `ESP_OTA_IMG_ABORTED` is possible:

- Get the last invalid application partition `esp_ota_get_last_invalid_partition()`.
- Pass the received partition to `esp_ota_set_boot_partition()`, this will update the otadata.
- Restart `esp_restart()`. The bootloader will boot the specified application.

To determine if self-tests should be run during startup of an application, call the `esp_ota_get_state_partition()` function. If result is `ESP_OTA_IMG_PENDING_VERIFY` then self-testing and subsequent confirmation of operability is required.

Where the States Are Set A brief description of where the states are set:

- `ESP_OTA_IMG_VALID` state is set by `esp_ota_mark_app_valid_cancel_rollback()` function.
- `ESP_OTA_IMG_UNDEFINED` state is set by `esp_ota_set_boot_partition()` function if `CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE` option is not enabled.
- `ESP_OTA_IMG_NEW` state is set by `esp_ota_set_boot_partition()` function if `CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE` option is enabled.
- `ESP_OTA_IMG_INVALID` state is set by `esp_ota_mark_app_invalid_rollback_and_reboot()` function.
- `ESP_OTA_IMG_ABORTED` state is set if there was no confirmation of the application operability and occurs reboots (if `CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE` option is enabled).
- `ESP_OTA_IMG_PENDING_VERIFY` state is set in a bootloader if `CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE` option is enabled and selected app has `ESP_OTA_IMG_NEW` state.

Anti-rollback

Anti-rollback prevents rollback to application with security version lower than one programmed in eFuse of chip.

This function works if set `CONFIG_BOOTLOADER_APP_ANTI_ROLLBACK` option. In the bootloader, when selecting a bootable application, an additional security version check is added which is on the chip and in the application image. The version in the bootable firmware must be greater than or equal to the version in the chip.

`CONFIG_BOOTLOADER_APP_ANTI_ROLLBACK` and `CONFIG_BOOTLOADER_APP_ROLLBACK_ENABLE` options are used together. In this case, rollback is possible only on the security version which is equal or higher than the version in the chip.

A Typical Anti-rollback Scheme Is

- New firmware released with the elimination of vulnerabilities with the previous version of security.
- After the developer makes sure that this firmware is working. He can increase the security version and release a new firmware.
- Download new application.
- To make it bootable, run the function `esp_ota_set_boot_partition()`. If the security version of the new application is smaller than the version in the chip, the new application will be erased. Update to new firmware is not possible.
- Reboot.
- In the bootloader, an application with a security version greater than or equal to the version in the chip will be selected. If otadata is in the initial state, and one firmware was loaded via a serial channel, whose secure version is higher than the chip, then the secure version of efuse will be immediately updated in the bootloader.
- New application booted. Then the application should perform diagnostics of the operation and if it is completed successfully, you should call `esp_ota_mark_app_valid_cancel_rollback()` function to mark the running application with the `ESP_OTA_IMG_VALID` state and update the secure version on chip. Note that if was called `esp_ota_mark_app_invalid_rollback_and_reboot()` function a rollback may not happen as the device may not have any bootable apps. It will then return `ESP_ERR_OTA_ROLLBACK_FAILED` error and stay in the `ESP_OTA_IMG_PENDING_VERIFY` state.
- The next update of app is possible if a running app is in the `ESP_OTA_IMG_VALID` state.

Recommendation:

If you want to avoid the download/erase overhead in case of the app from the server has security version lower than the running app, you have to get `new_app_info.secure_version` from the first package of an image and compare it with the secure version of efuse. Use `esp_efuse_check_secure_version(new_app_info.secure_version)` function if it is true then continue downloading otherwise abort.

```
....
bool image_header_was_checked = false;
while (1) {
    int data_read = esp_http_client_read(client, ota_write_data, BUFFSIZE);
    ...
    if (data_read > 0) {
        if (image_header_was_checked == false) {
            esp_app_desc_t new_app_info;
            if (data_read > sizeof(esp_image_header_t) + sizeof(esp_image_segment_
↪header_t) + sizeof(esp_app_desc_t)) {
                // check current version with downloading
                if (esp_efuse_check_secure_version(new_app_info.secure_version) ==_
↪false) {
                    ESP_LOGE(TAG, "This a new app can not be downloaded due to a_
↪secure version is lower than stored in efuse.");
                    http_cleanup(client);
                    task_fatal_error();
                }

                image_header_was_checked = true;

                esp_ota_begin(update_partition, OTA_SIZE_UNKNOWN, &update_handle);
            }
            esp_ota_write( update_handle, (const void *)ota_write_data, data_read);
        }
    }
}
....
```

Restrictions:

- The number of bits in the `secure_version` field is limited to 16 bits. This means that only 16 times you can do an anti-rollback. You can reduce the length of this efuse field using [CONFIG_BOOTLOADER_APP_SEC_VER_SIZE_EFUSE_FIELD](#) option.
- Factory and Test partitions are not supported in anti rollback scheme and hence partition table should not have partition with SubType set to factory or test.

security_version:

- In application image it is stored in `esp_app_desc` structure. The number is set [CONFIG_BOOTLOADER_APP_SECURE_VERSION](#).

Secure OTA Updates Without Secure Boot

The verification of signed OTA updates can be performed even without enabling hardware secure boot. This can be achieved by setting [CONFIG_SECURE_SIGNED_APPS_NO_SECURE_BOOT](#) and [CONFIG_SECURE_SIGNED_ON_UPDATE_NO_SECURE_BOOT](#)

OTA Tool otatool.py

The component `app_update` provides a tool [app_update/otatool.py](#) for performing OTA partition-related operations on a target device. The following operations can be performed using the tool:

- read contents of otadata partition (read_otadata)
- erase otadata partition, effectively resetting device to factory app (erase_otadata)
- switch OTA partitions (switch_ota_partition)
- erasing OTA partition (erase_ota_partition)
- write to OTA partition (write_ota_partition)
- read contents of OTA partition (read_ota_partition)

The tool can either be imported and used from another Python script or invoked from shell script for users wanting to perform operation programmatically. This is facilitated by the tool's Python API and command-line interface, respectively.

Python API Before anything else, make sure that the `otatool` module is imported.

```
import sys
import os

idf_path = os.environ["IDF_PATH"] # get value of IDF_PATH from environment
otatool_dir = os.path.join(idf_path, "components", "app_update") # otatool.py_
↳ lives in $IDF_PATH/components/app_update

sys.path.append(otatool_dir) # this enables Python to find otatool module
from otatool import * # import all names inside otatool module
```

The starting point for using the tool's Python API to do is create a `OtatoolTarget` object:

```
# Create a partool.py target device connected on serial port /dev/ttyUSB1
target = OtatoolTarget("/dev/ttyUSB1")
```

The created object can now be used to perform operations on the target device:

```
# Erase otadata, resetting the device to factory app
target.erase_otadata()

# Erase contents of OTA app slot 0
target.erase_ota_partition(0)

# Switch boot partition to that of app slot 1
target.switch_ota_partition(1)

# Read OTA partition 'ota_3' and save contents to a file named 'ota_3.bin'
target.read_ota_partition("ota_3", "ota_3.bin")
```

The OTA partition to operate on is specified using either the app slot number or the partition name.

More information on the Python API is available in the docstrings for the tool.

Command-line Interface The command-line interface of `otatool.py` has the following structure:

```
otatool.py [command-args] [subcommand] [subcommand-args]

- command-args - these are arguments that are needed for executing the main_
↳ command (partool.py), mostly pertaining to the target device
- subcommand - this is the operation to be performed
- subcommand-args - these are arguments that are specific to the chosen operation
```

```
# Erase otadata, resetting the device to factory app
otatool.py --port "/dev/ttyUSB1" erase_otadata

# Erase contents of OTA app slot 0
otatool.py --port "/dev/ttyUSB1" erase_ota_partition --slot 0
```

(continues on next page)

(continued from previous page)

```
# Switch boot partition to that of app slot 1
otatool.py --port "/dev/ttyUSB1" switch_ota_partition --slot 1

# Read OTA partition 'ota_3' and save contents to a file named 'ota_3.bin'
otatool.py --port "/dev/ttyUSB1" read_ota_partition --name=ota_3 --output=ota_3.bin
```

More information can be obtained by specifying `--help` as argument:

```
# Display possible subcommands and show main command argument descriptions
otatool.py --help

# Show descriptions for specific subcommand arguments
otatool.py [subcommand] --help
```

See Also

- [Partition Tables](#)
- [Partitions API](#)
- [SPI Flash API](#)
- [ESP HTTPS OTA](#)

Application Example

End-to-end example of OTA firmware update workflow: [system/ota](#).

API Reference

Header File

- [components/app_update/include/esp_ota_ops.h](#)
- This header file can be included with:

```
#include "esp_ota_ops.h"
```

- This header file is a part of the API provided by the `app_update` component. To declare that your component depends on `app_update`, add the following to your `CMakeLists.txt`:

```
REQUIRES app_update
```

or

```
PRIV_REQUIRES app_update
```

Functions

const [esp_app_desc_t](#) ***esp_ota_get_app_description** (void)

Return `esp_app_desc` structure. This structure includes app version.

Return description for running app.

Note: This API is present for backward compatibility reasons. Alternative function with the same functionality is `esp_app_get_description`

Returns Pointer to `esp_app_desc` structure.

int esp_ota_get_app_elf_sha256 (char *dst, size_t size)

Fill the provided buffer with SHA256 of the ELF file, formatted as hexadecimal, null-terminated. If the buffer size is not sufficient to fit the entire SHA256 in hex plus a null terminator, the largest possible number of bytes will be written followed by a null.

Note: This API is present for backward compatibility reasons. Alternative function with the same functionality is `esp_app_get_elf_sha256`

Parameters

- **dst** -- Destination buffer
- **size** -- Size of the buffer

Returns Number of bytes written to dst (including null terminator)

esp_err_t **esp_ota_begin** (const *esp_partition_t* *partition, size_t image_size, *esp_ota_handle_t* *out_handle)

Commence an OTA update writing to the specified partition.

The specified partition is erased to the specified image size.

If image size is not yet known, pass `OTA_SIZE_UNKNOWN` which will cause the entire partition to be erased.

On success, this function allocates memory that remains in use until `esp_ota_end()` is called with the returned handle.

Note: If the rollback option is enabled and the running application has the `ESP_OTA_IMG_PENDING_VERIFY` state then it will lead to the `ESP_ERR_OTA_ROLLBACK_INVALID_STATE` error. Confirm the running app before to run download a new app, use `esp_ota_mark_app_valid_cancel_rollback()` function for it (this should be done as early as possible when you first download a new application).

Parameters

- **partition** -- Pointer to info for partition which will receive the OTA update. Required.
- **image_size** -- Size of new OTA app image. Partition will be erased in order to receive this size of image. If 0 or `OTA_SIZE_UNKNOWN`, the entire partition is erased.
- **out_handle** -- On success, returns a handle which should be used for subsequent `esp_ota_write()` and `esp_ota_end()` calls.

Returns

- `ESP_OK`: OTA operation commenced successfully.
- `ESP_ERR_INVALID_ARG`: partition or out_handle arguments were NULL, or partition doesn't point to an OTA app partition.
- `ESP_ERR_NO_MEM`: Cannot allocate memory for OTA operation.
- `ESP_ERR_OTA_PARTITION_CONFLICT`: Partition holds the currently running firmware, cannot update in place.
- `ESP_ERR_NOT_FOUND`: Partition argument not found in partition table.
- `ESP_ERR_OTA_SELECT_INFO_INVALID`: The OTA data partition contains invalid data.
- `ESP_ERR_INVALID_SIZE`: Partition doesn't fit in configured flash size.
- `ESP_ERR_FLASH_OP_TIMEOUT` or `ESP_ERR_FLASH_OP_FAIL`: Flash write failed.
- `ESP_ERR_OTA_ROLLBACK_INVALID_STATE`: If the running app has not confirmed state. Before performing an update, the application must be valid.

esp_err_t **esp_ota_write** (*esp_ota_handle_t* handle, const void *data, size_t size)

Write OTA update data to partition.

This function can be called multiple times as data is received during the OTA operation. Data is written sequentially to the partition.

Parameters

- **handle** -- Handle obtained from `esp_ota_begin`
- **data** -- Data buffer to write

- **size** -- Size of data buffer in bytes.

Returns

- ESP_OK: Data was written to flash successfully, or size = 0
- ESP_ERR_INVALID_ARG: handle is invalid.
- ESP_ERR_OTA_VALIDATE_FAILED: First byte of image contains invalid app image magic byte.
- ESP_ERR_FLASH_OP_TIMEOUT or ESP_ERR_FLASH_OP_FAIL: Flash write failed.
- ESP_ERR_OTA_SELECT_INFO_INVALID: OTA data partition has invalid contents

esp_err_t **esp_ota_write_with_offset** (*esp_ota_handle_t* handle, const void *data, size_t size, uint32_t offset)

Write OTA update data to partition at an offset.

This function can write data in non-contiguous manner. If flash encryption is enabled, data should be 16 bytes aligned.

Note: While performing OTA, if the packets arrive out of order, `esp_ota_write_with_offset()` can be used to write data in non-contiguous manner. Use of `esp_ota_write_with_offset()` in combination with `esp_ota_write()` is not recommended.

Parameters

- **handle** -- Handle obtained from `esp_ota_begin`
- **data** -- Data buffer to write
- **size** -- Size of data buffer in bytes
- **offset** -- Offset in flash partition

Returns

- ESP_OK: Data was written to flash successfully.
- ESP_ERR_INVALID_ARG: handle is invalid.
- ESP_ERR_OTA_VALIDATE_FAILED: First byte of image contains invalid app image magic byte.
- ESP_ERR_FLASH_OP_TIMEOUT or ESP_ERR_FLASH_OP_FAIL: Flash write failed.
- ESP_ERR_OTA_SELECT_INFO_INVALID: OTA data partition has invalid contents

esp_err_t **esp_ota_end** (*esp_ota_handle_t* handle)

Finish OTA update and validate newly written app image.

Note: After calling `esp_ota_end()`, the handle is no longer valid and any memory associated with it is freed (regardless of result).

Parameters **handle** -- Handle obtained from `esp_ota_begin()`.

Returns

- ESP_OK: Newly written OTA app image is valid.
- ESP_ERR_NOT_FOUND: OTA handle was not found.
- ESP_ERR_INVALID_ARG: Handle was never written to.
- ESP_ERR_OTA_VALIDATE_FAILED: OTA image is invalid (either not a valid app image, or - if secure boot is enabled - signature failed to verify.)
- ESP_ERR_INVALID_STATE: If flash encryption is enabled, this result indicates an internal error writing the final encrypted bytes to flash.

esp_err_t **esp_ota_abort** (*esp_ota_handle_t* handle)

Abort OTA update, free the handle and memory associated with it.

Parameters **handle** -- obtained from `esp_ota_begin()`.

Returns

- ESP_OK: Handle and its associated memory is freed successfully.
- ESP_ERR_NOT_FOUND: OTA handle was not found.

esp_err_t **esp_ota_set_boot_partition** (const *esp_partition_t* *partition)

Configure OTA data for a new boot partition.

Note: If this function returns ESP_OK, calling esp_restart() will boot the newly configured app partition.

Parameters **partition** -- Pointer to info for partition containing app image to boot.

Returns

- ESP_OK: OTA data updated, next reboot will use specified partition.
- ESP_ERR_INVALID_ARG: partition argument was NULL or didn't point to a valid OTA partition of type "app".
- ESP_ERR_OTA_VALIDATE_FAILED: Partition contained invalid app image. Also returned if secure boot is enabled and signature validation failed.
- ESP_ERR_NOT_FOUND: OTA data partition not found.
- ESP_ERR_FLASH_OP_TIMEOUT or ESP_ERR_FLASH_OP_FAIL: Flash erase or write failed.

const *esp_partition_t* ***esp_ota_get_boot_partition** (void)

Get partition info of currently configured boot app.

If esp_ota_set_boot_partition() has been called, the partition which was set by that function will be returned.

If esp_ota_set_boot_partition() has not been called, the result is usually the same as esp_ota_get_running_partition(). The two results are not equal if the configured boot partition does not contain a valid app (meaning that the running partition will be an app that the bootloader chose via fallback).

If the OTA data partition is not present or not valid then the result is the first app partition found in the partition table. In priority order, this means: the factory app, the first OTA app slot, or the test app partition.

Note that there is no guarantee the returned partition is a valid app. Use esp_image_verify(ESP_IMAGE_VERIFY, ...) to verify if the returned partition contains a bootable image.

Returns Pointer to info for partition structure, or NULL if partition table is invalid or a flash read operation failed. Any returned pointer is valid for the lifetime of the application.

const *esp_partition_t* ***esp_ota_get_running_partition** (void)

Get partition info of currently running app.

This function is different to esp_ota_get_boot_partition() in that it ignores any change of selected boot partition caused by esp_ota_set_boot_partition(). Only the app whose code is currently running will have its partition information returned.

The partition returned by this function may also differ from esp_ota_get_boot_partition() if the configured boot partition is somehow invalid, and the bootloader fell back to a different app partition at boot.

Returns Pointer to info for partition structure, or NULL if no partition is found or flash read operation failed. Returned pointer is valid for the lifetime of the application.

const *esp_partition_t* ***esp_ota_get_next_update_partition** (const *esp_partition_t* *start_from)

Return the next OTA app partition which should be written with a new firmware.

Call this function to find an OTA app partition which can be passed to esp_ota_begin().

Finds next partition round-robin, starting from the current running partition.

Parameters **start_from** -- If set, treat this partition info as describing the current running partition. Can be NULL, in which case esp_ota_get_running_partition() is used to find the currently running partition. The result of this function is never the same as this argument.

Returns Pointer to info for partition which should be updated next. NULL result indicates invalid OTA data partition, or that no eligible OTA app slot partition was found.

esp_err_t **esp_ota_get_partition_description** (const *esp_partition_t* *partition, *esp_app_desc_t* *app_desc)

Returns esp_app_desc structure for app partition. This structure includes app version.

Returns a description for the requested app partition.

Parameters

- **partition** -- [in] Pointer to app partition. (only app partition)
- **app_desc** -- [out] Structure of info about app.

Returns

- ESP_OK Successful.
- ESP_ERR_NOT_FOUND app_desc structure is not found. Magic word is incorrect.
- ESP_ERR_NOT_SUPPORTED Partition is not application.
- ESP_ERR_INVALID_ARG Arguments is NULL or if partition's offset exceeds partition size.
- ESP_ERR_INVALID_SIZE Read would go out of bounds of the partition.
- or one of error codes from lower-level flash driver.

esp_err_t **esp_ota_get_bootloader_description** (const *esp_partition_t* *bootloader_partition, *esp_bootloader_desc_t* *desc)

Returns the description structure of the bootloader.

Parameters

- **bootloader_partition** -- [in] Pointer to bootloader partition. If NULL, then the current bootloader is used (the default location).
offset = CONFIG_BOOTLOADER_OFFSET_IN_FLASH,
size = CONFIG_PARTITION_TABLE_OFFSET - CONFIG_BOOTLOADER_OFFSET_IN_FLASH,
- **desc** -- [out] Structure of info about bootloader.

Returns

- ESP_OK Successful.
- ESP_ERR_NOT_FOUND Description structure is not found in the bootloader image. Magic byte is incorrect.
- ESP_ERR_INVALID_ARG Arguments is NULL.
- ESP_ERR_INVALID_SIZE Read would go out of bounds of the partition.
- or one of error codes from lower-level flash driver.

uint8_t **esp_ota_get_app_partition_count** (void)

Returns number of ota partitions provided in partition table.

Returns

- Number of OTA partitions

esp_err_t **esp_ota_mark_app_valid_cancel_rollback** (void)

This function is called to indicate that the running app is working well.

Returns

- ESP_OK: if successful.

esp_err_t **esp_ota_mark_app_invalid_rollback_and_reboot** (void)

This function is called to roll back to the previously workable app with reboot.

If rollback is successful then device will reset else API will return with error code. Checks applications on a flash drive that can be booted in case of rollback. If the flash does not have at least one app (except the running app) then rollback is not possible.

Returns

- ESP_FAIL: if not successful.
- ESP_ERR_OTA_ROLLBACK_FAILED: The rollback is not possible due to flash does not have any apps.

const *esp_partition_t* ***esp_ota_get_last_invalid_partition** (void)

Returns last partition with invalid state (ESP_OTA_IMG_INVALID or ESP_OTA_IMG_ABORTED).

Returns partition.

esp_err_t **esp_ota_get_state_partition** (const *esp_partition_t* *partition, *esp_ota_img_states_t* *ota_state)

Returns state for given partition.

Parameters

- **partition** -- [in] Pointer to partition.
- **ota_state** -- [out] state of partition (if this partition has a record in otadata).

Returns

- ESP_OK: Successful.
- ESP_ERR_INVALID_ARG: partition or ota_state arguments were NULL.
- ESP_ERR_NOT_SUPPORTED: partition is not ota.
- ESP_ERR_NOT_FOUND: Partition table does not have otadata or state was not found for given partition.

esp_err_t **esp_ota_erase_last_boot_app_partition** (void)

Erase previous boot app partition and corresponding otadata select for this partition.

When current app is marked to as valid then you can erase previous app partition.

Returns

- ESP_OK: Successful, otherwise ESP_ERR.

bool **esp_ota_check_rollback_is_possible** (void)

Checks applications on the slots which can be booted in case of rollback.

These applications should be valid (marked in otadata as not UNDEFINED, INVALID or ABORTED and crc is good) and be able booted, and secure_version of app >= secure_version of efuse (if anti-rollback is enabled).

Returns

- True: Returns true if the slots have at least one app (except the running app).
- False: The rollback is not possible.

esp_err_t **esp_ota_revoke_secure_boot_public_key** (*esp_ota_secure_boot_public_key_index_t* index)

Revokes the signature digest denoted by the given index. This should be called in the application only after the rollback logic otherwise the device may end up in unrecoverable state.

Relevant for Secure boot v2 on ESP32-S2, ESP32-S3, ESP32-C3, ESP32-C6, ESP32-H2 where up to 3 key digests can be stored (Key #N-1, Key #N, Key #N+1). When a key used to sign an app is invalidated, an OTA update is to be sent with an app signed with at least one of the other two keys which has not been revoked already. After successfully booting the OTA app should call this function to revoke Key #N-1.

Parameters **index** -- - The index of the signature block to be revoked

Returns

- ESP_OK: If revocation is successful.
- ESP_ERR_INVALID_ARG: If the index of the public key to be revoked is incorrect.
- ESP_FAIL: If secure boot v2 has not been enabled.

Macros

OTA_SIZE_UNKNOWN

Used for esp_ota_begin() if new image size is unknown

OTA_WITH_SEQUENTIAL_WRITES

Used for esp_ota_begin() if new image size is unknown and erase can be done in incremental manner (assuming write operation is in continuous sequence)

ESP_ERR_OTA_BASE

Base error code for ota_ops api

ESP_ERR_OTA_PARTITION_CONFLICT

Error if request was to write or erase the current running partition

ESP_ERR_OTA_SELECT_INFO_INVALID

Error if OTA data partition contains invalid content

ESP_ERR_OTA_VALIDATE_FAILED

Error if OTA app image is invalid

ESP_ERR_OTA_SMALL_SEC_VER

Error if the firmware has a secure version less than the running firmware.

ESP_ERR_OTA_ROLLBACK_FAILED

Error if flash does not have valid firmware in passive partition and hence rollback is not possible

ESP_ERR_OTA_ROLLBACK_INVALID_STATE

Error if current active firmware is still marked in pending validation state (ESP_OTA_IMG_PENDING_VERIFY), essentially first boot of firmware image post upgrade and hence firmware upgrade is not possible

Type Definitions

```
typedef uint32_t esp_ota_handle_t
```

Opaque handle for an application OTA update.

esp_ota_begin() returns a handle which is then used for subsequent calls to esp_ota_write() and esp_ota_end().

Enumerations

```
enum esp_ota_secure_boot_public_key_index_t
```

Secure Boot V2 public key indexes.

Values:

```
enumerator SECURE_BOOT_PUBLIC_KEY_INDEX_0
```

Points to the 0th index of the Secure Boot v2 public key

```
enumerator SECURE_BOOT_PUBLIC_KEY_INDEX_1
```

Points to the 1st index of the Secure Boot v2 public key

```
enumerator SECURE_BOOT_PUBLIC_KEY_INDEX_2
```

Points to the 2nd index of the Secure Boot v2 public key

Debugging OTA Failure




Fig. 32: How to Debug When OTA Fails (click to enlarge)

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.9.25 Power Management

Overview

Power management algorithm included in ESP-IDF can adjust the advanced peripheral bus (APB) frequency, CPU frequency, and put the chip into Light-sleep mode to run an application at smallest possible power consumption, given the requirements of application components.

Application components can express their requirements by creating and acquiring power management locks.

For example:

- Driver for a peripheral clocked from APB can request the APB frequency to be set to 80 MHz while the peripheral is used.
- RTOS can request the CPU to run at the highest configured frequency while there are tasks ready to run.
- A peripheral driver may need interrupts to be enabled, which means it has to request disabling Light-sleep.

Since requesting higher APB or CPU frequencies or disabling Light-sleep causes higher current consumption, please keep the usage of power management locks by components to a minimum.

Configuration

Power management can be enabled at compile time, using the option `CONFIG_PM_ENABLE`.

Enabling power management features comes at the cost of increased interrupt latency. Extra latency depends on a number of factors, such as the CPU frequency, single/dual core mode, whether or not frequency switch needs to be done. Minimum extra latency is 0.2 us (when the CPU frequency is 240 MHz and frequency scaling is not enabled). Maximum extra latency is 40 us (when frequency scaling is enabled, and a switch from 40 MHz to 80 MHz is performed on interrupt entry).

Dynamic frequency scaling (DFS) and automatic Light-sleep can be enabled in an application by calling the function `esp_pm_configure()`. Its argument is a structure defining the frequency scaling settings, `esp_pm_config_t`. In this structure, three fields need to be initialized:

- `max_freq_mhz`: Maximum CPU frequency in MHz, i.e., the frequency used when the `ESP_PM_CPU_FREQ_MAX` lock is acquired. This field is usually set to the default CPU frequency.
- `min_freq_mhz`: Minimum CPU frequency in MHz, i.e., the frequency used when only the `ESP_PM_APB_FREQ_MAX` lock is acquired. This field can be set to the XTAL frequency value, or the XTAL frequency divided by an integer. Note that 10 MHz is the lowest frequency at which the default REF_TICK clock of 1 MHz can be generated.
- `light_sleep_enable`: Whether the system should automatically enter Light-sleep when no locks are acquired (`true/false`).
Alternatively, if you enable the option `CONFIG_PM_DFS_INIT_AUTO` in menuconfig, the maximum CPU frequency will be determined by the `CONFIG_ESP_DEFAULT_CPU_FREQ_MHZ` setting, and the minimum CPU frequency will be locked to the XTAL frequency.

Note: Automatic Light-sleep is based on FreeRTOS Tickless Idle functionality. If automatic Light-sleep is requested while the option `CONFIG_FREERTOS_USE_TICKLESS_IDLE` is not enabled in menuconfig, `esp_pm_configure()` will return the error `ESP_ERR_NOT_SUPPORTED`.

Note: In Light-sleep, peripherals are clock gated, and interrupts (from GPIOs and internal peripherals) will not be generated. A wakeup source described in the [Sleep Modes](#) documentation can be used to trigger wakeup from the Light-sleep state.

Power Management Locks

Applications have the ability to acquire/release locks in order to control the power management algorithm. When an application acquires a lock, the power management algorithm operation is restricted in a way described below. When the lock is released, such restrictions are removed.

Power management locks have acquire/release counters. If the lock has been acquired a number of times, it needs to be released the same number of times to remove associated restrictions.

ESP32-P4 supports three types of locks described in the table below.

Lock	Description
ESP_PM_CPU_FREQ_MAX	Requests CPU frequency to be at the maximum value set with esp_pm_configure() . For ESP32-P4, this value can be set to Not updated yet.
ESP_PM_APB_FREQ_MAX	Requests the APB frequency to be at the maximum supported value. For ESP32-P4, this is 80 MHz.
ESP_PM_NO_LIGHT_SLEEP	Disables automatic switching to Light-sleep.

ESP32-P4 Power Management Algorithm

The table below shows how CPU and APB frequencies will be switched if dynamic frequency scaling is enabled. You can specify the maximum CPU frequency with either [esp_pm_configure\(\)](#) or [CONFIG_ESP_DEFAULT_CPU_FREQ_MHZ](#).

TO BE UPDATED IDF-7672

If none of the locks are acquired, and Light-sleep is enabled in a call to [esp_pm_configure\(\)](#), the system will go into Light-sleep mode. The duration of Light-sleep will be determined by:

- FreeRTOS tasks blocked with finite timeouts
- Timers registered with [High resolution timer](#) APIs

Light-sleep duration is chosen to wake up the chip before the nearest event (task being unblocked, or timer elapses).

To skip unnecessary wake-up, you can consider initializing an `esp_timer` with the `skip_unhandled_events` option as `true`. Timers with this flag will not wake up the system and it helps to reduce consumption.

Dynamic Frequency Scaling and Peripheral Drivers

When DFS is enabled, the APB frequency can be changed multiple times within a single RTOS tick. The APB frequency change does not affect the operation of some peripherals, while other peripherals may have issues. For example, Timer Group peripheral timers keeps counting, however, the speed at which they count changes proportionally to the APB frequency.

Peripheral clock sources such as `REF_TICK`, `XTAL`, `RC_FAST` (i.e., `RTC_8M`), their frequencies will not be influenced by APB frequency. And therefore, to ensure the peripheral behaves consistently during DFS, it is recommended to select one of these clocks as the peripheral clock source. For more specific guidelines, please refer to the "Power Management" section of each peripheral's "API Reference > Peripherals API" page.

Currently, the following peripheral drivers are aware of DFS and use the `ESP_PM_APB_FREQ_MAX` lock for the duration of the transaction:

- SPI master

- I2C
- I2S (If the APLL clock is used, then it will use the ESP_PM_NO_LIGHT_SLEEP lock)
- SDMMC

The following drivers hold the ESP_PM_APB_FREQ_MAX lock while the driver is enabled:

- **SPI slave:** between calls to `spi_slave_initialize()` and `spi_slave_free()`.
- **GPTimer:** between calls to `gptimer_enable()` and `gptimer_disable()`.
- **Ethernet:** between calls to `esp_eth_driver_install()` and `esp_eth_driver_uninstall()`.
- **WiFi:** between calls to `esp_wifi_start()` and `esp_wifi_stop()`. If modem sleep is enabled, the lock will be released for the periods of time when radio is disabled.
- **PCNT:** between calls to `pcnt_unit_enable()` and `pcnt_unit_disable()`.
- **MCPWM:** between calls to `mcpwm_timer_enable()` and `mcpwm_timer_disable()`, as well as `mcpwm_capture_timer_enable()` and `mcpwm_capture_timer_disable()`.

The following peripheral drivers are not aware of DFS yet. Applications need to acquire/release locks themselves, when necessary:

- The legacy PCNT driver
- The legacy timer group driver
- The legacy MCPWM driver

Light-sleep Peripheral Power Down

API Reference

Header File

- `components/esp_pm/include/esp_pm.h`
- This header file can be included with:

```
#include "esp_pm.h"
```

- This header file is a part of the API provided by the `esp_pm` component. To declare that your component depends on `esp_pm`, add the following to your CMakeLists.txt:

```
REQUIRES esp_pm
```

or

```
PRIV_REQUIRES esp_pm
```

Functions

`esp_err_t esp_pm_configure` (const void *config)

Set implementation-specific power management configuration.

Parameters `config` -- pointer to implementation-specific configuration structure (e.g. `esp_pm_config_esp32`)

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if the configuration values are not correct
- ESP_ERR_NOT_SUPPORTED if certain combination of values is not supported, or if CONFIG_PM_ENABLE is not enabled in sdkconfig

`esp_err_t esp_pm_get_configuration` (void *config)

Get implementation-specific power management configuration.

Parameters `config` -- pointer to implementation-specific configuration structure (e.g. `esp_pm_config_esp32`)

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if the pointer is null

esp_err_t **esp_pm_lock_create** (*esp_pm_lock_type_t* lock_type, int arg, const char *name, *esp_pm_lock_handle_t* *out_handle)

Initialize a lock handle for certain power management parameter.

When lock is created, initially it is not taken. Call esp_pm_lock_acquire to take the lock.

This function must not be called from an ISR.

Parameters

- **lock_type** -- Power management constraint which the lock should control
- **arg** -- argument, value depends on lock_type, see esp_pm_lock_type_t
- **name** -- arbitrary string identifying the lock (e.g. "wifi" or "spi"). Used by the esp_pm_dump_locks function to list existing locks. May be set to NULL. If not set to NULL, must point to a string which is valid for the lifetime of the lock.
- **out_handle** -- [out] handle returned from this function. Use this handle when calling esp_pm_lock_delete, esp_pm_lock_acquire, esp_pm_lock_release. Must not be NULL.

Returns

- ESP_OK on success
- ESP_ERR_NO_MEM if the lock structure can not be allocated
- ESP_ERR_INVALID_ARG if out_handle is NULL or type argument is not valid
- ESP_ERR_NOT_SUPPORTED if CONFIG_PM_ENABLE is not enabled in sdkconfig

esp_err_t **esp_pm_lock_acquire** (*esp_pm_lock_handle_t* handle)

Take a power management lock.

Once the lock is taken, power management algorithm will not switch to the mode specified in a call to esp_pm_lock_create, or any of the lower power modes (higher numeric values of 'mode').

The lock is recursive, in the sense that if esp_pm_lock_acquire is called a number of times, esp_pm_lock_release has to be called the same number of times in order to release the lock.

This function may be called from an ISR.

This function is not thread-safe w.r.t. calls to other esp_pm_lock_* functions for the same handle.

Parameters **handle** -- handle obtained from esp_pm_lock_create function

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if the handle is invalid
- ESP_ERR_NOT_SUPPORTED if CONFIG_PM_ENABLE is not enabled in sdkconfig

esp_err_t **esp_pm_lock_release** (*esp_pm_lock_handle_t* handle)

Release the lock taken using esp_pm_lock_acquire.

Call to this functions removes power management restrictions placed when taking the lock.

Locks are recursive, so if esp_pm_lock_acquire is called a number of times, esp_pm_lock_release has to be called the same number of times in order to actually release the lock.

This function may be called from an ISR.

This function is not thread-safe w.r.t. calls to other esp_pm_lock_* functions for the same handle.

Parameters **handle** -- handle obtained from esp_pm_lock_create function

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if the handle is invalid
- ESP_ERR_INVALID_STATE if lock is not acquired
- ESP_ERR_NOT_SUPPORTED if CONFIG_PM_ENABLE is not enabled in sdkconfig

esp_err_t **esp_pm_lock_delete** (*esp_pm_lock_handle_t* handle)

Delete a lock created using `esp_pm_lock`.

The lock must be released before calling this function.

This function must not be called from an ISR.

Parameters **handle** -- handle obtained from `esp_pm_lock_create` function

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if the handle argument is NULL
- ESP_ERR_INVALID_STATE if the lock is still acquired
- ESP_ERR_NOT_SUPPORTED if CONFIG_PM_ENABLE is not enabled in sdkconfig

esp_err_t **esp_pm_dump_locks** (FILE *stream)

Dump the list of all locks to stderr

This function dumps debugging information about locks created using `esp_pm_lock_create` to an output stream.

This function must not be called from an ISR. If `esp_pm_lock_acquire/release` are called while this function is running, inconsistent results may be reported.

Parameters **stream** -- stream to print information to; use stdout or stderr to print to the console; use `fmemopen/open_memstream` to print to a string buffer.

Returns

- ESP_OK on success
- ESP_ERR_NOT_SUPPORTED if CONFIG_PM_ENABLE is not enabled in sdkconfig

Structures

struct **esp_pm_config_t**

Power management config.

Pass a pointer to this structure as an argument to `esp_pm_configure` function.

Public Members

int **max_freq_mhz**

Maximum CPU frequency, in MHz

int **min_freq_mhz**

Minimum CPU frequency to use when no locks are taken, in MHz

bool **light_sleep_enable**

Enter light sleep when no locks are taken

Type Definitions

typedef *esp_pm_config_t* **esp_pm_config_esp32_t**

backward compatibility newer chips no longer require this typedef

typedef *esp_pm_config_t* **esp_pm_config_esp32s2_t**

typedef *esp_pm_config_t* **esp_pm_config_esp32s3_t**

typedef *esp_pm_config_t* **esp_pm_config_esp32c3_t**

```
typedef esp_pm_config_t esp_pm_config_esp32c2_t
```

```
typedef esp_pm_config_t esp_pm_config_esp32c6_t
```

```
typedef struct esp_pm_lock *esp_pm_lock_handle_t
```

Opaque handle to the power management lock.

Enumerations

```
enum esp_pm_lock_type_t
```

Power management constraints.

Values:

```
enumerator ESP_PM_CPU_FREQ_MAX
```

Require CPU frequency to be at the maximum value set via `esp_pm_configure`. Argument is unused and should be set to 0.

```
enumerator ESP_PM_APB_FREQ_MAX
```

Require APB frequency to be at the maximum value supported by the chip. Argument is unused and should be set to 0.

```
enumerator ESP_PM_NO_LIGHT_SLEEP
```

Prevent the system from going into light sleep. Argument is unused and should be set to 0.

2.9.26 POSIX Threads Support

Overview

ESP-IDF is based on FreeRTOS but offers a range of POSIX-compatible APIs that allow easy porting of third-party code. This includes support for common parts of the POSIX Threads `pthread` API.

POSIX Threads are implemented in ESP-IDF as wrappers around equivalent FreeRTOS features. The runtime memory or performance overhead of using the `pthread` API is quite low, but not every feature available in either `pthread` or FreeRTOS is available via the ESP-IDF `pthread` support.

`Pthreads` can be used in ESP-IDF by including standard `pthread.h` header, which is included in the toolchain `libc`. An additional ESP-IDF specific header, `esp_pthread.h`, provides additional non-POSIX APIs for using some ESP-IDF features with `pthread`.

C++ Standard Library implementations for `std::thread`, `std::mutex`, `std::condition_variable`, etc., are realized using `pthread` (via GCC `libstdc++`). Therefore, restrictions mentioned here also apply to the equivalent C++ standard library functionality.

RTOS Integration

Unlike many operating systems using POSIX Threads, ESP-IDF is a real-time operating system with a real-time scheduler. This means that a thread will only stop running if a higher priority task is ready to run, the thread blocks on an OS synchronization structure like a `mutex`, or the thread calls any of the functions `sleep`, `vTaskDelay()`, or `usleep`.

Note: When calling a standard libc or C++ sleep function, such as `usleep` defined in `unistd.h`, the task will only block and yield the core if the sleep time is longer than *one FreeRTOS tick period*. If the time is shorter, the thread will busy-wait instead of yielding to another RTOS task.

By default, all POSIX Threads have the same RTOS priority, but it is possible to change this by calling a *custom API*.

Standard Features

The following standard APIs are implemented in ESP-IDF.

Refer to [standard POSIX Threads documentation](#), or `pthread.h`, for details about the standard arguments and behaviour of each function. Differences or limitations compared to the standard APIs are noted below.

Thread APIs

- **`pthread_create()`**
 - The `attr` argument is supported for setting stack size and detach state only. Other attribute fields are ignored.
 - Unlike FreeRTOS task functions, the `start_routine` function is allowed to return. A detached type thread is automatically deleted if the function returns. The default joinable type thread will be suspended until `pthread_join()` is called on it.
- `pthread_join()`
- `pthread_detach()`
- `pthread_exit()`
- `sched_yield()`
- **`pthread_self()`**
 - An assert will fail if this function is called from a FreeRTOS task which is not a pthread.
- `pthread_equal()`

Thread Attributes

- `pthread_attr_init()`
- **`pthread_attr_destroy()`**
 - This function does not need to free any resources and instead resets the `attr` structure to defaults. The implementation is the same as `pthread_attr_init()`.
- `pthread_attr_getstacksize() / pthread_attr_setstacksize()`
- `pthread_attr_getdetachstate() / pthread_attr_setdetachstate()`

Once

- `pthread_once()`

Static initializer constant `PTHREAD_ONCE_INIT` is supported.

Note: This function can be called from tasks created using either pthread or FreeRTOS APIs.

Mutexes POSIX Mutexes are implemented as FreeRTOS Mutex Semaphores (normal type for "fast" or "error check" mutexes, and Recursive type for "recursive" mutexes). This means that they have the same priority inheritance behavior as mutexes created with `xSemaphoreCreateMutex()`.

- `pthread_mutex_init()`
- `pthread_mutex_destroy()`
- `pthread_mutex_lock()`
- `pthread_mutex_timedlock()`

- `pthread_mutex_trylock()`
- `pthread_mutex_unlock()`
- `pthread_mutexattr_init()`
- `pthread_mutexattr_destroy()`
- `pthread_mutexattr_gettype()` / `pthread_mutexattr_settype()`

Static initializer constant `PTHREAD_MUTEX_INITIALIZER` is supported, but the non-standard static initializer constants for other mutex types are not supported.

Note: These functions can be called from tasks created using either pthread or FreeRTOS APIs.

Condition Variables

- `pthread_cond_init()`
 - The `attr` argument is not implemented and is ignored.
- `pthread_cond_destroy()`
- `pthread_cond_signal()`
- `pthread_cond_broadcast()`
- `pthread_cond_wait()`
- `pthread_cond_timedwait()`

Static initializer constant `PTHREAD_COND_INITIALIZER` is supported.

- The resolution of `pthread_cond_timedwait()` timeouts is the RTOS tick period (see [CONFIG_FREERTOS_HZ](#)). Timeouts may be delayed up to one tick period after the requested timeout.

Note: These functions can be called from tasks created using either pthread or FreeRTOS APIs.

Semaphores In ESP-IDF, POSIX **unnamed** semaphores are implemented. The accessible API is described below. It implements [semaphores as specified in the POSIX standard](#), unless specified otherwise.

- `sem_init()`
- `sem_destroy()`
 - `pshared` is ignored. Semaphores can always be shared between FreeRTOS tasks.
- `sem_post()`
 - If the semaphore has a value of `SEM_VALUE_MAX` already, `-1` is returned and `errno` is set to `EAGAIN`.
- `sem_wait()`
- `sem_trywait()`
- `sem_timedwait()`
 - The time value passed by `abstime` will be rounded up to the next FreeRTOS tick.
 - The actual timeout happens after the tick that the time was rounded to and before the following tick.
 - It is possible, though unlikely, that the task is preempted directly after the timeout calculation, delaying the timeout of the following blocking operating system call by the duration of the preemption.
- `sem_getvalue()`

Read/Write Locks The following API functions of the POSIX reader-writer locks specification are implemented:

- `pthread_rwlock_init()`
 - The `attr` argument is not implemented and is ignored.
- `pthread_rwlock_destroy()`
- `pthread_rwlock_rdlock()`
- `pthread_rwlock_tryrdlock()`
- `pthread_rwlock_wrlock()`
- `pthread_rwlock_trywrlock()`
- `pthread_rwlock_unlock()`

The static initializer constant `PTHREAD_RWLOCK_INITIALIZER` is supported.

Note: These functions can be called from tasks created using either pthread or FreeRTOS APIs.

Thread-Specific Data

- `pthread_key_create()`
 - The `destr_function` argument is supported and will be called if a thread function exits normally, calls `pthread_exit()`, or if the underlying task is deleted directly using the FreeRTOS function `vTaskDelete()`.
- `pthread_key_delete()`
- `pthread_setspecific()` / `pthread_getspecific()`

Note: These functions can be called from tasks created using either pthread or FreeRTOS APIs. When calling these functions from tasks created using FreeRTOS APIs, [CONFIG_FREERTOS_TLSP_DELETION_CALLBACKS](#) config option must be enabled to ensure the thread-specific data is cleaned up before the task is deleted.

Note: There are other options for thread local storage in ESP-IDF, including options with higher performance. See [Thread Local Storage](#).

Not Implemented

The `pthread.h` header is a standard header and includes additional APIs and features which are not implemented in ESP-IDF. These include:

- `pthread_cancel()` returns `ENOSYS` if called.
- `pthread_condattr_init()` returns `ENOSYS` if called.

Other POSIX Threads functions (not listed here) are not implemented and will produce either a compiler or a linker error if referenced from an ESP-IDF application. If you identify a useful API that you would like to see implemented in ESP-IDF, please open a [feature request on GitHub](#) with the details.

ESP-IDF Extensions

The API `esp_pthread_set_cfg()` defined in the `esp_pthreads.h` header offers custom extensions to control how subsequent calls to `pthread_create()` behaves. Currently, the following configuration can be set:

- Default stack size of new threads, if not specified when calling `pthread_create()` (overrides [CONFIG_PTHREAD_TASK_STACK_SIZE_DEFAULT](#)).
- RTOS priority of new threads (overrides [CONFIG_PTHREAD_TASK_PRIO_DEFAULT](#)).
- Core affinity / core pinning of new threads (overrides [CONFIG_PTHREAD_TASK_CORE_DEFAULT](#)).
- FreeRTOS task name for new threads (overrides [CONFIG_PTHREAD_TASK_NAME_DEFAULT](#))

This configuration is scoped to the calling thread (or FreeRTOS task), meaning that `esp_pthread_set_cfg()` can be called independently in different threads or tasks. If the `inherit_cfg` flag is set in the current configuration then any new thread created will inherit the creator's configuration (if that thread calls `pthread_create()` recursively), otherwise the new thread will have the default configuration.

Examples

- [system/pthread](#) demonstrates using the pthreads API to create threads.
- [cxx/pthread](#) demonstrates using C++ Standard Library functions with threads.

API Reference

Header File

- [components/pthread/include/esp_pthread.h](#)
- This header file can be included with:

```
#include "esp_pthread.h"
```

- This header file is a part of the API provided by the `pthread` component. To declare that your component depends on `pthread`, add the following to your `CMakeLists.txt`:

```
REQUIRES pthread
```

or

```
PRIV_REQUIRES pthread
```

Functions

[esp_pthread_cfg_t](#) **esp_pthread_get_default_config** (void)

Creates a default pthread configuration based on the values set via `menuconfig`.

Returns A default configuration structure.

[esp_err_t](#) **esp_pthread_set_cfg** (const [esp_pthread_cfg_t](#) *cfg)

Configure parameters for creating pthread.

This API allows you to configure how the subsequent `pthread_create()` call will behave. This call can be used to setup configuration parameters like stack size, priority, configuration inheritance etc.

If the 'inherit' flag in the configuration structure is enabled, then the same configuration is also inherited in the thread subtree.

Note: Passing non-NULL attributes to `pthread_create()` will override the `stack_size` parameter set using this API

Parameters `cfg` -- The pthread config parameters

Returns

- `ESP_OK` if configuration was successfully set
- `ESP_ERR_NO_MEM` if out of memory
- `ESP_ERR_INVALID_ARG` if `stack_size` is less than `PTHREAD_STACK_MIN`

[esp_err_t](#) **esp_pthread_get_cfg** ([esp_pthread_cfg_t](#) *p)

Get current pthread creation configuration.

This will retrieve the current configuration that will be used for creating threads.

Parameters `p` -- Pointer to the pthread config structure that will be updated with the currently configured parameters

Returns

- `ESP_OK` if the configuration was available
- `ESP_ERR_NOT_FOUND` if a configuration wasn't previously set

[esp_err_t](#) **esp_pthread_init** (void)

Initialize pthread library.

Structures

struct **esp_pthread_cfg_t**

pthread configuration structure that influences pthread creation

Public Members

size_t **stack_size**

The stack size of the pthread.

size_t **prio**

The thread's priority.

bool **inherit_cfg**

Inherit this configuration further.

const char ***thread_name**

The thread name.

int **pin_to_core**

The core id to pin the thread to. Has the same value range as xCoreId argument of xTaskCreatePinnedToCore.

Macros

PTHREAD_STACK_MIN

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

2.9.27 Random Number Generation

ESP32-P4 contains a hardware random number generator, values from it can be obtained using the APIs *esp_random()* and *esp_fill_random()*.

The hardware RNG produces true random numbers under any of the following conditions:

- RF subsystem is enabled (i.e., Wi-Fi or Bluetooth are enabled).
- An internal entropy source has been enabled by calling *bootloader_random_enable()* and not yet disabled by calling *bootloader_random_disable()*.
- While the ESP-IDF *Second Stage Bootloader* is running. This is because the default ESP-IDF bootloader implementation calls *bootloader_random_enable()* when the bootloader starts, and *bootloader_random_disable()* before executing the app.

When any of these conditions are true, samples of physical noise are continuously mixed into the internal hardware RNG state to provide entropy. Consult the **ESP32-P4 Technical Reference Manual > Random Number Generator (RNG)** [PDF] chapter for more details.

If none of the above conditions are true, the output of the RNG should be considered pseudo-random only.

Startup

During startup, ESP-IDF bootloader temporarily enables a non-RF entropy source (internal reference voltage noise) that provides entropy for any first boot key generation. However, after the app starts executing then normally only pseudo-random numbers are available until Wi-Fi or Bluetooth are initialized.

To re-enable the entropy source temporarily during app startup, or for an application that does not use Wi-Fi or Bluetooth, call the function *bootloader_random_enable()* to re-enable the internal entropy source. The

function `bootloader_random_disable()` must be called to disable the entropy source again before using ADC, Wi-Fi or Bluetooth.

Note: The entropy source enabled during the boot process by the ESP-IDF Second Stage Bootloader seeds the internal RNG state with some entropy. However, the internal hardware RNG state is not large enough to provide a continuous stream of true random numbers. This is why a continuous entropy source must be enabled whenever true random numbers are required.

Note: If an application requires a source of true random numbers but it is not possible to permanently enable a hardware entropy source, consider using a strong software DRBG implementation such as the mbedTLS CTR-DRBG or HMAC-DRBG, with an initial seed of entropy from hardware RNG true random numbers.

Secondary Entropy

ESP32-P4 RNG contains a secondary entropy source, based on sampling an asynchronous 8 MHz internal oscillator (see the Technical Reference Manual for details). This entropy source is always enabled in ESP-IDF and continuously mixed into the RNG state by hardware. In testing, this secondary entropy source was sufficient to pass the [Dieharder](#) random number test suite without the main entropy source enabled (test input was created by concatenating short samples from a continuously resetting ESP32-P4). However, it is currently only guaranteed that true random numbers are produced when the main entropy source is also enabled as described above.

API Reference

Header File

- `components/esp_hw_support/include/esp_random.h`
- This header file can be included with:

```
#include "esp_random.h"
```

Functions

uint32_t **esp_random** (void)

Get one random 32-bit word from hardware RNG.

If Wi-Fi or Bluetooth are enabled, this function returns true random numbers. In other situations, if true random numbers are required then consult the ESP-IDF Programming Guide "Random Number Generation" section for necessary prerequisites.

This function automatically busy-waits to ensure enough external entropy has been introduced into the hardware RNG state, before returning a new random number. This delay is very short (always less than 100 CPU cycles).

Returns Random value between 0 and UINT32_MAX

void **esp_fill_random** (void *buf, size_t len)

Fill a buffer with random bytes from hardware RNG.

Note: This function is implemented via calls to `esp_random()`, so the same constraints apply.

Parameters

- **buf** -- Pointer to buffer to fill with random numbers.
- **len** -- Length of buffer in bytes

Header File

- `components/bootloader_support/include/bootloader_random.h`
- This header file can be included with:

```
#include "bootloader_random.h"
```

- This header file is a part of the API provided by the `bootloader_support` component. To declare that your component depends on `bootloader_support`, add the following to your `CMakeLists.txt`:

```
REQUIRES bootloader_support
```

or

```
PRIV_REQUIRES bootloader_support
```

Functions

void **bootloader_random_enable** (void)

Enable an entropy source for RNG if RF subsystem is disabled.

The exact internal entropy source mechanism depends on the chip in use but all SoCs use the SAR ADC to continuously mix random bits (an internal noise reading) into the HWRNG. Consult the SoC Technical Reference Manual for more information.

Can also be called from app code, if true random numbers are required without initialized RF subsystem. This might be the case in early startup code of the application when the RF subsystem has not started yet or if the RF subsystem should not be enabled for power saving.

Consult ESP-IDF Programming Guide "Random Number Generation" section for details.

Warning: This function is not safe to use if any other subsystem is accessing the RF subsystem or the ADC at the same time!

void **bootloader_random_disable** (void)

Disable entropy source for RNG.

Disables internal entropy source. Must be called after `bootloader_random_enable()` and before RF subsystem features, ADC, or I2S (ESP32 only) are initialized.

Consult the ESP-IDF Programming Guide "Random Number Generation" section for details.

void **bootloader_fill_random** (void *buffer, size_t length)

Fill buffer with 'length' random bytes.

Note: If this function is being called from app code only, and never from the bootloader, then it's better to call `esp_fill_random()`.

Parameters

- **buffer** -- Pointer to buffer
- **length** -- This many bytes of random data will be copied to buffer

getrandom()

A compatible version of the Linux `getrandom()` function is also provided for ease of porting:

```
#include <sys/random.h>

ssize_t getrandom(void *buf, size_t buflen, unsigned int flags);
```

This function is implemented by calling [esp_fill_random\(\)](#) internally.

The `flags` argument is ignored, this function is always non-blocking but the strength of any random numbers is dependent on the same conditions described above.

Return value is -1 (with `errno` set to `EFAULT`) if the `buf` argument is `NULL`, and equal to `buflen` otherwise.

getentropy()

A compatible version of the Linux `getentropy()` function is also provided for ease of porting:

```
#include <unistd.h>

int getentropy(void *buffer, size_t length);
```

This function is implemented by calling `getrandom()` internally.

Strength of any random numbers is dependent on the same conditions described above.

Return value is 0 on success and -1 otherwise with `errno` set to:

- `EFAULT` if the `buffer` argument is `NULL`.
- `EIO` if the `length` is more than 256.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

2.9.28 Sleep Modes

Overview

ESP32-P4 supports two major power saving modes: Light-sleep and Deep-sleep. According to the features used by an application, there are some sub sleep modes. See [Sleep Modes](#) for these sleep modes and sub sleep modes. Additionally, there are some power-down options that can be configured to further reduce the power consumption. See [Power-down Options](#) for more details.

There are several wakeup sources in the sleep modes. These sources can also be combined so that the chip will wake up when any of the sources are triggered. [Wakeup Sources](#) describes these wakeup sources and configuration APIs in detail.

The configuration of power-down options and wakeup sources are optional. They can be configured at any moment before entering the sleep modes.

Then the application can call sleep start APIs to enter one of the sleep modes. See [Entering Sleep](#) for more details. When the wakeup condition is met, the application is awoken from sleep. See [Checking Sleep Wakeup Cause](#) on how to get the wakeup cause, and [Disable Sleep Wakeup Source](#) on how to handle the wakeup sources after wakeup.

Sleep Modes

In Light-sleep mode, the digital peripherals, most of the RAM, and CPUs are clock-gated and their supply voltage is reduced. Upon exit from Light-sleep, the digital peripherals, RAM, and CPUs resume operation and their internal states are preserved.

In Deep-sleep mode, the CPUs, most of the RAM, and all digital peripherals that are clocked from APB_CLK are powered off. The only parts of the chip that remain powered on are:

- RTC controller
- RTC FAST memory

Wi-Fi and Sleep Modes In Deep-sleep and Light-sleep modes, the wireless peripherals are powered down. Before entering Deep-sleep or Light-sleep modes, applications must disable Wi-Fi using the appropriate calls (`esp_wifi_stop()`). Wi-Fi connections are not maintained in Deep-sleep or Light-sleep mode, even if these functions are not called.

If Wi-Fi connections need to be maintained, enable Wi-Fi Modem-sleep mode and automatic Light-sleep feature (see [Power Management APIs](#)). This will allow the system to wake up from sleep automatically when required by the Wi-Fi driver, thereby maintaining a connection to the AP.

Wakeup Sources

Wakeup sources can be enabled using `esp_sleep_enable_X_wakeup` APIs. Wakeup sources are not disabled after wakeup, you can disable them using `esp_sleep_disable_wakeup_source()` API if you do not need them any more. See [Disable Sleep Wakeup Source](#).

Following are the wakeup sources supported on ESP32-P4.

Timer The RTC controller has a built-in timer which can be used to wake up the chip after a predefined amount of time. Time is specified at microsecond precision, but the actual resolution depends on the clock source selected for RTC_SLOW_CLK.

RTC peripherals or RTC memories do not need to be powered on during sleep in this wakeup mode.

`esp_sleep_enable_timer_wakeup()` function can be used to enable sleep wakeup using a timer.

GPIO Wakeup (Light-sleep Only) One more method of wakeup from external inputs is available in Light-sleep mode. With this wakeup source, each pin can be individually configured to trigger wakeup on high or low level using `gpio_wakeup_enable()` function. This wakeup source can be used with any IO (RTC or digital).

`esp_sleep_enable_gpio_wakeup()` function can be used to enable this wakeup source.

Warning: Before entering Light-sleep mode, check if any GPIO pin to be driven is part of the VDD_SPI power domain. If so, this power domain must be configured to remain ON during sleep.

For example, on ESP32-WROOM-32 board, GPIO16 and GPIO17 are linked to VDD_SPI power domain. If they are configured to remain high during Light-sleep, the power domain should be configured to remain powered ON. This can be done with `esp_sleep_pd_config()`:

```
esp_sleep_pd_config(ESP_PD_DOMAIN_VDDSDIO, ESP_PD_OPTION_ON);
```

Note: In Light-sleep mode, if you set Kconfig option `CONFIG_PM_POWER_DOWN_PERIPHERAL_IN_LIGHT_SLEEP`, to continue using `gpio_wakeup_enable()` for GPIO wakeup, you need to first call `rtc_gpio_init()` and `rtc_gpio_set_direction()`, setting the RTCIO to input mode.

Alternatively, you can use `esp_deep_sleep_enable_gpio_wakeup()` directly in that condition for GPIO wakeup, because the digital IO power domain is being powered off, where the situation is the same as entering Deep-sleep.

UART Wakeup (Light-sleep Only) When ESP32-P4 receives UART input from external devices, it is often necessary to wake up the chip when input data is available. The UART peripheral contains a feature which allows waking up the chip from Light-sleep when a certain number of positive edges on RX pin are seen. This number of positive edges can be set using `uart_set_wakeup_threshold()` function. Note that the character which triggers wakeup (and any characters before it) will not be received by the UART after wakeup. This means that the external device typically needs to send an extra character to the ESP32-P4 to trigger wakeup before sending the data.

`esp_sleep_enable_uart_wakeup()` function can be used to enable this wakeup source.

After waking-up from UART, you should send some extra data through the UART port in Active mode, so that the internal wakeup indication signal can be cleared. Otherwise, the next UART wake-up would trigger with two less rising edges than the configured threshold value.

Disable Sleep Wakeup Source Previously configured wakeup sources can be disabled later using `esp_sleep_disable_wakeup_source()` API. This function deactivates trigger for the given wakeup source. Additionally, it can disable all triggers if the argument is `ESP_SLEEP_WAKEUP_ALL`.

Power-down Options

The application can force specific powerdown modes for RTC peripherals and RTC memories. In Deep-sleep mode, we can also isolate some IOs to further reduce current consumption.

Power-down of RTC Peripherals and Memories By default, `esp_deep_sleep_start()` and `esp_light_sleep_start()` functions power down all RTC power domains which are not needed by the enabled wakeup sources. To override this behaviour, `esp_sleep_pd_config()` function is provided.

In ESP32-P4, there is only RTC FAST memory, so if some variables in the program are marked by `RTC_DATA_ATTR`, `RTC_SLOW_ATTR` or `RTC_FAST_ATTR` attributes, all of them go to RTC FAST memory. It will be kept powered on by default. This can be overridden using `esp_sleep_pd_config()` function, if desired.

Power-down of Flash By default, to avoid potential issues, `esp_light_sleep_start()` function does **not** power down flash. To be more specific, it takes time to power down the flash and during this period the system may be woken up, which then actually powers up the flash before this flash could be powered down completely. As a result, there is a chance that the flash may not work properly.

So, in theory, it is ok if you only wake up the system after the flash is completely powered down. However, in reality, the flash power-down period can be hard to predict (for example, this period can be much longer when you add filter capacitors to the flash's power supply circuit) and uncontrollable (for example, the asynchronous wake-up signals make the actual sleep time uncontrollable).

Warning: If a filter capacitor is added to your flash power supply circuit, please do everything possible to avoid powering down flash.

Therefore, it is recommended not to power down flash when using ESP-IDF. For power-sensitive applications, it is recommended to use Kconfig option `CONFIG_ESP_SLEEP_FLASH_LEAKAGE_WORKAROUND` to reduce the power consumption of the flash during Light-sleep, instead of powering down the flash.

It is worth mentioning that PSRAM has a similar Kconfig option `CONFIG_ESP_SLEEP_PSRAM_LEAKAGE_WORKAROUND`.

However, for those who have fully understood the risk and are still willing to power down the flash to further reduce the power consumption, please check the following mechanisms:

- Setting Kconfig option `CONFIG_ESP_SLEEP_POWER_DOWN_FLASH` only powers down the flash when the RTC timer is the only wake-up source **and** the sleep time is longer than the flash power-down period.
- Calling `esp_sleep_pd_config(ESP_PD_DOMAIN_VDDSDIO, ESP_PD_OPTION_OFF)` powers down flash when the RTC timer is not enabled as a wakeup source **or** the sleep time is longer than the flash power-down period.

Note:

- ESP-IDF does not provide any mechanism that can power down the flash in all conditions when Light-sleep.
 - `esp_deep_sleep_start()` function forces power down flash regardless of user configuration.
-

Configuring IOs (Deep-sleep Only) Some ESP32-P4 IOs have internal pullups or pulldowns, which are enabled by default. If an external circuit drives this pin in Deep-sleep mode, current consumption may increase due to current flowing through these pullups and pulldowns.

To isolate a pin to prevent extra current draw, call `rtc_gpio_isolate()` function.

For example, on ESP32-WROVER module, GPIO12 is pulled up externally, and it also has an internal pulldown in the ESP32 chip. This means that in Deep-sleep, some current flows through these external and internal resistors, increasing Deep-sleep current above the minimal possible value.

Add the following code before `esp_deep_sleep_start()` to remove such extra current:

```
rtc_gpio_isolate(GPIO_NUM_12);
```

Entering Sleep

`esp_light_sleep_start()` or `esp_deep_sleep_start()` functions can be used to enter Light-sleep or Deep-sleep modes correspondingly. After that, the system configures the parameters of RTC controller according to the requested wakeup sources and power-down options.

It is also possible to enter sleep modes with no wakeup sources configured. In this case, the chip will be in sleep modes indefinitely until external reset is applied.

UART Output Handling Before entering sleep mode, `esp_deep_sleep_start()` will flush the contents of UART FIFOs.

When entering Light-sleep mode using `esp_light_sleep_start()`, UART FIFOs will not be flushed. Instead, UART output will be suspended, and remaining characters in the FIFO will be sent out after wakeup from Light-sleep.

Checking Sleep Wakeup Cause

`esp_sleep_get_wakeup_cause()` function can be used to check which wakeup source has triggered wakeup from sleep mode.

Application Example

- `protocols/sntp`: the implementation of basic functionality of Deep-sleep, where ESP module is periodically waken up to retrieve time from NTP server.
- `wifi/power_save`: the usage of Wi-Fi Modem-sleep mode and automatic Light-sleep feature to maintain Wi-Fi connections.
- `system/deep_sleep`: the usage of Deep-sleep wakeup triggered by timer.

API Reference

Header File

- `components/esp_hw_support/include/esp_sleep.h`
- This header file can be included with:

```
#include "esp_sleep.h"
```

Functions

esp_err_t **esp_sleep_disable_wakeup_source** (*esp_sleep_source_t* source)

Disable wakeup source.

This function is used to deactivate wake up trigger for source defined as parameter of the function.

See docs/sleep-modes.rst for details.

Note: This function does not modify wake up configuration in RTC. It will be performed in `esp_deep_sleep_start/esp_light_sleep_start` function.

Parameters **source** -- number of source to disable of type `esp_sleep_source_t`

Returns

- `ESP_OK` on success
- `ESP_ERR_INVALID_STATE` if trigger was not active

esp_err_t **esp_sleep_enable_timer_wakeup** (*uint64_t* time_in_us)

Enable wakeup by timer.

Note: The valid `time_in_us` value depends on the bit width of the `lp_timer/rtc_timer` counter and the current slow clock source selection (Refer RTC clock source configuration in `menuconfig`). Valid values should be positive values less than `RTC slow clock period * (2 ^ RTC timer bitwidth)`.

Parameters **time_in_us** -- time before wakeup, in microseconds

Returns

- `ESP_OK` on success
- `ESP_ERR_INVALID_ARG` if value is out of range.

bool **esp_sleep_is_valid_wakeup_gpio** (*gpio_num_t* gpio_num)

Returns true if a GPIO number is valid for use as wakeup source.

Note: For SoCs with RTC IO capability, this can be any valid RTC IO input pin.

Parameters **gpio_num** -- Number of the GPIO to test for wakeup source capability

Returns True if this GPIO number will be accepted as a sleep wakeup source.

esp_err_t **esp_deep_sleep_enable_gpio_wakeup** (*uint64_t* gpio_pin_mask,
esp_deepsleep_gpio_wake_up_mode_t mode)

Enable wakeup using specific gpio pins.

This function enables an IO pin to wake up the chip from deep sleep.

Note: 1.This function does not modify pin configuration. The pins are configured inside `esp_deep_sleep_start`, immediately before entering sleep mode. 2.This function is also applicable to waking up the lightsleep when the peripheral power domain is powered off, see `PM_POWER_DOWN_PERIPHERAL_IN_LIGHT_SLEEP` in `menuconfig`.

Note: You don't need to worry about pull-up or pull-down resistors before using this function because the `ESP_SLEEP_GPIO_ENABLE_INTERNAL_RESISTORS` option is enabled by default. It will automatically set pull-up or pull-down resistors internally in `esp_deep_sleep_start` based on the wakeup mode. However, when using external pull-up or pull-down resistors, please be sure to disable the `ESP_SLEEP_GPIO_ENABLE_INTERNAL_RESISTORS` option, as the combination of internal and external resistors may cause interference. BTW, when you use low level to wake up the chip, we strongly recommend you to add external resistors (pull-up).

Parameters

- **gpio_pin_mask** -- Bit mask of GPIO numbers which will cause wakeup. Only GPIOs which have RTC functionality (pads that powered by `VDD3P3_RTC`) can be used in this bit map.
- **mode** -- Select logic function used to determine wakeup condition:
 - `ESP_GPIO_WAKEUP_GPIO_LOW`: wake up when the gpio turn to low.
 - `ESP_GPIO_WAKEUP_GPIO_HIGH`: wake up when the gpio turn to high.

Returns

- `ESP_OK` on success
- `ESP_ERR_INVALID_ARG` if the mask contains any invalid deep sleep wakeup pin or wakeup mode is invalid

[*esp_err_t*](#) **esp_sleep_enable_gpio_wakeup** (void)

Enable wakeup from light sleep using GPIOs.

Each GPIO supports wakeup function, which can be triggered on either low level or high level. Unlike `EXT0` and `EXT1` wakeup sources, this method can be used both for all IOs: RTC IOs and digital IOs. It can only be used to wakeup from light sleep though.

To enable wakeup, first call `gpio_wakeup_enable`, specifying gpio number and wakeup level, for each GPIO which is used for wakeup. Then call this function to enable wakeup feature.

Note: 1. On ESP32, GPIO wakeup source can not be used together with touch or ULP wakeup sources.

- If `PM_POWER_DOWN_PERIPHERAL_IN_LIGHT_SLEEP` is enabled (if target supported), this API is unavailable since the GPIO module is powered down during sleep. You can use `esp_deep_sleep_enable_gpio_wakeup` instead, or use `EXT1` wakeup source by `esp_sleep_enable_ext1_wakeup_io` to achieve the same function. (Only GPIOs which have RTC functionality can be used)
-

Returns

- `ESP_OK` on success
- `ESP_ERR_INVALID_STATE` if wakeup triggers conflict

[*esp_err_t*](#) **esp_sleep_enable_uart_wakeup** (int uart_num)

Enable wakeup from light sleep using UART.

Use `uart_set_wakeup_threshold` function to configure UART wakeup threshold.

Wakeup from light sleep takes some time, so not every character sent to the UART can be received by the application.

Note: 1. ESP32 does not support wakeup from UART2.

- a. If PM_POWER_DOWN_PERIPHERAL_IN_LIGHT_SLEEP is enabled (if target supported), this API is unavailable since the UART module is powered down during sleep.
-

Parameters `uart_num` -- UART port to wake up from

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if wakeup from given UART is not supported

esp_err_t **esp_sleep_enable_bt_wakeup** (void)

Enable wakeup by bluetooth.

Returns

- ESP_OK on success
- ESP_ERR_NOT_SUPPORTED if wakeup from bluetooth is not supported

esp_err_t **esp_sleep_disable_bt_wakeup** (void)

Disable wakeup by bluetooth.

Returns

- ESP_OK on success
- ESP_ERR_NOT_SUPPORTED if wakeup from bluetooth is not supported

esp_err_t **esp_sleep_enable_wifi_wakeup** (void)

Enable wakeup by WiFi MAC.

Returns

- ESP_OK on success

esp_err_t **esp_sleep_disable_wifi_wakeup** (void)

Disable wakeup by WiFi MAC.

Returns

- ESP_OK on success

esp_err_t **esp_sleep_enable_wifi_beacon_wakeup** (void)

Enable beacon wakeup by WiFi MAC, it will wake up the system into modem state.

Returns

- ESP_OK on success

esp_err_t **esp_sleep_disable_wifi_beacon_wakeup** (void)

Disable beacon wakeup by WiFi MAC.

Returns

- ESP_OK on success

uint64_t **esp_sleep_get_ext1_wakeup_status** (void)

Get the bit mask of GPIOs which caused wakeup (ext1)

If wakeup was caused by another source, this function will return 0.

Returns bit mask, if GPIO_n caused wakeup, BIT(n) will be set

uint64_t **esp_sleep_get_gpio_wakeup_status** (void)

Get the bit mask of GPIOs which caused wakeup (gpio)

If wakeup was caused by another source, this function will return 0.

Returns bit mask, if GPIO_n caused wakeup, BIT(n) will be set

esp_err_t **esp_sleep_pd_config** (*esp_sleep_pd_domain_t* domain, *esp_sleep_pd_option_t* option)

Set power down mode for an RTC power domain in sleep mode.

If not set using this API, all power domains default to ESP_PD_OPTION_AUTO.

Parameters

- **domain** -- power domain to configure
- **option** -- power down option (ESP_PD_OPTION_OFF, ESP_PD_OPTION_ON, or ESP_PD_OPTION_AUTO)

Returns

- ESP_OK on success
- ESP_ERR_INVALID_ARG if either of the arguments is out of range

esp_err_t **esp_deep_sleep_try_to_start** (void)

Enter deep sleep with the configured wakeup options.

The reason for the rejection can be such as a short sleep time.

Note: In general, the function does not return, but if the sleep is rejected, then it returns from it.

Returns

- No return - If the sleep is not rejected.
- ESP_ERR_SLEEP_REJECT sleep request is rejected(wakeup source set before the sleep request)

void **esp_deep_sleep_start** (void)

Enter deep sleep with the configured wakeup options.

Note: The function does not do a return (no rejection). Even if wakeup source set before the sleep request it goes to deep sleep anyway.

esp_err_t **esp_light_sleep_start** (void)

Enter light sleep with the configured wakeup options.

Returns

- ESP_OK on success (returned after wakeup)
- ESP_ERR_SLEEP_REJECT sleep request is rejected(wakeup source set before the sleep request)
- ESP_ERR_SLEEP_TOO_SHORT_SLEEP_DURATION after deducting the sleep flow overhead, the final sleep duration is too short to cover the minimum sleep duration of the chip, when rtc timer wakeup source enabled

esp_err_t **esp_deep_sleep_try** (uint64_t time_in_us)

Enter deep-sleep mode.

The device will automatically wake up after the deep-sleep time. Upon waking up, the device calls deep sleep wake stub, and then proceeds to load application.

Call to this function is equivalent to a call to esp_deep_sleep_enable_timer_wakeup followed by a call to esp_deep_sleep_start.

The reason for the rejection can be such as a short sleep time.

Note: In general, the function does not return, but if the sleep is rejected, then it returns from it.

Parameters `time_in_us` -- deep-sleep time, unit: microsecond

Returns

- No return - If the sleep is not rejected.
- `ESP_ERR_SLEEP_REJECT` sleep request is rejected(wakeup source set before the sleep request)

void **esp_deep_sleep** (uint64_t time_in_us)

Enter deep-sleep mode.

The device will automatically wake up after the deep-sleep time. Upon waking up, the device calls deep sleep wake stub, and then proceeds to load application.

Call to this function is equivalent to a call to `esp_deep_sleep_enable_timer_wakeup` followed by a call to `esp_deep_sleep_start`.

Note: The function does not do a return (no rejection).. Even if wakeup source set before the sleep request it goes to deep sleep anyway.

Parameters `time_in_us` -- deep-sleep time, unit: microsecond

esp_err_t **esp_deep_sleep_register_hook** (*esp_deep_sleep_cb_t* new_dslp_cb)

Register a callback to be called from the deep sleep prepare.

Warning: deepsleep callbacks should without parameters, and MUST NOT, UNDER ANY CIRCUMSTANCES, CALL A FUNCTION THAT MIGHT BLOCK.

Parameters `new_dslp_cb` -- Callback to be called

Returns

- `ESP_OK`: Callback registered to the deepsleep `misc_modules_sleep_prepare`
- `ESP_ERR_NO_MEM`: No more hook space for register the callback

void **esp_deep_sleep_deregister_hook** (*esp_deep_sleep_cb_t* old_dslp_cb)

Unregister an deepsleep callback.

Parameters `old_dslp_cb` -- Callback to be unregistered

esp_sleep_wakeup_cause_t **esp_sleep_get_wakeup_cause** (void)

Get the wakeup source which caused wakeup from sleep.

Note: !!! This API will only return one wakeup source. If multiple wakeup sources wake up at the same time, the wakeup source information may be lost.

Returns cause of wake up from last sleep (deep sleep or light sleep)

uint32_t **esp_sleep_get_wakeup_causes** (void)

Get all wakeup sources bitmap which caused wakeup from sleep.

Returns The bitmap of the wakeup sources of the last wakeup from sleep. (deep sleep or light sleep)

void **esp_wake_deep_sleep** (void)

Default stub to run on wake from deep sleep.

Allows for executing code immediately on wake from sleep, before the software bootloader or ESP-IDF app has started up.

This function is weak-linked, so you can implement your own version to run code immediately when the chip wakes from sleep.

See docs/deep-sleep-stub.rst for details.

void **esp_set_deep_sleep_wake_stub** (*esp_deep_sleep_wake_stub_fn_t* new_stub)

Install a new stub at runtime to run on wake from deep sleep.

If implementing esp_wake_deep_sleep() then it is not necessary to call this function.

However, it is possible to call this function to substitute a different deep sleep stub. Any function used as a deep sleep stub must be marked RTC_IRAM_ATTR, and must obey the same rules given for esp_wake_deep_sleep().

void **esp_set_deep_sleep_wake_stub_default_entry** (void)

Set wake stub entry to default esp_wake_stub_entry

esp_deep_sleep_wake_stub_fn_t **esp_get_deep_sleep_wake_stub** (void)

Get current wake from deep sleep stub.

Returns Return current wake from deep sleep stub, or NULL if no stub is installed.

void **esp_default_wake_deep_sleep** (void)

The default esp-idf-provided esp_wake_deep_sleep() stub.

See docs/deep-sleep-stub.rst for details.

void **esp_deep_sleep_disable_rom_logging** (void)

Disable logging from the ROM code after deep sleep.

Using LSB of RTC_STORE4.

esp_err_t **esp_sleep_cpu_retention_init** (void)

CPU Power down initialize.

Returns

- ESP_OK on success
- ESP_ERR_NO_MEM not enough retention memory

esp_err_t **esp_sleep_cpu_retention_deinit** (void)

CPU Power down de-initialize.

Release system retention memory.

Returns

- ESP_OK on success

void **esp_sleep_config_gpio_isolate** (void)

Configure to isolate all GPIO pins in sleep state.

void **esp_sleep_enable_gpio_switch** (bool enable)

Enable or disable GPIO pins status switching between slept status and waked status.

Parameters **enable** -- decide whether to switch status or not

Macros

ESP_PD_DOMAIN_RTC8M

Type Definitions

typedef void (***esp_deep_sleep_cb_t**)(void)

typedef [*esp_sleep_source_t*](#) **esp_sleep_wakeup_cause_t**

typedef void (***esp_deep_sleep_wake_stub_fn_t**)(void)

Function type for stub to run on wake from sleep.

Enumerations

enum **esp_deepsleep_gpio_wake_up_mode_t**

Logic function used for EXT1 wakeup mode.

Values:

enumerator **ESP_GPIO_WAKEUP_GPIO_LOW**

enumerator **ESP_GPIO_WAKEUP_GPIO_HIGH**

enum **esp_sleep_pd_domain_t**

Power domains which can be powered down in sleep mode.

Values:

enumerator **ESP_PD_DOMAIN_XTAL**

XTAL oscillator.

enumerator **ESP_PD_DOMAIN_XTAL32K**

External 32 kHz XTAL oscillator.

enumerator **ESP_PD_DOMAIN_RC32K**

Internal 32 kHz RC oscillator.

enumerator **ESP_PD_DOMAIN_RC_FAST**

Internal Fast oscillator.

enumerator **ESP_PD_DOMAIN_CPU**

CPU core.

enumerator **ESP_PD_DOMAIN_VDDSDIO**

VDD_SDIO.

enumerator **ESP_PD_DOMAIN_MODEM**

MODEM, includes WiFi, Bluetooth and IEEE802.15.4.

enumerator **ESP_PD_DOMAIN_TOP**

SoC TOP.

enumerator **ESP_PD_DOMAIN_MAX**

Number of domains.

enum **esp_sleep_pd_option_t**

Power down options.

Values:

enumerator **ESP_PD_OPTION_OFF**

Power down the power domain in sleep mode.

enumerator **ESP_PD_OPTION_ON**

Keep power domain enabled during sleep mode.

enumerator **ESP_PD_OPTION_AUTO**

Keep power domain enabled in sleep mode, if it is needed by one of the wakeup options. Otherwise power it down.

enum **esp_sleep_source_t**

Sleep wakeup cause.

Values:

enumerator **ESP_SLEEP_WAKEUP_UNDEFINED**

In case of deep sleep, reset was not caused by exit from deep sleep.

enumerator **ESP_SLEEP_WAKEUP_ALL**

Not a wakeup cause, used to disable all wakeup sources with `esp_sleep_disable_wakeup_source`.

enumerator **ESP_SLEEP_WAKEUP_EXT0**

Wakeup caused by external signal using RTC_IO.

enumerator **ESP_SLEEP_WAKEUP_EXT1**

Wakeup caused by external signal using RTC_CNTL.

enumerator **ESP_SLEEP_WAKEUP_TIMER**

Wakeup caused by timer.

enumerator **ESP_SLEEP_WAKEUP_TOUCHPAD**

Wakeup caused by touchpad.

enumerator **ESP_SLEEP_WAKEUP_ULP**

Wakeup caused by ULP program.

enumerator **ESP_SLEEP_WAKEUP_GPIO**

Wakeup caused by GPIO (light sleep only on ESP32, S2 and S3)

enumerator **ESP_SLEEP_WAKEUP_UART**

Wakeup caused by UART0 (light sleep only)

enumerator **ESP_SLEEP_WAKEUP_UART1**

Wakeup caused by UART1 (light sleep only)

enumerator **ESP_SLEEP_WAKEUP_UART2**

Wakeup caused by UART2 (light sleep only)

enumerator **ESP_SLEEP_WAKEUP_WIFI**

Wakeup caused by WIFI (light sleep only)

enumerator **ESP_SLEEP_WAKEUP_COCPU**

Wakeup caused by COCPU int.

enumerator **ESP_SLEEP_WAKEUP_COCPU_TRAP_TRIG**

Wakeup caused by COCPU crash.

enumerator **ESP_SLEEP_WAKEUP_BT**

Wakeup caused by BT (light sleep only)

enum **esp_sleep_mode_t**

Sleep mode.

Values:

enumerator **ESP_SLEEP_MODE_LIGHT_SLEEP**

light sleep mode

enumerator **ESP_SLEEP_MODE_DEEP_SLEEP**

deep sleep mode

enum **[anonymous]**

Values:

enumerator **ESP_ERR_SLEEP_REJECT**

enumerator **ESP_ERR_SLEEP_TOO_SHORT_SLEEP_DURATION**

2.9.29 SoC Capabilities

This section lists the macro definitions of the ESP32-P4's SoC hardware capabilities. These macros are commonly used by conditional-compilation directives (e.g., `#if`) in ESP-IDF to determine which hardware-dependent features are supported, thus control what portions of code are compiled.

Warning: These macro definitions are currently not considered to be part of the public API, and may be changed in a breaking manner (see [ESP-IDF Versions](#) for more details).

API Reference

Header File

- [components/soc/esp32p4/include/soc/soc_caps.h](#)
- This header file can be included with:

```
#include "soc/soc_caps.h"
```

Macros

SOC_ANA_CMPR_SUPPORTED

SOC_UART_SUPPORTED

SOC_GDMA_SUPPORTED

SOC_AHB_GDMA_SUPPORTED

SOC_AXI_GDMA_SUPPORTED

SOC_GPTIMER_SUPPORTED

SOC_PCNT_SUPPORTED

SOC_MCPWM_SUPPORTED

SOC_ETM_SUPPORTED

SOC_PARLIO_SUPPORTED

SOC_ASYNC_MEMCPY_SUPPORTED

SOC_SUPPORTS_SECURE_DL_MODE

SOC_EFUSE_KEY_PURPOSE_FIELD

SOC_EFUSE_SUPPORTED

SOC_RTC_FAST_MEM_SUPPORTED

SOC_RTC_MEM_SUPPORTED

SOC_RMT_SUPPORTED

SOC_I2S_SUPPORTED

SOC_GPSPI_SUPPORTED

SOC_LEDC_SUPPORTED

SOC_I2C_SUPPORTED

SOC_SYSTIMER_SUPPORTED

SOC_MPI_SUPPORTED

SOC_HMAC_SUPPORTED

SOC_DIG_SIGN_SUPPORTED

SOC_ECC_SUPPORTED

SOC_ECC_EXTENDED_MODES_SUPPORTED

SOC_ECDSA_SUPPORTED

SOC_FLASH_ENC_SUPPORTED

SOC_SECURE_BOOT_SUPPORTED

SOC_LP_GPIO_MATRIX_SUPPORTED

SOC_LP_PERIPHERALS_SUPPORTED

SOC_SPIRAM_SUPPORTED

SOC_PSRAM_DMA_CAPABLE

SOC_SDMMC_HOST_SUPPORTED

SOC_WDT_SUPPORTED

SOC_SPI_FLASH_SUPPORTED

SOC_XTAL_SUPPORT_40M

SOC_AES_SUPPORT_DMA

SOC_AES_GDMA

SOC_AES_SUPPORT_AES_128

SOC_AES_SUPPORT_AES_256

SOC_ADC_DIG_SUPPORTED_UNIT (UNIT)

< SAR ADC Module

SOC_ADC_PERIPH_NUM

SOC_ADC_CHANNEL_NUM (PERIPH_NUM)

SOC_ADC_MAX_CHANNEL_NUM

SOC_ADC_ATTEN_NUM

Digital

SOC_ADC_DIGI_CONTROLLER_NUM

SOC_ADC_PATT_LEN_MAX

Two pattern tables, each contains 4 items. Each item takes 1 byte

SOC_ADC_DIGI_MAX_BITWIDTH

SOC_ADC_DIGI_MIN_BITWIDTH

SOC_ADC_DIGI_IIR_FILTER_NUM

SOC_ADC_DIGI_MONITOR_NUM

SOC_ADC_DIGI_RESULT_BYTES

SOC_ADC_DIGI_DATA_BYTES_PER_CONV

$F_{\text{sample}} = F_{\text{digi\_con}} / 2 / \text{interval}$. $F_{\text{digi\_con}} = 5\text{M}$ for now. $30 \leq \text{interval} \leq 4095$

SOC_ADC_SAMPLE_FREQ_THRES_HIGH

SOC_ADC_SAMPLE_FREQ_THRES_LOW

RTC

SOC_ADC_RTC_MIN_BITWIDTH

SOC_ADC_RTC_MAX_BITWIDTH

Calibration

SOC_ADC_CALIBRATION_V1_SUPPORTED

support HW offset calibration version 1

SOC_APB_BACKUP_DMA

SOC_BROWNOUT_RESET_SUPPORTED

SOC_SHARED_IDCACHE_SUPPORTED

SOC_CACHE_WRITEBACK_SUPPORTED

SOC_CACHE_FREEZE_SUPPORTED

SOC_CACHE_INTERNAL_MEM_VIA_L1CACHE

SOC_CPU_CORES_NUM

SOC_CPU_INTR_NUM

SOC_CPU_HAS_FLEXIBLE_INTC

SOC_INT_PLIC_SUPPORTED

SOC_INT_CLIC_SUPPORTED

SOC_INT_HW_NESTED_SUPPORTED

SOC_BRANCH_PREDICTOR_SUPPORTED

SOC_CPU_HAS_FPU

SOC_CPU_HAS_FPU_EXT_ILL_BUG

SOC_CPU_COPROC_NUM

SOC_HP_CPU_HAS_MULTIPLE_CORES

SOC_CPU_BREAKPOINTS_NUM

SOC_CPU_WATCHPOINTS_NUM

SOC_CPU_WATCHPOINT_MAX_REGION_SIZE

SOC_CPU_HAS_PMA

SOC_CPU_IDRAM_SPLIT_USING_PMP

SOC_DS_SIGNATURE_MAX_BIT_LEN

The maximum length of a Digital Signature in bits.

SOC_DS_KEY_PARAM_MD_IV_LENGTH

Initialization vector (IV) length for the RSA key parameter message digest (MD) in bytes.

SOC_DS_KEY_CHECK_MAX_WAIT_US

Maximum wait time for DS parameter decryption key. If overdue, then key error. See TRM DS chapter for more details

SOC_AHB_GDMA_VERSION

SOC_GDMA_SUPPORT_CRC

SOC_GDMA_NUM_GROUPS_MAX

SOC_GDMA_PAIRS_PER_GROUP_MAX

SOC_AXI_GDMA_SUPPORT_PSRAM

SOC_ETM_GROUPS

SOC_ETM_CHANNELS_PER_GROUP

SOC_GPIO_PORT

SOC_GPIO_PIN_COUNT

SOC_GPIO_SUPPORT_PIN_HYS_FILTER

SOC_GPIO_SUPPORT_ETM

SOC_GPIO_SUPPORT_RTC_INDEPENDENT

SOC_GPIO_SUPPORT_DEEPSLEEP_WAKEUP

SOC_GPIO_VALID_GPIO_MASK

SOC_GPIO_VALID_OUTPUT_GPIO_MASK

SOC_GPIO_IN_RANGE_MAX

SOC_GPIO_OUT_RANGE_MAX

SOC_GPIO_DEEP_SLEEP_WAKE_VALID_GPIO_MASK

SOC_GPIO_VALID_DIGITAL_IO_PAD_MASK

SOC_GPIO_SUPPORT_FORCE_HOLD

SOC_GPIO_SUPPORT_HOLD_SINGLE_IO_IN_DSLP

SOC_RTCIO_PIN_COUNT

SOC_RTCIO_INPUT_OUTPUT_SUPPORTED

SOC_RTCIO_HOLD_SUPPORTED

SOC_RTCIO_WAKE_SUPPORTED

SOC_RTCIO_EDGE_WAKE_SUPPORTED

SOC_DEDIC_GPIO_OUT_CHANNELS_NUM

8 outward channels on each CPU core

SOC_DEDIC_GPIO_IN_CHANNELS_NUM

8 inward channels on each CPU core

SOC_DEDIC_PERIPH_ALWAYS_ENABLE

The dedicated GPIO (a.k.a. fast GPIO) is featured by some customized CPU instructions, which is always enabled

SOC_ANA_CMPR_NUM

SOC_ANA_CMPR_CAN_DISTINGUISH_EDGE

SOC_ANA_CMPR_SUPPORT_ETM

SOC_I2C_NUM

SOC_I2C_FIFO_LEN

I2C hardware FIFO depth

SOC_I2C_CMD_REG_NUM

Number of I2C command registers

SOC_I2C_SUPPORT_SLAVE

SOC_I2C_SUPPORT_HW_FSM_RST

SOC_I2C_SUPPORT_HW_CLR_BUS

SOC_I2C_SUPPORT_XTAL

SOC_I2C_SUPPORT_RTC

SOC_I2C_SUPPORT_10BIT_ADDR

SOC_I2C_SLAVE_SUPPORT_BROADCAST

SOC_I2C_SLAVE_SUPPORT_I2CRAM_ACCESS

SOC_I2C_SLAVE_SUPPORT_SLAVE_UNMATCH

SOC_I2S_NUM

SOC_I2S_HW_VERSION_2

SOC_I2S_SUPPORTS_XTAL

SOC_I2S_SUPPORTS_APLL

SOC_I2S_SUPPORTS_PCM

SOC_I2S_SUPPORTS_PDM

SOC_I2S_SUPPORTS_PDM_TX

SOC_I2S_SUPPORTS_PDM_RX

SOC_I2S_SUPPORTS_PDM_RX_HP_FILTER

SOC_I2S_SUPPORTS_TX_SYNC_CNT

SOC_I2S_SUPPORTS_TDM

SOC_I2S_PDM_MAX_TX_LINES

SOC_I2S_PDM_MAX_RX_LINES

SOC_I2S_TDM_FULL_DATA_WIDTH

No limitation to data bit width when using multiple slots

SOC_LEDC_SUPPORT_PLL_DIV_CLOCK

SOC_LEDC_SUPPORT_XTAL_CLOCK

SOC_LEDC_CHANNEL_NUM

SOC_LEDC_TIMER_BIT_WIDTH

SOC_LEDC_GAMMA_CURVE_FADE_SUPPORTED

SOC_LEDC_GAMMA_CURVE_FADE_RANGE_MAX

SOC_LEDC_SUPPORT_FADE_STOP

SOC_LEDC_FADE_PARAMS_BIT_WIDTH

SOC_MMU_PAGE_SIZE_CONFIGURABLE

SOC_MMU_PERIPH_NUM

SOC_MMU_LINEAR_ADDRESS_REGION_NUM

SOC_MMU_DI_VADDR_SHARED

D/I vaddr are shared

SOC_MPU_CONFIGURABLE_REGIONS_SUPPORTED

SOC_MPU_MIN_REGION_SIZE

SOC_MPU_REGIONS_MAX_NUM

SOC_MPU_REGION_RO_SUPPORTED

SOC_MPU_REGION_WO_SUPPORTED

SOC_PCNT_GROUPS

SOC_PCNT_UNITS_PER_GROUP

SOC_PCNT_CHANNELS_PER_UNIT

SOC_PCNT_THRES_POINT_PER_UNIT

SOC_PCNT_SUPPORT_RUNTIME_THRES_UPDATE

SOC_PCNT_SUPPORT_CLEAR_SIGNAL

Support clear signal input

SOC_RMT_GROUPS

One RMT group

SOC_RMT_TX_CANDIDATES_PER_GROUP

Number of channels that capable of Transmit in each group

SOC_RMT_RX_CANDIDATES_PER_GROUP

Number of channels that capable of Receive in each group

SOC_RMT_CHANNELS_PER_GROUP

Total 8 channels

SOC_RMT_MEM_WORDS_PER_CHANNEL

Each channel owns 48 words memory (1 word = 4 Bytes)

SOC_RMT_SUPPORT_RX_PINGPONG

Support Ping-Pong mode on RX path

SOC_RMT_SUPPORT_RX_DEMODULATION

Support signal demodulation on RX path (i.e. remove carrier)

SOC_RMT_SUPPORT_ASYNC_STOP

Support stop transmission asynchronously

SOC_RMT_SUPPORT_TX_LOOP_COUNT

Support transmit specified number of cycles in loop mode

SOC_RMT_SUPPORT_TX_LOOP_AUTO_STOP

Hardware support of auto-stop in loop mode

SOC_RMT_SUPPORT_TX_SYNCHRO

Support coordinate a group of TX channels to start simultaneously

SOC_RMT_SUPPORT_TX_CARRIER_DATA_ONLY

TX carrier can be modulated to data phase only

SOC_RMT_SUPPORT_XTAL

Support set XTAL clock as the RMT clock source

SOC_RMT_SUPPORT_DMA

RMT peripheral can connect to DMA channel

SOC_MCPWM_GROUPS

2 MCPWM groups on the chip (i.e., the number of independent MCPWM peripherals)

SOC_MCPWM_TIMERS_PER_GROUP

The number of timers that each group has.

SOC_MCPWM_OPERATORS_PER_GROUP

The number of operators that each group has.

SOC_MCPWM_COMPARATORS_PER_OPERATOR

The number of comparators that each operator has.

SOC_MCPWM_EVENT_COMPARATORS_PER_OPERATOR

The number of event comparators that each operator has.

SOC_MCPWM_GENERATORS_PER_OPERATOR

The number of generators that each operator has.

SOC_MCPWM_TRIGGERS_PER_OPERATOR

The number of triggers that each operator has.

SOC_MCPWM_GPIO_FAULTS_PER_GROUP

The number of fault signal detectors that each group has.

SOC_MCPWM_CAPTURE_TIMERS_PER_GROUP

The number of capture timers that each group has.

SOC_MCPWM_CAPTURE_CHANNELS_PER_TIMER

The number of capture channels that each capture timer has.

SOC_MCPWM_GPIO_SYNCHROS_PER_GROUP

The number of GPIO synchros that each group has.

SOC_MCPWM_SWSYNC_CAN_PROPAGATE

Software sync event can be routed to its output.

SOC_MCPWM_SUPPORT_ETM

Support ETM (Event Task Matrix)

SOC_MCPWM_SUPPORT_EVENT_COMPARATOR

Support event comparator (based on ETM)

SOC_MCPWM_CAPTURE_CLK_FROM_GROUP

Capture timer shares clock with other PWM timers.

SOC_PARLIO_GROUPS

Number of parallel IO peripherals

SOC_PARLIO_TX_UNITS_PER_GROUP

number of TX units in each group

SOC_PARLIO_RX_UNITS_PER_GROUP

number of RX units in each group

SOC_PARLIO_TX_UNIT_MAX_DATA_WIDTH

Number of data lines of the TX unit

SOC_PARLIO_RX_UNIT_MAX_DATA_WIDTH

Number of data lines of the RX unit

SOC_PARLIO_TX_SIZE_BY_DMA

Transaction length is controlled by DMA instead of indicated by register

SOC_MPI_MEM_BLOCKS_NUM**SOC_MPI_OPERATIONS_NUM****SOC_RSA_MAX_BIT_LEN**

SOC_SDMMC_USE_IOMUX

Card detect, write protect, interrupt use GPIO Matrix on all chips. Slot 0 clock/cmd/data pins use IOMUX
Slot 1 clock/cmd/data pins use GPIO Matrix

SOC_SDMMC_USE_GPIO_MATRIX**SOC_SDMMC_NUM_SLOTS****SOC_SDMMC_DELAY_PHASE_NUM****SOC_SHA_DMA_MAX_BUFFER_SIZE****SOC_SHA_SUPPORT_DMA****SOC_SHA_SUPPORT_RESUME****SOC_SHA_GDMA****SOC_SHA_SUPPORT_SHA1****SOC_SHA_SUPPORT_SHA224****SOC_SHA_SUPPORT_SHA256****SOC_ECDSA_SUPPORT_EXPORT_PUBKEY****SOC_ECDSA_USES_MPI****SOC_SDM_GROUPS****SOC_SDM_CHANNELS_PER_GROUP****SOC_SDM_CLK_SUPPORT_PLL_F80M****SOC_SDM_CLK_SUPPORT_XTAL****SOC_SPI_PERIPH_NUM****SOC_SPI_PERIPH_CS_NUM** (i)**SOC_SPI_MAX_CS_NUM****SOC_SPI_MAXIMUM_BUFFER_SIZE****SOC_SPI_SLAVE_SUPPORT_SEG_TRANS**

SOC_SPI_SUPPORT_DDRCLK

SOC_SPI_SUPPORT_CD_SIG

SOC_SPI_SUPPORT_OCT

SOC_SPI_SUPPORT_CLK_XTAL

SOC_SPI_PERIPH_SUPPORT_MULTILINE_MODE (host_id)

SOC_MEMSPI_IS_INDEPENDENT

SOC_SPI_MAX_PRE_DIVIDER

SOC_SPI_MEM_SUPPORT_AUTO_WAIT_IDLE

SOC_SPI_MEM_SUPPORT_AUTO_RESUME

SOC_SPI_MEM_SUPPORT_IDLE_INTR

SOC_SPI_MEM_SUPPORT_SW_SUSPEND

SOC_SPI_MEM_SUPPORT_CHECK_SUS

SOC_SPI_MEM_SUPPORT_WRAP

SOC_MEMSPI_SRC_FREQ_80M_SUPPORTED

SOC_MEMSPI_SRC_FREQ_40M_SUPPORTED

SOC_MEMSPI_SRC_FREQ_20M_SUPPORTED

SOC_SYSTIMER_COUNTER_NUM

SOC_SYSTIMER_ALARM_NUM

SOC_SYSTIMER_BIT_WIDTH_LO

SOC_SYSTIMER_BIT_WIDTH_HI

SOC_SYSTIMER_FIXED_DIVIDER

SOC_SYSTIMER_SUPPORT_RC_FAST

SOC_SYSTIMER_INT_LEVEL

SOC_SYSTIMER_ALARM_MISS_COMPENSATE

SOC_LP_TIMER_BIT_WIDTH_LO

SOC_LP_TIMER_BIT_WIDTH_HI

SOC_TIMER_GROUPS

SOC_TIMER_GROUP_TIMERS_PER_GROUP

SOC_TIMER_GROUP_COUNTER_BIT_WIDTH

SOC_TIMER_GROUP_SUPPORT_XTAL

SOC_TIMER_GROUP_SUPPORT_RC_FAST

SOC_TIMER_GROUP_TOTAL_TIMERS

SOC_TIMER_SUPPORT_ETM

SOC_MWDT_SUPPORT_XTAL

SOC_TWAI_CONTROLLER_NUM

SOC_TWAI_CLK_SUPPORT_XTAL

SOC_TWAI_BRP_MIN

SOC_TWAI_BRP_MAX

SOC_TWAI_SUPPORTS_RX_STATUS

SOC_EFUSE_DIS_DOWNLOAD_ICACHE

SOC_EFUSE_DIS_PAD_JTAG

SOC_EFUSE_DIS_USB_JTAG

SOC_EFUSE_DIS_DIRECT_BOOT

SOC_EFUSE_SOFT_DIS_JTAG

SOC_SECURE_BOOT_V2_RSA

SOC_SECURE_BOOT_V2_ECC

SOC_EFUSE_SECURE_BOOT_KEY_DIGESTS

SOC_EFUSE_REVOKE_BOOT_KEY_DIGESTS

SOC_SUPPORT_SECURE_BOOT_REVOKE_KEY

SOC_FLASH_ENCRYPTED_XTS_AES_BLOCK_MAX

SOC_FLASH_ENCRYPTION_XTS_AES

SOC_FLASH_ENCRYPTION_XTS_AES_128

SOC_UART_NUM

SOC_UART_HP_NUM

SOC_UART_LP_NUM

SOC_UART_FIFO_LEN

The UART hardware FIFO length

SOC_LP_UART_FIFO_LEN

The LP UART hardware FIFO length

SOC_UART_BITRATE_MAX

Max bit rate supported by UART

SOC_UART_SUPPORT_PLL_F80M_CLK

Support PLL_F80M as the clock source

SOC_UART_SUPPORT_RTC_CLK

Support RTC clock as the clock source

SOC_UART_SUPPORT_XTAL_CLK

Support XTAL clock as the clock source

SOC_UART_SUPPORT_WAKEUP_INT

Support UART wakeup interrupt

SOC_UART_HAS_LP_UART

Support LP UART

SOC_UART_SUPPORT_FSM_TX_WAIT_SEND

SOC_COEX_HW_PTI

SOC_PHY_DIG_REGS_MEM_SIZE

SOC_WIFI_LIGHT_SLEEP_CLK_WIDTH

SOC_PM_SUPPORT_WIFI_WAKEUP

SOC_PM_SUPPORT_CPU_PD

SOC_PM_SUPPORT_MODEM_PD

SOC_PM_SUPPORT_XTAL32K_PD

SOC_PM_SUPPORT_RC32K_PD

SOC_PM_SUPPORT_RC_FAST_PD

SOC_PM_SUPPORT_VDDSDIO_PD

SOC_PM_SUPPORT_TOP_PD

SOC_PM_SUPPORT_DEEPSLEEP_CHECK_STUB_ONLY

Supports CRC only the stub code in RTC memory

SOC_PM_CPU_RETENTION_BY_SW

SOC_PM_PAU_LINK_NUM

SOC_CLK_RC_FAST_SUPPORT_CALIBRATION

SOC_MODEM_CLOCK_IS_INDEPENDENT

SOC_CLK_APLL_SUPPORTED

Support Audio PLL

SOC_CLK_XTAL32K_SUPPORTED

Support to connect an external low frequency crystal

SOC_CLK_OSC_SLOW_SUPPORTED

Support to connect an external oscillator, not a crystal

SOC_CLK_RC32K_SUPPORTED

Support an internal 32kHz RC oscillator

SOC_PERIPH_CLK_CTRL_SHARED

Peripheral clock control (e.g. set clock source) is shared between various peripherals

SOC_TEMPERATURE_SENSOR_SUPPORT_FAST_RC

`SOC_TEMPERATURE_SENSOR_SUPPORT_XTAL`

`SOC_MEM_TCM_SUPPORTED`

2.9.30 System Time

Overview

ESP32-P4 uses two hardware timers for the purpose of keeping system time. System time can be kept by using either one or both of the hardware timers depending on the application's purpose and accuracy requirements for system time. The two hardware timers are:

- **RTC timer:** This timer allows time keeping in various sleep modes, and can also persist time keeping across any resets (with the exception of power-on resets which reset the RTC timer). The frequency deviation depends on the [RTC Timer Clock Sources](#) and affects the accuracy only in sleep modes, in which case the time will be measured at 6.6667 μ s resolution.
- **High-resolution timer:** This timer is not available in sleep modes and will not persist over a reset, but has greater accuracy. The timer uses the APB_CLK clock source (typically 80 MHz), which has a frequency deviation of less than ± 10 ppm. Time will be measured at 1 μ s resolution.

The possible combinations of hardware timers used to keep system time are listed below:

- RTC and high-resolution timer (default)
- RTC
- High-resolution timer
- None

It is recommended that users stick to the default option as it provides the highest accuracy. However, users can also select a different setting via the [CONFIG_NEWLIB_TIME_SYSCALL](#) configuration option.

RTC Timer Clock Sources

The RTC timer has the following clock sources:

- `Internal 150 kHz RC oscillator (default):` Features the lowest Deep-sleep current consumption and no dependence on any external components. However, the frequency stability of this clock source is affected by temperature fluctuations, so time may drift in both Deep-sleep and Light-sleep modes.
- `External 32 kHz crystal:` Requires a 32 kHz crystal to be connected to the XTAL_32K_P and XTAL_32K_N pins. This source provides a better frequency stability at the expense of a slightly higher (by 1 μ A) Deep-sleep current consumption.
- `External 32 kHz oscillator at XTAL_32K_P pin:` Allows using 32 kHz clock generated by an external circuit. The external clock signal must be connected to the XTAL_32K_P pin. The amplitude should be less than 1.2 V for sine wave signal and less than 1 V for square wave signal. Common mode voltage should be in the range of $0.1 < V_{cm} < 0.5 \times V_{amp}$, where V_{amp} stands for signal amplitude. In this case, the XTAL_32K_P pin cannot be used as a GPIO pin.

The choice depends on your requirements for system time accuracy and power consumption in sleep modes. To modify the RTC clock source, set [CONFIG_RTC_CLK_SRC](#) in project configuration.

More details about the wiring requirements for the external crystal or external oscillator, please refer to [ESP32-P4 Hardware Design Guidelines](#).

Get Current Time

To get the current time, use the POSIX function `gettimeofday()`. Additionally, you can use the following standard C library functions to obtain time and manipulate it:

```
gettimeofday
time
asctime
clock
ctime
difftime
gmtime
localtime
mktime
strftime
adjtime*
```

To stop smooth time adjustment and update the current time immediately, use the POSIX function `settimeofday()`.

If you need to obtain time with one second resolution, use the following code snippet:

```
time_t now;
char strftime_buf[64];
struct tm timeinfo;

time(&now);
// Set timezone to China Standard Time
setenv("TZ", "CST-8", 1);
tzset();

localtime_r(&now, &timeinfo);
strftime(strftime_buf, sizeof(strftime_buf), "%c", &timeinfo);
ESP_LOGI(TAG, "The current date/time in Shanghai is: %s", strftime_buf);
```

If you need to obtain time with one microsecond resolution, use the code snippet below:

```
struct timeval tv_now;
gettimeofday(&tv_now, NULL);
int64_t time_us = (int64_t)tv_now.tv_sec * 1000000L + (int64_t)tv_now.tv_usec;
```

SNTP Time Synchronization

To set the current time, you can use the POSIX functions `settimeofday()` and `adjtime()`. They are used internally in the lwIP SNTP library to set current time when a response from the NTP server is received. These functions can also be used separately from the lwIP SNTP library.

Some lwIP APIs, including SNTP functions, are not thread safe, so it is recommended to use [esp_netif component](#) when interacting with SNTP module.

To initialize a particular SNTP server and also start the SNTP service, simply create a default SNTP server configuration with a particular server name, then call [esp_netif_sntp_init\(\)](#) to register that server and start the SNTP service.

```
esp_sntp_config_t config = ESP_NETIF_SNTP_DEFAULT_CONFIG("pool.ntp.org");
esp_netif_sntp_init(&config);
```

This code automatically performs time synchronization once a reply from the SNTP server is received. Sometimes it is useful to wait until the time gets synchronized, [esp_netif_sntp_sync_wait\(\)](#) can be used for this purpose:

```
if (esp_netif_sntp_sync_wait(pdMS_TO_TICKS(10000)) != ESP_OK) {
    printf("Failed to update system time within 10s timeout");
}
```

To configure multiple NTP servers (or use more advanced settings, such as DHCP provided NTP servers), please refer to the detailed description of *SNTP API* in *esp_netif* documentation.

The lwIP SNTP library could work in one of the following sync modes:

- *SNTP_SYNC_MODE_IMMED* (default): Updates system time immediately upon receiving a response from the SNTP server after using `settimeofday()`.
- *SNTP_SYNC_MODE_SMOOTH*: Updates time smoothly by gradually reducing time error using the function `adjtime()`. If the difference between the SNTP response time and system time is more than 35 minutes, update system time immediately by using `settimeofday()`.

If you want to choose the *SNTP_SYNC_MODE_SMOOTH* mode, please set the `esp_sntp_config::smooth` to `true` in the SNTP configuration struct. Otherwise (and by default) the *SNTP_SYNC_MODE_IMMED* mode will be used.

For setting a callback function that is called when time gets synchronized, use the `esp_sntp_config::sync_cb` field in the configuration struct.

An application with this initialization code periodically synchronizes the time. The time synchronization period is determined by *CONFIG_LWIP_SNTP_UPDATE_DELAY* (the default value is one hour). To modify the variable, set *CONFIG_LWIP_SNTP_UPDATE_DELAY* in project configuration.

A code example that demonstrates the implementation of time synchronization based on the lwIP SNTP library is provided in the `protocols/sntp` directory.

Note that it is also possible to use lwIP API directly, but care must be taken to thread safety. Here we list the thread-safe APIs:

- *sntp_set_time_sync_notification_cb()* can be used to set a callback function that notifies of the time synchronization process.
- *sntp_get_sync_status()* and *sntp_set_sync_status()* can be used to get/set time synchronization status.
- *sntp_set_sync_mode()* can be used to set the synchronization mode.
- *esp_sntp_setoperatingmode()* sets the preferred operating mode.:cpp:enumerator:*ESP_SNTP_OPMODE_POLL* and *esp_sntp_init()* initializes SNTP module.
- *esp_sntp_setservername()* configures one SNTP server.

Timezones

To set the local timezone, use the following POSIX functions:

1. Call `setenv()` to set the `TZ` environment variable to the correct value based on the device location. The format of the time string is the same as described in the [GNU libc documentation](#) (although the implementation is different).
2. Call `tzset()` to update C library runtime data for the new timezone.

Once these steps are completed, call the standard C library function `localtime()`, and it returns the correct local time taking into account the timezone offset and daylight saving time.

Year 2036 and 2038 Overflow Issues

SNTP/NTP 2036 Overflow SNTP/NTP timestamps are represented as 64-bit unsigned fixed point numbers, where the first 32 bits represent the integer part, and the last 32 bits represent the fractional part. The 64-bit unsigned fixed point number represents the number of seconds since 00:00 on 1st of January 1900, thus SNTP/NTP times will overflow in the year 2036.

To address this issue, lifetime of the SNTP/NTP timestamps has been extended by convention by using the MSB (bit 0 by convention) of the integer part to indicate time ranges between years 1968 to 2104 (see [RFC2030](#) for

more details). This convention is implemented in lwIP library SNTP module. Therefore SNTP-related functions in ESP-IDF are future-proof until year 2104.

Unix Time 2038 Overflow Unix time (type `time_t`) was previously represented as a 32-bit signed integer, leading to an overflow in year 2038 (i.e., [Y2K38 issue](#)). To address the Y2K38 issue, ESP-IDF uses a 64-bit signed integer to represent `time_t` starting from release v5.0, thus deferring `time_t` overflow for another 292 billion years.

API Reference

Header File

- [components/lwip/include/apps/esp_sntp.h](#)
- This header file can be included with:

```
#include "esp_sntp.h"
```

- This header file is a part of the API provided by the `lwip` component. To declare that your component depends on `lwip`, add the following to your `CMakeLists.txt`:

```
REQUIRES lwip
```

or

```
PRIV_REQUIRES lwip
```

Functions

void **sntp_sync_time** (struct timeval *tv)

This function updates the system time.

This is a weak-linked function. It is possible to replace all SNTP update functionality by placing a `sntp_sync_time()` function in the app firmware source. If the default implementation is used, calling `sntp_set_sync_mode()` allows the time synchronization mode to be changed to instant or smooth. If a callback function is registered via `sntp_set_time_sync_notification_cb()`, it will be called following time synchronization.

Parameters `tv` -- Time received from SNTP server.

void **sntp_set_sync_mode** ([sntp_sync_mode_t](#) sync_mode)

Set the sync mode.

Modes allowed: `SNTP_SYNC_MODE_IMMED` and `SNTP_SYNC_MODE_SMOOTH`.

Parameters `sync_mode` -- Sync mode.

[sntp_sync_mode_t](#) **sntp_get_sync_mode** (void)

Get set sync mode.

Returns `SNTP_SYNC_MODE_IMMED`: Update time immediately.
`SNTP_SYNC_MODE_SMOOTH`: Smooth time updating.

[sntp_sync_status_t](#) **sntp_get_sync_status** (void)

Get status of time sync.

After the update is completed, the status will be returned as `SNTP_SYNC_STATUS_COMPLETED`. After that, the status will be reset to `SNTP_SYNC_STATUS_RESET`. If the update operation is not completed yet, the status will be `SNTP_SYNC_STATUS_RESET`. If a smooth mode was chosen and the synchronization is still continuing (adjtime works), then it will be `SNTP_SYNC_STATUS_IN_PROGRESS`.

Returns `SNTP_SYNC_STATUS_RESET`: Reset status. `SNTP_SYNC_STATUS_COMPLETED`: Time is synchronized. `SNTP_SYNC_STATUS_IN_PROGRESS`: Smooth time sync in progress.

void **sntp_set_sync_status** (*sntp_sync_status_t* sync_status)

Set status of time sync.

Parameters **sync_status** -- status of time sync (see sntp_sync_status_t)

void **sntp_set_time_sync_notification_cb** (*sntp_sync_time_cb_t* callback)

Set a callback function for time synchronization notification.

Parameters **callback** -- a callback function

void **sntp_set_sync_interval** (uint32_t interval_ms)

Set the sync interval of SNTP operation.

Note: SNTPv4 RFC 4330 enforces a minimum sync interval of 15 seconds. This sync interval will be used in the next attempt update time through SNTP. To apply the new sync interval call the sntp_restart() function, otherwise, it will be applied after the last interval expired.

Parameters **interval_ms** -- The sync interval in ms. It cannot be lower than 15 seconds, otherwise 15 seconds will be set.

uint32_t **sntp_get_sync_interval** (void)

Get the sync interval of SNTP operation.

Returns the sync interval

bool **sntp_restart** (void)

Restart SNTP.

Returns True - Restart False - SNTP was not initialized yet

void **esp_sntp_setoperatingmode** (*esp_sntp_operatingmode_t* operating_mode)

Sets SNTP operating mode. The mode has to be set before init.

Parameters **operating_mode** -- Desired operating mode

void **esp_sntp_init** (void)

Init and start SNTP service.

void **esp_sntp_stop** (void)

Stops SNTP service.

void **esp_sntp_setserver** (u8_t idx, const ip_addr_t *addr)

Sets SNTP server address.

Parameters

- **idx** -- Index of the server
- **addr** -- IP address of the server

void **esp_sntp_setservername** (u8_t idx, const char *server)

Sets SNTP hostname.

Parameters

- **idx** -- Index of the server
- **server** -- Name of the server

const char ***esp_sntp_getservername** (u8_t idx)

Gets SNTP server name.

Parameters **idx** -- Index of the server

Returns Name of the server

const ip_addr_t ***esp_sntp_getserver** (u8_t idx)

Get SNTP server IP.

Parameters **idx** -- Index of the server

Returns IP address of the server

bool **esp_sntp_enabled** (void)

Checks if sntp is enabled.

Returns true if sntp module is enabled

uint8_t **esp_sntp_getreachability** (uint8_t idx)

Gets the server reachability shift register as described in RFC 5905.

Parameters **idx** -- Index of the SNTP server

Returns reachability shift register

esp_sntp_operatingmode_t **esp_sntp_getoperatingmode** (void)

Get the configured operating mode.

Returns operating mode enum

static inline void **sntp_setoperatingmode** (u8_t operating_mode)

if not build within lwip, provide translating inlines, that will warn about thread safety

static inline void **sntp_servermode_dhcp** (int set_servers_from_dhcp)

static inline void **sntp_setservername** (u8_t idx, const char *server)

static inline void **sntp_init** (void)

static inline const char ***sntp_getservername** (u8_t idx)

static inline const ip_addr_t ***sntp_getserver** (u8_t idx)

static inline uint8_t **sntp_getreachability** (uint8_t idx)

static inline *esp_sntp_operatingmode_t* **sntp_getoperatingmode** (void)

Macros

esp_sntp_sync_time

Aliases for esp_sntp prefixed API (inherently thread safe)

esp_sntp_set_sync_mode

esp_sntp_get_sync_mode

esp_sntp_get_sync_status

esp_sntp_set_sync_status

esp_sntp_set_time_sync_notification_cb

esp_sntp_set_sync_interval

esp_sntp_get_sync_interval

esp_sntp_restart

SNTP_OPMODE_POLL

Type Definitions

typedef void (***sntp_sync_time_cb_t**)(struct timeval *tv)

SNTP callback function for notifying about time sync event.

Param tv Time received from SNTP server.

Enumerations

enum **sntp_sync_mode_t**

SNTP time update mode.

Values:

enumerator **SNTP_SYNC_MODE_IMMED**

Update system time immediately when receiving a response from the SNTP server.

enumerator **SNTP_SYNC_MODE_SMOOTH**

Smooth time updating. Time error is gradually reduced using adjtime function. If the difference between SNTP response time and system time is large (more than 35 minutes) then update immediately.

enum **sntp_sync_status_t**

SNTP sync status.

Values:

enumerator **SNTP_SYNC_STATUS_RESET**

enumerator **SNTP_SYNC_STATUS_COMPLETED**

enumerator **SNTP_SYNC_STATUS_IN_PROGRESS**

enum **esp_sntp_operatingmode_t**

SNTP operating modes per lwip SNTP module.

Values:

enumerator **ESP_SNTP_OPMODE_POLL**

enumerator **ESP_SNTP_OPMODE_LISTENONLY**

2.9.31 Asynchronous Memory Copy

Overview

ESP32-P4 has a DMA engine which can help to offload internal memory copy operations from the CPU in an asynchronous way.

The async memcpy API wraps all DMA configurations and operations. The signature of [*esp_async_memcpy\(\)*](#) is almost the same as the standard libc `memcpy` function.

The DMA allows multiple memory copy requests to be queued up before the first one is completed, which allows overlap of computation and memory copy. Moreover, it is still possible to know the exact time when a memory copy request is completed by registering an event callback.

If the async memcopy is constructed upon the AXI GDMA, it is also possible to copy data from/to PSRAM with a proper alignment.

Configure and Install Driver

There are several ways to install the async memcopy driver, depending on the underlying DMA engine:

- `esp_async_memcopy_install_gdma_ahb()` is used to install the async memcopy driver based on the AHB GDMA engine.
- `esp_async_memcopy_install_gdma_axi()` is used to install the async memcopy driver based on the AXI GDMA engine.
- `esp_async_memcopy_install()` is a generic API to install the async memcopy driver with a default DMA engine. If the SoC has the CP DMA engine, the default DMA engine is CP DMA. Otherwise, the default DMA engine is AHB GDMA.

Driver configuration is described in `async_memcopy_config_t`:

- `backlog`: This is used to configure the maximum number of memory copy transactions that can be queued up before the first one is completed. If this field is set to zero, then the default value 4 will be applied.
- `sram_trans_align`: Declare SRAM alignment for both data address and copy size, set to zero if the data has no restriction in alignment. If set to a quadruple value (i.e., 4X), the driver will enable the burst mode internally, which is helpful for some performance related application.
- `psram_trans_align`: Declare PSRAM alignment for both data address and copy size. User has to give it a valid value (only 16, 32, 64 are supported) if the destination of memcopy is located in PSRAM. The default alignment (i.e., 16) will be applied if it is set to zero. Internally, the driver configures the size of block used by DMA to access PSRAM, according to the alignment.
- `flags`: This is used to enable some special driver features.

```
async_memcopy_config_t config = ASYNC_MEMCOPY_DEFAULT_CONFIG();
// update the maximum data stream supported by underlying DMA engine
config.backlog = 8;
async_memcopy_handle_t driver = NULL;
ESP_ERROR_CHECK(esp_async_memcopy_install(&config, &driver)); // install driver
↳with default DMA engine
```

Send Memory Copy Request

`esp_async_memcopy()` is the API to send memory copy request to DMA engine. It must be called after driver is installed successfully. This API is thread safe, so it can be called from different tasks.

Different from the libc version of `memcpy`, you can optionally pass a callback to `esp_async_memcopy()`, so that you can be notified when the memory copy is finished. Note that the callback is executed in the ISR context, please make sure you will not call any blocking functions in the callback.

The prototype of the callback function is `async_memcopy_isr_cb_t`. The callback function should only return true if it wakes up a high priority task by RTOS APIs like `xSemaphoreGiveFromISR()`.

```
// Callback implementation, running in ISR context
static bool my_async_memcopy_cb(async_memcopy_handle_t mcp_hdl, async_memcopy_event_t
↳*event, void *cb_args)
{
    SemaphoreHandle_t sem = (SemaphoreHandle_t)cb_args;
    BaseType_t high_task_wakeup = pdFALSE;
    xSemaphoreGiveFromISR(semphr, &high_task_wakeup); // high_task_wakeup set to
↳pdTRUE if some high priority task unblocked
```

(continues on next page)

(continued from previous page)

```

    return high_task_wakeup == pdTRUE;
}

// Create a semaphore used to report the completion of async memcpy
SemaphoreHandle_t semphr = xSemaphoreCreateBinary();

// Called from user's context
ESP_ERROR_CHECK(esp_async_memcpy(driver_handle, to, from, copy_len, my_async_
↪memcpy_cb, my_semaphore));
// Do something else here
xSemaphoreTake(my_semaphore, portMAX_DELAY); // Wait until the buffer copy is done

```

Uninstall Driver

`esp_async_memcpy_uninstall()` is used to uninstall asynchronous memcpy driver. It is not necessary to uninstall the driver after each memcpy operation. If you know your application will not use this driver anymore, then this API can recycle the memory and other hardware resources for you.

API Reference

Header File

- [components/esp_hw_support/include/esp_async_memcpy.h](#)
- This header file can be included with:

```
#include "esp_async_memcpy.h"
```

Functions

`esp_err_t esp_async_memcpy_install_gdma_ahb` (const [async_memcpy_config_t](#) *config, [async_memcpy_handle_t](#) *mcp)

Install async memcpy driver, with AHB-GDMA as the backend.

Parameters

- **config** -- [in] Configuration of async memcpy
- **mcp** -- [out] Returned driver handle

Returns

- ESP_OK: Install async memcpy driver successfully
- ESP_ERR_INVALID_ARG: Install async memcpy driver failed because of invalid argument
- ESP_ERR_NO_MEM: Install async memcpy driver failed because out of memory
- ESP_FAIL: Install async memcpy driver failed because of other error

`esp_err_t esp_async_memcpy_install_gdma_axi` (const [async_memcpy_config_t](#) *config, [async_memcpy_handle_t](#) *mcp)

Install async memcpy driver, with AXI-GDMA as the backend.

Parameters

- **config** -- [in] Configuration of async memcpy
- **mcp** -- [out] Returned driver handle

Returns

- ESP_OK: Install async memcpy driver successfully
- ESP_ERR_INVALID_ARG: Install async memcpy driver failed because of invalid argument
- ESP_ERR_NO_MEM: Install async memcpy driver failed because out of memory
- ESP_FAIL: Install async memcpy driver failed because of other error

esp_err_t **esp_async_memcpy_install** (const *async_memcpy_config_t* *config, *async_memcpy_handle_t* *mcp)

Install async memcpy driver with the default DMA backend.

Note: On chip with CPDMA support, CPDMA is the default choice. On chip with AHB-GDMA support, AHB-GDMA is the default choice.

Parameters

- **config** -- [in] Configuration of async memcpy
- **mcp** -- [out] Returned driver handle

Returns

- ESP_OK: Install async memcpy driver successfully
- ESP_ERR_INVALID_ARG: Install async memcpy driver failed because of invalid argument
- ESP_ERR_NO_MEM: Install async memcpy driver failed because out of memory
- ESP_FAIL: Install async memcpy driver failed because of other error

esp_err_t **esp_async_memcpy_uninstall** (*async_memcpy_handle_t* mcp)

Uninstall async memcpy driver.

Parameters **mcp** -- [in] Handle of async memcpy driver that returned from `esp_async_memcpy_install`

Returns

- ESP_OK: Uninstall async memcpy driver successfully
- ESP_ERR_INVALID_ARG: Uninstall async memcpy driver failed because of invalid argument
- ESP_FAIL: Uninstall async memcpy driver failed because of other error

esp_err_t **esp_async_memcpy** (*async_memcpy_handle_t* mcp, void *dst, void *src, size_t n, *async_memcpy_isr_cb_t* cb_isr, void *cb_args)

Send an asynchronous memory copy request.

Note: The callback function is invoked in interrupt context, never do blocking jobs in the callback.

Parameters

- **mcp** -- [in] Handle of async memcpy driver that returned from `esp_async_memcpy_install`
- **dst** -- [in] Destination address (copy to)
- **src** -- [in] Source address (copy from)
- **n** -- [in] Number of bytes to copy
- **cb_isr** -- [in] Callback function, which got invoked in interrupt context. Set to NULL can bypass the callback.
- **cb_args** -- [in] User defined argument to be passed to the callback function

Returns

- ESP_OK: Send memory copy request successfully
- ESP_ERR_INVALID_ARG: Send memory copy request failed because of invalid argument
- ESP_FAIL: Send memory copy request failed because of other error

Structures

struct **async_memcpy_event_t**

Async memory copy event data.

Public Members

void ***data**

Event data

struct **async_memcpy_config_t**

Type of async memcpy configuration.

Public Members

uint32_t **backlog**

Maximum number of transactions that can be prepared in the background

size_t **sram_trans_align**

DMA transfer alignment (both in size and address) for SRAM memory

size_t **psram_trans_align**

DMA transfer alignment (both in size and address) for PSRAM memory

uint32_t **flags**

Extra flags to control async memcpy feature

Macros

ASYNC_MEMCPY_DEFAULT_CONFIG()

Default configuration for async memcpy.

Type Definitions

typedef struct async_memcpy_context_t ***async_memcpy_handle_t**

Async memory copy driver handle.

typedef bool (***async_memcpy_isr_cb_t**)(*async_memcpy_handle_t* mcp_hdl, *async_memcpy_event_t* *event, void *cb_args)

Type of async memcpy interrupt callback function.

Note: User can call OS primitives (semaphore, mutex, etc) in the callback function. Keep in mind, if any OS primitive wakes high priority task up, the callback should return true.

Param mcp_hdl Handle of async memcpy

Param event Event object, which contains related data, reserved for future

Param cb_args User defined arguments, passed from esp_async_memcpy function

Return Whether a high priority task is woken up by the callback function

2.9.32 Watchdogs

Overview

ESP-IDF supports multiple types of watchdogs:

- Interrupt Watchdog Timer (IWDT)
- Task Watchdog Timer (TWDT)

The Interrupt Watchdog is responsible for ensuring that ISRs (Interrupt Service Routines) are not blocked for a prolonged period of time. The TWDT is responsible for detecting instances of tasks running without yielding for a prolonged period.

The various watchdog timers can be enabled using the [Project Configuration Menu](#). However, the TWDT can also be enabled during runtime.

Interrupt Watchdog Timer (IWDT)

The purpose of the IWDT is to ensure that interrupt service routines (ISRs) are not blocked from running for a prolonged period of time (i.e., the IWDT timeout period). Preventing ISRs from running in a timely manner is undesirable as it can increase ISR latency, and also prevent task switching (as task switching is executed from an ISR). The things that can block ISRs from running include:

- Disabling interrupts
- Critical Sections (also disables interrupts)
- Other same/higher priority ISRs which block same/lower priority ISRs from running

The IWDT utilizes the watchdog timer in Timer Group 1 as its underlying hardware timer and leverages the FreeRTOS tick interrupt on each CPU to feed the watchdog timer. If the tick interrupt on a particular CPU is not run at within the IWDT timeout period, it is indicative that something is blocking ISRs from being run on that CPU (see the list of reasons above).

When the IWDT times out, the default action is to invoke the panic handler and display the panic reason as `Interrupt wdt timeout on CPU0` or `Interrupt wdt timeout on CPU1` (as applicable). Depending on the panic handler's configured behavior (see [CONFIG_ESP_SYSTEM_PANIC](#)), users can then debug the source of the IWDT timeout (via the backtrace, OpenOCD, gdbstub etc) or simply reset the chip (which may be preferred in a production environment).

If for whatever reason the panic handler is unable to run after an IWDT timeout, the IWDT has a second stage timeout that will hard-reset the chip (i.e., a system reset).

Configuration

- The IWDT is enabled by default via the [CONFIG_ESP_INT_WDT](#) option.
- The IWDT's timeout is configured by setting the [CONFIG_ESP_INT_WDT_TIMEOUT_MS](#) option.
 - Note that the default timeout is higher if PSRAM support is enabled, as a critical section or interrupt routine that accesses a large amount of PSRAM takes longer to complete in some circumstances.
 - The timeout should always be at least twice longer than the period between FreeRTOS ticks (see [CONFIG_FREERTOS_HZ](#)).

Tuning If you find the IWDT timeout is triggered because an interrupt or critical section is running longer than the timeout period, consider rewriting the code:

- Critical sections should be made as short as possible. Any non-critical code/computation should be placed outside the critical section.
- Interrupt handlers should also perform the minimum possible amount of computation. Users can consider deferring any computation to a task by having the ISR push data to a task using queues.

Neither critical sections or interrupt handlers should ever block waiting for another event to occur. If changing the code to reduce the processing time is not possible or desirable, it is possible to increase the [CONFIG_ESP_INT_WDT_TIMEOUT_MS](#) setting instead.

Task Watchdog Timer (TWDT)

The Task Watchdog Timer (TWDT) is used to monitor particular tasks, ensuring that they are able to execute within a given timeout period. The TWDT primarily watches the Idle Tasks of each CPU, however any task can subscribe to be watched by the TWDT. By watching the Idle Tasks of each CPU, the TWDT can detect instances of tasks running for a prolonged period of time without yielding. This can be an indicator of poorly written code that spins on a peripheral, or a task that is stuck in an infinite loop.

The TWDT is built around the Hardware Watchdog Timer in Timer Group 0. When a timeout occurs, an interrupt is triggered.

Users can define the function `esp_task_wdt_isr_user_handler` in the user code, in order to receive the timeout event and extend the default behavior.

Usage The following functions can be used to watch tasks using the TWDT:

- `esp_task_wdt_init()` to initialize the TWDT and subscribe the idle tasks.
- `esp_task_wdt_add()` subscribes other tasks to the TWDT.
- Once subscribed, `esp_task_wdt_reset()` should be called from the task to feed the TWDT.
- `esp_task_wdt_delete()` unsubscribes a previously subscribed task.
- `esp_task_wdt_deinit()` unsubscribes the idle tasks and deinitializes the TWDT.

In the case where applications need to watch at a more granular level (i.e., ensure that a particular functions/stub/code-path is called), the TWDT allows subscription of `users`.

- `esp_task_wdt_add_user()` to subscribe an arbitrary user of the TWDT. This function returns a user handle to the added user.
- `esp_task_wdt_reset_user()` must be called using the user handle in order to prevent a TWDT timeout.
- `esp_task_wdt_delete_user()` unsubscribes an arbitrary user of the TWDT.

Configuration The default timeout period for the TWDT is set using config item `CONFIG_ESP_TASK_WDT_TIMEOUT_S`. This should be set to at least as long as you expect any single task needs to monopolize the CPU (for example, if you expect the app will do a long intensive calculation and should not yield to other tasks). It is also possible to change this timeout at runtime by calling `esp_task_wdt_init()`.

Note: Erasing large flash areas can be time consuming and can cause a task to run continuously, thus triggering a TWDT timeout. The following two methods can be used to avoid this:

- Increase `CONFIG_ESP_TASK_WDT_TIMEOUT_S` in menuconfig for a larger watchdog timeout period.
- You can also call `esp_task_wdt_init()` to increase the watchdog timeout period before erasing a large flash area.

For more information, you can refer to [SPI Flash API](#).

The following config options control TWDT configuration. They are all enabled by default:

- `CONFIG_ESP_TASK_WDT_EN` - enables TWDT feature. If this option is disabled, TWDT cannot be used, even if initialized at runtime.
- `CONFIG_ESP_TASK_WDT_INIT` - the TWDT is initialized automatically during startup. If this option is disabled, it is still possible to initialize the Task WDT at runtime by calling `esp_task_wdt_init()`.
- `CONFIG_ESP_TASK_WDT_CHECK_IDLE_TASK_CPU0` - CPU0 Idle task is subscribed to the TWDT during startup. If this option is disabled, it is still possible to subscribe the idle task by calling `esp_task_wdt_init()` again.
- `CONFIG_ESP_TASK_WDT_CHECK_IDLE_TASK_CPU1` - CPU1 Idle task is subscribed to the TWDT during startup.

Note: On a TWDT timeout the default behaviour is to simply print a warning and a backtrace before continuing running the app. If you want a timeout to cause a panic and a system reset then this can be configured through `CONFIG_ESP_TASK_WDT_PANIC`.

JTAG & Watchdogs

While debugging using OpenOCD, the CPUs are halted every time a breakpoint is reached. However if the watchdog timers continue to run when a breakpoint is encountered, they will eventually trigger a reset making it very difficult to debug code. Therefore OpenOCD will disable the hardware timers of both the interrupt and task watchdogs at every breakpoint. Moreover, OpenOCD will not reenale them upon leaving the breakpoint. This means that interrupt watchdog and task watchdog functionality will essentially be disabled. No warnings or panics from either watchdogs will be generated when the ESP32-P4 is connected to OpenOCD via JTAG.

API Reference

Task Watchdog A full example using the Task Watchdog is available in esp-idf: [system/task_watchdog](#)

Header File

- `components/esp_system/include/esp_task_wdt.h`
- This header file can be included with:

```
#include "esp_task_wdt.h"
```

Functions

`esp_err_t esp_task_wdt_init` (const `esp_task_wdt_config_t` *config)

Initialize the Task Watchdog Timer (TWDT)

This function configures and initializes the TWDT. This function will subscribe the idle tasks if configured to do so. For other tasks, users can subscribe them using `esp_task_wdt_add()` or `esp_task_wdt_add_user()`. This function won't start the timer if no task have been registered yet.

Note: `esp_task_wdt_init()` must only be called after the scheduler is started. Moreover, it must not be called by multiple tasks simultaneously.

Parameters `config` -- [in] Configuration structure

Returns

- `ESP_OK`: Initialization was successful
- `ESP_ERR_INVALID_STATE`: Already initialized
- Other: Failed to initialize TWDT

`esp_err_t esp_task_wdt_reconfigure` (const `esp_task_wdt_config_t` *config)

Reconfigure the Task Watchdog Timer (TWDT)

The function reconfigures the running TWDT. It must already be initialized when this function is called.

Note: `esp_task_wdt_reconfigure()` must not be called by multiple tasks simultaneously.

Parameters `config` -- [in] Configuration structure

Returns

- `ESP_OK`: Reconfiguring was successful

- ESP_ERR_INVALID_STATE: TWDT not initialized yet
- Other: Failed to initialize TWDT

esp_err_t **esp_task_wdt_deinit** (void)

Deinitialize the Task Watchdog Timer (TWDT)

This function will deinitialize the TWDT, and unsubscribe any idle tasks. Calling this function whilst other tasks are still subscribed to the TWDT, or when the TWDT is already deinitialized, will result in an error code being returned.

Note: `esp_task_wdt_deinit()` must not be called by multiple tasks simultaneously.

Returns

- ESP_OK: TWDT successfully deinitialized
- Other: Failed to deinitialize TWDT

esp_err_t **esp_task_wdt_add** (*TaskHandle_t* task_handle)

Subscribe a task to the Task Watchdog Timer (TWDT)

This function subscribes a task to the TWDT. Each subscribed task must periodically call `esp_task_wdt_reset()` to prevent the TWDT from elapsing its timeout period. Failure to do so will result in a TWDT timeout.

Parameters **task_handle** -- Handle of the task. Input NULL to subscribe the current running task to the TWDT

Returns

- ESP_OK: Successfully subscribed the task to the TWDT
- Other: Failed to subscribe task

esp_err_t **esp_task_wdt_add_user** (const char *user_name, *esp_task_wdt_user_handle_t* *user_handle_ret)

Subscribe a user to the Task Watchdog Timer (TWDT)

This function subscribes a user to the TWDT. A user of the TWDT is usually a function that needs to run periodically. Each subscribed user must periodically call `esp_task_wdt_reset_user()` to prevent the TWDT from elapsing its timeout period. Failure to do so will result in a TWDT timeout.

Parameters

- **user_name** -- [in] String to identify the user
- **user_handle_ret** -- [out] Handle of the user

Returns

- ESP_OK: Successfully subscribed the user to the TWDT
- Other: Failed to subscribe user

esp_err_t **esp_task_wdt_reset** (void)

Reset the Task Watchdog Timer (TWDT) on behalf of the currently running task.

This function will reset the TWDT on behalf of the currently running task. Each subscribed task must periodically call this function to prevent the TWDT from timing out. If one or more subscribed tasks fail to reset the TWDT on their own behalf, a TWDT timeout will occur.

Returns

- ESP_OK: Successfully reset the TWDT on behalf of the currently running task
- Other: Failed to reset

esp_err_t **esp_task_wdt_reset_user** (*esp_task_wdt_user_handle_t* user_handle)

Reset the Task Watchdog Timer (TWDT) on behalf of a user.

This function will reset the TWDT on behalf of a user. Each subscribed user must periodically call this function to prevent the TWDT from timing out. If one or more subscribed users fail to reset the TWDT on their own behalf, a TWDT timeout will occur.

Parameters **user_handle** -- [in] User handle

- ESP_OK: Successfully reset the TWDT on behalf of the user
- Other: Failed to reset

esp_err_t **esp_task_wdt_delete** (*TaskHandle_t* task_handle)

Unsubscribes a task from the Task Watchdog Timer (TWDT)

This function will unsubscribe a task from the TWDT. After being unsubscribed, the task should no longer call `esp_task_wdt_reset()`.

Parameters **task_handle** -- [in] Handle of the task. Input NULL to unsubscribe the current running task.

Returns

- ESP_OK: Successfully unsubscribed the task from the TWDT
- Other: Failed to unsubscribe task

esp_err_t **esp_task_wdt_delete_user** (*esp_task_wdt_user_handle_t* user_handle)

Unsubscribes a user from the Task Watchdog Timer (TWDT)

This function will unsubscribe a user from the TWDT. After being unsubscribed, the user should no longer call `esp_task_wdt_reset_user()`.

Parameters **user_handle** -- [in] User handle

Returns

- ESP_OK: Successfully unsubscribed the user from the TWDT
- Other: Failed to unsubscribe user

esp_err_t **esp_task_wdt_status** (*TaskHandle_t* task_handle)

Query whether a task is subscribed to the Task Watchdog Timer (TWDT)

This function will query whether a task is currently subscribed to the TWDT, or whether the TWDT is initialized.

Parameters **task_handle** -- [in] Handle of the task. Input NULL to query the current running task.

Returns :

- ESP_OK: The task is currently subscribed to the TWDT
- ESP_ERR_NOT_FOUND: The task is not subscribed
- ESP_ERR_INVALID_STATE: TWDT was never initialized

void **esp_task_wdt_isr_user_handler** (void)

User ISR callback placeholder.

This function is called by `task_wdt_isr` function (ISR for when TWDT times out). It can be defined in user code to handle TWDT events.

Note: It has the same limitations as the interrupt function. Do not use ESP_LOGx functions inside.

esp_err_t **esp_task_wdt_print_triggered_tasks** (*task_wdt_msg_handler* msg_handler, void *opaque, int *cpus_fail)

Prints or retrieves information about tasks/users that triggered the Task Watchdog Timeout.

This function provides various operations to handle tasks/users that did not reset the Task Watchdog in time. It can print detailed information about these tasks/users, such as their names, associated CPUs, and whether they have been reset. Additionally, it can retrieve the total length of the printed information or the CPU affinity of the failing tasks.

Note:

- If `msg_handler` is not provided, the information will be printed to console using `ESP_EARLY_LOGE`.

- If `msg_handler` is provided, the function will send the printed information to the provided message handler function.
 - If `cpus_fail` is provided, the function will store the CPU affinity of the failing tasks in the provided integer.
 - During the execution of this function, logging is allowed in critical sections, as TWDT timeouts are considered fatal errors.
-

Parameters

- **`msg_handler`** -- [in] Optional message handler function that will be called for each printed line.
- **`opaque`** -- [in] Optional pointer to opaque data that will be passed to the message handler function.
- **`cpus_fail`** -- [out] Optional pointer to an integer where the CPU affinity of the failing tasks will be stored.

Returns

- `ESP_OK`: The function executed successfully.
- `ESP_FAIL`: No triggered tasks were found, and thus no information was printed or retrieved.

Structures

struct **`esp_task_wdt_config_t`**

Task Watchdog Timer (TWDT) configuration structure.

Public Members

uint32_t **`timeout_ms`**

TWDT timeout duration in milliseconds

uint32_t **`idle_core_mask`**

Bitmask of the core whose idle task should be subscribed on initialization where $1 \ll i$ means that core i 's idle task will be monitored by the TWDT

bool **`trigger_panic`**

Trigger panic when timeout occurs

Type Definitions

typedef struct esp_task_wdt_user_handle_s **`*esp_task_wdt_user_handle_t`**

Task Watchdog Timer (TWDT) user handle.

typedef void (**`*task_wdt_msg_handler`**)(void *opaque, const char *msg)

Code examples for this API section are provided in the [system](#) directory of ESP-IDF examples.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Chapter 3

Hardware Reference

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

Chapter 4

API Guides

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.1 Application Level Tracing Library

4.1.1 Overview

ESP-IDF provides a useful feature for program behavior analysis: application level tracing. It is implemented in the corresponding library and can be enabled in menuconfig. This feature allows to transfer arbitrary data between host and ESP32-P4 via JTAG, UART, or USB interfaces with small overhead on program execution. It is possible to use JTAG and UART interfaces simultaneously. The UART interface is mostly used for connection with SEGGER SystemView tool (see [SystemView](#)).

Developers can use this library to send application-specific state of execution to the host and receive commands or other types of information from the opposite direction at runtime. The main use cases of this library are:

1. Collecting application-specific data. See [Application Specific Tracing](#).
2. Lightweight logging to the host. See [Logging to Host](#).
3. System behavior analysis. See [System Behavior Analysis with SEGGER SystemView](#).
4. Source code coverage. See [Gcov \(Source Code Coverage\)](#).

Tracing components used when working over JTAG interface are shown in the figure below.

4.1.2 Modes of Operation

The library supports two modes of operation:

Post-mortem mode: This is the default mode. The mode does not need interaction with the host side. In this mode, tracing module does not check whether the host has read all the data from *HW UP BUFFER*, but directly overwrites old data with the new ones. This mode is useful when only the latest trace data is interesting to the user, e.g., for analyzing program's behavior just before the crash. The host can read the data later on upon user request, e.g., via special OpenOCD command in case of working via JTAG interface.

Streaming mode: Tracing module enters this mode when the host connects to ESP32-P4. In this mode, before writing new data to *HW UP BUFFER*, the tracing module checks that whether there is enough space in it and if




Fig. 1: Tracing Components Used When Working Over JTAG

necessary, waits for the host to read data and free enough memory. Maximum waiting time is controlled via timeout values passed by users to corresponding API routines. So when application tries to write data to the trace buffer using the finite value of the maximum waiting time, it is possible that this data will be dropped. This is especially true for tracing from time critical code (ISRs, OS scheduler code, etc.) where infinite timeouts can lead to system malfunction. In order to avoid loss of such critical data, developers can enable additional data buffering via menuconfig option `CONFIG_APPTRACE_PENDING_DATA_SIZE_MAX`. This macro specifies the size of data which can be buffered in above conditions. The option can also help to overcome situation when data transfer to the host is temporarily slowed down, e.g., due to USB bus congestions. But it will not help when the average bitrate of the trace data stream exceeds the hardware interface capabilities.

4.1.3 Configuration Options and Dependencies

Using of this feature depends on two components:

1. **Host side:** Application tracing is done over JTAG, so it needs OpenOCD to be set up and running on host machine. For instructions on how to set it up, please see [JTAG Debugging](#) for details.
2. **Target side:** Application tracing functionality can be enabled in menuconfig. Please go to Component config>Application Level Tracing menu, which allows selecting destination for the trace data (hardware interface for transport: JTAG or/and UART). Choosing any of the destinations automatically enables the `CONFIG_APPTRACE_ENABLE` option. For UART interfaces, users have to define baud rate, TX and RX pins numbers, and additional UART-related parameters.

Note: In order to achieve higher data rates and minimize the number of dropped packets, it is recommended to optimize the setting of JTAG clock frequency, so that it is at maximum and still provides stable operation of JTAG. See [Optimize JTAG Speed](#).

There are two additional menuconfig options not mentioned above:

1. *Threshold for flushing last trace data to host on panic* (`CONFIG_APPTRACE_POSTMORTEM_FLUSH_THRESH`). This option is necessary due to the nature of working over JTAG. In this mode, trace data is exposed to the

host in 16 KB blocks. In post-mortem mode, when one block is filled, it is exposed to the host and the previous one becomes unavailable. In other words, the trace data is overwritten in 16 KB granularity. On panic, the latest data from the current input block is exposed to the host and the host can read them for post-analysis. System panic may occur when a very small amount of data are not exposed to the host yet. In this case, the previous 16 KB of collected data will be lost and the host will see the latest, but very small piece of the trace. It can be insufficient to diagnose the problem. This menuconfig option allows avoiding such situations. It controls the threshold for flushing data in case of apanic. For example, users can decide that it needs no less than 512 bytes of the recent trace data, so if there is less than 512 bytes of pending data at the moment of panic, they will not be flushed and will not overwrite the previous 16 KB. The option is only meaningful in post-mortem mode and when working over JTAG.

2. *Timeout for flushing last trace data to host on panic* ([CONFIG_APPTRACE_ONPANIC_HOST_FLUSH_TMO](#)). The option is only meaningful in streaming mode and it controls the maximum time that the tracing module will wait for the host to read the last data in case of panic.
3. *UART RX/TX ring buffer size* ([CONFIG_APPTRACE_UART_TX_BUFF_SIZE](#)). The size of the buffer depends on the amount of data transferred through the UART.
4. *UART TX message size* ([CONFIG_APPTRACE_UART_TX_MSG_SIZE](#)). The maximum size of the single message to transfer.

4.1.4 How to Use This Library

This library provides APIs for transferring arbitrary data between the host and ESP32-P4. When enabled in menuconfig, the target application tracing module is initialized automatically at the system startup, so all what the user needs to do is to call corresponding APIs to send, receive or flush the data.

Application Specific Tracing

In general, users should decide what type of data should be transferred in every direction and how these data must be interpreted (processed). The following steps must be performed to transfer data between the target and the host:

1. On the target side, users should implement algorithms for writing trace data to the host. Piece of code below shows an example on how to do this.

```
#include "esp_app_trace.h"
...
char buf[] = "Hello World!";
esp_err_t res = esp_apptrace_write(ESP_APPTRACE_DEST_TRAX, buf,
↪ strlen(buf), ESP_APPTRACE_TMO_INFINITE);
if (res != ESP_OK) {
    ESP_LOGE(TAG, "Failed to write data to host!");
    return res;
}
```

`esp_apptrace_write()` function uses `memcpy` to copy user data to the internal buffer. In some cases, it can be more optimal to use `esp_apptrace_buffer_get()` and `esp_apptrace_buffer_put()` functions. They allow developers to allocate buffer and fill it themselves. The following piece of code shows how to do this.

```
#include "esp_app_trace.h"
...
int number = 10;
char *ptr = (char *)esp_apptrace_buffer_get(ESP_APPTRACE_DEST_TRAX, 32,
↪ 100/*tmo in us*/);
if (ptr == NULL) {
    ESP_LOGE(TAG, "Failed to get buffer!");
    return ESP_FAIL;
}
sprintf(ptr, "Here is the number %d", number);
esp_err_t res = esp_apptrace_buffer_put(ESP_APPTRACE_DEST_TRAX, ptr,
↪ 100/*tmo in us*/);
```

(continues on next page)

(continued from previous page)

```

if (res != ESP_OK) {
    /* in case of error host tracing tool (e.g., OpenOCD) will report
    ↪ incomplete user buffer */
    ESP_LOGE(TAG, "Failed to put buffer!");
    return res;
}

```

Also according to his needs, the user may want to receive data from the host. Piece of code below shows an example on how to do this.

```

#include "esp_app_trace.h"
...
char buf[32];
char down_buf[32];
size_t sz = sizeof(buf);

/* config down buffer */
esp_apptrace_down_buffer_config(down_buf, sizeof(down_buf));
/* check for incoming data and read them if any */
esp_err_t res = esp_apptrace_read(ESP_APPTRACE_DEST_TRAX, buf, &sz, 0/
    ↪ *do not wait*/);
if (res != ESP_OK) {
    ESP_LOGE(TAG, "Failed to read data from host!");
    return res;
}
if (sz > 0) {
    /* we have data, process them */
    ...
}

```

esp_apptrace_read() function uses memcpy to copy host data to user buffer. In some cases it can be more optimal to use esp_apptrace_down_buffer_get() and esp_apptrace_down_buffer_put() functions. They allow developers to occupy chunk of read buffer and process it in-place. The following piece of code shows how to do this.

```

#include "esp_app_trace.h"
...
char down_buf[32];
uint32_t *number;
size_t sz = 32;

/* config down buffer */
esp_apptrace_down_buffer_config(down_buf, sizeof(down_buf));
char *ptr = (char *)esp_apptrace_down_buffer_get(ESP_APPTRACE_DEST_
    ↪ TRAX, &sz, 100/*tmo in us*/);
if (ptr == NULL) {
    ESP_LOGE(TAG, "Failed to get buffer!");
    return ESP_FAIL;
}
if (sz > 4) {
    number = (uint32_t *)ptr;
    printf("Here is the number %d", *number);
} else {
    printf("No data");
}
esp_err_t res = esp_apptrace_down_buffer_put(ESP_APPTRACE_DEST_TRAX,
    ↪ ptr, 100/*tmo in us*/);
if (res != ESP_OK) {
    /* in case of error host tracing tool (e.g., OpenOCD) will report
    ↪ incomplete user buffer */
    ESP_LOGE(TAG, "Failed to put buffer!");
    return res;
}

```

(continues on next page)

(continued from previous page)

}

2. The next step is to build the program image and download it to the target as described in the [Getting Started Guide](#).
3. Run OpenOCD (see [JTAG Debugging](#)).
4. Connect to OpenOCD telnet server. It can be done using the following command in terminal `telnet <oocd_host> 4444`. If telnet session is opened on the same machine which runs OpenOCD, you can use `localhost` as `<oocd_host>` in the command above.
5. Start trace data collection using special OpenOCD command. This command will transfer tracing data and redirect them to the specified file or socket (currently only files are supported as trace data destination). For description of the corresponding commands, see [OpenOCD Application Level Tracing Commands](#).
6. The final step is to process received data. Since the format of data is defined by users, the processing stage is out of the scope of this document. Good starting points for data processor are python scripts in `$IDF_PATH/tools/esp_app_trace: appttrace_proc.py` (used for feature tests) and `logtrace_proc.py` (see more details in section [Logging to Host](#)).

OpenOCD Application Level Tracing Commands *HW UP BUFFER* is shared between user data blocks and the filling of the allocated memory is performed on behalf of the API caller (in task or ISR context). In multithreading environment, it can happen that the task/ISR which fills the buffer is preempted by another high priority task/ISR. So it is possible that the user data preparation process is not completed at the moment when that chunk is read by the host. To handle such conditions, the tracing module prepends all user data chunks with header which contains the allocated user buffer size (2 bytes) and the length of the actually written data (2 bytes). So the total length of the header is 4 bytes. OpenOCD command which reads trace data reports error when it reads incomplete user data chunk, but in any case, it puts the contents of the whole user chunk (including unfilled area) to the output file.

Below is the description of available OpenOCD application tracing commands.

Note: Currently, OpenOCD does not provide commands to send arbitrary user data to the target.

Command usage:

```
esp appttrace [start <options>] | [stop] | [status] | [dump <cores_num> <outfile>]
```

Sub-commands:

start Start tracing (continuous streaming).
stop Stop tracing.
status Get tracing status.
dump Dump all data from (post-mortem dump).

Start command syntax:

```
start <outfile> [poll_period] [trace_size] [stop_tmo] [wait4halt  
[skip_size]]]
```

outfile Path to file to save data from both CPUs. This argument should have the following format: `file://path/to/file`.

poll_period Data polling period (in ms) for available trace data. If greater than 0, then command runs in non-blocking mode. By default, 1 ms.

trace_size Maximum size of data to collect (in bytes). Tracing is stopped after specified amount of data is received. By default, -1 (trace size stop trigger is disabled).

stop_tmo Idle timeout (in sec). Tracing is stopped if there is no data for specified period of time. By default, -1 (disable this stop trigger). Optionally set it to value longer than longest pause between tracing commands from target.

wait4halt If 0, start tracing immediately, otherwise command waits for the target to be halted (after reset, by breakpoint etc.) and then automatically resumes it and starts tracing. By default, 0.

skip_size Number of bytes to skip at the start. By default, 0.

Note: If `poll_period` is 0, OpenOCD telnet command line will not be available until tracing is stopped. You must stop it manually by resetting the board or pressing Ctrl+C in OpenOCD window (not one with the telnet session). Another option is to set `trace_size` and wait until this size of data is collected. At this point, tracing stops automatically.

Command usage examples:

1. Collect 2048 bytes of tracing data to the file `trace.log`. The file will be saved in the `openocd-esp32` directory.

```
esp apptrace start file://trace.log 1 2048 5 0 0
```

The tracing data will be retrieved and saved in non-blocking mode. This process will stop automatically after 2048 bytes are collected, or if no data are available for more than 5 seconds.

Note: Tracing data is buffered before it is made available to OpenOCD. If you see "Data timeout!" message, then it is likely that the target is not sending enough data to empty the buffer to OpenOCD before the timeout. Either increase the timeout or use the function `esp_apptrace_flush()` to flush the data on specific intervals.

2. Retrieve tracing data indefinitely in non-blocking mode.

```
esp apptrace start file://trace.log 1 -1 -1 0 0
```

There is no limitation on the size of collected data and there is no data timeout set. This process may be stopped by issuing `esp apptrace stop` command on OpenOCD telnet prompt, or by pressing Ctrl+C in OpenOCD window.

3. Retrieve tracing data and save them indefinitely.

```
esp apptrace start file://trace.log 0 -1 -1 0 0
```

OpenOCD telnet command line prompt will not be available until tracing is stopped. To stop tracing, press Ctrl+C in the OpenOCD window.

4. Wait for the target to be halted. Then resume the target's operation and start data retrieval. Stop after collecting 2048 bytes of data:

```
esp apptrace start file://trace.log 0 2048 -1 1 0
```

To configure tracing immediately after reset, use the OpenOCD `reset halt` command.

Logging to Host

ESP-IDF implements a useful feature: logging to the host via application level tracing library. This is a kind of semihosting when all `ESP_LOGx` calls send strings to be printed to the host instead of UART. This can be useful because "printing to host" eliminates some steps performed when logging to UART. Most part of the work is done on the host.

By default, ESP-IDF's logging library uses `vprintf`-like function to write formatted output to dedicated UART. In general, it involves the following steps:

1. Format string is parsed to obtain type of each argument.
2. According to its type, every argument is converted to string representation.
3. Format string combined with converted arguments is sent to UART.

Though the implementation of the `vprintf`-like function can be optimized to a certain level, all steps above have to be performed in any case and every step takes some time (especially item 3). So it frequently occurs that with additional log added to the program to identify the problem, the program behavior is changed and the problem cannot be reproduced. And in the worst cases, the program cannot work normally at all and ends up with an error or even hangs.

Possible ways to overcome this problem are to use higher UART bitrates (or another faster interface) and/or to move string formatting procedure to the host.

The application level tracing feature can be used to transfer log information to the host using `esp_apptrace_vprintf` function. This function does not perform full parsing of the format string and arguments. Instead, it just calculates the number of arguments passed and sends them along with the format string address to the host. On the host, log data is processed and printed out by a special Python script.

Limitations Current implementation of logging over JTAG has some limitations:

1. No support for tracing from `ESP_EARLY_LOGx` macros.
2. No support for printf arguments whose size exceeds 4 bytes (e.g., `double` and `uint64_t`).
3. Only strings from the `.rodata` section are supported as format strings and arguments.
4. The maximum number of printf arguments is 256.

How To Use It In order to use logging via trace module, users need to perform the following steps:

1. On the target side, the special vprintf-like function `esp_apptrace_vprintf()` needs to be installed. It sends log data to the host. An example is `esp_log_set_vprintf(esp_apptrace_vprintf);`. To send log data to UART again, use `esp_log_set_vprintf(vprintf);`.
2. Follow instructions in items 2-5 in [Application Specific Tracing](#).
3. To print out collected log records, run the following command in terminal: `$IDF_PATH/tools/esp_app_trace/logtrace_proc.py /path/to/trace/file /path/to/program/elf/file`.

Log Trace Processor Command Options Command usage:

```
logtrace_proc.py [-h] [--no-errors] <trace_file> <elf_file>
```

Positional arguments:

trace_file Path to log trace file.

elf_file Path to program ELF file.

Optional arguments:

-h, --help Show this help message and exit.

--no-errors, -n Do not print errors.

System Behavior Analysis with SEGGER SystemView

Another useful ESP-IDF feature built on top of application tracing library is the system level tracing which produces traces compatible with SEGGER SystemView tool (see [SystemView](#)). SEGGER SystemView is a real-time recording and visualization tool that allows to analyze runtime behavior of an application. It is possible to view events in real-time through the UART interface.

How To Use It Support for this feature is enabled by Component config>Application Level Tracing>FreeRTOS SystemView Tracing ([CONFIG_APPTRACE_SV_ENABLE](#)) menuconfig option. There are several other options enabled under the same menu:

1. SystemView destination. Select the destination interface: JTAG or UART. In case of UART, it will be possible to connect SystemView application to the ESP32-P4 directly and receive data in real-time.
2. ESP32-P4 timer to use as SystemView timestamp source: ([CONFIG_APPTRACE_SV_TS_SOURCE](#)) selects the source of timestamps for SystemView events. In the single core mode, timestamps are generated using ESP32-P4 internal cycle counter running at maximum 240 Mhz (about 4 ns granularity). In the dual-core mode, external timer working at 40 Mhz is used, so the timestamp granularity is 25 ns.
3. Individually enabled or disabled collection of SystemView events ([CONFIG_APPTRACE_SV_EVT_XXX](#)):
 - Trace Buffer Overflow Event
 - ISR Enter Event
 - ISR Exit Event
 - ISR Exit to Scheduler Event

- Task Start Execution Event
- Task Stop Execution Event
- Task Start Ready State Event
- Task Stop Ready State Event
- Task Create Event
- Task Terminate Event
- System Idle Event
- Timer Enter Event
- Timer Exit Event

ESP-IDF has all the code required to produce SystemView compatible traces, so users can just configure necessary project options (see above), build, download the image to target, and use OpenOCD to collect data as described in the previous sections.

4. Select Pro or App CPU in menuconfig options `Component config>Application Level Tracing>FreeRTOS SystemView Tracing` to trace over the UART interface in real-time.

OpenOCD SystemView Tracing Command Options Command usage:

```
esp sysview [start <options>] | [stop] | [status]
```

Sub-commands:

start Start tracing (continuous streaming).

stop Stop tracing.

status Get tracing status.

Start command syntax:

```
start <outfile1> [outfile2] [poll_period [trace_size [stop_tmo]]]
```

outfile1 Path to file to save data from PRO CPU. This argument should have the following format: `file://path/to/file`.

outfile2 Path to file to save data from APP CPU. This argument should have the following format: `file://path/to/file`.

poll_period Data polling period (in ms) for available trace data. If greater than 0, then command runs in non-blocking mode. By default, 1 ms.

trace_size Maximum size of data to collect (in bytes). Tracing is stopped after specified amount of data is received. By default, -1 (trace size stop trigger is disabled).

stop_tmo Idle timeout (in sec). Tracing is stopped if there is no data for specified period of time. By default, -1 (disable this stop trigger).

Note: If `poll_period` is 0, OpenOCD telnet command line will not be available until tracing is stopped. You must stop it manually by resetting the board or pressing Ctrl+C in the OpenOCD window (not the one with the telnet session). Another option is to set `trace_size` and wait until this size of data is collected. At this point, tracing stops automatically.

Command usage examples:

1. Collect SystemView tracing data to files `pro-cpu.SVData` and `app-cpu.SVData`. The files will be saved in `openocd-esp32` directory.

```
esp sysview start file://pro-cpu.SVData file://app-cpu.SVData
```

The tracing data will be retrieved and saved in non-blocking mode. To stop this process, enter `esp sysview stop` command on OpenOCD telnet prompt, optionally pressing Ctrl+C in the OpenOCD window.

2. Retrieve tracing data and save them indefinitely.

```
esp sysview start file://pro-cpu.SVData file://app-cpu.SVData 0 -1 -1
```

OpenOCD telnet command line prompt will not be available until tracing is stopped. To stop tracing, press Ctrl+C in the OpenOCD window.

Data Visualization After trace data are collected, users can use a special tool to visualize the results and inspect behavior of the program.

Unfortunately, SystemView does not support tracing from multiple cores. So when tracing from ESP32-P4 with JTAG interfaces in the dual-core mode, two files are generated: one for PRO CPU and another for APP CPU. Users can load each file into separate instances of the tool. For tracing over UART, users can select `Component config > Application Level Tracing > FreeRTOS SystemView Tracing` in `menuconfig Pro or App` to choose which CPU has to be traced.

It is uneasy and awkward to analyze data for every core in separate instance of the tool. Fortunately, there is an Eclipse plugin called *Impulse* which can load several trace files, thus making it possible to inspect events from both cores in one view. Also, this plugin has no limitation of 1,000,000 events as compared to the free version of SystemView.

Good instructions on how to install, configure, and visualize data in Impulse from one core can be found [here](#).

Note: ESP-IDF uses its own mapping for SystemView FreeRTOS events IDs, so users need to replace the original file mapping `$SYSVIEW_INSTALL_DIR/Description/SYSVIEW_FreeRTOS.txt` with `$IDF_PATH/tools/esp_app_trace/SYSVIEW_FreeRTOS.txt`. Also, contents of that ESP-IDF-specific file should be used when configuring SystemView serializer using the above link.

Configure Impulse for Dual Core Traces After installing Impulse and ensuring that it can successfully load trace files for each core in separate tabs, users can add special Multi Adapter port and load both files into one view. To do this, users need to do the following steps in Eclipse:

1. Open the Signal Ports view. Go to `Windows > Show View > Other` menu. Find the Signal Ports view in Impulse folder and double-click it.
2. In the Signal Ports view, right-click Ports and select `Add > New Multi Adapter Port`.
3. In the open dialog box, click Add and select `New Pipe/File`.
4. In the open dialog box, select `SystemView Serializer` as Serializer and set path to PRO CPU trace file. Click OK.
5. Repeat the steps 3-4 for APP CPU trace file.
6. Double-click the created port. View for this port should open.
7. Click the `Start/Stop Streaming` button. Data should be loaded.
8. Use the `Zoom Out`, `Zoom In` and `Zoom Fit` buttons to inspect data.
9. For settings measurement cursors and other features, please see [Impulse documentation](#)).

Note: If you have problems with visualization (no data is shown or strange behaviors of zoom action are observed), you can try to delete current signal hierarchy and double-click on the necessary file or port. Eclipse will ask you to create a new signal hierarchy.

Gcov (Source Code Coverage)

Basics of Gcov and Gcovr Source code coverage is data indicating the count and frequency of every program execution path that has been taken within a program's runtime. [Gcov](#) is a GCC tool that, when used in concert with the compiler, can generate log files indicating the execution count of each line of a source file. The [Gcovr](#) tool is a utility for managing Gcov and generating summarized code coverage results.

Generally, using Gcov to compile and run programs on the host will undergo these steps:

1. Compile the source code using GCC with the `--coverage` option enabled. This will cause the compiler to generate a `.gcno` notes files during compilation. The notes files contain information to reconstruct execution path block graphs and map each block to source code line numbers. Each source file compiled with the `--coverage` option should have their own `.gcno` file of the same name (e.g., a `main.c` will generate a `main.gcno` when compiled).
2. Execute the program. During execution, the program should generate `.gda` data files. These data files contain the counts of the number of times an execution path was taken. The program will generate a `.gda` file for each source file compiled with the `--coverage` option (e.g., `main.c` will generate a `main.gda`).

3. Gcov or Gcovr can be used to generate a code coverage based on the `.gcno`, `.gcda`, and source files. Gcov will generate a text-based coverage report for each source file in the form of a `.gcov` file, whilst Gcovr will generate a coverage report in HTML format.

Gcov and Gcovr in ESP-IDF Using Gcov in ESP-IDF is complicated due to the fact that the program is running remotely from the host (i.e., on the target). The code coverage data (i.e., the `.gcda` files) is initially stored on the target itself. OpenOCD is then used to dump the code coverage data from the target to the host via JTAG during runtime. Using Gcov in ESP-IDF can be split into the following steps.

1. *Setting Up a Project for Gcov*
2. *Dumping Code Coverage Data*
3. *Generating Coverage Report*

Setting Up a Project for Gcov

Compiler Option In order to obtain code coverage data in a project, one or more source files within the project must be compiled with the `--coverage` option. In ESP-IDF, this can be achieved at the component level or the individual source file level:

- To cause all source files in a component to be compiled with the `--coverage` option, you can add `target_compile_options(${COMPONENT_LIB} PRIVATE --coverage)` to the `CMakeLists.txt` file of the component.
- To cause a select number of source files (e.g., `source1.c` and `source2.c`) in the same component to be compiled with the `--coverage` option, you can add `set_source_files_properties(source1.c source2.c PROPERTIES COMPILE_FLAGS --coverage)` to the `CMakeLists.txt` file of the component.

When a source file is compiled with the `--coverage` option (e.g., `gcov_example.c`), the compiler will generate the `gcov_example.gcno` file in the project's build directory.

Project Configuration Before building a project with source code coverage, make sure that the following project configuration options are enabled by running `idf.py menuconfig`.

- Enable the application tracing module by selecting `Trace Memory` for the [*CONFIG_APPTTRACE_DESTINATION1*](#) option.
- Enable Gcov to the host via the [*CONFIG_APPTTRACE_GCOV_ENABLE*](#).

Dumping Code Coverage Data Once a project has been compiled with the `--coverage` option and flashed onto the target, code coverage data will be stored internally on the target (i.e., in trace memory) whilst the application runs. The process of transferring code coverage data from the target to the host is known as dumping.

The dumping of coverage data is done via OpenOCD (see [JTAG Debugging](#) on how to setup and run OpenOCD). A dump is triggered by issuing commands to OpenOCD, therefore a telnet session to OpenOCD must be opened to issue such commands (run `telnet localhost 4444`). Note that GDB could be used instead of telnet to issue commands to OpenOCD, however all commands issued from GDB will need to be prefixed as `mon <occd_command>`.

When the target dumps code coverage data, the `.gcda` files are stored in the project's build directory. For example, if `gcov_example_main.c` of the `main` component is compiled with the `--coverage` option, then dumping the code coverage data would generate a `gcov_example_main.gcda` in `build/esp-idf/main/CMakeFiles/__idf_main.dir/gcov_example_main.c.gcda`. Note that the `.gcno` files produced during compilation are also placed in the same directory.

The dumping of code coverage data can be done multiple times throughout an application's lifetime. Each dump will simply update the `.gcda` file with the newest code coverage information. Code coverage data is accumulative, thus the newest data will contain the total execution count of each code path over the application's entire lifetime.

ESP-IDF supports two methods of dumping code coverage data from the target to the host:

- Instant Run-Time Dumpgit

- Hard-coded Dump

Instant Run-Time Dump An Instant Run-Time Dump is triggered by calling the `ESP32-P4 gcov OpenOCD` command (via a telnet session). Once called, OpenOCD will immediately preempt the ESP32-P4's current state and execute a built-in ESP-IDF Gcov debug stub function. The debug stub function will handle the dumping of data to the host. Upon completion, the ESP32-P4 will resume its current state.

Hard-coded Dump A Hard-coded Dump is triggered by the application itself by calling `esp_gcov_dump()` from somewhere within the application. When called, the application will halt and wait for OpenOCD to connect and retrieve the code coverage data. Once `esp_gcov_dump()` is called, the host must execute the `esp gcov dump` OpenOCD command (via a telnet session). The `esp gcov dump` command will cause OpenOCD to connect to the ESP32-P4, retrieve the code coverage data, then disconnect from the ESP32-P4, thus allowing the application to resume. Hard-coded Dumps can also be triggered multiple times throughout an application's lifetime.

Hard-coded dumps are useful if code coverage data is required at certain points of an application's lifetime by placing `esp_gcov_dump()` where necessary (e.g., after application initialization, during each iteration of an application's main loop).

GDB can be used to set a breakpoint on `esp_gcov_dump()`, then call `mon esp gcov dump` automatically via the use a `gdbinit` script (see Using GDB from [Command Line](#)).

The following GDB script will add a breakpoint at `esp_gcov_dump()`, then call the `mon esp gcov dump` OpenOCD command.

```
b esp_gcov_dump
commands
mon esp gcov dump
end
```

Note: Note that all OpenOCD commands should be invoked in GDB as: `mon <occd_command>`.

Generating Coverage Report Once the code coverage data has been dumped, the `.gcno`, `.gcda` and the source files can be used to generate a code coverage report. A code coverage report is simply a report indicating the number of times each line in a source file has been executed.

Both Gcov and Gcovr can be used to generate code coverage reports. Gcov is provided along with the Xtensa toolchain, whilst Gcovr may need to be installed separately. For details on how to use Gcov or Gcovr, refer to [Gcov documentation](#) and [Gcovr documentation](#).

Adding Gcovr Build Target to Project To make report generation more convenient, users can define additional build targets in their projects such that the report generation can be done with a single build command.

Add the following lines to the `CMakeLists.txt` file of your project.

```
include($ENV{IDF_PATH}/tools/cmake/gcov.cmake)
idf_create_coverage_report(${CMAKE_CURRENT_BINARY_DIR}/coverage_report)
idf_clean_coverage_report(${CMAKE_CURRENT_BINARY_DIR}/coverage_report)
```

The following commands can now be used:

- `cmake --build build/ --target gcovr-report` will generate an HTML coverage report in `$(BUILD_DIR_BASE)/coverage_report/html` directory.
- `cmake --build build/ --target cov-data-clean` will remove all coverage data files.

4.2 Application Startup Flow

This note explains various steps which happen before `app_main` function of an ESP-IDF application is called.

The high level view of startup process is as follows:

1. *First Stage Bootloader* in ROM loads second-stage bootloader image to RAM (IRAM & DRAM) from flash offset 0x2000.
2. *Second Stage Bootloader* loads partition table and main app image from flash. Main app incorporates both RAM segments and read-only segments mapped via flash cache.
3. *Application Startup* executes. At this point, the second CPU and RTOS scheduler are started, which then run the `main_task`, leading to the execution of `app_main`.

This process is explained in detail in the following sections.

4.2.1 First Stage Bootloader

After SoC reset, PRO CPU will start running immediately, executing reset vector code, while APP CPU will be held in reset. During startup process, PRO CPU does all the initialization. APP CPU reset is de-asserted in the `call_start_cpu0` function of application startup code. Reset vector code is located in the mask ROM of the ESP32-P4 chip and cannot be modified.

Startup code called from the reset vector determines the boot mode by checking `GPIO_STRAP_REG` register for bootstrap pin states. Depending on the reset reason, the following takes place:

1. Reset from deep sleep: if the value in `RTC_CNTL_STORE6_REG` is non-zero, and CRC value of RTC memory in `RTC_CNTL_STORE7_REG` is valid, use `RTC_CNTL_STORE6_REG` as an entry point address and jump immediately to it. If `RTC_CNTL_STORE6_REG` is zero, or `RTC_CNTL_STORE7_REG` contains invalid CRC, or once the code called via `RTC_CNTL_STORE6_REG` returns, proceed with boot as if it was a power-on reset. **Note:** to run customized code at this point, a deep sleep stub mechanism is provided. Please see [deep sleep](#) documentation for this.
2. For power-on reset, software SoC reset, and watchdog SoC reset: check the `GPIO_STRAP_REG` register if a custom boot mode (such as UART Download Mode) is requested. If this is the case, this custom loader mode is executed from ROM. Otherwise, proceed with boot as if it was due to software CPU reset. Consult ESP32-P4 datasheet for a description of SoC boot modes and how to execute them.
3. For software CPU reset and watchdog CPU reset: configure SPI flash based on EFUSE values, and attempt to load the code from flash. This step is described in more detail in the next paragraphs.

Note: During normal boot modes the RTC watchdog is enabled when this happens, so if the process is interrupted or stalled then the watchdog will reset the SOC automatically and repeat the boot process. This may cause the SoC to strap into a new boot mode, if the strapping GPIOs have changed.

Second stage bootloader binary image is loaded from flash starting at address 0x2000. The 8 kB sector of flash before this address is reserved for the key manager for use with flash encryption (AES-XTS).

4.2.2 Second Stage Bootloader

In ESP-IDF, the binary image which resides at offset 0x2000 in flash is the second stage bootloader. Second stage bootloader source code is available in [components/bootloader](#) directory of ESP-IDF. Second stage bootloader is used in ESP-IDF to add flexibility to flash layout (using partition tables), and allow for various flows associated with flash encryption, secure boot, and over-the-air updates (OTA) to take place.

When the first stage bootloader is finished checking and loading the second stage bootloader, it jumps to the second stage bootloader entry point found in the binary image header.

Second stage bootloader reads the partition table found by default at offset 0x8000 (*configurable value*). See [partition tables](#) documentation for more information. The bootloader finds factory and OTA app partitions. If OTA app partitions are found in the partition table, the bootloader consults the `otadata` partition to determine which one should be booted. See [Over The Air Updates \(OTA\)](#) for more information.

For a full description of the configuration options available for the ESP-IDF bootloader, see [Bootloader](#).

For the selected partition, second stage bootloader reads the binary image from flash one segment at a time:

- For segments with load addresses in internal *IRAM (Instruction RAM)* or *DRAM (Data RAM)*, the contents are copied from flash to the load address.
- For segments which have load addresses in *DROM (Data Stored in flash)* or *IROM (Code Executed from flash)* regions, the flash MMU is configured to provide the correct mapping from the flash to the load address.

Once all segments are processed - meaning code is loaded and flash MMU is set up, second stage bootloader verifies the integrity of the application and then jumps to the application entry point found in the binary image header.

4.2.3 Application Startup

Application startup covers everything that happens after the app starts executing and before the `app_main` function starts running inside the main task. This is split into three stages:

- Port initialization of hardware and basic C runtime environment.
- System initialization of software services and FreeRTOS.
- Running the main task and calling `app_main`.

Note: Understanding all stages of ESP-IDF app initialization is often not necessary. To understand initialization from the application developer's perspective only, skip forward to [Running the Main Task](#).

Port Initialization

ESP-IDF application entry point is `call_start_cpu0` function found in [components/esp_system/port/cpu_start.c](#). This function is executed by the second stage bootloader, and never returns.

This port-layer initialization function initializes the basic C Runtime Environment ("CRT") and performs initial configuration of the SoC's internal hardware:

- Reconfigure CPU exceptions for the app (allowing app interrupt handlers to run, and causing *Fatal Errors* to be handled using the options configured for the app rather than the simpler error handler provided by ROM).
- If the option `CONFIG_BOOTLOADER_WDT_ENABLE` is not set then the RTC watchdog timer is disabled.
- Initialize internal memory (data & bss).
- Finish configuring the MMU cache.
- Enable PSRAM if configured.
- Set the CPU clocks to the frequencies configured for the project.
- If the app is configured to run on multiple cores, start the other core and wait for it to initialize as well (inside the similar "port layer" initialization function `call_start_cpu1`).

Once `call_start_cpu0` completes running, it calls the "system layer" initialization function `start_cpu0` found in [components/esp_system/startup.c](#). Other cores will also complete port-layer initialization and call `start_other_cores` found in the same file.

System Initialization

The main system initialization function is `start_cpu0`. By default, this function is weak-linked to the function `start_cpu0_default`. This means that it is possible to override this function to add some additional initialization steps.

The primary system initialization stage includes:

- Log information about this application (project name, *App Version*, etc.) if default log level enables this.
- Initialize the heap allocator (before this point all allocations must be static or on the stack).
- Initialize newlib component syscalls and time functions.
- Configure the brownout detector.
- Setup libc stdin, stdout, and stderr according to the *serial console configuration*.
- Perform any security-related checks, including burning efuses that should be burned for this configuration (including *permanently limiting ROM download modes*).
- Initialize SPI flash API support.
- Call global C++ constructors and any C functions marked with `__attribute__((constructor))`.

Secondary system initialization allows individual components to be initialized. If a component has an initialization function annotated with the `ESP_SYSTEM_INIT_FN` macro, it will be called as part of secondary initialization. Component initialization functions have priorities assigned to them to ensure the desired initialization order. The priorities are documented in `esp_system/system_init_fn.txt` and `ESP_SYSTEM_INIT_FN` definition in source code are checked against this file.

Running the Main Task

After all other components are initialized, the main task is created and the FreeRTOS scheduler starts running.

After doing some more initialization tasks (that require the scheduler to have started), the main task runs the application-provided function `app_main` in the firmware.

The main task that runs `app_main` has a fixed RTOS priority (one higher than the minimum) and a *configurable stack size*.

The main task core affinity is also configurable: `CONFIG_ESP_MAIN_TASK_AFFINITY`.

Unlike normal FreeRTOS tasks (or embedded C main functions), the `app_main` task is allowed to return. If this happens, the task is cleaned up and the system will continue running with other RTOS tasks scheduled normally. Therefore, it is possible to implement `app_main` as either a function that creates other application tasks and then returns, or as a main application task itself.

Second Core Startup

A similar but simpler startup process happens on the APP CPU:

When running system initialization, the code on PRO CPU sets the entry point for APP CPU, de-asserts APP CPU reset, and waits for a global flag to be set by the code running on APP CPU, indicating that it has started. Once this is done, APP CPU jumps to `call_start_cpu1` function in `components/esp_system/port/cpu_start.c`.

While PRO CPU does initialization in `start_cpu0` function, APP CPU runs `start_cpu_other_cores` function. Similar to `start_cpu0`, this function is weak-linked and defaults to the `start_cpu_other_cores_default` function but can be replaced with a different function by the application.

The `start_cpu_other_cores_default` function does some core-specific system initialization and then waits for the PRO CPU to start the FreeRTOS scheduler, at which point it executes `esp_startup_start_app_other_cores` which is another weak-linked function defaulting to `esp_startup_start_app_other_cores_default`.

By default `esp_startup_start_app_other_cores_default` does nothing but spin in a busy-waiting loop until the scheduler of the PRO CPU triggers an interrupt to start the RTOS scheduler on the APP CPU.

4.3 Bootloader

The ESP-IDF Software Bootloader performs the following functions:

1. Minimal initial configuration of internal modules;
2. Initialize [Flash Encryption](#) and/or [Secure](#) features, if configured;
3. Select the application partition to boot, based on the partition table and ota_data (if any);
4. Load this image to RAM (IRAM & DRAM) and transfer management to the image that was just loaded.

Bootloader is located at the address 0x2000 in the flash.

For a full description of the startup process including the ESP-IDF bootloader, see [Application Startup Flow](#).

4.3.1 Bootloader Compatibility

It is recommended to update to newer [versions of ESP-IDF](#): when they are released. The OTA (over the air) update process can flash new apps in the field but cannot flash a new bootloader. For this reason, the bootloader supports booting apps built from newer versions of ESP-IDF.

The bootloader does not support booting apps from older versions of ESP-IDF. When updating ESP-IDF manually on an existing product that might need to downgrade the app to an older version, keep using the older ESP-IDF bootloader binary as well.

Note: If testing an OTA update for an existing product in production, always test it using the same ESP-IDF bootloader binary that is deployed in production.

SPI Flash Configuration

Each ESP-IDF application or bootloader .bin file contains a header with [CONFIG_ESPTOOLPY_FLASHMODE](#), [CONFIG_ESPTOOLPY_FLASHFREQ](#), [CONFIG_ESPTOOLPY_FLASHSIZE](#) embedded in it. These are used to configure the SPI flash during boot.

The [First Stage Bootloader](#) in ROM reads the [Second Stage Bootloader](#) header information from flash and uses this information to load the rest of the [Second Stage Bootloader](#) from flash. However, at this time the system clock speed is lower than configured and not all flash modes are supported. When the [Second Stage Bootloader](#) then runs, it will reconfigure the flash using values read from the currently selected app binary's header (and NOT from the [Second Stage Bootloader](#) header). This allows an OTA update to change the SPI flash settings in use.

4.3.2 Log Level

The default bootloader log level is "Info". By setting the [CONFIG_BOOTLOADER_LOG_LEVEL](#) option, it is possible to increase or decrease this level. This log level is separate from the log level used in the app (see [Logging library](#)).

Reducing bootloader log verbosity can improve the overall project boot time by a small amount.

4.3.3 Factory Reset

Sometimes it is desirable to have a way for the device to fall back to a known-good state, in case of some problem with an update.

To roll back to the original "factory" device configuration and clear any user settings, configure the config item [CONFIG_BOOTLOADER_FACTORY_RESET](#) in the bootloader.

The factory reset mechanism allows the device to be factory reset in two ways:

- Clear one or more data partitions. The [CONFIG_BOOTLOADER_DATA_FACTORY_RESET](#) option allows users to specify which data partitions will be erased when the factory reset is executed. Users can specify the names of partitions as a comma-delimited list with optional spaces for readability. (Like this: `nvs, phy_init, nvs_custom`). Make sure that the names of partitions specified in the option are the same as those found in the partition table. Partitions of type "app" cannot be specified here.
- Boot from "factory" app partition. Enabling the [CONFIG_BOOTLOADER_OTA_DATA_ERASE](#) option will cause the device to boot from the default "factory" app partition after a factory reset (or if there is no factory app partition in the partition table then the default ota app partition is selected instead). This reset process involves erasing the OTA data partition which holds the currently selected OTA partition slot. The "factory" app partition slot (if it exists) is never updated via OTA, so resetting to this allows reverting to a "known good" firmware application.

Either or both of these configuration options can be enabled independently.

In addition, the following configuration options control the reset condition:

- [CONFIG_BOOTLOADER_NUM_PIN_FACTORY_RESET](#)- The input GPIO number used to trigger a factory reset. This GPIO must be pulled low or high (configurable) on reset to trigger this.
- [CONFIG_BOOTLOADER_HOLD_TIME_GPIO](#)- this is hold time of GPIO for reset/test mode (by default 5 seconds). The GPIO must be held continuously for this period of time after reset before a factory reset or test partition boot (as applicable) is performed.
- [CONFIG_BOOTLOADER_FACTORY_RESET_PIN_LEVEL](#) - configure whether a factory reset should trigger on a high or low level of the GPIO. If the GPIO has an internal pullup then this is enabled before the pin is sampled, consult the ESP32-P4 datasheet for details on pin internal pullups.

If an application needs to know if the factory reset has occurred, users can call the function `bootloader_common_get_rtc_retain_mem_factory_reset_state()`.

- If the status is read as true, the function will return the status, indicating that the factory reset has occurred. The function then resets the status to false for subsequent factory reset judgement.
- If the status is read as false, the function will return the status, indicating that the factory reset has not occurred, or the memory where this status is stored is invalid.

Note that this feature reserves some RTC FAST memory (the same size as the [CONFIG_BOOTLOADER_SKIP_VALIDATE_IN_DEEP_SLEEP](#) feature).

4.3.4 Boot from Test Firmware

It is possible to write a special firmware app for testing in production, and boot this firmware when needed. The project partition table will need a dedicated app partition entry for this testing app, type `app` and subtype `test` (see [Partition Tables](#)).

Implementing a dedicated test app firmware requires creating a totally separate ESP-IDF project for the test app (each project in ESP-IDF only builds one app). The test app can be developed and tested independently of the main project, and then integrated at production testing time as a pre-compiled `.bin` file which is flashed to the address of the main project's test app partition.

To support this functionality in the main project's bootloader, set the configuration item [CONFIG_BOOTLOADER_APP_TEST](#) and configure the following three items:

- [CONFIG_BOOTLOADER_NUM_PIN_APP_TEST](#) - GPIO number to boot test partition. The selected GPIO will be configured as an input with internal pull-up enabled. This GPIO must be pulled low or high (configurable) on reset to trigger this. Once the GPIO input is released and the device has been rebooted, the default boot sequence will be enabled again to boot the factory partition or any OTA app partition slot.
- [CONFIG_BOOTLOADER_HOLD_TIME_GPIO](#) - this is the hold time of GPIO for reset/test mode (by default 5 seconds). The GPIO must be held continuously for this period of time after reset before a factory reset or test partition boot (as applicable) is performed.
- [CONFIG_BOOTLOADER_APP_TEST_PIN_LEVEL](#) - configure whether a test partition boot should trigger on a high or low level of the GPIO. If the GPIO has an internal pull-up, then this is enabled before the pin is sampled. Consult the ESP32-P4 datasheet for details on pin internal pull-ups.

4.3.5 Rollback

Rollback and anti-rollback features must be configured in the bootloader as well.

Consult the [App Rollback](#) and [Anti-rollback](#) sections in the [OTA API reference document](#).

4.3.6 Watchdog

By default, the hardware RTC Watchdog timer remains running while the bootloader is running and will automatically reset the chip if no app has successfully started after 9 seconds.

- The timeout period can be adjusted by setting [CONFIG_BOOTLOADER_WDT_TIME_MS](#) and recompiling the bootloader.
- The app's behaviour can be adjusted so the RTC Watchdog remains enabled after app startup. The Watchdog would need to be explicitly reset (i.e., fed) by the app to avoid a reset. To do this, set the [CONFIG_BOOTLOADER_WDT_DISABLE_IN_USER_CODE](#) option, modify the app as needed, and then recompile the app.
- The RTC Watchdog can be disabled in the bootloader by disabling the [CONFIG_BOOTLOADER_WDT_ENABLE](#) setting and recompiling the bootloader. This is not recommended.

4.3.7 Bootloader Size

When enabling additional bootloader functions, including [Flash Encryption](#) or Secure Boot, and especially if setting a high [CONFIG_BOOTLOADER_LOG_LEVEL](#) level, then it is important to monitor the bootloader .bin file's size.

When using the default [CONFIG_PARTITION_TABLE_OFFSET](#) value 0x8000, the size limit is 0x8000 bytes.

If the bootloader binary is too large, then the bootloader build will fail with an error "Bootloader binary size [...] is too large for partition table offset". If the bootloader binary is flashed anyhow then the ESP32-P4 will fail to boot - errors will be logged about either invalid partition table or invalid bootloader checksum.

Options to work around this are:

- Set [bootloader compiler optimization](#) back to "Size" if it has been changed from this default value.
- Reduce [bootloader log level](#). Setting log level to Warning, Error or None all significantly reduce the final binary size (but may make it harder to debug).
- Set [CONFIG_PARTITION_TABLE_OFFSET](#) to a higher value than 0x8000, to place the partition table later in the flash. This increases the space available for the bootloader. If the [partition table](#) CSV file contains explicit partition offsets, they will need changing so no partition has an offset lower than `CONFIG_PARTITION_TABLE_OFFSET + 0x1000`. (This includes the default partition CSV files supplied with ESP-IDF.)

When Secure Boot V2 is enabled, there is also an absolute binary size limit of 64 KB (0x10000 bytes) (excluding the 4 KB signature), because the bootloader is first loaded into a fixed size buffer for verification.

4.3.8 Fast Boot from Deep-Sleep

The bootloader has the [CONFIG_BOOTLOADER_SKIP_VALIDATE_IN_DEEP_SLEEP](#) option which allows the wake-up time from Deep-sleep to be reduced (useful for reducing power consumption). This option is available when [CONFIG_SECURE_BOOT](#) option is disabled. Reduction of time is achieved due to the lack of image verification. During the first boot, the bootloader stores the address of the application being launched in the RTC FAST memory. And during the awakening, this address is used for booting without any checks, thus fast loading is achieved.

4.3.9 Custom Bootloader

The current bootloader implementation allows a project to extend it or modify it. There are two ways of doing it: by implementing hooks or by overriding it. Both ways are presented in [custom_bootloader](#) folder in ESP-IDF examples:

- *bootloader_hooks* which presents how to connect some hooks to the bootloader initialization
- *bootloader_override* which presents how to override the bootloader implementation

In the bootloader space, you cannot use the drivers and functions from other components. If necessary, then the required functionality should be placed in the project's *bootloader_components* directory (note that this will increase its size).

If the bootloader grows too large then it can collide with the partition table, which is flashed at offset 0x8000 by default. Increase the *partition table offset* value to place the partition table later in the flash. This increases the space available for the bootloader.

4.4 Build System

This document explains the implementation of the ESP-IDF build system and the concept of "components". Read this document if you want to know how to organize and build a new ESP-IDF project or component.

4.4.1 Overview

An ESP-IDF project can be seen as an amalgamation of a number of components. For example, for a web server that shows the current humidity, there could be:

- The ESP-IDF base libraries (libc, ROM bindings, etc)
- The Wi-Fi drivers
- A TCP/IP stack
- The FreeRTOS operating system
- A web server
- A driver for the humidity sensor
- Main code tying it all together

ESP-IDF makes these components explicit and configurable. To do that, when a project is compiled, the build system will look up all the components in the ESP-IDF directories, the project directories and (optionally) in additional custom component directories. It then allows the user to configure the ESP-IDF project using a text-based menu system to customize each component. After the components in the project are configured, the build system will compile the project.

Concepts

- A "project" is a directory that contains all the files and configuration to build a single "app" (executable), as well as additional supporting elements such as a partition table, data/filesystem partitions, and a bootloader.
- "Project configuration" is held in a single file called `sdkconfig` in the root directory of the project. This configuration file is modified via `idf.py menuconfig` to customize the configuration of the project. A single project contains exactly one project configuration.
- An "app" is an executable that is built by ESP-IDF. A single project will usually build two apps - a "project app" (the main executable, ie your custom firmware) and a "bootloader app" (the initial bootloader program which launches the project app).
- "components" are modular pieces of standalone code that are compiled into static libraries (.a files) and linked to an app. Some are provided by ESP-IDF itself, others may be sourced from other places.
- "Target" is the hardware for which an application is built. A full list of supported targets in your version of ESP-IDF can be seen by running `idf.py --list-targets`.

Some things are not part of the project:

- "ESP-IDF" is not part of the project. Instead, it is standalone, and linked to the project via the `IDF_PATH` environment variable which holds the path of the `esp-idf` directory. This allows the ESP-IDF framework to be decoupled from your project.

- The toolchain for compilation is not part of the project. The toolchain should be installed in the system command line PATH.

4.4.2 Using the Build System

idf.py

The `idf.py` command-line tool provides a front-end for easily managing your project builds. It manages the following tools:

- [CMake](#), which configures the project to be built
- [Ninja](#) which builds the project
- [esptool.py](#) for flashing the target.

You can read more about configuring the build system using `idf.py` [here](#).

Using CMake Directly

`idf.py` is a wrapper around [CMake](#) for convenience. However, you can also invoke CMake directly if you prefer.

When `idf.py` does something, it prints each command that it runs for easy reference. For example, the `idf.py build` command is the same as running these commands in a bash shell (or similar commands for Windows Command Prompt):

```
mkdir -p build
cd build
cmake .. -G Ninja # or 'Unix Makefiles'
ninja
```

In the above list, the `cmake` command configures the project and generates build files for use with the final build tool. In this case, the final build tool is [Ninja](#): running `ninja` actually builds the project.

It's not necessary to run `cmake` more than once. After the first build, you only need to run `ninja` each time. `ninja` will automatically re-invoke `cmake` if the project needs reconfiguration.

If using CMake with `ninja` or `make`, there are also targets for more of the `idf.py` sub-commands. For example, running `make menuconfig` or `ninja menuconfig` in the build directory will work the same as `idf.py menuconfig`.

Note: If you're already familiar with [CMake](#), you may find the ESP-IDF CMake-based build system unusual because it wraps a lot of CMake's functionality to reduce boilerplate. See [writing pure CMake components](#) for some information about writing more "CMake style" components.

Flashing with Ninja or Make It's possible to build and flash directly from `ninja` or `make` by running a target like:

```
ninja flash
```

Or:

```
make app-flash
```

Available targets are: `flash`, `app-flash` (app only), `bootloader-flash` (bootloader only).

When flashing this way, optionally set the `ESPPORT` and `ESPBAUD` environment variables to specify the serial port and baud rate. You can set environment variables in your operating system or IDE project. Alternatively, set them directly on the command line:

```
ESPSPORT=/dev/ttyUSB0 ninja flash
```

Note: Providing environment variables at the start of the command like this is Bash shell Syntax. It will work on Linux and macOS. It won't work when using Windows Command Prompt, but it will work when using Bash-like shells on Windows.

Or:

```
make -j3 app-flash ESPSPORT=COM4 ESPBAUD=2000000
```

Note: Providing variables at the end of the command line is make syntax, and works for make on all platforms.

Using CMake in an IDE

You can also use an IDE with CMake integration. The IDE will want to know the path to the project's `CMakeLists.txt` file. IDEs with CMake integration often provide their own build tools (CMake calls these "generators") to build the source files as part of the IDE.

When adding custom non-build steps like "flash" to the IDE, it is recommended to execute `idf.py` for these "special" commands.

For more detailed information about integrating ESP-IDF with CMake into an IDE, see [Build System Metadata](#).

Setting up the Python Interpreter

ESP-IDF works well with Python version 3.8+.

`idf.py` and other Python scripts will run with the default Python interpreter, i.e., `python`. You can switch to a different one like `python3 $IDF_PATH/tools/idf.py ...`, or you can set up a shell alias or another script to simplify the command.

If using CMake directly, running `cmake -D PYTHON=python3 ...` will cause CMake to override the default Python interpreter.

If using an IDE with CMake, setting the `PYTHON` value as a CMake cache override in the IDE UI will override the default Python interpreter.

To manage the Python version more generally via the command line, check out the tools [pyenv](#) or [virtualenv](#). These let you change the default Python version.

4.4.3 Example Project

An example project directory tree might look like this:

```
- myProject/
  - CMakeLists.txt
  - sdkconfig
  - bootloader_components/ - boot_component/ - CMakeLists.txt
                                - Kconfig
                                - src1.c
  - components/ - component1/ - CMakeLists.txt
                                    - Kconfig
                                    - src1.c
                        - component2/ - CMakeLists.txt
                                    - Kconfig
```

(continues on next page)

(continued from previous page)

```

- src1.c
- include/ - component2.h
- main/      - CMakeLists.txt
              - src1.c
              - src2.c
- build/

```

This example "myProject" contains the following elements:

- A top-level project CMakeLists.txt file. This is the primary file which CMake uses to learn how to build the project; and may set project-wide CMake variables. It includes the file [/tools/cmake/project.cmake](#) which implements the rest of the build system. Finally, it sets the project name and defines the project.
- "sdkconfig" project configuration file. This file is created/updated when `idf.py menuconfig` runs, and holds the configuration for all of the components in the project (including ESP-IDF itself). The `sdkconfig` file may or may not be added to the source control system of the project.
- Optional "bootloader_components" directory contains components that need to be compiled and linked inside the bootloader project. A project does not have to contain custom bootloader components of this kind, but it can be useful in case the bootloader needs to be modified to embed new features.
- Optional "components" directory contains components that are part of the project. A project does not have to contain custom components of this kind, but it can be useful for structuring reusable code or including third-party components that aren't part of ESP-IDF. Alternatively, `EXTRA_COMPONENT_DIRS` can be set in the top-level CMakeLists.txt to look for components in other places.
- "main" directory is a special component that contains source code for the project itself. "main" is a default name, the CMake variable `COMPONENT_DIRS` includes this component but you can modify this variable. See the [renaming main](#) section for more info. If you have a lot of source files in your project, we recommend grouping most into components instead of putting them all in "main".
- "build" directory is where the build output is created. This directory is created by `idf.py` if it doesn't already exist. CMake configures the project and generates interim build files in this directory. Then, after the main build process is run, this directory will also contain interim object files and libraries as well as final binary output files. This directory is usually not added to source control or distributed with the project source code.

Component directories each contain a component CMakeLists.txt file. This file contains variable definitions to control the build process of the component, and its integration into the overall project. See [Component CMakeLists Files](#) for more details.

Each component may also include a Kconfig file defining the [component configuration](#) options that can be set via `menuconfig`. Some components may also include `Kconfig.projbuild` and `project_include.cmake` files, which are special files for [overriding parts of the project](#).

4.4.4 Project CMakeLists File

Each project has a single top-level CMakeLists.txt file that contains build settings for the entire project. By default, the project CMakeLists can be quite minimal.

Minimal Example CMakeLists

Minimal project:

```

cmake_minimum_required(VERSION 3.16)
include($ENV{IDF_PATH}/tools/cmake/project.cmake)
project(myProject)

```

Mandatory Parts

The inclusion of these three lines, in the order shown above, is necessary for every project:

- `cmake_minimum_required(VERSION 3.16)` tells CMake the minimum version that is required to build the project. ESP-IDF is designed to work with CMake 3.16 or newer. This line must be the first line in the `CMakeLists.txt` file.
- `include($ENV{IDF_PATH}/tools/cmake/project.cmake)` pulls in the rest of the CMake functionality to configure the project, discover all the components, etc.
- `project(myProject)` creates the project itself, and specifies the project name. The project name is used for the final binary output files of the app - ie `myProject.elf`, `myProject.bin`. Only one project can be defined per `CMakeLists` file.

Optional Project Variables

These variables all have default values that can be overridden for custom behavior. Look in [/tools/cmake/project.cmake](#) for all of the implementation details.

- `COMPONENT_DIRS`: Directories to search for components. Defaults to `IDF_PATH/components`, `PROJECT_DIR/components`, and `EXTRA_COMPONENT_DIRS`. Override this variable if you don't want to search for components in these places.
- `EXTRA_COMPONENT_DIRS`: Optional list of additional directories to search for components. Paths can be relative to the project directory, or absolute.
- `COMPONENTS`: A list of component names to build into the project. Defaults to all components found in the `COMPONENT_DIRS` directories. Use this variable to "trim down" the project for faster build times. Note that any component which "requires" another component via the `REQUIRES` or `PRIV_REQUIRES` arguments on component registration will automatically have it added to this list, so the `COMPONENTS` list can be very short.
- `BOOTLOADER_IGNORE_EXTRA_COMPONENT`: A list of components, placed in `bootloader_components/`, that should be ignored by the bootloader compilation. Use this variable if a bootloader component needs to be included conditionally inside the project.

Any paths in these variables can be absolute paths, or set relative to the project directory.

To set these variables, use the `cmake set command` ie `set(VARIABLE "VALUE")`. The `set()` commands should be placed after the `cmake_minimum(...)` line but before the `include(...)` line.

Renaming main Component

The build system provides special treatment to the `main` component. It is a component that gets automatically added to the build provided that it is in the expected location, `PROJECT_DIR/main`. All other components in the build are also added as its dependencies, saving the user from hunting down dependencies and providing a build that works right out of the box. Renaming the `main` component causes the loss of these behind-the-scenes heavy lifting, requiring the user to specify the location of the newly renamed component and manually specify its dependencies. Specifically, the steps to renaming `main` are as follows:

1. Rename `main` directory.
2. Set `EXTRA_COMPONENT_DIRS` in the project `CMakeLists.txt` to include the renamed `main` directory.
3. Specify the dependencies in the renamed component's `CMakeLists.txt` file via `REQUIRES` or `PRIV_REQUIRES` arguments [on component registration](#).

Overriding Default Build Specifications

The build sets some global build specifications (compile flags, definitions, etc.) that gets used in compiling all sources from all components.

For example, one of the default build specifications set is the compile option `-Wextra`. Suppose a user wants to use override this with `-Wno-extra`, it should be done after `project()`:

```
cmake_minimum_required(VERSION 3.16)
include($ENV{IDF_PATH}/tools/cmake/project.cmake)
project(myProject)

idf_build_set_property(COMPILER_OPTIONS "-Wno-error" APPEND)
```

This ensures that the compile options set by the user won't be overridden by the default build specifications, since the latter are set inside `project()`.

4.4.5 Component CMakeLists Files

Each project contains one or more components. Components can be part of ESP-IDF, part of the project's own components directory, or added from custom component directories ([see above](#)).

A component is any directory in the `COMPONENT_DIRS` list which contains a `CMakeLists.txt` file.

Searching for Components

The list of directories in `COMPONENT_DIRS` is searched for the project's components. Directories in this list can either be components themselves (ie they contain a `CMakeLists.txt` file), or they can be top-level directories whose sub-directories are components.

When CMake runs to configure the project, it logs the components included in the build. This list can be useful for debugging the inclusion/exclusion of certain components.

Multiple Components with the Same Name

When ESP-IDF is collecting all the components to compile, the search precedence is as follows (from highest to lowest):

- Project components
- Components from `EXTRA_COMPONENT_DIRS`
- Project managed components, downloaded by the IDF Component Manager into `PROJECT_DIR/managed_components`, unless the IDF Component Manager is disabled.
- ESP-IDF components (`IDF_PATH/components`)

If two or more of these directories contain component sub-directories with the same name, the component with higher precedence is used. This allows, for example, overriding ESP-IDF components with a modified version by copying that component from the ESP-IDF components directory to the project components directory and then modifying it there. If used in this way, the ESP-IDF directory itself can remain untouched.

Note: If a component is overridden in an existing project by moving it to a new location, the project will not automatically see the new component path. Run `idf.py reconfigure` (or delete the project build folder) and then build again.

Minimal Component CMakeLists

The minimal component `CMakeLists.txt` file simply registers the component to the build system using `idf_component_register`:

```
idf_component_register(SRCS "foo.c" "bar.c"
                       INCLUDE_DIRS "include"
                       REQUIRES mbedtls)
```

- `SRCS` is a list of source files (`*.c`, `*.cpp`, `*.cc`, `*.S`). These source files will be compiled into the component library.
- `INCLUDE_DIRS` is a list of directories to add to the global include search path for any component which requires this component, and also the main source files.
- `REQUIRES` is not actually required, but it is very often required to declare what other components this component will use. See [component requirements](#).

A library with the name of the component will be built and linked to the final app.

Directories are usually specified relative to the `CMakeLists.txt` file itself, although they can be absolute.

There are other arguments that can be passed to `idf_component_register`. These arguments are discussed [here](#).

See [example component requirements](#) and [example component CMakeLists](#) for more complete component `CMakeLists.txt` examples.

Preset Component Variables

The following component-specific variables are available for use inside component `CMakeLists`, but should not be modified:

- `COMPONENT_DIR`: The component directory. Evaluates to the absolute path of the directory containing `CMakeLists.txt`. The component path cannot contain spaces. This is the same as the `CMAKE_CURRENT_SOURCE_DIR` variable.
- `COMPONENT_NAME`: Name of the component. Same as the name of the component directory.
- `COMPONENT_ALIAS`: Alias of the library created internally by the build system for the component.
- `COMPONENT_LIB`: Name of the library created internally by the build system for the component.

The following variables are set at the project level, but available for use in component `CMakeLists`:

- `CONFIG_*`: Each value in the project configuration has a corresponding variable available in `cmake`. All names begin with `CONFIG_`. [More information here](#).
- `ESP_PLATFORM`: Set to 1 when the `CMake` file is processed within the ESP-IDF build system.

Build/Project Variables

The following are some project/build variables that are available as build properties and whose values can be queried using `idf_build_get_property` from the component `CMakeLists.txt`:

- `PROJECT_NAME`: Name of the project, as set in project `CMakeLists.txt` file.
- `PROJECT_DIR`: Absolute path of the project directory containing the project `CMakeLists`. Same as the `CMAKE_SOURCE_DIR` variable.
- `COMPONENTS`: Names of all components that are included in this build, formatted as a semicolon-delimited `CMake` list.
- `IDF_VER`: Git version of ESP-IDF (produced by `git describe`)
- `IDF_VERSION_MAJOR`, `IDF_VERSION_MINOR`, `IDF_VERSION_PATCH`: Components of ESP-IDF version, to be used in conditional expressions. Note that this information is less precise than that provided by `IDF_VER` variable. `v4.0-dev-*`, `v4.0-beta1`, `v4.0-rc1` and `v4.0` will all have the same values of `IDF_VERSION_*` variables, but different `IDF_VER` values.
- `IDF_TARGET`: Name of the target for which the project is being built.
- `PROJECT_VER`: Project version.
 - If `CONFIG_APP_PROJECT_VER_FROM_CONFIG` option is set, the value of `CONFIG_APP_PROJECT_VER` will be used.
 - Else, if `PROJECT_VER` variable is set in project `CMakeLists.txt` file, its value will be used.
 - Else, if the `PROJECT_DIR/version.txt` exists, its contents will be used as `PROJECT_VER`.
 - Else, if the project is located inside a Git repository, the output of `git` description will be used.
 - Otherwise, `PROJECT_VER` will be "1".
- `EXTRA_PARTITION_SUBTYPES`: `CMake` list of extra partition subtypes. Each subtype description is a comma-separated string with `type_name`, `subtype_name`, `numeric_value` format. Components may add new subtypes by appending them to this list.

Other build properties are listed [here](#).

Controlling Component Compilation

To pass compiler options when compiling source files belonging to a particular component, use the `target_compile_options` function:

```
target_compile_options(${COMPONENT_LIB} PRIVATE -Wno-unused-variable)
```

To apply the compilation flags to a single source file, use the CMake `set_source_files_properties` command:

```
set_source_files_properties(mysrc.c
    PROPERTIES COMPILE_FLAGS
        -Wno-unused-variable
)
```

This can be useful if there is upstream code that emits warnings.

Note: CMake `set_source_files_properties` command is not applicable when the source files have been populated with help of the `SRC_DIRS` variable in `idf_component_register`. See [File Globbing & Incremental Builds](#) for more details.

When using these commands, place them after the call to `idf_component_register` in the component CMakeLists file.

4.4.6 Component Configuration

Each component can also have a `Kconfig` file, alongside `CMakeLists.txt`. This contains configuration settings to add to the configuration menu for this component.

These settings are found under the "Component Settings" menu when `menuconfig` is run.

To create a component `Kconfig` file, it is easiest to start with one of the `Kconfig` files distributed with ESP-IDF.

For an example, see [Adding conditional configuration](#).

4.4.7 Preprocessor Definitions

The ESP-IDF build system adds the following C preprocessor definitions on the command line:

- `ESP_PLATFORM`: Can be used to detect that build happens within ESP-IDF.
- `IDF_VER`: Defined to a git version string. E.g. `v2.0` for a tagged release or `v1.0-275-g0efaa4f` for an arbitrary commit.

4.4.8 Component Requirements

When compiling each component, the ESP-IDF build system recursively evaluates its dependencies. This means each component needs to declare the components that it depends on ("requires").

When Writing a Component

```
idf_component_register(...
    REQUIRES mbedtls
    PRIV_REQUIRES console spiffs)
```

- `REQUIRES` should be set to all components whose header files are `#included` from the *public* header files of this component.

- `PRIV_REQUIRES` should be set to all components whose header files are `#included` from *any source files* in this component, unless already listed in `REQUIRES`. Also, any component which is required to be linked in order for this component to function correctly.
- The values of `REQUIRES` and `PRIV_REQUIRES` should not depend on any configuration choices (`CONFIG_xxx` macros). This is because requirements are expanded before the configuration is loaded. Other component variables (like include paths or source files) can depend on configuration choices.
- Not setting either or both `REQUIRES` variables is fine. If the component has no requirements except for the *Common component requirements* needed for RTOS, libc, etc.

If a component only supports some target chips (values of `IDF_TARGET`) then it can specify `REQUIRED_IDF_TARGETS` in the `idf_component_register` call to express these requirements. In this case, the build system will generate an error if the component is included in the build, but does not support the selected target.

Note: In CMake terms, `REQUIRES` & `PRIV_REQUIRES` are approximate wrappers around the CMake functions `target_link_libraries(... PUBLIC ...)` and `target_link_libraries(... PRIVATE ...)`.

Example of Component Requirements

Imagine there is a `car` component, which uses the `engine` component, which uses the `spark_plug` component:

```
- autoProject/
  - CMakeLists.txt
  - components/ - car/ - CMakeLists.txt
                        - car.c
                        - car.h
    - engine/ - CMakeLists.txt
              - engine.c
              - include/ - engine.h
  - spark_plug/ - CMakeLists.txt
                - spark_plug.c
                - spark_plug.h
```

Car Component The `car.h` header file is the public interface for the `car` component. This header includes `engine.h` directly because it uses some declarations from this header:

```
/* car.h */
#include "engine.h"

#ifdef ENGINE_IS_HYBRID
#define CAR_MODEL "Hybrid"
#endif
```

And `car.c` includes `car.h` as well:

```
/* car.c */
#include "car.h"
```

This means the `car/CMakeLists.txt` file needs to declare that `car` requires `engine`:

```
idf_component_register(SRCS "car.c"
                      INCLUDE_DIRS "."
                      REQUIRES engine)
```

- `SRCS` gives the list of source files in the `car` component.
- `INCLUDE_DIRS` gives the list of public include directories for this component. Because the public interface is `car.h`, the directory containing `car.h` is listed here.

- `REQUIRES` gives the list of components required by the public interface of this component. Because `car.h` is a public header and includes a header from `engine`, we include `engine` here. This makes sure that any other component which includes `car.h` will be able to recursively include the required `engine.h` also.

Engine Component The engine component also has a public header file `include/engine.h`, but this header is simpler:

```
/* engine.h */
#define ENGINE_IS_HYBRID

void engine_start(void);
```

The implementation is in `engine.c`:

```
/* engine.c */
#include "engine.h"
#include "spark_plug.h"

...
```

In this component, `engine` depends on `spark_plug` but this is a private dependency. `spark_plug.h` is needed to compile `engine.c`, but not needed to include `engine.h`.

This means that the `engine/CMakeLists.txt` file can use `PRIV_REQUIRES`:

```
idf_component_register(SRCS "engine.c"
                      INCLUDE_DIRS "include"
                      PRIV_REQUIRES spark_plug)
```

As a result, source files in the `car` component don't need the `spark_plug` include directories added to their compiler search path. This can speed up compilation, and stops compiler command lines from becoming longer than necessary.

Spark Plug Component The `spark_plug` component doesn't depend on anything else. It has a public header file `spark_plug.h`, but this doesn't include headers from any other components.

This means that the `spark_plug/CMakeLists.txt` file doesn't need any `REQUIRES` or `PRIV_REQUIRES` clauses:

```
idf_component_register(SRCS "spark_plug.c"
                      INCLUDE_DIRS ".")
```

Source File Include Directories

Each component's source file is compiled with these include path directories, as specified in the passed arguments to `idf_component_register`:

```
idf_component_register(..
                      INCLUDE_DIRS "include"
                      PRIV_INCLUDE_DIRS "other")
```

- The current component's `INCLUDE_DIRS` and `PRIV_INCLUDE_DIRS`.
- The `INCLUDE_DIRS` belonging to all other components listed in the `REQUIRES` and `PRIV_REQUIRES` parameters (ie all the current component's public and private dependencies).
- Recursively, all of the `INCLUDE_DIRS` of those components `REQUIRES` lists (ie all public dependencies of this component's dependencies, recursively expanded).

Main Component Requirements

The component named `main` is special because it automatically requires all other components in the build. So it's not necessary to pass `REQUIRES` or `PRIV_REQUIRES` to this component. See [renaming main](#) for a description of what needs to be changed if no longer using the `main` component.

Common Component Requirements

To avoid duplication, every component automatically requires some "common" IDF components even if they are not mentioned explicitly. Headers from these components can always be included.

The list of common components is: `cxx`, `newlib`, `freertos`, `esp_hw_support`, `heap`, `log`, `soc`, `hal`, `esp_rom`, `esp_common`, `esp_system`, `xtensa/riscv`.

Including Components in the Build

- By default, every component is included in the build.
- If you set the `COMPONENTS` variable to a minimal list of components used directly by your project, then the build will expand to also include required components. The full list of components will be:
 - Components mentioned explicitly in `COMPONENTS`.
 - Those components' requirements (evaluated recursively).
 - The "common" components that every component depends on.
- Setting `COMPONENTS` to the minimal list of required components can significantly reduce compile times.

Circular Dependencies

It's possible for a project to contain Component A that requires (`REQUIRES` or `PRIV_REQUIRES`) Component B, and Component B that requires Component A. This is known as a dependency cycle or a circular dependency.

CMake will usually handle circular dependencies automatically by repeating the component library names twice on the linker command line. However this strategy doesn't always work, and the build may fail with a linker error about "Undefined reference to ...", referencing a symbol defined by one of the components inside the circular dependency. This is particularly likely if there is a large circular dependency, i.e., $A > B > C > D > A$.

The best solution is to restructure the components to remove the circular dependency. In most cases, a software architecture without circular dependencies has desirable properties of modularity and clean layering and will be more maintainable in the long term. However, removing circular dependencies is not always possible.

To bypass a linker error caused by a circular dependency, the simplest workaround is to increase the CMake [LINK_INTERFACE_MULTIPLICITY](#) property of one of the component libraries. This causes CMake to repeat this library and its dependencies more than two times on the linker command line.

For example:

```
set_property(TARGET ${COMPONENT_LIB} APPEND PROPERTY LINK_INTERFACE_MULTIPLICITY 3)
```

- This line should be placed after `idf_component_register` in the component CMakeLists.txt file.
- If possible, place this line in the component that creates the circular dependency by depending on a lot of other components. However, the line can be placed inside any component that is part of the cycle. Choosing the component that owns the source file shown in the linker error message, or the component that defines the symbol(s) mentioned in the linker error message, is a good place to start.
- Usually increasing the value to 3 (default is 2) is enough, but if this doesn't work then try increasing the number further.
- Adding this option will make the linker command line longer, and the linking stage slower.

Advanced Workaround: Undefined Symbols If only one or two symbols are causing a circular dependency, and all other dependencies are linear, then there is an alternative method to avoid linker errors: Specify the specific symbols required for the "reverse" dependency as undefined symbols at link time.

For example, if component A depends on component B but component B also needs to reference `reverse_ops` from component A (but nothing else), then you can add a line like the following to the component B CMakeLists.txt to resolve the cycle at link time:

```
# This symbol is provided by 'Component A' at link time
target_link_libraries(${COMPONENT_LIB} INTERFACE "-u reverse_ops")
```

- The `-u` argument means that the linker will always include this symbol in the link, regardless of dependency ordering.
- This line should be placed after `idf_component_register` in the component CMakeLists.txt file.
- If 'Component B' doesn't need to access any headers of 'Component A', only link to a few symbol(s), then this line can be used instead of any `REQUIRES` from B to A. This further simplifies the component structure in the build system.

See the [target_link_libraries](#) documentation for more information about this CMake function.

Requirements in the Build System Implementation

- Very early in the CMake configuration process, the script `expand_requirements.cmake` is run. This script does a partial evaluation of all component CMakeLists.txt files and builds a graph of component requirements (this *graph may have cycles*). The graph is used to generate a file `component_depends.cmake` in the build directory.
- The main CMake process then includes this file and uses it to determine the list of components to include in the build (internal `BUILD_COMPONENTS` variable). The `BUILD_COMPONENTS` variable is sorted so dependencies are listed first, however, as the component dependency graph has cycles this cannot be guaranteed for all components. The order should be deterministic given the same set of components and component dependencies.
- The value of `BUILD_COMPONENTS` is logged by CMake as "Component names: "
- Configuration is then evaluated for the components included in the build.
- Each component is included in the build normally and the CMakeLists.txt file is evaluated again to add the component libraries to the build.

Component Dependency Order The order of components in the `BUILD_COMPONENTS` variable determines other orderings during the build:

- Order that *Project_include.cmake* files are included in the project.
- Order that the list of header paths is generated for compilation (via `-I` argument). (Note that for a given component's source files, only that component's dependency's header paths are passed to the compiler.)

Adding Link-Time Dependencies The ESP-IDF CMake helper function `idf_component_add_link_dependency` adds a link-only dependency between one component and another. In almost all cases, it is better to use the `PRIV_REQUIRES` feature in `idf_component_register` to create a dependency. However, in some cases, it's necessary to add the link-time dependency of another component to this component, i.e., the reverse order to `PRIV_REQUIRES` (for example: *Overriding Default Chip Drivers*).

To make another component depend on this component at link time:

```
idf_component_add_link_dependency(FROM other_component)
```

Place this line after the line with `idf_component_register`.

It's also possible to specify both components by name:

```
idf_component_add_link_dependency(FROM other_component TO that_component)
```

4.4.9 Overriding Parts of the Project

Project_include.cmake

For components that have build requirements that must be evaluated before any component CMakeLists files are evaluated, you can create a file called `project_include.cmake` in the component directory. This CMake file is included when `project.cmake` is evaluating the entire project.

`project_include.cmake` files are used inside ESP-IDF, for defining project-wide build features such as `esptool.py` command line arguments and the bootloader "special app".

Unlike component `CMakeLists.txt` files, when including a `project_include.cmake` file the current source directory (`CMAKE_CURRENT_SOURCE_DIR` and working directory) is the project directory. Use the variable `COMPONENT_DIR` for the absolute directory of the component.

Note that `project_include.cmake` isn't necessary for the most common component uses, such as adding include directories to the project, or `LDFLAGS` to the final linking step. These values can be customized via the `CMakeLists.txt` file itself. See [Optional Project Variables](#) for details.

`project_include.cmake` files are included in the order given in `BUILD_COMPONENTS` variable (as logged by CMake). This means that a component's `project_include.cmake` file will be included after it's all dependencies' `project_include.cmake` files, unless both components are part of a dependency cycle. This is important if a `project_include.cmake` file relies on variables set by another component. See also [above](#).

Take great care when setting variables or targets in a `project_include.cmake` file. As the values are included in the top-level project CMake pass, they can influence or break functionality across all components!

KConfig.projbuild

This is an equivalent to `project_include.cmake` for [Component Configuration](#) KConfig files. If you want to include configuration options at the top level of `menuconfig`, rather than inside the "Component Configuration" sub-menu, then these can be defined in the `KConfig.projbuild` file alongside the `CMakeLists.txt` file.

Take care when adding configuration values in this file, as they will be included across the entire project configuration. Where possible, it's generally better to create a KConfig file for [Component Configuration](#).

`project_include.cmake` files are used inside ESP-IDF, for defining project-wide build features such as `esptool.py` command line arguments and the bootloader "special app".

Wrappers to Redefine or Extend Existing Functions

Thanks to the linker's wrap feature, it is possible to redefine or extend the behavior of an existing ESP-IDF function. To do so, you will need to provide the following CMake declaration in your project's `CMakeLists.txt` file:

```
target_link_libraries(${COMPONENT_LIB} INTERFACE "-Wl,--wrap=function_to_redefine")
```

Where `function_to_redefine` is the name of the function to redefine or extend. This option will let the linker replace all the calls to `function_to_redefine` functions in the binary libraries with calls to `__wrap_function_to_redefine` function. Thus, you must define this new symbol in your application.

The linker will provide a new symbol named `__real_function_to_redefine` which points to the former implementation of the function to redefine. It can be called from the new implementation, making it an extension of the former one.

This mechanism is shown in the example [build_system/wrappers](#). Check [examples/build_system/wrappers/README.md](#) for more details.

Override the Default Bootloader

Thanks to the optional `bootloader_components` directory present in your ESP-IDF project, it is possible to override the default ESP-IDF bootloader. To do so, a new `bootloader_components/main` component should be defined, which will make the project directory tree look like the following:

- **myProject/**
 - CMakeLists.txt
 - sdkconfig
 - **bootloader_components/ - main/ - CMakeLists.txt**
 - * Kconfig
 - * my_bootloader.c
 - **main/ - CMakeLists.txt**
 - * app_main.c
 - build/

Here the `my_bootloader.c` file becomes source code for the new bootloader, which means that it will need to perform all the required operations to set up and load the `main` application from flash.

It is also possible to conditionally replace the bootloader depending on a certain condition, such as the target for example. This can be achieved thanks to the `BOOTLOADER_IGNORE_EXTRA_COMPONENT` CMake variable. This list can be used to tell the ESP-IDF bootloader project to ignore and not compile the given components present in `bootloader_components`. For example, if one wants to use the default bootloader for ESP32 target, then `myProject/CMakeLists.txt` should look like the following:

```
include($ENV{IDF_PATH}/tools/cmake/project.cmake)

if(${IDF_TARGET} STREQUAL "esp32")
    set(BOOTLOADER_IGNORE_EXTRA_COMPONENT "main")
endif()

project(main)
```

It is important to note that this can also be used for any other bootloader components than `main`. In all cases, the prefix `bootloader_component` must not be specified.

See [custom_bootloader/bootloader_override](#) for an example of overriding the default bootloader.

4.4.10 Configuration-Only Components

Special components which contain no source files, only `Kconfig.projbuild` and `KConfig`, can have a one-line `CMakeLists.txt` file which calls the function `idf_component_register()` with no arguments specified. This function will include the component in the project build, but no library will be built *and* no header files will be added to any included paths.

4.4.11 Debugging CMake

For full details about [CMake](#) and CMake commands, see the [CMake v3.16 documentation](#).

Some tips for debugging the ESP-IDF CMake-based build system:

- When CMake runs, it prints quite a lot of diagnostic information including lists of components and component paths.
- Running `cmake -DDEBUG=1` will produce more verbose diagnostic output from the IDF build system.
- Running `cmake` with the `--trace` or `--trace-expand` options will give a lot of information about control flow. See the [cmake command line documentation](#).

When included from a project CMakeLists file, the `project.cmake` file defines some utility modules and global variables and then sets `IDF_PATH` if it was not set in the system environment.

It also defines an overridden custom version of the built-in `CMake project` function. This function is overridden to add all of the ESP-IDF specific project functionality.

Warning On Undefined Variables

By default, the function of warnings on undefined variables is disabled.

To enable this function, we can pass the `--warn-uninitialized` flag to `CMake` or pass the `--cmake-warn-uninitialized` flag to `idf.py` so it will print a warning if an undefined variable is referenced in the build. This can be very useful to find buggy CMake files.

Browse the [/tools/cmake/project.cmake](#) file and supporting functions in [/tools/cmake/](#) for more details.

4.4.12 Example Component CMakeLists

Because the build environment tries to set reasonable defaults that will work most of the time, component `CMakeLists.txt` can be very small or even empty (see [Minimal Component CMakeLists](#)). However, overriding [pre-set_component_variables](#) is usually required for some functionality.

Here are some more advanced examples of component CMakeLists files.

Adding Conditional Configuration

The configuration system can be used to conditionally compile some files depending on the options selected in the project configuration.

Kconfig:

```
config FOO_ENABLE_BAR
    bool "Enable the BAR feature."
    help
        This enables the BAR feature of the FOO component.
```

CMakeLists.txt:

```
set(srcs "foo.c" "more_foo.c")

if(CONFIG_FOO_ENABLE_BAR)
    list(APPEND srcs "bar.c")
endif()

idf_component_register(SRCS "${srcs}"
    ...)
```

This example makes use of the CMake `if` function and `list APPEND` function.

This can also be used to select or stub out an implementation, as such:

Kconfig:

```
config ENABLE_LCD_OUTPUT
    bool "Enable LCD output."
    help
        Select this if your board has an LCD.

config ENABLE_LCD_CONSOLE
    bool "Output console text to LCD"
    depends on ENABLE_LCD_OUTPUT
    help
        Select this to output debugging output to the LCD
```

(continues on next page)

(continued from previous page)

```
config ENABLE_LCD_PLOT
    bool "Output temperature plots to LCD"
    depends on ENABLE_LCD_OUTPUT
    help
        Select this to output temperature plots
```

CMakeLists.txt:

```
if(CONFIG_ENABLE_LCD_OUTPUT)
    set(srcs lcd-real.c lcd-spi.c)
else()
    set(srcs lcd-dummy.c)
endif()

# We need font if either console or plot is enabled
if(CONFIG_ENABLE_LCD_CONSOLE OR CONFIG_ENABLE_LCD_PLOT)
    list(APPEND srcs "font.c")
endif()

idf_component_register(SRCS "${srcs}"
    ...)
```

Conditions Which Depend on the Target

The current target is available to CMake files via `IDF_TARGET` variable.

In addition to that, if target `xyz` is used (`IDF_TARGET=xyz`), then Kconfig variable `CONFIG_IDF_TARGET_XYZ` will be set.

Note that component dependencies may depend on `IDF_TARGET` variable, but not on Kconfig variables. Also one can not use Kconfig variables in include statements in CMake files, but `IDF_TARGET` can be used in such context.

Source Code Generation

Some components will have a situation where a source file isn't supplied with the component itself but has to be generated from another file. Say our component has a header file that consists of the converted binary data of a BMP file, converted using a hypothetical tool called `bmp2h`. The header file is then included in as C source file called `graphics_lib.c`:

```
add_custom_command(OUTPUT logo.h
    COMMAND bmp2h -i ${COMPONENT_DIR}/logo.bmp -o log.h
    DEPENDS ${COMPONENT_DIR}/logo.bmp
    VERBATIM)

add_custom_target(logo DEPENDS logo.h)
add_dependencies(${COMPONENT_LIB} logo)

set_property(DIRECTORY "${COMPONENT_DIR}" APPEND PROPERTY
    ADDITIONAL_CLEAN_FILES logo.h)
```

This answer is adapted from the [CMake FAQ entry](#), which contains some other examples that will also work with ESP-IDF builds.

In this example, `logo.h` will be generated in the current directory (the build directory) while `logo.bmp` comes with the component and resides under the component path. Because `logo.h` is a generated file, it should be cleaned when the project is cleaned. For this reason, it is added to the [ADDITIONAL_CLEAN_FILES](#) property.

Note: If generating files as part of the project CMakeLists.txt file, not a component CMakeLists.txt, then use build property `PROJECT_DIR` instead of `${COMPONENT_DIR}` and `PROJECT_NAME` .elf instead of `COMPONENT_LIB`}).

If a source file from another component included `logo.h`, then `add_dependencies` would need to be called to add a dependency between the two components, to ensure that the component source files were always compiled in the correct order.

Embedding Binary Data

Sometimes you have a file with some binary or text data that you'd like to make available to your component, but you don't want to reformat the file as a C source.

You can specify argument `MBED_FILES` in the component registration, giving space-delimited names of the files to embed:

```
idf_component_register(...
    EMBED_FILES server_root_cert.der)
```

Or if the file is a string, you can use the variable `MBED_TXTFILES`. This will embed the contents of the text file as a null-terminated string:

```
idf_component_register(...
    EMBED_TXTFILES server_root_cert.pem)
```

The file's contents will be added to the `.rodata` section in flash, and are available via symbol names as follows:

```
extern const uint8_t server_root_cert_pem_start[] asm("_binary_server_root_cert_
↪pem_start");
extern const uint8_t server_root_cert_pem_end[]   asm("_binary_server_root_cert_
↪pem_end");
```

The names are generated from the full name of the file, as given in `EMBED_FILES`. Characters `/`, `.`, etc. are replaced with underscores. The `_binary` prefix in the symbol name is added by objcopy and is the same for both text and binary files.

To embed a file into a project, rather than a component, you can call the function `target_add_binary_data` like this:

```
target_add_binary_data(myproject.elf "main/data.bin" TEXT)
```

Place this line after the `project()` line in your project CMakeLists.txt file. Replace `myproject.elf` with your project name. The final argument can be `TEXT` to embed a null-terminated string, or `BINARY` to embed the content as-is.

For an example of using this technique, see the "main" component of the file_serving example [protocols/http_server/file_serving/main/CMakeLists.txt](https://docs.espressif.com/projects/esp-idf/en/latest/esp32/examples/networking/file_serving/main/CMakeLists.txt) - two files are loaded at build time and linked into the firmware.

It is also possible to embed a generated file:

```
add_custom_command(OUTPUT my_processed_file.bin
    COMMAND my_process_file_cmd my_unprocessed_file.bin)
target_add_binary_data(my_target "my_processed_file.bin" BINARY)
```

In the example above, `my_processed_file.bin` is generated from `my_unprocessed_file.bin` through some command `my_process_file_cmd`, then embedded into the target.

To specify a dependence on a target, use the `DEPENDS` argument:

```
add_custom_target(my_process COMMAND ...)
target_add_binary_data(my_target "my_embed_file.bin" BINARY_DEPENDS my_process)
```

The `DEPENDS` argument to `target_add_binary_data` ensures that the target executes first.

Code and Data Placements

ESP-IDF has a feature called linker script generation that enables components to define where its code and data will be placed in memory through linker fragment files. These files are processed by the build system, and is used to augment the linker script used for linking app binary. See [Linker Script Generation](#) for a quick start guide as well as a detailed discussion of the mechanism.

Fully Overriding the Component Build Process

Obviously, there are cases where all these recipes are insufficient for a certain component, for example when the component is basically a wrapper around another third-party component not originally intended to be compiled under this build system. In that case, it's possible to forego the ESP-IDF build system entirely by using a CMake feature called [ExternalProject](#). Example component CMakeLists:

```
# External build process for quirc, runs in source dir and
# produces libquirc.a
externalproject_add(quirc_build
    PREFIX ${COMPONENT_DIR}
    SOURCE_DIR ${COMPONENT_DIR}/quirc
    CONFIGURE_COMMAND ""
    BUILD_IN_SOURCE 1
    BUILD_COMMAND make CC=${CMAKE_C_COMPILER} libquirc.a
    INSTALL_COMMAND ""
)

# Add libquirc.a to the build process
add_library(quirc STATIC IMPORTED GLOBAL)
add_dependencies(quirc quirc_build)

set_target_properties(quirc PROPERTIES IMPORTED_LOCATION
    ${COMPONENT_DIR}/quirc/libquirc.a)
set_target_properties(quirc PROPERTIES INTERFACE_INCLUDE_DIRECTORIES
    ${COMPONENT_DIR}/quirc/lib)

set_directory_properties( PROPERTIES ADDITIONAL_CLEAN_FILES
    "${COMPONENT_DIR}/quirc/libquirc.a")
```

(The above CMakeLists.txt can be used to create a component named `quirc` that builds the `quirc` project using its own Makefile.)

- `externalproject_add` defines an external build system.
 - `SOURCE_DIR`, `CONFIGURE_COMMAND`, `BUILD_COMMAND` and `INSTALL_COMMAND` should always be set. `CONFIGURE_COMMAND` can be set to an empty string if the build system has no "configure" step. `INSTALL_COMMAND` will generally be empty for ESP-IDF builds.
 - Setting `BUILD_IN_SOURCE` means the build directory is the same as the source directory. Otherwise, you can set `BUILD_DIR`.
 - Consult the [ExternalProject](#) documentation for more details about `externalproject_add()`
- The second set of commands adds a library target, which points to the "imported" library file built by the external system. Some properties need to be set in order to add include directories and tell CMake where this file is.
- Finally, the generated library is added to `ADDITIONAL_CLEAN_FILES`. This means `make clean` will delete this library. (Note that the other object files from the build won't be deleted.)

ExternalProject Dependencies and Clean Builds CMake has some unusual behavior around external project builds:

- `ADDITIONAL_CLEAN_FILES` only works when "make" or "ninja" is used as the build system. If an IDE build system is used, it won't delete these files when cleaning.
- However, the `ExternalProject` configure & build commands will *always* be re-run after a clean is run.
- Therefore, there are two alternative recommended ways to configure the external build command:
 1. Have the external `BUILD_COMMAND` run a full clean compile of all sources. The build command will be run if any of the dependencies passed to `externalproject_add` with `DEPENDS` have changed, or if this is a clean build (ie any of `idf.py clean`, `ninja clean`, or `make clean` was run.)
 2. Have the external `BUILD_COMMAND` be an incremental build command. Pass the parameter `BUILD_ALWAYS 1` to `externalproject_add`. This means the external project will be built each time a build is run, regardless of dependencies. This is only recommended if the external project has correct incremental build behavior, and doesn't take too long to run.

The best of these approaches for building an external project will depend on the project itself, its build system, and whether you anticipate needing to frequently recompile the project.

4.4.13 Custom Sdkconfig Defaults

For example projects or other projects where you don't want to specify a full `sdkconfig` configuration, but you do want to override some key values from the ESP-IDF defaults, it is possible to create a file `sdkconfig.defaults` in the project directory. This file will be used when creating a new config from scratch, or when any new config value hasn't yet been set in the `sdkconfig` file.

To override the name of this file or to specify multiple files, set the `SDKCONFIG_DEFAULTS` environment variable or set `SDKCONFIG_DEFAULTS` in top-level `CMakeLists.txt`. File names that are not specified as full paths are resolved relative to current project's directory.

When specifying multiple files, use a semicolon as the list separator. Files listed first will be applied first. If a particular key is defined in multiple files, the definition in the latter file will override definitions from former files.

Some of the IDF examples include a `sdkconfig.ci` file. This is part of the continuous integration (CI) test framework and is ignored by the normal build process.

Target-dependent Sdkconfig Defaults

If and only if an `sdkconfig.defaults` file exists, the build system will also attempt to load defaults from an `sdkconfig.defaults.TARGET_NAME` file, where `TARGET_NAME` is the value of `IDF_TARGET`. For example, for `esp32` target, default settings will be taken from `sdkconfig.defaults` first, and then from `sd-kconfig.defaults.esp32`. If there are no generic default settings, an empty `sdkconfig.defaults` still needs to be created if the build system should recognize any additional target-dependent `sdkconfig.defaults.TARGET_NAME` files.

If `SDKCONFIG_DEFAULTS` is used to override the name of defaults file/files, the name of target-specific defaults file will be derived from `SDKCONFIG_DEFAULTS` value/values using the rule above. When there are multiple files in `SDKCONFIG_DEFAULTS`, target-specific file will be applied right after the file bringing it in, before all latter files in `SDKCONFIG_DEFAULTS`.

For example, if `SDKCONFIG_DEFAULTS="sdkconfig.defaults; sdkconfig_devkit1"`, and there is a file `sdkconfig.defaults.esp32` in the same folder, then the files will be applied in the following order: (1) `sd-kconfig.defaults` (2) `sd-kconfig.defaults.esp32` (3) `sd-kconfig_devkit1`.

4.4.14 Flash Arguments

There are some scenarios that we want to flash the target board without IDF. For this case we want to save the built binaries, `esptool.py` and `esptool write_flash` arguments. It's simple to write a script to save binaries and `esptool.py`.

After running a project build, the build directory contains binary output files (`.bin` files) for the project and also the following flashing data files:

- `flash_project_args` contains arguments to flash the entire project (app, bootloader, partition table, PHY data if this is configured).
- `flash_app_args` contains arguments to flash only the app.
- `flash_bootloader_args` contains arguments to flash only the bootloader.

You can pass any of these flasher argument files to `esptool.py` as follows:

```
python esptool.py --chip esp32 write_flash @build/flash_project_args
```

Alternatively, it is possible to manually copy the parameters from the argument file and pass them on the command line.

The build directory also contains a generated file `flasher_args.json` which contains project flash information, in JSON format. This file is used by `idf.py` and can also be used by other tools which need information about the project build.

4.4.15 Building the Bootloader

The bootloader is a special "subproject" inside `/components/bootloader/subproject`. It has its own project CMakeLists.txt file and builds separate .ELF and .BIN files to the main project. However, it shares its configuration and build directory with the main project.

The subproject is inserted as an external project from the top-level project, by the file `/components/bootloader/project_include.cmake`. The main build process runs CMake for the subproject, which includes discovering components (a subset of the main components) and generating a bootloader-specific config (derived from the main `sdkconfig`).

4.4.16 Writing Pure CMake Components

The ESP-IDF build system "wraps" CMake with the concept of "components", and helper functions to automatically integrate these components into a project build.

However, underneath the concept of "components" is a full CMake build system. It is also possible to make a component which is pure CMake.

Here is an example minimal "pure CMake" component CMakeLists file for a component named `json`:

```
add_library(json STATIC
cJSON/cJSON.c
cJSON/cJSON_Utils.c)

target_include_directories(json PUBLIC cJSON)
```

- This is actually an equivalent declaration to the IDF `json` component `/components/json/CMakeLists.txt`.
- This file is quite simple as there are not a lot of source files. For components with a large number of files, the globbing behavior of ESP-IDF's component logic can make the component CMakeLists style simpler.)
- Any time a component adds a library target with the component name, the ESP-IDF build system will automatically add this to the build, expose public include directories, etc. If a component wants to add a library target with a different name, dependencies will need to be added manually via CMake commands.

4.4.17 Using Third-Party CMake Projects with Components

CMake is used for a lot of open-source C and C++ projects —code that users can tap into for their applications. One of the benefits of having a CMake build system is the ability to import these third-party projects, sometimes even without modification! This allows for users to be able to get functionality that may not yet be provided by a component, or use another library for the same functionality.

Importing a library might look like this for a hypothetical library `foo` to be used in the `main` component:

```
# Register the component
idf_component_register(...)

# Set values of hypothetical variables that control the build of `foo`
set(FOO_BUILD_STATIC OFF)
set(FOO_BUILD_TESTS OFF)

# Create and import the library targets
add_subdirectory(foo)

# Publicly link `foo` to `main` component
target_link_libraries(main PUBLIC foo)
```

For an actual example, take a look at [build_system/cmake/import_lib](#). Take note that what needs to be done in order to import the library may vary. It is recommended to read up on the library's documentation for instructions on how to import it from other projects. Studying the library's CMakeLists.txt and build structure can also be helpful.

It is also possible to wrap a third-party library to be used as a component in this manner. For example, the [mbedtls](#) component is a wrapper for Espressif's fork of [mbedtls](#). See its [component CMakeLists.txt](#).

The CMake variable `ESP_PLATFORM` is set to 1 whenever the ESP-IDF build system is being used. Tests such as `if (ESP_PLATFORM)` can be used in generic CMake code if special IDF-specific logic is required.

Using ESP-IDF Components from External Libraries

The above example assumes that the external library `foo` (or `tinyclib` in the case of the `import_lib` example) doesn't need to use any ESP-IDF APIs apart from common APIs such as `libc`, `libstdc++`, etc. If the external library needs to use APIs provided by other ESP-IDF components, this needs to be specified in the external CMakeLists.txt file by adding a dependency on the library target `idf::<componentname>`.

For example, in the `foo/CMakeLists.txt` file:

```
add_library(foo bar.c fizz.cpp buzz.cpp)

if(ESP_PLATFORM)
    # On ESP-IDF, bar.c needs to include esp_flash.h from the spi_flash component
    target_link_libraries(foo PRIVATE idf::spi_flash)
endif()
```

4.4.18 Using Prebuilt Libraries with Components

Another possibility is that you have a prebuilt static library (`.a` file), built by some other build process.

The ESP-IDF build system provides a utility function `add_prebuilt_library` for users to be able to easily import and use prebuilt libraries:

```
add_prebuilt_library(target_name lib_path [REQUIRES req1 req2 ...] [PRIV_REQUIRES_
↪req1 req2 ...])
```

where:

- `target_name`- name that can be used to reference the imported library, such as when linking to other targets
- `lib_path`- path to prebuilt library; may be an absolute or relative path to the component directory

Optional arguments `REQUIRES` and `PRIV_REQUIRES` specify dependency on other components. These have the same meaning as the arguments for `idf_component_register`.

Take note that the prebuilt library must have been compiled for the same target as the consuming project. Configuration relevant to the prebuilt library must also match. If not paid attention to, these two factors may contribute to subtle bugs in the app.

For an example, take a look at [build_system/cmake/import_prebuilt](#).

4.4.19 Using ESP-IDF in Custom CMake Projects

ESP-IDF provides a template CMake project for easily creating an application. However, in some instances the user might already have an existing CMake project or may want to create a custom one. In these cases it is desirable to be able to consume IDF components as libraries to be linked to the user's targets (libraries/executables).

It is possible to do so by using the *build system APIs provided* by `tools/cmake/idf.cmake`. For example:

```
cmake_minimum_required(VERSION 3.16)
project(my_custom_app C)

# Include CMake file that provides ESP-IDF CMake build system APIs.
include($ENV{IDF_PATH}/tools/cmake/idf.cmake)

# Include ESP-IDF components in the build, may be thought as an equivalent of
# add_subdirectory() but with some additional processing and magic for ESP-IDF
↳ build
# specific build processes.
idf_build_process(esp32)

# Create the project executable and plainly link the newlib component to it using
# its alias, idf::newlib.
add_executable(${CMAKE_PROJECT_NAME}.elf main.c)
target_link_libraries(${CMAKE_PROJECT_NAME}.elf idf::newlib)

# Let the build system know what the project executable is to attach more targets,
↳ dependencies, etc.
idf_build_executable(${CMAKE_PROJECT_NAME}.elf)
```

The example in `build_system/cmake/idf_as_lib` demonstrates the creation of an application equivalent to `hello world application` using a custom CMake project.

4.4.20 ESP-IDF CMake Build System API

Idf-build-commands

```
idf_build_get_property(var property [GENERATOR_EXPRESSION])
```

Retrieve a *build property* property and store it in *var* accessible from the current scope. Specifying *GENERATOR_EXPRESSION* will retrieve the generator expression string for that property, instead of the actual value, which can be used with CMake commands that support generator expressions.

```
idf_build_set_property(property val [APPEND])
```

Set a *build property* property with value *val*. Specifying *APPEND* will append the specified value to the current value of the property. If the property does not previously exist or it is currently empty, the specified value becomes the first element/member instead.

```
idf_build_component(component_dir [component_source])
```

Present a directory *component_dir* that contains a component to the build system. Relative paths are converted to absolute paths with respect to current directory.

An optional *component_source* argument can be specified to indicate the source of the component. (default: "project_components")

This argument determines the overriding priority for components with the same name. For detailed information, see *Multiple Components with the Same Name*.

This argument supports the following values (from highest to lowest priority):

- "project_components" - project components

- "project_extra_components" - components from EXTRA_COMPONENT_DIRS
- "project_managed_components" - custom project dependencies managed by the IDF Component Manager
- "idf_components" - ESP-IDF built-in components, typically under [/components](#)

For instance, if a component named "json" is present as both "idf_components", and "project_components", the component as "project_components" takes precedence over the one as "idf_components".

Warning: All calls to this command must be performed before *idf_build_process*. This command does not guarantee that the component will be processed during build (see the *COMPONENTS* argument description for *idf_build_process*).

```
idf_build_process(target
    [PROJECT_DIR project_dir]
    [PROJECT_VER project_ver]
    [PROJECT_NAME project_name]
    [SDKCONFIG sdkconfig]
    [SDKCONFIG_DEFAULTS sdkconfig_defaults]
    [BUILD_DIR build_dir]
    [COMPONENTS component1 component2 ...])
```

Performs the bulk of the behind-the-scenes magic for including ESP-IDF components such as component configuration, libraries creation, dependency expansion and resolution. Among these functions, perhaps the most important from a user's perspective is the libraries creation by calling each component's *idf_component_register*. This command creates the libraries for each component, which are accessible using aliases in the form *idf::component_name*. These aliases can be used to link the components to the user's own targets, either libraries or executables.

The call requires the target chip to be specified with *target* argument. Optional arguments for the call include:

- PROJECT_DIR - directory of the project; defaults to CMAKE_SOURCE_DIR
- PROJECT_NAME - name of the project; defaults to CMAKE_PROJECT_NAME
- PROJECT_VER - version/revision of the project; defaults to "1"
- SDKCONFIG - output path of generated sdkconfig file; defaults to PROJECT_DIR/sdkconfig or CMAKE_SOURCE_DIR/sdkconfig depending if PROJECT_DIR is set
- SDKCONFIG_DEFAULTS - list of files containing default config to use in the build (list must contain full paths); defaults to empty. For each value *filename* in the list, the config from file *filename.target*, if it exists, is also loaded.
- BUILD_DIR - directory to place ESP-IDF build-related artifacts, such as generated binaries, text files, components; defaults to CMAKE_BINARY_DIR
- COMPONENTS - select components to process among the components known by the build system (added via *idf_build_component*). This argument is used to trim the build. Other components are automatically added if they are required in the dependency chain, i.e., the public and private requirements of the components in this list are automatically added, and in turn the public and private requirements of those requirements, so on and so forth. If not specified, all components known to the build system are processed.

```
idf_build_executable(executable)
```

Specify the executable *executable* for ESP-IDF build. This attaches additional targets such as dependencies related to flashing, generating additional binary files, etc. Should be called after *idf_build_process*.

```
idf_build_get_config(var config [GENERATOR_EXPRESSION])
```

Get the value of the specified config. Much like build properties, specifying *GENERATOR_EXPRESSION* will retrieve the generator expression string for that config, instead of the actual value, which can be used with CMake commands that support generator expressions. Actual config values are only known after call to *idf_build_process*, however.

Idf-build-properties

These are properties that describe the build. Values of build properties can be retrieved by using the build command `idf_build_get_property`. For example, to get the Python interpreter used for the build:

```
idf_build_get_property(python PYTHON)
message(STATUS "The Python interpreter is: ${python}")
```

- `BUILD_DIR` - build directory; set from `idf_build_process BUILD_DIR` argument
- `BUILD_COMPONENTS` - list of components included in the build; set by `idf_build_process`
- `BUILD_COMPONENT_ALIASES` - list of library alias of components included in the build; set by `idf_build_process`
- `C_COMPILE_OPTIONS` - compile options applied to all components' C source files
- `COMPILE_OPTIONS` - compile options applied to all components' source files, regardless of it being C or C++
- `COMPILE_DEFINITIONS` - compile definitions applied to all component source files
- `CXX_COMPILE_OPTIONS` - compile options applied to all components' C++ source files
- `DEPENDENCIES_LOCK` - lock file path used in component manager. The default value is *dependencies.lock* under the project path.
- `EXECUTABLE` - project executable; set by call to `idf_build_executable`
- `EXECUTABLE_NAME` - name of project executable without extension; set by call to `idf_build_executable`
- `EXECUTABLE_DIR` - path containing the output executable
- `IDF_COMPONENT_MANAGER` - the component manager is enabled by default, but if this property is set to 0 it was disabled by the `IDF_COMPONENT_MANAGER` environment variable
- `IDF_PATH` - ESP-IDF path; set from `IDF_PATH` environment variable, if not, inferred from the location of `idf.cmake`
- `IDF_TARGET` - target chip for the build; set from the required target argument for `idf_build_process`
- `IDF_VER` - ESP-IDF version; set from either a version file or the Git revision of the `IDF_PATH` repository
- `INCLUDE_DIRECTORIES` - include directories for all component source files
- `KCONFIGS` - list of Kconfig files found in components in build; set by `idf_build_process`
- `KCONFIG_PROJBUILDS` - list of Kconfig.projbuild files found in components in build; set by `idf_build_process`
- `PROJECT_NAME` - name of the project; set from `idf_build_process PROJECT_NAME` argument
- `PROJECT_DIR` - directory of the project; set from `idf_build_process PROJECT_DIR` argument
- `PROJECT_VER` - version of the project; set from `idf_build_process PROJECT_VER` argument
- `PYTHON` - Python interpreter used for the build; set from `PYTHON` environment variable if available, if not "python" is used
- `SDKCONFIG` - full path to output config file; set from `idf_build_process SDKCONFIG` argument
- `SDKCONFIG_DEFAULTS` - list of files containing default config to use in the build; set from `idf_build_process SDKCONFIG_DEFAULTS` argument
- `SDKCONFIG_HEADER` - full path to C/C++ header file containing component configuration; set by `idf_build_process`
- `SDKCONFIG_CMAKE` - full path to CMake file containing component configuration; set by `idf_build_process`
- `SDKCONFIG_JSON` - full path to JSON file containing component configuration; set by `idf_build_process`
- `SDKCONFIG_JSON_MENUS` - full path to JSON file containing config menus; set by `idf_build_process`

Idf-component-commands

```
idf_component_get_property(var component property [GENERATOR_EXPRESSION])
```

Retrieve a specified *component's component property*, *property* and store it in *var* accessible from the current scope. Specifying *GENERATOR_EXPRESSION* will retrieve the generator expression string for that property, instead of the actual value, which can be used with CMake commands that support generator expressions.

```
idf_component_set_property(component property val [APPEND])
```

Set a specified *component's component property*, *property* with value *val*. Specifying *APPEND* will append the specified value to the current value of the property. If the property does not previously exist or it is currently empty, the specified value becomes the first element/member instead.

```
idf_component_register([[SRCS src1 src2 ...] | [[SRC_DIRS dir1 dir2 ...] [EXCLUDE_
↪SRCS src1 src2 ...]]
    [INCLUDE_DIRS dir1 dir2 ...]
    [PRIV_INCLUDE_DIRS dir1 dir2 ...]
    [REQUIRES component1 component2 ...]
    [PRIV_REQUIRES component1 component2 ...]
    [LDFRAGMENTS ldfragment1 ldfragment2 ...]
    [REQUIRED_IDF_TARGETS target1 target2 ...]
    [EMBED_FILES file1 file2 ...]
    [EMBED_TXTFILES file1 file2 ...]
    [KCONFIG kconfig]
    [KCONFIG_PROJBUILD kconfig_projbuild]
    [WHOLE_ARCHIVE])
```

Register a component to the build system. Much like the `project()` CMake command, this should be called from the component's `CMakeLists.txt` directly (not through a function or macro) and is recommended to be called before any other command. Here are some guidelines on what commands can **not** be called before `idf_component_register`:

- commands that are not valid in CMake script mode
- custom commands defined in `project_include.cmake`
- build system API commands except `idf_build_get_property`; although consider whether the property may not have been set yet

Commands that set and operate on variables are generally okay to call before `idf_component_register`.

The arguments for `idf_component_register` include:

- **SRCS** - component source files used for creating a static library for the component; if not specified, component is treated as a config-only component and an interface library is created instead.
- **SRC_DIRS**, **EXCLUDE_SRCS** - used to glob source files (.c, .cpp, .S) by specifying directories, instead of specifying source files manually via **SRCS**. Note that this is subject to the *limitations of globbing in CMake*. Source files specified in **EXCLUDE_SRCS** are removed from the globbed files.
- **INCLUDE_DIRS** - paths, relative to the component directory, which will be added to the include search path for all other components which require the current component
- **PRIV_INCLUDE_DIRS** - directory paths, must be relative to the component directory, which will be added to the include search path for this component's source files only
- **REQUIRES** - public component requirements for the component
- **PRIV_REQUIRES** - private component requirements for the component; ignored on config-only components
- **LDFRAGMENTS** - component linker fragment files
- **REQUIRED_IDF_TARGETS** - specify the only target the component supports
- **KCONFIG** - override the default Kconfig file
- **KCONFIG_PROJBUILD** - override the default Kconfig.projbuild file
- **WHOLE_ARCHIVE** - if specified, the component library is surrounded by `-Wl,--whole-archive`, `-Wl,--no-whole-archive` when linked. This has the same effect as setting **WHOLE_ARCHIVE** component property.

The following are used for *embedding data into the component*, and is considered as source files when determining if a component is config-only. This means that even if the component does not specify source files, a static library is still created internally for the component if it specifies either:

- **EMBED_FILES** - binary files to be embedded in the component
- **EMBED_TXTFILES** - text files to be embedded in the component

Idf-component-properties

These are properties that describe a component. Values of component properties can be retrieved by using the build command `idf_component_get_property`. For example, to get the directory of the `freertos` component:

```
idf_component_get_property(dir freertos COMPONENT_DIR)
message(STATUS "The 'freertos' component directory is: ${dir}")
```

- **COMPONENT_ALIAS** - alias for **COMPONENT_LIB** used for linking the component to external targets; set by `idf_build_component` and alias library itself is created by `idf_component_register`
- **COMPONENT_DIR** - component directory; set by `idf_build_component`
- **COMPONENT_OVERRIDEN_DIR** - contains the directory of the original component if *this component overrides another component*
- **COMPONENT_LIB** - name for created component static/interface library; set by `idf_build_component` and library itself is created by `idf_component_register`
- **COMPONENT_NAME** - name of the component; set by `idf_build_component` based on the component directory name
- **COMPONENT_TYPE** - type of the component, whether **LIBRARY** or **CONFIG_ONLY**. A component is of type **LIBRARY** if it specifies source files or embeds a file
- **COMPONENT_SOURCE** - source of the component, one of "idf_components", "project_managed_components", "project_components", "project_extra_components". This is used to determine the override precedence of components with the same name.
- **EMBED_FILES** - list of files to embed in component; set from `idf_component_register` **EMBED_FILES** argument
- **EMBED_TXTFILES** - list of text files to embed in component; set from `idf_component_register` **EMBED_TXTFILES** argument
- **INCLUDE_DIRS** - list of component include directories; set from `idf_component_register` **INCLUDE_DIRS** argument
- **KCONFIG** - component Kconfig file; set by `idf_build_component`
- **KCONFIG_PROJBUILD** - component Kconfig.projbuild; set by `idf_build_component`
- **LDFRAGMENTS** - list of component linker fragment files; set from `idf_component_register` **LDFRAGMENTS** argument
- **MANAGED_PRIV_REQUIRES** - list of private component dependencies added by the IDF component manager from dependencies in `idf_component.yml` manifest file
- **MANAGED_REQUIRES** - list of public component dependencies added by the IDF component manager from dependencies in `idf_component.yml` manifest file
- **PRIV_INCLUDE_DIRS** - list of component private include directories; set from `idf_component_register` **PRIV_INCLUDE_DIRS** on components of type **LIBRARY**
- **PRIV_REQUIRES** - list of private component dependencies; set from value of `idf_component_register` **PRIV_REQUIRES** argument and dependencies in `idf_component.yml` manifest file
- **REQUIRED_IDF_TARGETS** - list of targets the component supports; set from `idf_component_register` **EMBED_TXTFILES** argument
- **REQUIRES** - list of public component dependencies; set from value of `idf_component_register` **REQUIRES** argument and dependencies in `idf_component.yml` manifest file
- **SRCS** - list of component source files; set from **SRCS** or **SRC_DIRS/EXCLUDE_SRCS** argument of `idf_component_register`
- **WHOLE_ARCHIVE** - if this property is set to **TRUE** (or any boolean "true" CMake value: 1, ON, YES, Y), the component library is surrounded by `-Wl, --whole-archive, -Wl, --no-whole-archive` when linked. This can be used to force the linker to include every object file into the executable, even if the object file doesn't resolve any references from the rest of the application. This is commonly used when a component contains plugins or modules which rely on link-time registration. This property is **FALSE** by default. It can be set to **TRUE** from the component CMakeLists.txt file.

4.4.21 File Globbing & Incremental Builds

The preferred way to include source files in an ESP-IDF component is to list them manually via **SRCS** argument to `idf_component_register`:

```
idf_component_register(SRCS library/a.c library/b.c platform/platform.c
                      ...)
```

This preference reflects the [CMake best practice](#) of manually listing source files. This could, however, be inconvenient when there are lots of source files to add to the build. The ESP-IDF build system provides an alternative way for specifying source files using `SRC_DIRS`:

```
idf_component_register(SRC_DIRS library platform
                      ...)
```

This uses globbing behind the scenes to find source files in the specified directories. Be aware, however, that if a new source file is added and this method is used, then CMake won't know to automatically re-run and this file won't be added to the build.

The trade-off is acceptable when you're adding the file yourself, because you can trigger a clean build or run `idf.py reconfigure` to manually re-run [CMake](#). However, the problem gets harder when you share your project with others who may check out a new version using a source control tool like Git...

For components which are part of ESP-IDF, we use a third party Git CMake integration module ([/tools/cmake/third_party/GetGitRevisionDescription.cmake](#)) which automatically re-runs CMake any time the repository commit changes. This means if you check out a new ESP-IDF version, CMake will automatically re-run.

For project components (not part of ESP-IDF), there are a few different options:

- If keeping your project file in Git, ESP-IDF will automatically track the Git revision and re-run CMake if the revision changes.
- If some components are kept in a third git repository (not the project repository or ESP-IDF repository), you can add a call to the `git_describe` function in a component CMakeLists file in order to automatically trigger re-runs of CMake when the Git revision changes.
- If not using Git, remember to manually run `idf.py reconfigure` whenever a source file may change.
- To avoid this problem entirely, use `SRCS` argument to `idf_component_register` to list all source files in project components.

The best option will depend on your particular project and its users.

4.4.22 Build System Metadata

For integration into IDEs and other build systems, when CMake runs the build process generates a number of metadata files in the `build/` directory. To regenerate these files, run `cmake` or `idf.py reconfigure` (or any other `idf.py build` command).

- `compile_commands.json` is a standard format JSON file which describes every source file which is compiled in the project. A CMake feature generates this file, and many IDEs know how to parse it.
- `project_description.json` contains some general information about the ESP-IDF project, configured paths, etc.
- `flasher_args.json` contains `esptool.py` arguments to flash the project's binary files. There are also `flash_*_args` files which can be used directly with `esptool.py`. See [Flash arguments](#).
- `CMakeCache.txt` is the CMake cache file which contains other information about the CMake process, toolchain, etc.
- `config/sdkconfig.json` is a JSON-formatted version of the project configuration values.
- `config/kconfig_menus.json` is a JSON-formatted version of the menus shown in `menuconfig`, for use in external IDE UIs.

JSON Configuration Server

A tool called `kconfserver` is provided to allow IDEs to easily integrate with the configuration system logic. `kconfserver` is designed to run in the background and interact with a calling process by reading and writing JSON over process stdin & stdout.

You can run `kconfserver` from a project via `idf.py confserver` or `ninja kconfserver`, or a similar target triggered from a different build generator.

For more information about `kconfserver`, see the [esp-idf-kconfig documentation](#).

4.4.23 Build System Internals

Build Scripts

The listfiles for the ESP-IDF build system reside in [/tools/cmake](#). The modules which implement core build system functionality are as follows:

- `build.cmake` - Build related commands i.e., build initialization, retrieving/setting build properties, build processing.
- `component.cmake` - Component related commands i.e., adding components, retrieving/setting component properties, registering components.
- `kconfig.cmake` - Generation of configuration files (`sdkconfig`, `sdkconfig.h`, `sdkconfig.cmake`, etc.) from Kconfig files.
- `ldgen.cmake` - Generation of final linker script from linker fragment files.
- `target.cmake` - Setting build target and toolchain file.
- `utilities.cmake` - Miscellaneous helper commands.

Aside from these files, there are two other important CMake scripts in [/tools/cmake](#):

- `idf.cmake` - Sets up the build and includes the core modules listed above. Included in CMake projects in order to access ESP-IDF build system functionality.
- `project.cmake` - Includes `idf.cmake` and provides a custom `project()` command that takes care of all the heavy lifting of building an executable. Included in the top-level `CMakeLists.txt` of standard ESP-IDF projects.

The rest of the files in [/tools/cmake](#) are support or third-party scripts used in the build process.

Build Process

This section describes the standard ESP-IDF application build process. The build process can be broken down roughly into four phases:



Fig. 2: ESP-IDF Build System Process

Initialization This phase sets up necessary parameters for the build.

- **Upon inclusion of `idf.cmake` in `project.cmake`, the following steps are performed:**
 - Set `IDF_PATH` from environment variable or inferred from path to `project.cmake` included in the top-level `CMakeLists.txt`.
 - Add [/tools/cmake](#) to `CMAKE_MODULE_PATH` and include core modules plus the various helper/third-party scripts.
 - Set build tools/executables such as default Python interpreter.
 - Get ESP-IDF git revision and store as `IDF_VER`.
 - Set global build specifications i.e., compile options, compile definitions, include directories for all components in the build.
 - Add components in [components](#) to the build.
- **The initial part of the custom `project()` command performs the following steps:**

- Set `IDF_TARGET` from environment variable or CMake cache and the corresponding `CMAKE_TOOLCHAIN_FILE` to be used.
- Add components in `EXTRA_COMPONENT_DIRS` to the build.
- Prepare arguments for calling command `idf_build_process()` from variables such as `COMPONENTS/EXCLUDE_COMPONENTS`, `SDKCONFIG`, `SDKCONFIG_DEFAULTS`.

The call to `idf_build_process()` command marks the end of this phase.

Enumeration

This phase builds a final list of components to be processed in the build, and is performed in the first half of `idf_build_process()`.

- Retrieve each component's public and private requirements. A child process is created which executes each component's `CMakeLists.txt` in script mode. The values of `idf_component_register` `REQUIRES` and `PRIV_REQUIRES` argument is returned to the parent build process. This is called early expansion. The variable `CMAKE_BUILD_EARLY_EXPANSION` is defined during this step.
- Recursively include components based on public and private requirements.

Processing

This phase processes the components in the build, and is the second half of `idf_build_process()`.

- Load project configuration from `sdkconfig` file and generate an `sdkconfig.cmake` and `sdkconfig.h` header. These define configuration variables/macros that are accessible from the build scripts and C/C++ source/header files, respectively.
- Include each component's `project_include.cmake`.
- Add each component as a subdirectory, processing its `CMakeLists.txt`. The component `CMakeLists.txt` calls the registration command, `idf_component_register` which adds source files, include directories, creates component library, links dependencies, etc.

Finalization

This phase is everything after `idf_build_process()`.

- Create executable and link the component libraries to it.
- Generate project metadata files such as `project_description.json` and display relevant information about the project built.

Browse [/tools/cmake/project.cmake](#) for more details.

4.4.24 Migrating from ESP-IDF GNU Make System

Some aspects of the CMake-based ESP-IDF build system are very similar to the older GNU Make-based system. The developer needs to provide values the include directories, source files etc. There is a syntactical difference, however, as the developer needs to pass these as arguments to the registration command, `idf_component_register`.

Automatic Conversion Tool

An automatic project conversion tool is available in `tools/cmake/convert_to_cmake.py` in ESP-IDF v4.x releases. The script was removed in v5.0 because of its `make` build system dependency.

No Longer Available in CMake

Some features are significantly different or removed in the CMake-based system. The following variables no longer exist in the CMake-based build system:

- `COMPONENT_BUILD_DIR`: Use `CMAKE_CURRENT_BINARY_DIR` instead.
- `COMPONENT_LIBRARY`: Defaulted to `$(COMPONENT_NAME).a`, but the library name could be overridden by the component. The name of the component library can no longer be overridden by the component.
- `CC`, `LD`, `AR`, `OBJCOPY`: Full paths to each tool from the gcc xtensa cross-toolchain. Use `CMAKE_C_COMPILER`, `CMAKE_C_LINK_EXECUTABLE`, `CMAKE_OBJCOPY`, etc instead. [Full list here](#).
- `HOSTCC`, `HOSTLD`, `HOSTAR`: Full names of each tool from the host native toolchain. These are no longer provided, external projects should detect any required host toolchain manually.
- `COMPONENT_ADD_LDFLAGS`: Used to override linker flags. Use the CMake [target_link_libraries](#) command instead.
- `COMPONENT_ADD_LINKER_DEPS`: List of files that linking should depend on. [target_link_libraries](#) will usually infer these dependencies automatically. For linker scripts, use the provided custom CMake function `target_linker_scripts`.
- `COMPONENT_SUBMODULES`: No longer used, the build system will automatically enumerate all submodules in the ESP-IDF repository.
- `COMPONENT_EXTRA_INCLUDES`: Used to be an alternative to `COMPONENT_PRIV_INCLUDEDIRS` for absolute paths. Use `PRIV_INCLUDE_DIRS` argument to `idf_component_register` for all cases now (can be relative or absolute).
- `COMPONENT_OBJS`: Previously, component sources could be specified as a list of object files. Now they can be specified as a list of source files via `SRCS` argument to `idf_component_register`.
- `COMPONENT_OBJEXCLUDE`: Has been replaced with `EXCLUDE_SRCS` argument to `idf_component_register`. Specify source files (as absolute paths or relative to component directory), instead.
- `COMPONENT_EXTRA_CLEAN`: Set property `ADDITIONAL_CLEAN_FILES` instead but note [CMake has some restrictions around this functionality](#).
- `COMPONENT_OWNBUILDTARGET` & `COMPONENT_OWNCLEANTARGET`: Use CMake [ExternalProject](#) instead. See [Fully Overriding the Component Build Process](#) for full details.
- `COMPONENT_CONFIG_ONLY`: Call `idf_component_register` without any arguments instead. See [Configuration-Only Components](#).
- `CFLAGS`, `CPPFLAGS`, `CXXFLAGS`: Use equivalent CMake commands instead. See [Controlling Component Compilation](#).

No Default Values

Unlike in the legacy Make-based build system, the following have no default values:

- Source directories (`COMPONENT_SRCDIRS` variable in Make, `SRC_DIRS` argument to `idf_component_register` in CMake)
- Include directories (`COMPONENT_ADD_INCLUDEDIRS` variable in Make, `INCLUDE_DIRS` argument to `idf_component_register` in CMake)

No Longer Necessary

- In the legacy Make-based build system, it is required to also set `COMPONENT_SRCDIRS` if `COMPONENT_SRCS` is set. In CMake, the equivalent is not necessary i.e., specifying `SRC_DIRS` to `idf_component_register` if `SRCS` is also specified (in fact, `SRCS` is ignored if `SRC_DIRS` is specified).

Flashing from Make

`make flash` and similar targets still work to build and flash. However, project `sdkconfig` no longer specifies serial port and baud rate. Environment variables can be used to override these. See [Flashing with Ninja or Make](#) for more details.

4.5 Core Dump

4.5.1 Overview

A core dump is a set of software state information that is automatically saved by the panic handler when a fatal error occurs. Core dumps are useful for conducting post-mortem analysis of the software's state at the moment of failure. ESP-IDF provides support for generating core dumps.

A core dump contains snapshots of all tasks in the system at the moment of failure, where each snapshot includes a task's control block (TCB) and stack. By analyzing the task snapshots, it is possible to find out what task, at what instruction (line of code), and what call stack of that task lead to the crash. It is also possible to dump the contents of variables on demand, provided those variables are assigned special core dump attributes.

Core dump data is saved to a core dump file according to a particular format, see [Core dump internals](#) for more details. However, ESP-IDF's `idf.py` command provides special subcommands to decode and analyze the core dump file.

4.5.2 Configurations

Destination

The [CONFIG_ESP_COREDUMP_TO_FLASH_OR_UART](#) option enables or disables core dump, and selects the core dump destination if enabled. When a crash occurs, the generated core dump file can either be saved to flash, or output to a connected host over UART.

Format & Size

The [CONFIG_ESP_COREDUMP_DATA_FORMAT](#) option controls the format of the core dump file, namely ELF format or Binary format.

The ELF format contains extended features and allows more information regarding erroneous tasks and crashed software to be saved. However, using the ELF format causes the core dump file to be larger. This format is recommended for new software designs and is flexible enough to be extended in future revisions to save more information.

The Binary format is kept for compatibility reasons. Binary format core dump files are smaller while provide better performance.

The [CONFIG_ESP_COREDUMP_MAX_TASKS_NUM](#) option configures the number of task snapshots saved by the core dump. Crashed task registers and the stack are always saved, regardless of this configuration option. Other tasks are included in order of their priority (starting with the highest-priority ready task).

Core dump data integrity checking is supported via the `Components > Core dump > Core dump data integrity check` option.

Reserved Stack Size

Core dump routines run from a separate stack due to core dump itself needing to parse and save all other task stacks. The [CONFIG_ESP_COREDUMP_STACK_SIZE](#) option controls the size of the core dump's stack in number of bytes.

Setting this option to 0 bytes will cause the core dump routines to run from the ISR stack, thus saving a bit of memory. Setting the option greater than zero will cause a separate stack to be instantiated.

Note: If a separate stack is used, the recommended stack size should be larger than 800 bytes to ensure that the core dump routines themselves do not cause a stack overflow.

4.5.3 Core Dump to Flash

When the core dump file is saved to flash, the file is saved to a special core dump partition in flash. Specifying the core dump partition will reserve space on the flash chip to store the core dump file.

The core dump partition is automatically declared when using the default partition table provided by ESP-IDF. However, when using a custom partition table, you need to declare the core dump partition, as illustrated below:

```
# Name,   Type, SubType, Offset,  Size
# Note: if you have increased the bootloader size, make sure to update the offsets.
# ↳to avoid overlap
nvs,      data, nvs,      0x9000, 0x6000
phy_init, data, phy,      0xf000, 0x1000
factory,  app,  factory,  0x10000, 1M
coredump, data, coredump,,      64K
```

Important: If *Flash Encryption* is enabled on the device, please add an encrypted flag to the core dump partition declaration.

```
coredump, data, coredump,,      64K, encrypted
```

There are no special requirements for the partition name. It can be chosen according to the application's needs, but the partition type should be `data` and the sub-type should be `coredump`. Also, when choosing partition size, note that the core dump file introduces a constant overhead of 20 bytes and a per-task overhead of 12 bytes. This overhead does not include the size of TCB and stack for every task. So the partition size should be at least $20 + \text{max tasks number} \times (12 + \text{TCB size} + \text{max task stack size})$ bytes.

An example of the generic command to analyze core dump from flash is:

```
idf.py coredump-info
```

or

```
idf.py coredump-debug
```

4.5.4 Core Dump to UART

When the core dump file is output to UART, the output file is Base64-encoded. The `CONFIG_ESP_COREDUMP_DECODE` option allows for selecting whether the output file is automatically decoded by the ESP-IDF monitor or kept encoded for manual decoding.

Automatic Decoding

If `CONFIG_ESP_COREDUMP_DECODE` is set to automatically decode the UART core dump, ESP-IDF monitor will automatically decode the data, translate any function addresses to source code lines, and display it in the monitor. The output to ESP-IDF monitor would resemble the following output:

The `CONFIG_ESP_COREDUMP_UART_DELAY` allows for an optional delay to be added before the core dump file is output to UART.

```
=====
===== ESP32 CORE DUMP START =====

Crashed task handle: 0x3ffafba0, name: 'main', GDB name: 'process 1073413024'
Crashed task is not in the interrupt context
Panic reason: abort() was called at PC 0x400d66b9 on core 0
```

(continues on next page)

(continued from previous page)

```

===== CURRENT THREAD REGISTERS =====
exccause      0x1d (StoreProhibitedCause)
excaddr       0x0
epc1          0x40084013
epc2          0x0
...
===== CURRENT THREAD STACK =====
#0  0x4008110d in panic_abort (details=0x3ffb4f0b "abort() was called at PC
↳0x400d66b9 on core 0") at /builds/espressif/esp-idf/components/esp_system/panic.
↳c:472
#1  0x4008510c in esp_system_abort (details=0x3ffb4f0b "abort() was called at PC
↳0x400d66b9 on core 0") at /builds/espressif/esp-idf/components/esp_system/port/
↳esp_system_chip.c:93
...
===== THREADS INFO =====
  Id   Target Id      Frame
* 1    process 1073413024 0x4008110d in panic_abort (details=0x3ffb4f0b "abort()
↳was called at PC 0x400d66b9 on core 0") at /builds/espressif/esp-idf/components/
↳esp_system/panic.c:472
  2    process 1073413368 vPortTaskWrapper (pxCode=0x0, pvParameters=0x0) at /
↳builds/espressif/esp-idf/components/freertos/FreeRTOS-Kernel/portable/xtensa/
↳port.c:133
...
      TCB          NAME PRI0 C/B  STACK USED/FREE
-----
0x3ffaeba0        main      1/1    368/3724
0x3ffaefcf8       IDLE0      0/0    288/1240
0x3ffaef50       IDLE1      0/0    416/1108
...
===== THREAD 1 (TCB: 0x3ffaeba0, name: 'main') =====
#0  0x4008110d in panic_abort (details=0x3ffb4f0b "abort() was called at PC
↳0x400d66b9 on core 0") at /builds/espressif/esp-idf/components/esp_system/panic.
↳c:472
#1  0x4008510c in esp_system_abort (details=0x3ffb4f0b "abort() was called at PC
↳0x400d66b9 on core 0") at /builds/espressif/esp-idf/components/esp_system/port/
↳esp_system_chip.c:93
...
===== THREAD 2 (TCB: 0x3ffaefcf8, name: 'IDLE0')
↳=====
#0  vPortTaskWrapper (pxCode=0x0, pvParameters=0x0) at /builds/espressif/esp-idf/
↳components/freertos/FreeRTOS-Kernel/portable/xtensa/port.c:133
#1  0x40000000 in ?? ()
...
===== ALL MEMORY REGIONS =====
Name      Address      Size  Attrs
...
.iram0.vectors 0x40080000 0x403 R XA
.iram0.text 0x40080404 0xb8ab R XA
.dram0.data 0x3ffb0000 0x2114 RW A
...
===== ESP32 CORE DUMP END =====
=====

```

Manual Decoding

If you set `CONFIG_ESP_COREDUMP_DECODE` to no decoding, then the raw Base64-encoded body of core dump is output to UART between the following header and footer of the UART output:

```
===== CORE DUMP START =====  
<body of Base64-encoded core dump, save it to file on disk>  
===== CORE DUMP END =====
```

It is advised to manually save the core dump text body to a file. The `CORE DUMP START` and `CORE DUMP END` lines must not be included in a core dump text file. The saved text can then be decoded using the following command:

```
idf.py coredump-info -c </path/to/saved/base64/text>
```

or

```
idf.py coredump-debug -c </path/to/saved/base64/text>
```

4.5.5 Core Dump Commands

ESP-IDF provides special commands to help to retrieve and analyze core dumps:

- `idf.py coredump-info` - prints crashed task's registers, call stack, list of available tasks in the system, memory regions, and contents of memory stored in core dump (TCBs and stacks).
- `idf.py coredump-debug` - creates core dump ELF file and runs GDB debug session with this file. You can examine memory, variables, and task states manually. Note that since not all memory is saved in the core dump, only the values of variables allocated on the stack are meaningful.

4.5.6 ROM Functions in Backtraces

It is possible that at the moment of a crash, some tasks and/or the crashed task itself have one or more ROM functions in their call stacks. Since ROM is not part of the program ELF, it is impossible for GDB to parse such call stacks due to GDB analyzing functions' prologues to decode backtraces. Thus, call stack parsing will break with an error message upon the first ROM function that is encountered.

To overcome this issue, the [ROM ELF](#) provided by Espressif is loaded automatically by ESP-IDF monitor based on the target and its revision. More details about ROM ELFs can be found in [esp-rom-elfs](#).

4.5.7 Dumping Variables on Demand

Sometimes you want to read the last value of a variable to understand the root cause of a crash. Core dump supports retrieving variable data over GDB by applying special attributes to declared variables.

Supported Notations and RAM Regions

- `COREDUMP_DRAM_ATTR` places the variable into the DRAM area, which is included in the dump.
- `COREDUMP_RTC_ATTR` places the variable into the RTC area, which is included in the dump.
- `COREDUMP_RTC_FAST_ATTR` places the variable into the RTC_FAST area, which is included in the dump.

Example

1. In *Project Configuration Menu*, enable *COREDUMP TO FLASH*, then save and exit.
2. In your project, create a global variable in the DRAM area, such as:

```
// uint8_t global_var;  
COREDUMP_DRAM_ATTR uint8_t global_var;
```

3. In the main application, set the variable to any value and `assert(0)` to cause a crash.

```
global_var = 25;
assert(0);
```

4. Build, flash, and run the application on a target device and wait for the dumping information.
5. Run the command below to start core dumping in GDB, where `PORT` is the device USB port:

```
idf.py coredump-debug
```

6. In GDB shell, type `p global_var` to get the variable content:

```
(gdb) p global_var
$1 = 25 '\031'
```

4.5.8 Running `idf.py coredump-info` and `idf.py coredump-debug`

`idf.py coredump-info --help` and `idf.py coredump-debug --help` commands can be used to get more details on usage.

Related Documents

Anatomy of Core Dump Image

A core dump file's format can be configured to use the ELF format, or a legacy binary format. The ELF format is recommended for all new designs as it provides more information regarding the software's state at the moment the crash occurs, e.g., CPU registers and memory contents.

The memory state embeds a snapshot of all tasks mapped in the memory space of the program. The CPU state contains register values when the core dump has been generated. The core dump file uses a subset of the ELF structures to register this information.

Loadable ELF segments are used to store the process' memory state, while ELF notes (`ELF.PT_NOTE`) are used to store the process' metadata (e.g., PID, registers, signal etc). In particular, the CPU's status is stored in a note with a special name and type (`CORE, NT_PRSTATUS` type).

Here is an overview of the core dump layout:

Note: The format of the image file shown in the above pictures represents the current version of the image and can be changed in future releases.

Overview of Implementation The figure below describes some basic aspects related to the implementation of the core dump:

Note: The diagram above hides some details and represents the current implementation of the core dump which can be changed later.

4.6 C++ Support

ESP-IDF is primarily written in C and provides C APIs. However, ESP-IDF supports development of applications in C++. This document covers various topics relevant to C++ development.

The following C++ features are supported:




Fig. 3: Core Dump ELF Image Format




Fig. 4: Core Dump Binary Image Format




Fig. 5: Core Dump Implementation Overview

- *Exception Handling*
- *Multithreading*
- *Runtime Type Information (RTTI)*
- *Thread Local Storage* (thread_local keyword)
- All C++ features implemented by GCC, except for some *Limitations*. See [GCC documentation](#) for details on features implemented by GCC.

4.6.1 esp-idf-cxx Component

esp-idf-cxx component provides higher-level C++ APIs for some of the ESP-IDF features. This component is available from the [ESP-IDF Component Registry](#).

4.6.2 C++ Language Standard

By default, ESP-IDF compiles C++ code with C++23 language standard with GNU extensions (-std=gnu++23).

To compile the source code of a certain component using a different language standard, set the desired compiler flag in the component's CMakeLists.txt file:

```
idf_component_register( ... )
target_compile_options(${COMPONENT_LIB} PRIVATE -std=gnu++11)
```

Use PUBLIC instead of PRIVATE if the public header files of the component also need to be compiled with the same language standard.

4.6.3 Multithreading

C++ threads, mutexes, and condition variables are supported. C++ threads are built on top of pthreads, which in turn wrap FreeRTOS tasks.

See [cxx/pthread](#) for an example of creating threads in C++.

Note: The destructor of `std::jthread` can only safely be called from a task that has been created by *Thread APIs* or by the *C++ threading library API*.

4.6.4 Exception Handling

Support for C++ Exceptions in ESP-IDF is disabled by default, but can be enabled using the [CONFIG_COMPILER_CXX_EXCEPTIONS](#) option.

If an exception is thrown, but there is no `catch` block, the program is terminated by the `abort` function, and the backtrace is printed. See [Fatal Errors](#) for more information about backtraces.

C++ Exceptions should **only** be used for exceptional cases, i.e., something happening unexpectedly and occurs rarely, such as events that happen less frequently than 1/100 times. **Do not** use them for control flow (see also the section about resource usage below). For more information on how to use C++ Exceptions, see the [ISO C++ FAQ](#) and [CPP Core Guidelines](#).

See [cxx/exceptions](#) for an example of C++ exception handling.

C++ Exception Handling and Resource Usage

Enabling exception handling normally increases application binary size by a few KB.

Additionally, it may be necessary to reserve some amount of RAM for the exception emergency memory pool. Memory from this pool is used if it is not possible to allocate an exception object from the heap.

The amount of memory in the emergency pool can be set using the [CONFIG_COMPILER_CXX_EXCEPTIONS_EMG_POOL_SIZE](#) variable.

Some additional stack memory (around 200 bytes) is also used if and only if a C++ Exception is actually thrown, because it requires calling some functions from the top of the stack to initiate exception handling.

The run time of code using C++ exceptions depends on what actually happens at run time.

- If no exception is thrown, the code tends to be somewhat faster since there is no need to check error codes.
- If an exception is thrown, the run time of the code that handles exceptions is orders of magnitude slower than code returning an error code.

If an exception is thrown, the run time of the code that unwinds the stack is orders of magnitude slower than code returning an error code. The significance of the increased run time will depend on the application's requirements and implementation of error handling (e.g., requiring user input or messaging to a cloud). As a result, exception-throwing code should never be used in real-time critical code paths.

4.6.5 Runtime Type Information (RTTI)

Support for RTTI in ESP-IDF is disabled by default, but can be enabled using [CONFIG_COMPILER_CXX_RTTI](#) option.

Enabling this option compiles all C++ files with RTTI support enabled, which allows using `dynamic_cast` conversion and `typeid` operator. Enabling this option typically increases the binary size by tens of kB.

See [cxx/rtti](#) for an example of using RTTI in ESP-IDF.

4.6.6 Developing in C++

The following sections provide tips on developing ESP-IDF applications in C++.

Combining C and C++ Code

When an application is developed using both C and C++, it is important to understand the concept of [language linkage](#).

In order for a C++ function to be callable from C code, it has to be both **declared** and **defined** with C linkage (`extern "C"`):

```
// declaration in the .h file:
#ifdef __cplusplus
extern "C" {
#endif

void my_cpp_func(void);

#ifdef __cplusplus
}
#endif

// definition in a .cpp file:
extern "C" void my_cpp_func(void) {
    // ...
}
```

In order for a C function to be callable from C++, it has to be **declared** with C linkage:

```
// declaration in .h file:
#ifdef __cplusplus
extern "C" {
#endif

void my_c_func(void);

#ifdef __cplusplus
}
#endif

// definition in a .c file:
void my_c_func(void) {
    // ...
}
```

Defining app_main in C++

ESP-IDF expects the application entry point, `app_main`, to be defined with C linkage. When `app_main` is defined in a .cpp source file, it has to be designated as `extern "C"`:

```
extern "C" void app_main()
{
}
```

Designated Initializers

Many of the ESP-IDF components use [Configuration Structures](#) as arguments to the initialization functions. ESP-IDF examples written in C routinely use [designated initializers](#) to fill these structures in a readable and a maintainable way.

C and C++ languages have different rules with regards to the designated initializers. For example, C++23 (currently the default in ESP-IDF) does not support out-of-order designated initialization, nested designated initialization, mixing of designated initializers and regular initializers, and designated initialization of arrays. Therefore, when porting ESP-IDF C examples to C++, some changes to the structure initializers may be necessary. See the [C++ aggregate initialization reference](#) for more details.

iostream

`iostream` functionality is supported in ESP-IDF, with a couple of caveats:

1. Normally, ESP-IDF build process eliminates the unused code. However, in the case of `iostreams`, simply including `<iostream>` header in one of the source files significantly increases the binary size by about 200 kB.
2. By default, ESP-IDF uses a simple non-blocking implementation of the standard input stream (`stdin`). To get the usual behavior of `std::cin`, the application has to initialize the UART driver and enable the blocking mode as shown in [common_components/protocol_examples_common/stdin_out.c](#).

4.6.7 Limitations

- Linker script generator does not support function level placements for functions with C++ linkage.
- Various section attributes (such as `IRAM_ATTR`) are ignored when used with template functions.
- Vtables are placed into Flash and are not accessible when the flash cache is disabled. Therefore, virtual function calls should be avoided in [IRAM-Safe Interrupt Handlers](#). Placement of Vtables cannot be adjusted using the linker script generator, yet.
- C++ filesystem (`std::filesystem`) features are not supported.

4.6.8 What to Avoid

Do not use `setjmp/longjmp` in C++. `longjmp` blindly jumps up the stack without calling any destructors, easily introducing undefined behavior and memory leaks. Use C++ exceptions instead, they guarantee correctly calling destructors. If you cannot use C++ exceptions, use alternatives (except `setjmp/longjmp` themselves) such as simple return codes.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.7 Current Consumption Measurement of Modules

You may want to know the current consumption of a [module](#) in deep-sleep mode, [other power-saving modes](#), and active mode to develop some applications sensitive to power consumption. This section introduces how to measure the current consumption of a module running such an application.

4.7.1 Notes to Measurement

Can We Use a Development Board?

How to Choose an Appropriate Ammeter?

In the [deep_sleep](#) example, the module will be woken up every 20 seconds. In deep-sleep mode, the current in the module is just several microamps (μA), while in active mode, the current is in the order of milliamps (mA). The

high dynamic current range makes accurate measurement difficult. Ordinary ammeters cannot dynamically switch the measurement range fast enough.

Additionally, ordinary ammeters have a relatively high internal resistance, resulting in a significant voltage drop. This may cause the module to enter an unstable state, as it is powered by a voltage smaller than the minimum required voltage supply.

Therefore, an ammeter suitable for measuring current in deep-sleep mode should have low internal resistance and, ideally, switch current ranges dynamically. We recommend two options: the [Joulescope ammeter](#) and the [Power Profiler Kit II from Nordic](#).

Joulescope Ammeter The Joulescope ammeter combines high-speed sampling and rapid dynamic current range switching to provide accurate and seamless current and energy measurements, even for devices with rapidly varying current consumption. Joulescope accurately measures electrical current over nine orders of magnitude from amps down to nanoamps. This wide range allows for accurate and precise current measurements for devices. Additionally, Joulescope has a total voltage drop of 25 mV at 1 A, which keeps the module running normally. These two features make Joulescope a perfect option for measuring the module switching between deep-sleep mode and wake-up mode.

Joulescope has no display screen. You need to connect it to a PC to visualize the current waveforms of the measured module. For specific instructions, please follow the documentation provided by the manufacturer.

Nordic Power Profiler Kit II The Nordic Power Profiler Kit II has an advanced analog measurement unit with a high dynamic measurement range. This allows for accurate power consumption measurements for the entire range typically seen in low-power embedded applications, all the way from single μ As to 1 A. The resolution varies between 100 nA and 1 mA, depending on the measurement range, and is high enough to detect small spikes often seen in low-power optimized systems.

4.7.2 Hardware Connection

To measure the power consumption of a bare module, you need an [ESP-Prog](#) to flash the [deep_sleep](#) example to the module and power the module during measurement, a suitable ammeter (here we use the Joulescope ammeter), a computer, and of course a bare module with necessary jumper wires. For the connection, please refer to the following figure.

Please connect the pins of **UART TX**, **UART RX**, **SPI Boot**, **Enable**, and **Ground** on the measured module with corresponding pins on ESP-Prog, and connect the **VPROG** pin on ESP-Prog with the **IN+** port on the Joulescope ammeter and connect its **OUT+** port with the **Power supply (3V3)** pin on the measured module. For the specific names of these pins in different modules, please refer to the list below.

Table 1: Pin Names of Modules Based on ESP32-P4 Chip

Function of Module Pin	Pin Name
UART TX	TXD0
UART RX	RXD0
SPI Boot	Not updated
Enable	EN
Power Supply	3V3
Ground	GND

For details of the pin names, please refer to the [datasheet of specific module](#).

4.7.3 Measurement Steps

ESP32-S3-WROOM-1 is used as an example in the measurement, and other modules can be measured similarly. For the specific current consumption of chips in different modes, please refer to the Current Consumption subsection in the corresponding [chip datasheet](#).

You can refer to the following steps to measure the current in deep-sleep mode.




Fig. 6: Hardware Connection (click to enlarge)

- Connect the aforementioned devices according to the hardware connection.
- Flash the `deep_sleep` example to the module. For details, please refer to [Start a Project on Linux and macOS](#) for a computer with Linux or macOS system or [Start a Project on Windows](#) for a computer with Windows system.
- By default, the module will be woken up every 20 seconds (you can change the timing by modifying the code of this example). To check if the example runs as expected, you can monitor the module operation by running `idf.py -p PORT monitor` (please replace PORT with your serial port name).
- Open the Joulescope software to see the current waveform as shown in the image below.

From the waveforms, you can obtain that the current of the module in deep-sleep mode is 8.14 μA . In addition, you can also see the current of the module in active mode, which is about 23.88 mA. The waveforms also show that the average power consumption during deep-sleep mode is 26.85 μW , and the average power consumption during active mode is 78.32 mW.



Fig. 7: Current Waveform of ESP32-S3-WROOM-1 (click to enlarge)

The figure below shows the total power consumption of one cycle is 6.37 mW.

By referring to these power consumption in different modes, you can estimate the power consumption of your applications and choose the appropriate power source.




Fig. 8: Power Consumption of ESP32-S3-WROOM-1 (click to enlarge)

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.8 Deep Sleep Wake Stubs

ESP32-P4 supports running a "deep sleep wake stub" when coming out of deep sleep. This function runs immediately as soon as the chip wakes up - before any normal initialisation, bootloader, or ESP-IDF code has run. After the wake stub runs, the SoC can go back to sleep or continue to start ESP-IDF normally.

Deep sleep wake stub code is loaded into "RTC Fast Memory" and any data which it uses must also be loaded into RTC memory. RTC memory regions hold their contents during deep sleep.

4.8.1 Rules for Wake Stubs

Wake stub code must be carefully written:

- As the SoC has freshly woken from sleep, most of the peripherals are in reset states. The SPI flash is unmapped.
- The wake stub code can only call functions implemented in ROM or loaded into RTC Fast Memory (see below.)
- The wake stub code can only access data loaded in RTC memory. All other RAM will be uninitialised and have random contents. The wake stub can use other RAM for temporary storage, but the contents will be overwritten when the SoC goes back to sleep or starts ESP-IDF.
- RTC memory must include any read-only data (.rodata) used by the stub.
- Data in RTC memory is initialised whenever the SoC restarts, except when waking from deep sleep. When waking from deep sleep, the values which were present before going to sleep are kept.
- Wake stub code is a part of the main esp-idf app. During normal running of esp-idf, functions can call the wake stub functions or access RTC memory. It is as if these were regular parts of the app.

4.8.2 Implementing A Stub

The wake stub in esp-idf is called `esp_wake_deep_sleep()`. This function runs whenever the SoC wakes from deep sleep. There is a default version of this function provided in esp-idf, but the default function is weak-linked so if your app contains a function named `esp_wake_deep_sleep()` then this will override the default.

If supplying a custom wake stub, the first thing it does should be to call `esp_default_wake_deep_sleep()`.

It is not necessary to implement `esp_wake_deep_sleep()` in your app in order to use deep sleep. It is only necessary if you want to have special behaviour immediately on wake.

If you want to swap between different deep sleep stubs at runtime, it is also possible to do this by calling the `esp_set_deep_sleep_wake_stub()` function. This is not necessary if you only use the default `esp_wake_deep_sleep()` function.

All of these functions are declared in the `esp_sleep.h` header under `components/esp32p4`.

4.8.3 Loading Code Into RTC Memory

Wake stub code must be resident in RTC Fast Memory. This can be done in one of two ways.

The first way is to use the `RTC_IRAM_ATTR` attribute to place a function into RTC memory:

```
void RTC_IRAM_ATTR esp_wake_deep_sleep(void) {  
    esp_default_wake_deep_sleep();  
    // Add additional functionality here  
}
```

The second way is to place the function into any source file whose name starts with `rtc_wake_stub`. Files names `rtc_wake_stub*` have their contents automatically put into RTC memory by the linker.

The first way is simpler for very short and simple code, or for source files where you want to mix "normal" and "RTC" code. The second way is simpler when you want to write longer pieces of code for RTC memory.

4.8.4 Loading Data Into RTC Memory

Data used by stub code must be resident in RTC memory.

Specifying this data can be done in one of two ways:

The first way is to use the `RTC_DATA_ATTR` and `RTC_RODATA_ATTR` to specify any data (writeable or read-only, respectively) which should be loaded into RTC memory:

```
RTC_DATA_ATTR int wake_count;

void RTC_IRAM_ATTR esp_wake_deep_sleep(void) {
    esp_default_wake_deep_sleep();
    static RTC_RODATA_ATTR const char fmt_str[] = "Wake count %d\n";
    esp_rom_printf(fmt_str, wake_count++);
}
```

The attributes `RTC_FAST_ATTR` and `RTC_SLOW_ATTR` can be used to specify data that will be force placed into `RTC_FAST` and `RTC_SLOW` memory respectively, but for ESP32-P4 there is only RTC fast memory, so both attributes will map to this region.

Unfortunately, any string constants used in this way must be declared as arrays and marked with `RTC_RODATA_ATTR`, as shown in the example above.

The second way is to place the data into any source file whose name starts with `rtc_wake_stub`.

For example, the equivalent example in `rtc_wake_stub_counter.c`:

```
int wake_count;

void RTC_IRAM_ATTR esp_wake_deep_sleep(void) {
    esp_default_wake_deep_sleep();
    esp_rom_printf("Wake count %d\n", wake_count++);
}
```

The second way is a better option if you need to use strings, or write other more complex code.

To reduce wake-up time use the `CONFIG_BOOTLOADER_SKIP_VALIDATE_IN_DEEP_SLEEP` Kconfig option, see more information in [Fast boot from Deep Sleep](#).

4.8.5 CRC Check For Wake Stubs

During deep sleep, only the wake stubs area of RTC Fast memory is validated with CRC. When ESP32-P4 wakes up from deep sleep, the wake stubs area is validated again. If the validation passes, the wake stubs code will be executed. Otherwise, the normal initialization, bootloader, and esp-idf codes will be executed.

Note: When the `CONFIG_ESP_SYSTEM_ALLOW_RTC_FAST_MEM_AS_HEAP` option is enabled, all the RTC fast memory except the wake stubs area is added to the heap.

4.8.6 Example

ESP-IDF provides an example to show how to implement the Deep-sleep wake stub.

- [system/deep_sleep_wake_stub](#)

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

4.9 Error Handling

4.9.1 Overview

Identifying and handling run-time errors is important for developing robust applications. There can be multiple kinds of run-time errors:

- Recoverable errors:
 - Errors indicated by functions through return values (error codes)
 - C++ exceptions, thrown using `throw` keyword
- Unrecoverable (fatal) errors:
 - Failed assertions (using `assert` macro and equivalent methods, see [Assertions](#)) and `abort()` calls.
 - CPU exceptions: access to protected regions of memory, illegal instruction, etc.
 - System level checks: watchdog timeout, cache access error, stack overflow, stack smashing, heap corruption, etc.

This guide explains ESP-IDF error handling mechanisms related to recoverable errors, and provides some common error handling patterns.

For instructions on diagnosing unrecoverable errors, see [Fatal Errors](#).

4.9.2 Error Codes

The majority of ESP-IDF-specific functions use `esp_err_t` type to return error codes. `esp_err_t` is a signed integer type. Success (no error) is indicated with `ESP_OK` code, which is defined as zero.

Various ESP-IDF header files define possible error codes using preprocessor defines. Usually these defines start with `ESP_ERR_` prefix. Common error codes for generic failures (out of memory, timeout, invalid argument, etc.) are defined in `esp_err.h` file. Various components in ESP-IDF may define additional error codes for specific situations.

For the complete list of error codes, see [Error Code Reference](#).

4.9.3 Converting Error Codes to Error Messages

For each error code defined in ESP-IDF components, `esp_err_t` value can be converted to an error code name using `esp_err_to_name()` or `esp_err_to_name_r()` functions. For example, passing `0x101` to `esp_err_to_name()` will return "ESP_ERR_NO_MEM" string. Such strings can be used in log output to make it easier to understand which error has happened.

Additionally, `esp_err_to_name_r()` function will attempt to interpret the error code as a [standard POSIX error code](#), if no matching `ESP_ERR_` value is found. This is done using `strerror_r` function. POSIX error codes (such as `ENOENT`, `ENOMEM`) are defined in `errno.h` and are typically obtained from `errno` variable. In ESP-IDF this variable is thread-local: multiple FreeRTOS tasks have their own copies of `errno`. Functions which set `errno` only modify its value for the task they run in.

This feature is enabled by default, but can be disabled to reduce application binary size. See [CONFIG_ESP_ERR_TO_NAME_LOOKUP](#). When this feature is disabled, `esp_err_to_name()` and

`esp_err_to_name_r()` are still defined and can be called. In this case, `esp_err_to_name()` will return UNKNOWN_ERROR, and `esp_err_to_name_r()` will return Unknown error 0XXXXX(YYYYY), where 0XXXXX and YYYYY are the hexadecimal and decimal representations of the error code, respectively.

4.9.4 ESP_ERROR_CHECK Macro

`ESP_ERROR_CHECK` macro serves similar purpose as `assert`, except that it checks `esp_err_t` value rather than a `bool` condition. If the argument of `ESP_ERROR_CHECK` is not equal `ESP_OK`, then an error message is printed on the console, and `abort()` is called.

Error message will typically look like this:

```
ESP_ERROR_CHECK failed: esp_err_t 0x107 (ESP_ERR_TIMEOUT) at 0x400d1fdf

file: "/Users/user/esp/example/main/main.c" line 20
func: app_main
expression: sdmmc_card_init(host, &card)

Backtrace: 0x40086e7c:0x3ffb4ff0 0x40087328:0x3ffb5010 0x400d1fdf:0x3ffb5030_
↳ 0x400d0816:0x3ffb5050
```

Note: If *ESP-IDF monitor* is used, addresses in the backtrace will be converted to file names and line numbers.

- The first line mentions the error code as a hexadecimal value, and the identifier used for this error in source code. The latter depends on `CONFIG_ESP_ERR_TO_NAME_LOOKUP` option being set. Address in the program where error has occurred is printed as well.
- Subsequent lines show the location in the program where `ESP_ERROR_CHECK` macro was called, and the expression which was passed to the macro as an argument.
- Finally, backtrace is printed. This is part of panic handler output common to all fatal errors. See *Fatal Errors* for more information about the backtrace.

4.9.5 ESP_ERROR_CHECK_WITHOUT_ABORT Macro

`ESP_ERROR_CHECK_WITHOUT_ABORT` macro serves similar purpose as `ESP_ERROR_CHECK`, except that it will not call `abort()`.

4.9.6 ESP_RETURN_ON_ERROR Macro

`ESP_RETURN_ON_ERROR` macro checks the error code, if the error code is not equal `ESP_OK`, it prints the message and returns.

4.9.7 ESP_GOTO_ON_ERROR Macro

`ESP_GOTO_ON_ERROR` macro checks the error code, if the error code is not equal `ESP_OK`, it prints the message, sets the local variable `ret` to the code, and then exits by jumping to `goto_tag`.

4.9.8 ESP_RETURN_ON_FALSE Macro

`ESP_RETURN_ON_FALSE` macro checks the condition, if the condition is not equal `true`, it prints the message and returns with the supplied `err_code`.

4.9.9 ESP_GOTO_ON_FALSE Macro

`ESP_GOTO_ON_FALSE` macro checks the condition, if the condition is not equal *true*, it prints the message, sets the local variable *ret* to the supplied *err_code*, and then exits by jumping to *goto_tag*.

4.9.10 CHECK MACROS Examples

Some examples:

```
static const char* TAG = "Test";

esp_err_t test_func(void)
{
    esp_err_t ret = ESP_OK;

    ESP_ERROR_CHECK(x); // err message_
    ↪printed if `x` is not `ESP_OK`, and then `abort()``.
    ESP_ERROR_CHECK_WITHOUT_ABORT(x); // err message_
    ↪printed if `x` is not `ESP_OK`, without `abort()``.
    ESP_RETURN_ON_ERROR(x, TAG, "fail reason 1"); // err message_
    ↪printed if `x` is not `ESP_OK`, and then function returns with code `x`.
    ESP_GOTO_ON_ERROR(x, err, TAG, "fail reason 2"); // err message_
    ↪printed if `x` is not `ESP_OK`, `ret` is set to `x`, and then jumps to `err`.
    ESP_RETURN_ON_FALSE(a, err_code, TAG, "fail reason 3"); // err message_
    ↪printed if `a` is not `true`, and then function returns with code `err_code`.
    ESP_GOTO_ON_FALSE(a, err_code, err, TAG, "fail reason 4"); // err message_
    ↪printed if `a` is not `true`, `ret` is set to `err_code`, and then jumps to_
    ↪`err`.

err:
    // clean up
    return ret;
}
```

Note: If the option `CONFIG_COMPILER_OPTIMIZATION_CHECKS_SILENT` in Kconfig is enabled, the err message will be discarded, while the other action works as is.

The `ESP_RETURN_XX` and `ESP_GOTO_xx` macros cannot be called from ISR. While there are `xx_ISR` versions for each of them, e.g., `ESP_RETURN_ON_ERROR_ISR`, these macros could be used in ISR.

4.9.11 Error Handling Patterns

1. Attempt to recover. Depending on the situation, we may try the following methods:

- retry the call after some time;
- attempt to de-initialize the driver and re-initialize it again;
- fix the error condition using an out-of-band mechanism (e.g reset an external peripheral which is not responding).

Example:

```
esp_err_t err;
do {
    err = sdio_slave_send_queue(addr, len, arg, timeout);
    // keep retrying while the sending queue is full
} while (err == ESP_ERR_TIMEOUT);
if (err != ESP_OK) {
    // handle other errors
}
```

2. Propagate the error to the caller. In some middleware components this means that a function must exit with the same error code, making sure any resource allocations are rolled back.

Example:

```
sdmmc_card_t* card = calloc(1, sizeof(sdmmc_card_t));
if (card == NULL) {
    return ESP_ERR_NO_MEM;
}
esp_err_t err = sdmmc_card_init(host, &card);
if (err != ESP_OK) {
    // Clean up
    free(card);
    // Propagate the error to the upper layer (e.g., to notify the user).
    // Alternatively, application can define and return custom error code.
    return err;
}
```

3. Convert into unrecoverable error, for example using `ESP_ERROR_CHECK`. See [ESP_ERROR_CHECK macro](#) section for details.

Terminating the application in case of an error is usually undesirable behavior for middleware components, but is sometimes acceptable at application level.

Many ESP-IDF examples use `ESP_ERROR_CHECK` to handle errors from various APIs. This is not the best practice for applications, and is done to make example code more concise.

Example:

```
ESP_ERROR_CHECK(spi_bus_initialize(host, bus_config, dma_chan));
```

4.9.12 C++ Exceptions

See [Exception Handling](#).

4.10 Support for External RAM

4.10.1 Introduction

ESP32-P4 has a few hundred kilobytes of internal RAM, residing on the same die as the rest of the chip components. It can be insufficient for some purposes, so ESP32-P4 has the ability to use up to 64 MB of virtual addresses for external PSRAM (Pseudostatic RAM) memory. The external memory is incorporated in the memory map and, with certain restrictions, is usable in the same way as internal data RAM.

4.10.2 Hardware

ESP32-P4 supports PSRAM connected in parallel with the SPI flash chip. While ESP32-P4 is capable of supporting several types of RAM chips, ESP-IDF currently only supports Espressif branded PSRAM chips (e.g., ESP-PSRAM32, ESP-PSRAM64, etc).

Note: Some PSRAM chips are 1.8 V devices and some are 3.3 V. The working voltage of the PSRAM chip must match the working voltage of the flash component. Consult the datasheet for your PSRAM chip and ESP32-P4 device to find out the working voltages. For a 1.8 V PSRAM chip, make sure to either set the MTDI pin to a high signal level on bootup, or program ESP32-P4 eFuses to always use the VDD_SIO level of 1.8 V. Not doing this can damage the PSRAM and/or flash chip.

Note: Espressif produces both modules and system-in-package chips that integrate compatible PSRAM and flash and are ready to mount on a product PCB. Consult the Espressif website for more information. If you are using a custom PSRAM chip, ESP-IDF SDK might not be compatible with it.

For specific details about connecting the SoC or module pins to an external PSRAM chip, consult the SoC or module datasheet.

4.10.3 Configuring External RAM

ESP-IDF fully supports the use of external RAM in applications. Once the external RAM is initialized at startup, ESP-IDF can be configured to integrate the external RAM in several ways:

- *Integrate RAM into the ESP32-P4 Memory Map*
- *Add External RAM to the Capability Allocator*
- *Provide External RAM via malloc() (default)*
- *Allow .bss Segment to Be Placed in External Memory*
- *Allow .noinit Segment to Be Placed in External Memory*

Integrate RAM into the ESP32-P4 Memory Map

Select this option by choosing `Integrate RAM into memory map` from [*CONFIG_SPIRAM_USE*](#).

This is the most basic option for external RAM integration. Most likely, you will need another, more advanced option.

During the ESP-IDF startup, external RAM is mapped into the data virtual address space. The address space is dynamically allocated. The length will be the minimum length between the PSRAM size and the available data virtual address space size.

Applications can manually place data in external memory by creating pointers to this region. So if an application uses external memory, it is responsible for all management of the external RAM: coordinating buffer usage, preventing corruption, etc.

It is recommended to access the PSRAM by ESP-IDF heap memory allocator (see next chapter).

Add External RAM to the Capability Allocator

Select this option by choosing `Make RAM allocatable using heap_caps_malloc(..., MALLOC_CAP_SPIRAM)` from [*CONFIG_SPIRAM_USE*](#).

When enabled, memory is mapped to data virtual address space and also added to the *capabilities-based heap memory allocator* using `MALLOC_CAP_SPIRAM`.

To allocate memory from external RAM, a program should call `heap_caps_malloc(size, MALLOC_CAP_SPIRAM)`. After use, this memory can be freed by calling the normal `free()` function.

Provide External RAM via malloc()

Select this option by choosing `Make RAM allocatable using malloc() as well` from [*CONFIG_SPIRAM_USE*](#). This is the default option.

In this case, memory is added to the capability allocator as described for the previous option. However, it is also added to the pool of RAM that can be returned by the standard `malloc()` function.

This allows any application to use the external RAM without having to rewrite the code to use `heap_caps_malloc(..., MALLOC_CAP_SPIRAM)`.

An additional configuration item, [*CONFIG_SPIRAM_MALLOC_ALWAYSINTERNAL*](#), can be used to set the size threshold when a single allocation should prefer external memory:

- When allocating a size less than or equal to the threshold, the allocator will try internal memory first.
- When allocating a size larger than the threshold, the allocator will try external memory first.

If a suitable block of preferred internal/external memory is not available, the allocator will try the other type of memory.

Because some buffers can only be allocated in internal memory, a second configuration item [*CONFIG_SPIRAM_MALLOC_RESERVE_INTERNAL*](#) defines a pool of internal memory which is reserved for *only* explicitly internal allocations (such as memory for DMA use). Regular `malloc()` will not allocate from this pool. The [*MALLOC_CAP_DMA*](#) and `MALLOC_CAP_INTERNAL` flags can be used to allocate memory from this pool.

Allow .bss Segment to Be Placed in External Memory

Enable this option by checking [*CONFIG_SPIRAM_ALLOW_BSS_SEG_EXTERNAL_MEMORY*](#).

If enabled, the region of the data virtual address space where the PSRAM is mapped to will be used to store zero-initialized data (BSS segment) from the lwIP, net80211, libpp, wpa_supplicant and bluedroid ESP-IDF libraries.

Additional data can be moved from the internal BSS segment to external RAM by applying the macro `EXT_RAM_BSS_ATTR` to any static declaration (which is not initialized to a non-zero value).

It is also possible to place the BSS section of a component or a library to external RAM using linker fragment scheme `extram_bss`.

This option reduces the internal static memory used by the BSS segment.

Remaining external RAM can also be added to the capability heap allocator using the method shown above.

Allow .noinit Segment to Be Placed in External Memory

Enable this option by checking [*CONFIG_SPIRAM_ALLOW_NOINIT_SEG_EXTERNAL_MEMORY*](#). If enabled, the region of the data virtual address space where the PSRAM is mapped to will be used to store non-initialized data. The values placed in this segment will not be initialized or modified even during startup or restart.

By applying the macro `EXT_RAM_NOINIT_ATTR`, data could be moved from the internal NOINIT segment to external RAM. Remaining external RAM can still be added to the capability heap allocator using the method shown above, [Add External RAM to the Capability Allocator](#).

4.10.4 Restrictions

External RAM use has the following restrictions:

- When flash cache is disabled (for example, if the flash is being written to), the external RAM also becomes inaccessible. Any read operations from or write operations to it will lead to an illegal cache access exception. This is also the reason why ESP-IDF does not by default allocate any task stacks in external RAM (see below).
- External RAM uses the same cache region as the external flash. This means that frequently accessed variables in external RAM can be read and modified almost as quickly as in internal RAM. However, when accessing large chunks of data (> 32 KB), the cache can be insufficient, and speeds will fall back to the access speed of the external RAM. Moreover, accessing large chunks of data can "push out" cached flash, possibly making the execution of code slower afterwards.
- In general, external RAM will not be used as task stack memory. [*xTaskCreate\(\)*](#) and similar functions will always allocate internal memory for stack and task TCBs.

The option [*CONFIG_SPIRAM_ALLOW_STACK_EXTERNAL_MEMORY*](#) can be used to allow placing task stacks into external memory. In these cases [*xTaskCreateStatic\(\)*](#) must be used to specify a task stack buffer allocated from external memory, otherwise task stacks will still be allocated from internal memory.

4.10.5 Failure to Initialize

By default, failure to initialize external RAM will cause the ESP-IDF startup to abort. This can be disabled by enabling the config item [CONFIG_SPIRAM_IGNORE_NOTFOUND](#).

4.10.6 Encryption

It is possible to enable automatic encryption for data stored in external RAM. When this is enabled any data read and written through the cache will automatically be encrypted or decrypted by the external memory encryption hardware.

This feature is enabled whenever flash encryption is enabled. For more information on how to enable and how it works see [Flash Encryption](#).

4.11 Fatal Errors

4.11.1 Overview

In certain situations, the execution of the program can not be continued in a well-defined way. In ESP-IDF, these situations include:

- CPU Exceptions: Illegal Instruction, Load/Store Alignment Error, Load/Store Prohibited error.
- System level checks and safeguards:
 - [Interrupt watchdog](#) timeout
 - [Task watchdog](#) timeout (only fatal if [CONFIG_ESP_TASK_WDT_PANIC](#) is set)
 - Cache access error
 - Brownout detection event
 - Stack overflow
 - Stack smashing protection check
 - Heap integrity check
 - Undefined behavior sanitizer (UBSAN) checks
- Failed assertions, via `assert`, `configASSERT` and similar macros.

This guide explains the procedure used in ESP-IDF for handling these errors, and provides suggestions on troubleshooting the errors.

4.11.2 Panic Handler

Every error cause listed in the [Overview](#) will be handled by the *panic handler*.

The panic handler will start by printing the cause of the error to the console. For CPU exceptions, the message will be similar to

```
Guru Meditation Error: Core 0 panic'ed (Illegal instruction). Exception_
↳was unhandled.
```

For some of the system level checks (interrupt watchdog, cache access error), the message will be similar to

```
Guru Meditation Error: Core 0 panic'ed (Cache error). Exception was_
↳unhandled.
```

In all cases, the error cause will be printed in parentheses. See [Guru Meditation Errors](#) for a list of possible error causes.

Subsequent behavior of the panic handler can be set using [CONFIG_ESP_SYSTEM_PANIC](#) configuration choice. The available options are:

- Print registers and reboot (`CONFIG_ESP_SYSTEM_PANIC_PRINT_REBOOT`) —default option.
This will print register values at the point of the exception, print the backtrace, and restart the chip.
- Print registers and halt (`CONFIG_ESP_SYSTEM_PANIC_PRINT_HALT`)
Similar to the above option, but halt instead of rebooting. External reset is required to restart the program.
- Silent reboot (`CONFIG_ESP_SYSTEM_PANIC_SILENT_REBOOT`)
Do not print registers or backtrace, restart the chip immediately.
- Invoke GDB Stub (`CONFIG_ESP_SYSTEM_PANIC_GDBSTUB`)
Start GDB server which can communicate with GDB over console UART port. This option will only provide read-only debugging or post-mortem debugging. See [GDB Stub](#) for more details.

Note: The `CONFIG_ESP_SYSTEM_PANIC_GDBSTUB` choice in the configuration option `CONFIG_ESP_SYSTEM_PANIC` is only available when the component `esp_gdbstub` is included in the build.

The behavior of the panic handler is affected by three other configuration options.

- If `CONFIG_ESP_DEBUG_OCDAWARE` is enabled (which is the default), the panic handler will detect whether a JTAG debugger is connected. If it is, execution will be halted and control will be passed to the debugger. In this case, registers and backtrace are not dumped to the console, and GDBStub / Core Dump functions are not used.
- If the *Core Dump* feature is enabled, then the system state (task stacks and registers) will be dumped to either Flash or UART, for later analysis.
- If `CONFIG_ESP_PANIC_HANDLER_IRAM` is disabled (disabled by default), the panic handler code is placed in flash memory, not IRAM. This means that if ESP-IDF crashes while flash cache is disabled, the panic handler will automatically re-enable flash cache before running GDB Stub or Core Dump. This adds some minor risk, if the flash cache status is also corrupted during the crash.
If this option is enabled, the panic handler code (including required UART functions) is placed in IRAM, and hence will decrease the usable memory space in SRAM. But this may be necessary to debug some complex issues with crashes while flash cache is disabled (for example, when writing to SPI flash) or when flash cache is corrupted when an exception is triggered.
- If `CONFIG_ESP_SYSTEM_PANIC_REBOOT_DELAY_SECONDS` is enabled (disabled by default) and set to a number higher than 0, the panic handler will delay the reboot for that amount of time in seconds. This can help if the tool used to monitor serial output does not provide a possibility to stop and examine the serial output. In that case, delaying the reboot will allow users to examine and debug the panic handler output (backtrace, etc.) for the duration of the delay. After the delay, the device will reboot. The reset reason is preserved.

The following diagram illustrates the panic handler behavior:

4.11.3 Register Dump and Backtrace

Unless the `CONFIG_ESP_SYSTEM_PANIC_SILENT_REBOOT` option is enabled, the panic handler prints some of the CPU registers, and the backtrace, to the console

```
Core 0 register dump:
MEPC   : 0x420048b4  RA      : 0x420048b4  SP      : 0x3fc8f2f0  GP      : 0x3fc8a600
TP     : 0x3fc8a2ac  T0     : 0x40057fa6  T1     : 0x0000000f  T2     : 0x00000000
S0/FP  : 0x00000000  S1     : 0x00000000  A0     : 0x00000001  A1     : 0x00000001
A2     : 0x00000064  A3     : 0x00000004  A4     : 0x00000001  A5     : 0x00000000
A6     : 0x42001fd6  A7     : 0x00000000  S2     : 0x00000000  S3     : 0x00000000
S4     : 0x00000000  S5     : 0x00000000  S6     : 0x00000000  S7     : 0x00000000
S8     : 0x00000000  S9     : 0x00000000  S10    : 0x00000000  S11    : 0x00000000
T3     : 0x00000000  T4     : 0x00000000  T5     : 0x00000000  T6     : 0x00000000
```

(continues on next page)




Fig. 9: Panic Handler Flowchart (click to enlarge)

(continued from previous page)

```

MSTATUS : 0x00001881  MTVEC   : 0x40380001  MCAUSE   : 0x00000007  MTVAL    : 0x00000000
MHARTID  : 0x00000000

```

The register values printed are the register values in the exception frame, i.e., values at the moment when the CPU exception or another fatal error has occurred.

A Register dump is not printed if the panic handler has been executed as a result of an `abort()` call.

If *IDF Monitor* is used, Program Counter values will be converted to code locations (function name, file name, and line number), and the output will be annotated with additional lines:

```

Core 0 register dump:
MEPC   : 0x420048b4  RA      : 0x420048b4  SP      : 0x3fc8f2f0  GP      : 0x3fc8a600
0x420048b4: app_main at /Users/user/esp/example/main/hello_world_main.c:20
0x420048b4: app_main at /Users/user/esp/example/main/hello_world_main.c:20

TP      : 0x3fc8a2ac  T0      : 0x40057fa6  T1      : 0x0000000f  T2      : 0x00000000
S0/FP   : 0x00000000  S1      : 0x00000000  A0      : 0x00000001  A1      : 0x00000001
A2      : 0x00000064  A3      : 0x00000004  A4      : 0x00000001  A5      : 0x00000001
A6      : 0x42001fd6  A7      : 0x00000000  S2      : 0x00000000  S3      : 0x00000000
0x42001fd6: uart_write at /Users/user/esp/esp-idf/components/vfs/vfs_uart.c:201

S4      : 0x00000000  S5      : 0x00000000  S6      : 0x00000000  S7      : 0x00000000
S8      : 0x00000000  S9      : 0x00000000  S10     : 0x00000000  S11     : 0x00000000
T3      : 0x00000000  T4      : 0x00000000  T5      : 0x00000000  T6      : 0x00000000
MSTATUS : 0x00001881  MTVEC   : 0x40380001  MCAUSE   : 0x00000007  MTVAL    : 0x00000000
MHARTID  : 0x00000000

```

Moreover, *IDF Monitor* is also capable of generating and printing a backtrace thanks to the stack dump provided by the board in the panic handler. The output looks like this:

```

Backtrace:
0x42006686 in bar (ptr=ptr@entry=0x0) at ../main/hello_world_main.c:18
18      *ptr = 0x42424242;
#0  0x42006686 in bar (ptr=ptr@entry=0x0) at ../main/hello_world_main.c:18
#1  0x42006692 in foo () at ../main/hello_world_main.c:22
#2  0x420066ac in app_main () at ../main/hello_world_main.c:28
#3  0x42015ece in main_task (args=<optimized out>) at /Users/user/esp/components/freertos/port/port_common.c:142
#4  0x403859b8 in vPortEnterCritical () at /Users/user/esp/components/freertos/port/riscv/port.c:130
#5  0x00000000 in ?? ()
Backtrace stopped: frame did not save the PC

```

While the backtrace above is very handy, it requires the user to use *IDF Monitor*. Thus, in order to generate and print a backtrace while using another monitor program, it is possible to activate `CONFIG_ESP_SYSTEM_USE_EH_FRAME` option from the menuconfig.

This option will let the compiler generate DWARF information for each function of the project. Then, when a CPU exception occurs, the panic handler will parse these data and determine the backtrace of the task that failed. The

output looks like this:

```
Backtrace: 0x42009e9a:0x3fc92120 0x42009ea6:0x3fc92120 0x42009ec2:0x3fc92130_
↳0x42024620:0x3fc92150 0x40387d7c:0x3fc92160 0xfffffffffe:0x3fc92170
```

These PC:SP pairs represent the PC (Program Counter) and SP (Stack Pointer) for each stack frame of the current task.

The main benefit of the `CONFIG_ESP_SYSTEM_USE_EH_FRAME` option is that the backtrace is generated by the board itself (without the need for *IDF Monitor*). However, the option's drawback is that it results in an increase of the compiled binary's size (ranging from 20% to 100% increase in size). Furthermore, this option causes debug information to be included within the compiled binary. Therefore, users are strongly advised not to enable this option in mass/final production builds.

To find the location where a fatal error has happened, look at the lines which follow the "Backtrace" line. Fatal error location is the top line, and subsequent lines show the call stack.

4.11.4 GDB Stub

If the `CONFIG_ESP_SYSTEM_PANIC_GDBSTUB` option is enabled, the panic handler will not reset the chip when a fatal error happens. Instead, it will start a GDB remote protocol server, commonly referred to as GDB Stub. When this happens, a GDB instance running on the host computer can be instructed to connect to the ESP32-P4 UART port.

If *IDF Monitor* is used, GDB is started automatically when a GDB Stub prompt is detected on the UART. The output looks like this:

```
Entering gdb stub now.
$T0b#e6GNU gdb (crosstool-NG crosstool-ng-1.22.0-80-gff1f415) 7.10
Copyright (C) 2015 Free Software Foundation, Inc.
License GPLv3+: GNU GPL version 3 or later <http://gnu.org/licenses/gpl.html>
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law. Type "show copying"
and "show warranty" for details.
This GDB was configured as "--host=x86_64-build_apple-darwin16.3.0 --
↳target=riscv32-esp-elf".
Type "show configuration" for configuration details.
For bug reporting instructions, please see:
<http://www.gnu.org/software/gdb/bugs/>.
Find the GDB manual and other documentation resources online at:
<http://www.gnu.org/software/gdb/documentation/>.
For help, type "help".
Type "apropos word" to search for commands related to "word"...
Reading symbols from /Users/user/esp/example/build/example.elf...done.
Remote debugging using /dev/cu.usbserial-31301
0x400e1b41 in app_main ()
    at /Users/user/esp/example/main/main.cpp:36
36      *((int*) 0) = 0;
(gdb)
```

The GDB prompt can be used to inspect CPU registers, local and static variables, and arbitrary locations in memory. It is not possible to set breakpoints, change the PC, or continue execution. To reset the program, exit GDB and perform an external reset: Ctrl-T Ctrl-R in *IDF Monitor*, or using the external reset button on the development board.

4.11.5 RTC Watchdog Timeout

The RTC watchdog is used in the startup code to keep track of execution time and it also helps to prevent a lock-up caused by an unstable power source. It is enabled by default (see `CONFIG_BOOTLOADER_WDT_ENABLE`). If the execution time is exceeded, the RTC watchdog will restart the system. In this case, the ROM bootloader will print a message with the RTC Watchdog Timeout reason for the reboot.

```
rst:0x10 (LP_WDT_SYS)
```

The RTC watchdog covers the execution time from the first stage bootloader (ROM bootloader) to application startup. It is initially set in the ROM bootloader, then configured in the bootloader with the [CONFIG_BOOTLOADER_WDT_TIME_MS](#) option (9000 ms by default). During the application initialization stage, it is reconfigured because the source of the slow clock may have changed, and finally disabled right before the `app_main()` call. There is an option [CONFIG_BOOTLOADER_WDT_DISABLE_IN_USER_CODE](#) which prevents the RTC watchdog from being disabled before `app_main`. Instead, the RTC watchdog remains active and must be fed periodically in your application's code.

The RTC watchdog is also used by the system [panic handler](#) to protect the system from hanging during a panic. The RTC watchdog is reconfigured in the panic handler to have a timeout of 10 seconds. If the panic handler takes longer than 10 seconds to execute, the system will be reset by the RTC watchdog.

4.11.6 Guru Meditation Errors

This section explains the meaning of different error causes, printed in parens after the `Guru Meditation Error: Core panic'ed` message.

Note: See the [Guru Meditation Wikipedia article](#) for historical origins of "Guru Meditation".

Illegal instruction

This CPU exception indicates that the instruction which was executed was not a valid instruction. The most common reasons for this error include:

- FreeRTOS task function has returned. In FreeRTOS, if a task function needs to terminate, it should call [vTaskDelete\(\)](#) and delete itself, instead of returning.
- Failure to read next instruction from SPI flash. This usually happens if:
 - Application has reconfigured the SPI flash pins as some other function (GPIO, UART, etc.). Consult the Hardware Design Guidelines and the datasheet for the chip or module for details about the SPI flash pins.
 - Some external device has accidentally been connected to the SPI flash pins, and has interfered with communication between ESP32-P4 and SPI flash.
- In C++ code, exiting from a non-void function without returning a value is considered to be an undefined behavior. When optimizations are enabled, the compiler will often omit the epilogue in such functions. This most often results in an Illegal instruction exception. By default, ESP-IDF build system enables `-Werror=return-type` which means that missing return statements are treated as compile time errors. However if the application project disables compiler warnings, this issue might go undetected and the Illegal instruction exception will occur at run time.

Instruction Address Misaligned

This CPU exception indicates that the address of the instruction to execute is not 2-byte aligned.

Instruction Access Fault, Load Access Fault, Store Access Fault

This CPU exception happens when application attempts to execute, read from or write to an invalid memory location. The address which was written/read is found in `MTVAL` register in the register dump. If this address is zero, it usually means that application attempted to dereference a NULL pointer. If this address is close to zero, it usually means that application attempted to access member of a structure, but the pointer to the structure was NULL. If this address is something else (garbage value, not in `0x3fxxxxxx - 0x6xxxxxxx` range), it likely means that the pointer used to access the data was either not initialized or was corrupted.

Breakpoint

This CPU exception happens when the instruction `EBREAK` is executed. See also [FreeRTOS End of Stack Watchpoint](#).

Load Address Misaligned, Store Address Misaligned

Application has attempted to read or write memory location, and address alignment did not match load/store size. For example, 32-bit load can only be done from 4-byte aligned address, and 16-bit load can only be done from a 2-byte aligned address.

Interrupt Watchdog Timeout on CPU0/CPU1

Indicates that an interrupt watchdog timeout has occurred. See [Watchdogs](#) for more information.

Cache error

In some situations, ESP-IDF will temporarily disable access to external SPI Flash and SPI RAM via caches. For example, this happens when `spi_flash` APIs are used to read/write/erase/mmap regions of SPI Flash. In these situations, tasks are suspended, and interrupt handlers not registered with `ESP_INTR_FLAG_IRAM` are disabled. Make sure that any interrupt handlers registered with this flag have all the code and data in IRAM/DRAM. Refer to the [SPI flash API documentation](#) for more details.

4.11.7 Other Fatal Errors

Corrupt Heap

ESP-IDF's heap implementation contains a number of run-time checks of the heap structure. Additional checks ("Heap Poisoning") can be enabled in `menuconfig`. If one of the checks fails, a message similar to the following will be printed:

```
CORRUPT HEAP: Bad tail at 0x3ffe270a. Expected 0xbaad5678 got 0xbaac5678
assertion "head != NULL" failed: file "/Users/user/esp/esp-idf/components/heap/
↳multi_heap_poisoning.c", line 201, function: multi_heap_free
abort() was called at PC 0x400dca43 on core 0
```

Consult [Heap Memory Debugging](#) documentation for further information.

Stack overflow

FreeRTOS End of Stack Watchpoint ESP-IDF provides a custom FreeRTOS stack overflow detecting mechanism based on watchpoints. Every time FreeRTOS switches task context, one of the watchpoints is set to watch the last 32 bytes of stack.

Generally, this may cause the watchpoint to be triggered up to 28 bytes earlier than expected. The value 32 is chosen because it is larger than the stack canary size in FreeRTOS (20 bytes). Adopting this approach ensures that the watchpoint triggers before the stack canary is corrupted, not after.

Note: Not every stack overflow is guaranteed to trigger the watchpoint. It is possible that the task writes to memory beyond the stack canary location, in which case the watchpoint will not be triggered.

If watchpoint triggers, the message will be similar to:

```
Guru Meditation Error: Core  0 panic'ed (Breakpoint). Exception was unhandled.
```

This feature can be enabled by using the `CONFIG_FREERTOS_WATCHPOINT_END_OF_STACK` option.

FreeRTOS Stack Checks See [CONFIG_FREERTOS_CHECK_STACKOVERFLOW](#)

Stack Smashing

Stack smashing protection (based on GCC `-fstack-protector*` flags) can be enabled in ESP-IDF using [CONFIG_COMPILER_STACK_CHECK_MODE](#) option. If stack smashing is detected, message similar to the following will be printed:

```
Stack smashing protect failure!

abort() was called at PC 0x400d2138 on core 0

Backtrace: 0x4008e6c0:0x3ffc1780 0x4008e8b7:0x3ffc17a0 0x400d2138:0x3ffc17c0
↳ 0x400e79d5:0x3ffc17e0 0x400e79a7:0x3ffc1840 0x400e79df:0x3ffc18a0
↳ 0x400e2235:0x3ffc18c0 0x400e1916:0x3ffc18f0 0x400e19cd:0x3ffc1910
↳ 0x400e1a11:0x3ffc1930 0x400e1bb2:0x3ffc1950 0x400d2c44:0x3ffc1a80
0
```

The backtrace should point to the function where stack smashing has occurred. Check the function code for unbounded access to local arrays.

Undefined Behavior Sanitizer (UBSAN) Checks

Undefined behavior sanitizer (UBSAN) is a compiler feature which adds run-time checks for potentially incorrect operations, such as:

- overflows (multiplication overflow, signed integer overflow)
- shift base or exponent errors (e.g., shift by more than 32 bits)
- integer conversion errors

See [GCC documentation](#) of `-fsanitize=undefined` option for the complete list of supported checks.

Enabling UBSAN UBSAN is disabled by default. It can be enabled at file, component, or project level by adding the `-fsanitize=undefined` compiler option in the build system.

When enabling UBSAN for code which uses the SOC hardware register header files (`soc/xxx_reg.h`), it is recommended to disable shift-base sanitizer using `-fno-sanitize=shift-base` option. This is due to the fact that ESP-IDF register header files currently contain patterns which cause false positives for this specific sanitizer option.

To enable UBSAN at project level, add the following code at the end of the project's `CMakeLists.txt` file:

```
idf_build_set_property (COMPILE_OPTIONS "-fsanitize=undefined" "-fno-sanitize=shift-
↳base" APPEND)
```

Alternatively, pass these options through the `EXTRA_CFLAGS` and `EXTRA_CXXFLAGS` environment variables.

Enabling UBSAN results in significant increase of code and data size. Most applications, except for the trivial ones, will not fit into the available RAM of the microcontroller when UBSAN is enabled for the whole application. Therefore it is recommended that UBSAN is instead enabled for specific components under test.

To enable UBSAN for a specific component (`component_name`) from the project's `CMakeLists.txt` file, add the following code at the end of the file:

```
idf_component_get_property(lib component_name COMPONENT_LIB)
target_compile_options(${lib} PRIVATE "-fsanitize=undefined" "-fno-sanitize=shift-
↳base")
```

Note: See the build system documentation for more information about [build properties](#) and [component properties](#).

To enable UBSAN for a specific component (`component_name`) from `CMakeLists.txt` of the same component, add the following at the end of the file:

```
target_compile_options(${COMPONENT_LIB} PRIVATE "-fsanitize=undefined" "-fno-
↪sanitize=shift-base")
```

UBSAN Output When UBSAN detects an error, a message and the backtrace are printed, for example:

```
Undefined behavior of type out_of_bounds

Backtrace:0x4008b383:0x3ffcd8b0 0x4008c791:0x3ffcd8d0 0x4008c587:0x3ffcd8f0_
↪0x4008c6be:0x3ffcd950 0x400db74f:0x3ffcd970 0x400db99c:0x3ffcd9a0
```

When using *IDF Monitor*, the backtrace will be decoded to function names and source code locations, pointing to the location where the issue has happened (here it is `main.c:128`):

```
0x4008b383: panic_abort at /path/to/esp-idf/components/esp_system/panic.c:367

0x4008c791: esp_system_abort at /path/to/esp-idf/components/esp_system/system_api.
↪c:106

0x4008c587: __ubsan_default_handler at /path/to/esp-idf/components/esp_system/
↪ubsan.c:152

0x4008c6be: __ubsan_handle_out_of_bounds at /path/to/esp-idf/components/esp_system/
↪ubsan.c:223

0x400db74f: test_ub at main.c:128

0x400db99c: app_main at main.c:56 (discriminator 1)
```

The types of errors reported by UBSAN can be as follows:

Name	Meaning
<code>type_mismatch</code> , <code>type_mismatch_v1</code>	Incorrect pointer value: null, unaligned, not compatible with the given type.
<code>add_overflow</code> , <code>sub_overflow</code> , <code>mul_overflow</code> , <code>negate_overflow</code>	Integer overflow during addition, subtraction, multiplication, negation.
<code>divrem_overflow</code>	Integer division by 0 or <code>INT_MIN</code> .
<code>shift_out_of_bounds</code>	Overflow in left or right shift operators.
<code>out_of_bounds</code>	Access outside of bounds of an array.
<code>unreachable</code>	Unreachable code executed.
<code>missing_return</code>	Non-void function has reached its end without returning a value (C++ only).
<code>vla_bound_not_positive</code>	Size of variable length array is not positive.
<code>load_invalid_value</code>	Value of <code>bool</code> or <code>enum</code> (C++ only) variable is invalid (out of bounds).
<code>nonnull_arg</code>	Null argument passed to a function which is declared with a <code>nonnull</code> attribute.
<code>nonnull_return</code>	Null value returned from a function which is declared with <code>returns_nonnull</code> attribute.
<code>builtin_unreachable</code>	<code>__builtin_unreachable</code> function called.
<code>pointer_overflow</code>	Overflow in pointer arithmetic.

4.12 Hardware Abstraction

ESP-IDF provides a group of APIs for hardware abstraction. These APIs allow you to control peripherals at different levels of abstraction, giving you more flexibility compared to using only the ESP-IDF drivers to interact with hardware. ESP-IDF Hardware abstraction is likely to be useful for writing high-performance bare-metal drivers, or for attempting to port an ESP chip to another platform.

This guide is split into the following sections:

1. *Architecture*
2. *LL (Low Level) Layer*
3. *HAL (Hardware Abstraction Layer)*

Warning: Hardware abstraction API (excluding the driver and `xxx_types.h`) should be considered an experimental feature, thus cannot be considered public API. The hardware abstraction API does not adhere to the API name changing restrictions of ESP-IDF's versioning scheme. In other words, it is possible that Hardware Abstraction API may change in between non-major release versions.

Note: Although this document mainly focuses on hardware abstraction of peripherals, e.g., UART, SPI, I2C, certain layers of hardware abstraction extend to other aspects of hardware as well, e.g., some of the CPU's features are partially abstracted.

4.12.1 Architecture

Hardware abstraction in ESP-IDF is comprised of the following layers, ordered from low level of abstraction that is closer to hardware, to high level of abstraction that is further away from hardware.

- Low Level (LL) Layer
- Hardware Abstraction Layer (HAL)
- Driver Layers

The LL Layer, and HAL are entirely contained within the `hal` component. Each layer is dependent on the layer below it, i.e., driver depends on HAL, HAL depends on LL, LL depends on the register header files.

For a particular peripheral `xxx`, its hardware abstraction generally consists of the header files described in the table below. Files that are **Target Specific** have a separate implementation for each target, i.e., a separate copy for each chip. However, the `#include` directive is still target-independent, i.e., is the same for different targets, as the build system automatically includes the correct version of the header and source files.

Table 2: Hardware Abstraction Header Files

Include Directive	Target Specific	Description
<code>#include 'soc/xxx_caps.h'</code>	Y	This header contains a list of C macros specifying the various capabilities of the ESP32-P4's peripheral xxx. Hardware capabilities of a peripheral include things such as the number of channels, DMA support, hardware FIFO/buffer lengths, etc.
<code>#include "soc/xxx_struct.h"</code> <code>#include "soc/xxx_reg.h"</code>	Y	The two headers contain a representation of a peripheral's registers in C structure and C macro format respectively, allowing you to operate a peripheral at the register level via either of these two header files.
<code>#include "soc/xxx_pins.h"</code>	Y	If certain signals of a peripheral are mapped to a particular pin of the ESP32-P4, their mappings are defined in this header as C macros.
<code>#include "soc/xxx_periph.h"</code>	N	This header is mainly used as a convenience header file to automatically include <code>xxx_caps.h</code> , <code>xxx_struct.h</code> , and <code>xxx_reg.h</code> .
<code>#include "hal/xxx_types.h"</code>	N	This header contains type definitions and macros that are shared among the LL, HAL, and driver layers. Moreover, it is considered public API thus can be included by the application level. The shared types and definitions usually related to non-implementation specific concepts such as the following: • Protocol-related types/macros such as frames, modes, common bus speeds, etc. • Features/characteristics of an xxx peripheral that are likely to be present on any implementation (implementation-independent) such as channels, operating modes, signal amplification or attenuation intensities, etc.
<code>#include "hal/xxx_ll.h"</code>	Y	This header contains the Low Level (LL) Layer of hardware abstraction. LL Layer API are primarily used to abstract away register operations into readable functions.
<code>#include "hal/xxx_hal.h"</code>	Y	The Hardware Abstraction Layer (HAL) is used to abstract away peripheral operation steps into functions (e.g., reading a buffer, starting a transmission, handling an event, etc). The HAL is built on top of the LL Layer.
<code>#include "driver/xxx.h"</code>	N	The driver layer is the highest level of ESP-IDF's hardware abstraction. Driver layer API are meant to be called from ESP-IDF applications, and internally utilize OS primitives. Thus, driver layer API are event-driven, and can be used in a multi-threaded environment.

4.12.2 LL (Low Level) Layer

The primary purpose of the LL Layer is to abstract away register field access into more easily understandable functions. LL functions essentially translate various in/out arguments into the register fields of a peripheral in the form of get/set functions. All the necessary bit shifting, masking, offsetting, and endianness of the register fields should be handled by the LL functions.

```
//Inside xxx_ll.h

static inline void xxx_ll_set_baud_rate(xxx_dev_t *hw,
                                       xxx_ll_clk_src_t clock_source,
                                       uint32_t baud_rate) {
    uint32_t src_clk_freq = (source_clk == XXX_SCLK_APB) ? APB_CLK_FREQ : REF_CLK_FREQ;
    uint32_t clock_divider = src_clk_freq / baud;
    // Set clock select field
    hw->clk_div_reg.divider = clock_divider >> 4;
    // Set clock divider field
```

(continues on next page)

(continued from previous page)

```

hw->config.clk_sel = (source_clk == XXX_SCLK_APB) ? 0 : 1;
}

static inline uint32_t xxx_ll_get_rx_byte_count(xxx_dev_t *hw) {
    return hw->status_reg.rx_cnt;
}

```

The code snippet above illustrates typical LL functions for a peripheral `xxx`. LL functions typically have the following characteristics:

- All LL functions are defined as `static inline` so that there is minimal overhead when calling these functions due to compiler optimization. These functions are not guaranteed to be inlined by the compiler, so any LL function that is called when the cache is disabled (e.g., from an IRAM ISR context) should be marked with `__attribute__((always_inline))`.
- The first argument should be a pointer to a `xxx_dev_t` type. The `xxx_dev_t` type is a structure representing the peripheral's registers, thus the first argument is always a pointer to the starting address of the peripheral's registers. Note that in some cases where the peripheral has multiple channels with identical register layouts, `xxx_dev_t *hw` may point to the registers of a particular channel instead.
- LL functions should be short, and in most cases are deterministic. In other words, in the worst case, runtime of the LL function can be determined at compile time. Thus, any loops in LL functions should be finite bounded; however, there are currently a few exceptions to this rule.
- LL functions are not thread-safe, it is the responsibility of the upper layers (driver layer) to ensure that registers or register fields are not accessed concurrently.

4.12.3 HAL (Hardware Abstraction Layer)

The HAL layer models the operational process of a peripheral as a set of general steps, where each step has an associated function. For each step, the details of a peripheral's register implementation (i.e., which registers need to be set/read) are hidden (abstracted away) by the HAL. By modeling peripheral operation as a set of functional steps, any minor hardware implementation differences of the peripheral between different targets or chip versions can be abstracted away by the HAL (i.e., handled transparently). In other words, the HAL API for a particular peripheral remains mostly the same across multiple targets/chip versions.

The following HAL function examples are selected from the Watchdog Timer HAL as each function maps to one of the steps in a WDT's operation life cycle, thus illustrating how a HAL abstracts a peripheral's operation into functional steps.

```

// Initialize one of the WDTs
void wdt_hal_init(wdt_hal_context_t *hal, wdt_inst_t wdt_inst, uint32_t prescaler,
    ↪ bool enable_intr);

// Configure a particular timeout stage of the WDT
void wdt_hal_config_stage(wdt_hal_context_t *hal, wdt_stage_t stage, uint32_t
    ↪ timeout, wdt_stage_action_t behavior);

// Start the WDT
void wdt_hal_enable(wdt_hal_context_t *hal);

// Feed (i.e., reset) the WDT
void wdt_hal_feed(wdt_hal_context_t *hal);

// Handle a WDT timeout
void wdt_hal_handle_intr(wdt_hal_context_t *hal);

// Stop the WDT
void wdt_hal_disable(wdt_hal_context_t *hal);

// De-initialize the WDT
void wdt_hal_deinit(wdt_hal_context_t *hal);

```

HAL functions generally have the following characteristics:

- The first argument to a HAL function has the `xxx_hal_context_t *` type. The HAL context type is used to store information about a particular instance of the peripheral (i.e., the context instance). A HAL context is initialized by the `xxx_hal_init()` function and can store information such as the following:
 - The channel number of this instance
 - Pointer to the peripheral's (or channel's) registers (i.e., a `xxx_dev_t *` type)
 - Information about an ongoing transaction (e.g., pointer to DMA descriptor list in use)
 - Some configuration values for the instance (e.g., channel configurations)
 - Variables to maintain state information regarding the instance (e.g., a flag to indicate if the instance is waiting for transaction to complete)
- HAL functions should not contain any OS primitives such as queues, semaphores, mutexes, etc. All synchronization/concurrency should be handled at higher layers (e.g., the driver).
- Some peripherals may have steps that cannot be further abstracted by the HAL, thus end up being a direct wrapper (or macro) for an LL function.
- Some HAL functions may be placed in IRAM thus may carry an `IRAM_ATTR` or be placed in a separate `xxx_hal_iram.c` source file.

4.13 JTAG Debugging

This document provides a guide to installing OpenOCD for ESP32-P4 and debugging using GDB.

Note: You can also debug your ESP32-P4 without needing to setup JTAG or OpenOCD by using `idf.py monitor`. See: [IDF Monitor](#) and [CONFIG_ESP_SYSTEM_GDBSTUB_RUNTIME](#).

The document is structured as follows:

Introduction Introduction to the purpose of this guide.

How it Works? Description how ESP32-P4, JTAG interface, OpenOCD and GDB are interconnected and working together to enable debugging of ESP32-P4.

Selecting JTAG Adapter What are the criteria and options to select JTAG adapter hardware.

Setup of OpenOCD Procedure to install OpenOCD and verify that it is installed.

Configuring ESP32-P4 Target Configuration of OpenOCD software and setting up of JTAG adapter hardware, which together make up the debugging target.

Launching Debugger Steps to start up a debug session with GDB from [Eclipse](#) and from [Command Line](#).

Debugging Examples If you are not familiar with GDB, check this section for debugging examples provided from [Eclipse](#) as well as from [Command Line](#).

Building OpenOCD from Sources Procedure to build OpenOCD from sources for [Windows](#), [Linux](#) and [macOS](#) operating systems.

Tips and Quirks This section provides collection of tips and quirks related to JTAG debugging of ESP32-P4 with OpenOCD and GDB.

4.13.1 Introduction

Espressif has ported OpenOCD to support the ESP32-P4 processor and the multi-core FreeRTOS (which is the foundation of most ESP32-P4 apps). Additionally, some extra tools have been written to provide extra features that OpenOCD does not support natively.

This document provides a guide to installing OpenOCD for ESP32-P4 and debugging using GDB under Linux, Windows and macOS. Except for OS specific installation procedures, the s/w user interface and use procedures are the same across all supported operating systems.

Note: Screenshots presented in this document have been made for Eclipse Neon 3 running on Ubuntu 16.04 LTS. There may be some small differences in what a particular user interface looks like, depending on whether you are

using Windows, macOS or Linux and/or a different release of Eclipse.

4.13.2 How it Works?

The key software and hardware components that perform debugging of ESP32-P4 with OpenOCD over JTAG (Joint Test Action Group) interface is presented in the diagram below under the "Debugging With JTAG" label. These components include riscv32-esp-elf-gdb debugger, OpenOCD on chip debugger, and the JTAG adapter connected to ESP32-P4 target.




Fig. 10: JTAG debugging - overview diagram

Likewise, the "Application Loading and Monitoring" label indicates the key software and hardware components that allow an application to be compiled, built, and flashed to ESP32-P4, as well as to provide means to monitor diagnostic messages from ESP32-P4.

"Debugging With JTAG" and "Application Loading and Monitoring" is integrated under the [Eclipse IDE](#) in order to provide a quick and easy transition between writing/compiling/loading/debugging code. The Eclipse IDE (and the integrated debugging software) is available for Windows, Linux and macOS platforms. Depending on user preferences, both the debugger and `idf.py build` can also be used directly from terminal/command line, instead of Eclipse.

4.13.3 Selecting JTAG Adapter

If you decide to use separate JTAG adapter, look for one that is compatible with both the voltage levels on the ESP32-P4 as well as with the OpenOCD software. The JTAG port on the ESP32-P4 is an industry-standard JTAG port which lacks (and does not need) the TRST pin. The JTAG I/O pins all are powered from the VDD_3P3_RTC pin (which normally would be powered by a 3.3 V rail) so the JTAG adapter needs to be able to work with JTAG pins in that voltage range.

On the software side, OpenOCD supports a fair amount of JTAG adapters. See <https://openocd.org/doc/html/Debug-Adapter-Hardware.html> for an (unfortunately slightly incomplete) list of the adapters OpenOCD works with. This page lists SWD-compatible adapters as well; take note that the ESP32-P4 does not support SWD. JTAG adapters that are hardcoded to a specific product line, e.g., ST-LINK debugging adapters for STM32 families, will not work.

The minimal signalling to get a working JTAG connection are TDI, TDO, TCK, TMS and GND. Some JTAG debuggers also need a connection from the ESP32-P4 power line to a line called e.g., Vtar to set the working voltage.

SRST can optionally be connected to the CH_PD of the ESP32-P4, although for now, support in OpenOCD for that line is pretty minimal.

[ESP-Prog](#) is an example for using an external board for debugging by connecting it to the JTAG pins of ESP32-P4.

4.13.4 Setup of OpenOCD

If you have already set up ESP-IDF with CMake build system according to the [Getting Started Guide](#), then OpenOCD is already installed. After [setting up the environment](#) in your terminal, you should be able to run OpenOCD. Check this by executing the following command:

```
openocd --version
```

The output should be as follows (although the version may be more recent than listed here):

```
Open On-Chip Debugger v0.12.0-esp32-20240318 (2024-03-18-18:25)
Licensed under GNU GPL v2
For bug reports, read
    http://openocd.org/doc/doxygen/bugs.html
```

You may also verify that OpenOCD knows where its configuration scripts are located by printing the value of OPENOCD_SCRIPTS environment variable, by typing `echo $OPENOCD_SCRIPTS` (for Linux and macOS) or `echo %OPENOCD_SCRIPTS%` (for Windows). If a valid path is printed, then OpenOCD is set up correctly.

If any of these steps do not work, please go back to the [setting up the tools](#) section of the Getting Started Guide.

Note: It is also possible to build OpenOCD from source. Please refer to [Building OpenOCD from Sources](#) section for details.

4.13.5 Configuring ESP32-P4 Target

Once OpenOCD is installed, you can proceed to configuring the ESP32-P4 target (i.e ESP32-P4 board with JTAG interface). Configuring the target is split into the following three steps:

- [Configure and Connect JTAG Interface](#)
- [Run OpenOCD](#)
- [Upload Application for Debugging](#)

Configure and Connect JTAG Interface

This step depends on the JTAG and ESP32-P4 board you are using (see the two cases described below).

Configure Other JTAG Interfaces

For guidance about which JTAG interface to select when using OpenOCD with ESP32-P4, refer to the section [Selecting JTAG Adapter](#). Then follow the configuration steps below to get it working.

Configure Hardware

1. Identify all pins/signals on JTAG interface and ESP32-P4 board that should be connected to establish communication.

Table 3: ESP32-P4 pins and JTAG signals

ESP32-P4 Pin	JTAG Signal
MTDO / GPIO5	TDO
MTDI / GPIO3	TDI
MTCK / GPIO2	TCK
MTMS / GPIO4	TMS

2. Verify if ESP32-P4 pins used for JTAG communication are not connected to some other hardware that may disturb JTAG operation.
3. Connect identified pin/signals of ESP32-P4 and JTAG interface.

Configure Drivers You may need to install driver software to make JTAG work with computer. Refer to documentation of your JTAG adapter for related details.

On Linux, adding OpenOCD udev rules is required and is done by copying the [udev rules file](#) into the `/etc/udev/rules.d` directory.

Connect Connect JTAG interface to the computer. Power on ESP32-P4 and JTAG interface boards. Check if the JTAG interface is visible on the computer.

To carry on with debugging environment setup, proceed to section [Run OpenOCD](#).

Run OpenOCD

Once target is configured and connected to computer, you are ready to launch OpenOCD.

Open a terminal and set it up for using the ESP-IDF as described in the [setting up the environment](#) section of the Getting Started Guide. To run OpenOCD for a specific board, you must pass the board-specific configuration. The default configuration for the built project can be found in the `debug_arguments_openocd` field of the `build/project_description.json` file. There is an example to run OpenOCD (this command works on Windows, Linux, and macOS):

```
openocd -f board/esp32p4-builtin.cfg
```

Note: The files provided after `-f` above are specific for ESP32-P4 through built-in USB connection. You may need to provide different files depending on the hardware that is used. For guidance see [Configuration of OpenOCD for Specific Target](#).

You should now see similar output (this log is for ESP32-P4 through built-in USB connection):

```
user-name@computer-name:~/esp/esp-idf$ openocd -f board/esp32p4-builtin.cfg
Open On-Chip Debugger v0.12.0-esp32-20240318 (2024-03-18-18:25)
Licensed under GNU GPL v2
For bug reports, read
    http://openocd.org/doc/doxygen/bugs.html
Info : only one transport option; autoselecting 'jtag'
Info : esp_usb_jtag: VID set to 0x303a and PID to 0x1001
Info : esp_usb_jtag: capabilities descriptor set to 0x2000
Info : Listening on port 6666 for tcl connections
Info : Listening on port 4444 for telnet connections
Info : esp_usb_jtag: serial (60:55:F9:F9:04:AD)
Info : esp_usb_jtag: Device found. Base speed 24000KHz, div range 1 to 255
Info : clock speed 24000 kHz
Info : JTAG tap: esp32p4.tap0 tap/device found: 0x00012c25 (mfg: 0x612 (Espressif_
↳Systems), part: 0x0012, ver: 0x0)
Info : JTAG tap: esp32p4.tap1 tap/device found: 0x00012c25 (mfg: 0x612 (Espressif_
↳Systems), part: 0x0012, ver: 0x0)
```

(continues on next page)

(continued from previous page)

```

Info : [esp32p4.hp.cpu0] datacount=1 progbufsize=2
Info : [esp32p4.hp.cpu0] Core 0 made part of halt group 1.
Info : [esp32p4.hp.cpu0] Examined RISC-V core; found 2 harts
Info : [esp32p4.hp.cpu0] XLEN=32, misa=0x40901125
Info : [esp32p4.hp.cpu0] Examination succeed
Info : [esp32p4.hp.cpu1] datacount=1 progbufsize=2
Info : [esp32p4.hp.cpu1] Core 1 made part of halt group 1.
Info : [esp32p4.hp.cpu1] Examined RISC-V core; found 2 harts
Info : [esp32p4.hp.cpu1] XLEN=32, misa=0x40901125
Info : [esp32p4.hp.cpu1] Examination succeed
Info : [esp32p4.hp.cpu0] starting gdb server on 3333
Info : Listening on port 3333 for gdb connections

```

- If there is an error indicating permission problems, please see section on "Permissions delegation" in the OpenOCD README file located in the `~/esp/openocd-esp32` directory.
- In case there is an error in finding the configuration files, e.g., `Can't find board/esp32p4-builtin.cfg`, check if the `OPENOCD_SCRIPTS` environment variable is set correctly. This variable is used by OpenOCD to look for the files specified after the `-f` option. See [Setup of OpenOCD](#) section for details. Also check if the file is indeed under the provided path.
- If you see JTAG errors (e.g., `...all ones` or `...all zeroes`), please check your JTAG connections, whether other signals are connected to JTAG besides ESP32-P4's pins, and see if everything is powered on correctly.

Upload Application for Debugging

Build and upload your application to ESP32-P4 as usual, see [Step 5. First Steps on ESP-IDF](#).

Another option is to write application image to flash using OpenOCD via JTAG with commands like this:

```
openocd -f board/esp32p4-builtin.cfg -c "program_esp filename.bin 0x10000 verify_
↳exit"
```

OpenOCD flashing command `program_esp` has the following format:

```
program_esp <image_file> <offset> [verify] [reset] [exit] [compress] [en-
crypt] [no_clock_boost] [restore_clock] [no_skip_loaded]
```

- `image_file` - Path to program image file.
- `offset` - Offset in flash bank to write image.
- `verify` - Optional. Verify flash contents after writing.
- `reset` - Optional. Reset target after programing.
- `exit` - Optional. Finally exit OpenOCD.
- `compress` - Optional. Compress image file before programming.
- `encrypt` - Optional. Encrypt binary before writing to flash. Same functionality with `idf.py encrypt-flash`
- `no_clock_boost` - Optional. Disable setting target clock frequency to its maximum possible value before programming. Clock boost is enabled by default.
- `restore_clock` - Optional. Restore clock frequency to its initial value after programming. Disabled by default.
- `no_skip_loaded` - Optional. Do not check whether the binary is already loaded before flashing. Disabled by default.

Alternative Method: Using `program_esp_bins`

For convenience when working with ESP-IDF projects, OpenOCD provides an alternative command `program_esp_bins` that can flash multiple binaries in a single command by reading the build configuration from the `flasher_args.json` file generated during the ESP-IDF build process.

This method is particularly useful because:

- It automatically reads all the binary files and their flash addresses from the build output
- It handles encrypted partitions automatically based on the project configuration
- It eliminates the need to manually specify addresses for each binary (bootloader, partition table, application, etc.)

Basic usage:

```
openocd -f board/esp32-wrover-kit-3.3v.cfg -c "program_esp_bins build flasher_args.  
↪json verify exit"
```

Command Format The OpenOCD flashing command `program_esp_bins` has the following format:

```
program_esp_bins <build_dir> <json_file> [verify] [reset] [exit] [compress]  
[no_clock_boost] [restore_clock] [no_skip_loaded]
```

- `build_dir` - Path to the build directory containing the `flasher_args.json` file.
- `json_file` - Name of the JSON file containing flash configuration (typically `flasher_args.json`).
- Other optional parameters work the same as `program_esp` command. See [Upload Application for Debugging](#) section for details.

You are now ready to start application debugging. Follow the steps described in the section below.

4.13.6 Launching Debugger

The toolchain for ESP32-P4 features GNU Debugger, in short GDB. It is available with other toolchain programs under filename: `riscv32-esp-elf-gdb`. GDB can be called and operated directly from command line in a terminal. Another option is to call it from within IDE (like Eclipse, Visual Studio Code, etc.) and operate indirectly with help of GUI instead of typing commands in a terminal.

The options of using debugger are discussed under links below.

- [Eclipse](#)
- [Command Line](#)
- [Configuration for Visual Studio Code Debug](#)

It is recommended to first check if debugger works from [Command Line](#) and then move to using [Eclipse](#).

4.13.7 Debugging Examples

This section is intended for users not familiar with GDB. It presents example debugging session from [Eclipse](#) using simple application available under [get-started/blink](#) and covers the following debugging actions:

1. [Navigating Through the Code, Call Stack and Threads](#)
2. [Setting and Clearing Breakpoints](#)
3. [Halting the Target Manually](#)
4. [Stepping Through the Code](#)
5. [Checking and Setting Memory](#)
6. [Watching and Setting Program Variables](#)
7. [Setting Conditional Breakpoints](#)

Similar debugging actions are provided using GDB from [Command Line](#).

Note: [Debugging FreeRTOS Objects](#) is currently only available for command line debugging.

Before proceeding to examples, set up your ESP32-P4 target and load it with [get-started/blink](#).

4.13.8 Building OpenOCD from Sources

Please refer to separate documents listed below, that describe build process.

Building OpenOCD from Sources for Windows

Note: This document outlines how to build a binary of OpenOCD from its source files instead of downloading the pre-built binary. For a quick setup, users can download a pre-built binary of OpenOCD from [Espressif GitHub](#) instead of compiling it themselves (see [Setup of OpenOCD](#) for more details).

Note: All code snippets in this document are assumed to be running in an MSYS2 shell with the MINGW32 subsystem.

Install Dependencies Install packages that are required to compile OpenOCD:

```
pacman -S --noconfirm --needed autoconf automake git make \
mingw-w64-i686-gcc \
mingw-w64-i686-toolchain \
mingw-w64-i686-libtool \
mingw-w64-i686-pkg-config \
mingw-w64-cross-winpthreads-git \
p7zip
```

Download Sources of OpenOCD The sources for the ESP32-P4-enabled variant of OpenOCD are available from Espressif's GitHub under <https://github.com/espressif/openocd-esp32>. These source files can be pulled via Git using the following commands:

```
cd ~/esp
git clone --recursive https://github.com/espressif/openocd-esp32.git
```

The clone of sources should be now saved in ~/esp/openocd-esp32 directory.

Downloading libusb The libusb library is also required when building OpenOCD. The following commands will download a particular release of libusb and uncompress it to the current directory.

```
wget https://github.com/libusb/libusb/releases/download/v1.0.22/libusb-1.0.22.7z
7z x -olibusb ./libusb-1.0.22.7z
```

We now need to export the following variables such that the libusb library gets linked into the OpenOCD build.

```
export CPPFLAGS="$CPPFLAGS -I${PWD}/libusb/include/libusb-1.0"
export LDFLAGS="$LDFLAGS -L${PWD}/libusb/MingW32/.libs/dll"
```

Build OpenOCD The following commands will configure OpenOCD then build it.

```
cd ~/esp/openocd-esp32
export CPPFLAGS="$CPPFLAGS -D__USE_MINGW_ANSI_STDIO=1 -Wno-error"; export CFLAGS="
↪$CFLAGS -Wno-error"
./bootstrap
./configure --disable-doxxygen-pdf --enable-ftdi --enable-jlink --enable-ulink --
↪build=i686-w64-mingw32 --host=i686-w64-mingw32
```

(continues on next page)

(continued from previous page)

```
make
cp ../libusb/MinGW32/dll/libusb-1.0.dll ./src
cp /opt/i686-w64-mingw32/bin/libwinpthread-1.dll ./src
```

Once the build is completed, the OpenOCD binary will be placed in `~/esp/openocd-esp32/src/`.

You can then optionally call `make install`. This will copy the OpenOCD binary to a user specified location.

- This location can be specified when OpenOCD is configured, or by setting `export DESTDIR="/custom/install/dir"` before calling `make install`.
- If you have an existing OpenOCD (from e.g., another development platform), you may want to skip this call as your existing OpenOCD may get overwritten.

Note:

- Should an error occur, resolve it and try again until the command `make` works.
 - If there is a submodule problem from OpenOCD, please `cd` to the `openocd-esp32` directory and input `git submodule update --init`.
 - If the `./configure` is successfully run, information of enabled JTAG will be printed under OpenOCD configuration summary.
 - If the information of your device is not shown in the log, use `./configure` to enable it as described in `../openocd-esp32/doc/INSTALL.txt`.
 - For details concerning compiling OpenOCD, please refer to `openocd-esp32/README.Windows`.
 - Don't forget to copy `libusb-1.0.dll` and `libwinpthread-1.dll` into `OOCD_INSTALLDIR/bin` from `~/esp/openocd-esp32/src`.
-

Once `make` process is successfully completed, the executable of OpenOCD will be saved in `~/esp/openocd-esp32/src` directory.

Full Listing For greater convenience, all of commands called throughout the OpenOCD build process have been listed in the code snippet below. Users can copy this code snippet into a shell script then execute it:

```
pacman -S --noconfirm --needed autoconf automake git make mingw-w64-i686-gcc mingw-
↪w64-i686-toolchain mingw-w64-i686-libtool mingw-w64-i686-pkg-config mingw-w64-
↪cross-winpthreads-git p7zip
cd ~/esp
git clone --recursive https://github.com/espressif/openocd-esp32.git

wget https://github.com/libusb/libusb/releases/download/v1.0.22/libusb-1.0.22.7z
7z x -olibusb ./libusb-1.0.22.7z
export CPPFLAGS="$CPPFLAGS -I${PWD}/libusb/include/libusb-1.0"; export LDFLAGS="
↪$LDFLAGS -L${PWD}/libusb/MinGW32/.libs/dll"

export CPPFLAGS="$CPPFLAGS -D__USE_MINGW_ANSI_STDIO=1 -Wno-error"; export CFLAGS="
↪$CFLAGS -Wno-error"
cd ~/esp/openocd-esp32
./bootstrap
./configure --disable-doxxygen-pdf --enable-ftdi --enable-jlink --enable-ulink --
↪build=i686-w64-mingw32 --host=i686-w64-mingw32
make
cp ../libusb/MinGW32/dll/libusb-1.0.dll ./src
cp /opt/i686-w64-mingw32/bin/libwinpthread-1.dll ./src

# # optional
# export DESTDIR="$PWD"
# make install
# cp ./src/libusb-1.0.dll $DESTDIR/mingw32/bin
# cp ./src/libwinpthread-1.dll $DESTDIR/mingw32/bin
```

Next Steps To carry on with debugging environment setup, proceed to section [Configuring ESP32-P4 Target](#).

Building OpenOCD from Sources for Linux

The following instructions are alternative to downloading binary OpenOCD from [Espressif GitHub](#). To quickly setup the binary OpenOCD, instead of compiling it yourself, backup and proceed to section [Setup of OpenOCD](#).

Download Sources of OpenOCD The sources for the ESP32-P4-enabled variant of OpenOCD are available from Espressif GitHub under <https://github.com/espressif/openocd-esp32>. To download the sources, use the following commands:

```
cd ~/esp
git clone --recursive https://github.com/espressif/openocd-esp32.git
```

The clone of sources should be now saved in `~/esp/openocd-esp32` directory.

Install Dependencies Install packages that are required to compile OpenOCD.

Note: Install the following packages one by one, check if installation was successful and then proceed to the next package. Resolve reported problems before moving to the next step.

```
sudo apt-get install make
sudo apt-get install libtool
sudo apt-get install pkg-config
sudo apt-get install autoconf
sudo apt-get install automake
sudo apt-get install texinfo
sudo apt-get install libusb-1.0
```

Note:

- Version of pkg-config should be 0.2.3 or above.
 - Version of autoconf should be 2.6.4 or above.
 - Version of automake should be 1.9 or above.
 - When using USB-Blaster, ASIX Presto, OpenJTAG and FT2232 as adapters, drivers libFTDI and FTD2XX need to be downloaded and installed.
 - When using CMSIS-DAP, HIDAPI is needed.
-

Build OpenOCD Proceed with configuring and building OpenOCD:

```
cd ~/esp/openocd-esp32
./bootstrap
./configure
make
```

Optionally you can add `sudo make install` step at the end. Skip it, if you have an existing OpenOCD (from e.g., another development platform), as it may get overwritten.

Note:

- Should an error occur, resolve it and try again until the command `make` works.

- If there is a submodule problem from OpenOCD, please `cd` to the `openocd-esp32` directory and input `git submodule update --init`.
 - If the `./configure` is successfully run, information of enabled JTAG will be printed under OpenOCD configuration summary.
 - If the information of your device is not shown in the log, use `./configure` to enable it as described in `../openocd-esp32/doc/INSTALL.txt`.
 - For details concerning compiling OpenOCD, please refer to `openocd-esp32/README`.
-

Once `make` process is successfully completed, the executable of OpenOCD will be saved in `~/openocd-esp32/bin` directory.

Next Steps To carry on with debugging environment setup, proceed to section [Configuring ESP32-P4 Target](#).

Building OpenOCD from Sources for MacOS

The following instructions are alternative to downloading binary OpenOCD from [Espressif GitHub](#). To quickly setup the binary OpenOCD, instead of compiling it yourself, backup and proceed to section [Setup of OpenOCD](#).

Download Sources of OpenOCD The sources for the ESP32-P4-enabled variant of OpenOCD are available from Espressif GitHub under <https://github.com/espressif/openocd-esp32>. To download the sources, use the following commands:

```
cd ~/esp
git clone --recursive https://github.com/espressif/openocd-esp32.git
```

The clone of sources should be now saved in `~/esp/openocd-esp32` directory.

Install Dependencies Install packages that are required to compile OpenOCD using Homebrew:

```
brew install automake libtool libusb wget gcc@4.9 pkg-config
```

Build OpenOCD Proceed with configuring and building OpenOCD:

```
cd ~/esp/openocd-esp32
./bootstrap
./configure
make
```

Optionally you can add `sudo make install` step at the end. Skip it, if you have an existing OpenOCD (from e.g., another development platform), as it may get overwritten.

Note:

- Should an error occur, resolve it and try again until the command `make` works.
- Error Unknown command 'raggedright' may indicate that the required version of `texinfo` was not installed on your computer or installed but was not linked to your `PATH`. To resolve this issue make sure `texinfo` is installed and `PATH` is adjusted prior to the `./bootstrap` by running:

```
brew install texinfo
export PATH=/usr/local/opt/texinfo/bin:$PATH
```

- If there is a submodule problem from OpenOCD, please `cd` to the `openocd-esp32` directory and input `git submodule update --init`.

- If the `./configure` is successfully run, information of enabled JTAG will be printed under OpenOCD configuration summary.
 - If the information of your device is not shown in the log, use `./configure` to enable it as described in `../openocd-esp32/doc/INSTALL.txt`.
 - For details concerning compiling OpenOCD, please refer to `openocd-esp32/README.OSX`.
-

Once make process is successfully completed, the executable of OpenOCD will be saved in `~/esp/openocd-esp32/src/openocd` directory.

Next Steps To carry on with debugging environment setup, proceed to section [Configuring ESP32-P4 Target](#).

The examples of invoking OpenOCD in this document assume using pre-built binary distribution described in section [Setup of OpenOCD](#).

To use binaries build locally from sources, change the path to OpenOCD executable to `src/openocd` and set the `OPENOCD_SCRIPTS` environment variable so that OpenOCD can find the configuration files. For Linux and macOS:

```
cd ~/esp/openocd-esp32
export OPENOCD_SCRIPTS=$PWD/tcl
```

For Windows:

```
cd %USERPROFILE%\esp\openocd-esp32
set "OPENOCD_SCRIPTS=%CD%\tcl"
```

Example of invoking OpenOCD build locally from sources, for Linux and macOS:

```
src/openocd -f board/esp32p4-builtin.cfg
```

and Windows:

```
src\openocd -f board/esp32p4-builtin.cfg
```

4.13.9 Tips and Quirks

This section provides collection of links to all tips and quirks referred to from various parts of this guide.

Tips and Quirks

This section provides collection of all tips and quirks referred to from various parts of this guide.

Breakpoints and Watchpoints Available ESP32-P4 debugger supports 3 hardware implemented breakpoints and 64 software ones. Hardware breakpoints are implemented by ESP32-P4 chip's logic and can be set anywhere in the code: either in flash or IRAM program's regions. Additionally there are 2 types of software breakpoints implemented by OpenOCD: flash (up to 32) and IRAM (up to 32) breakpoints. Currently GDB can not set software breakpoints in flash. So until this limitation is removed those breakpoints have to be emulated by OpenOCD as hardware ones (see [below](#) for details). ESP32-P4 also supports 3 watchpoints, so 3 variables can be watched for change or read by the GDB command `watch myVariable`. Note that menuconfig option [CONFIG_FREERTOS_WATCHPOINT_END_OF_STACK](#) uses the last watchpoint and will not provide expected results, if you also try to use it within OpenOCD/GDB. See menuconfig's help for detailed description.

What Else Should I Know About Breakpoints? Emulating part of hardware breakpoints using software flash ones means that the GDB command `hb myFunction` which is invoked for function in flash will use pure hardware breakpoint if it is available otherwise one of the 32 software flash breakpoints is used. The same rule applies to `b myFunction`-like commands. In this case GDB will decide what type of breakpoint to set itself. If `myFunction` is resided in writable region (IRAM) software IRAM breakpoint will be used otherwise hardware or software flash breakpoint is used as it is done for `hb` command.

Flash Mappings vs SW Flash Breakpoints In order to set/clear software breakpoints in flash, OpenOCD needs to know their flash addresses. To accomplish conversion from the ESP32-P4 address space to the flash one, OpenOCD uses mappings of program's code regions resided in flash. Those mappings are kept in the image header which is prepended to program binary data (code and data segments) and is specific to every application image written to the flash. So to support software flash breakpoints OpenOCD should know where application image under debugging is resided in the flash. By default OpenOCD reads partition table at 0x8000 and uses mappings from the first found application image, but there can be the cases when it will not work, e.g., partition table is not at standard flash location or even there can be multiple images: one factory and two OTA and you may want to debug any of them. To cover all possible debugging scenarios OpenOCD supports special command which can be used to set arbitrary location of application image to debug. The command has the following format:

```
esp appimage_offset <offset>
```

Offset should be in hex format. To reset to the default behaviour you can specify `-1` as offset.

Note: Since GDB requests memory map from OpenOCD only once when connecting to it, this command should be specified in one of the TCL configuration files, or passed to OpenOCD via its command line. In the latter case command line should look like below:

```
openocd -f board/esp32p4-builtin.cfg -c "init; halt; esp appimage_offset 0x210000"
```

Another option is to execute that command via OpenOCD telnet session and then connect GDB, but it seems to be less handy.

Why Stepping with "next" Does Not Bypass Subroutine Calls? When stepping through the code with `next` command, GDB is internally setting a breakpoint ahead in the code to bypass the subroutine calls. If all 3 breakpoints are already set, this functionality will not work. If this is the case, delete breakpoints to have one "spare". With all breakpoints already used, stepping through the code with `next` command will work as like with `step` command and debugger will step inside subroutine calls.

Support Options for OpenOCD at Compile Time ESP-IDF has some support options for OpenOCD debugging which can be set at compile time:

- `CONFIG_ESP_DEBUG_OCDAWARE` is enabled by default. If a panic or unhandled exception is thrown and a JTAG debugger is connected (ie OpenOCD is running), ESP-IDF will break into the debugger.
- `CONFIG_FREERTOS_WATCHPOINT_END_OF_STACK` (disabled by default) sets watchpoint index 1 (the second of two) at the end of any task stack. This is the most accurate way to debug task stack overflows. Click the link for more details.

Please see the [project configuration menu](#) menu for more details on setting compile-time options.

FreeRTOS Support OpenOCD has explicit support for the ESP-IDF FreeRTOS. GDB can see FreeRTOS tasks as threads. Viewing them all can be done using the GDB `i threads` command, changing to a certain task is done with `thread n`, with `n` being the number of the thread. FreeRTOS detection can be disabled in target's configuration. For more details see [Configuration of OpenOCD for Specific Target](#).

GDB has a Python extension for FreeRTOS support. ESP-IDF automatically loads this module into GDB with the `idf.py gdb` command when the system requirements are met. See more details in [Debugging FreeRTOS Objects](#).

Optimize JTAG Speed In order to achieve higher data rates and minimize number of dropped packets it is recommended to optimize setting of JTAG clock frequency, so it is at maximum and still provides stable operation of JTAG. To do so use the following tips.

1. The upper limit of JTAG clock frequency is 20 MHz if CPU runs at 80 MHz, or 26 MHz if CPU runs at 160 MHz or 240 MHz.
2. Depending on particular JTAG adapter and the length of connecting cables, you may need to reduce JTAG frequency below 20 MHz or 26 MHz.
3. In particular reduce frequency, if you get DSR/DIR errors (and they do not relate to OpenOCD trying to read from a memory range without physical memory being present there).
4. ESP-WROVER-KIT operates stable at 20 MHz or 26 MHz.

What Is the Meaning of Debugger's Startup Commands? On startup, debugger is issuing sequence of commands to reset the chip and halt it at specific line of code. This sequence (shown below) is user defined to pick up at most convenient/appropriate line and start debugging.

- `set remote hardware-watchpoint-limit 3` —Restrict GDB to using available hardware watchpoints supported by the chip, 3 for ESP32-P4. For more information see <https://sourceware.org/gdb/onlinedocs/gdb/Remote-Configuration.html>.
- `mon reset halt` —reset the chip and keep the CPUs halted
- `maintenance flush register-cache` —monitor (`mon`) command can not inform GDB that the target state has changed. GDB will assume that whatever stack the target had before `mon reset halt` will still be valid. In fact, after reset the target state will change, and executing `maintenance flush register-cache` is a way to force GDB to get new state from the target.
- `thb app_main` —insert a temporary hardware breakpoint at `app_main`, put here another function name if required
- `c` —resume the program. It will then stop at breakpoint inserted at `app_main`.

Configuration of OpenOCD for Specific Target There are several kinds of OpenOCD configuration files (*.cfg). All configuration files are located in subdirectories of `share/openocd/scripts` directory of OpenOCD distribution (or `tcl/scripts` directory of the source repository). For the purposes of this guide, the most important ones are `board`, `interface` and `target`.

- `interface` configuration files describe the JTAG adapter. Examples of JTAG adapters are ESP-Prog and J-Link.
- `target` configuration files describe specific chips, or in some cases, modules.
- `board` configuration files are provided for development boards with a built-in JTAG adapter. Such files include an `interface` configuration file to choose the adapter, and `target` configuration file to choose the chip/module.

The following configuration files are available for ESP32-P4:

Table 4: OpenOCD configuration files for ESP32-P4

Name	Description
<code>board/esp32p4-builtin.cfg</code>	Board configuration file for ESP32-P4 through built-in USB, includes target and adapter configuration.
<code>board/esp32p4-ftdi.cfg</code>	Board configuration file for ESP32-P4 for via externally connected FTDI-based probe like ESP-Prog, includes target and adapter configuration.
<code>target/esp32p4.cfg</code>	ESP32-P4 target configuration file. Can be used together with one of the <code>interface/</code> configuration files.
<code>interface/esp_usb_jtag.cfg</code>	JTAG adapter configuration file for ESP32-P4.
<code>interface/ftdi/esp_ftdi.cfg</code>	JTAG adapter configuration file for ESP-Prog boards.

If you are using one of the boards which have a pre-defined configuration file, you only need to pass one `-f` argument to OpenOCD, specifying that file.

If you are using a board not listed here, you need to specify both the interface configuration file and target configuration file.

Custom Configuration Files OpenOCD configuration files are written in TCL, and include a variety of choices for customization and scripting. This can be useful for non-standard debugging situations. Please refer to [OpenOCD Manual](#) for the TCL scripting reference.

OpenOCD Configuration Variables The following variables can be optionally set before including the ESP-specific target configuration file. This can be done either in a custom configuration file, or from the command line.

The syntax for setting a variable in TCL is:

```
set VARIABLE_NAME value
```

To set a variable from the command line (replace the name of .cfg file with the correct file for your board):

```
openocd -c 'set VARIABLE_NAME value' -f board/esp-xxxxx-kit.cfg
```

It is important to set the variable before including the ESP-specific configuration file, otherwise the variable will not have effect. You can set multiple variables by repeating the `-c` option.

Table 5: Common ESP-related OpenOCD variables

Variable	Description
ESP_RTOS	Set to <code>none</code> to disable RTOS support. In this case, thread list will not be available in GDB. Can be useful when debugging FreeRTOS itself, and stepping through the scheduler code.
ESP_FLASH_SIZE	Set to 0 to disable Flash breakpoints support.
ESP_SEMIHOST_BASED	Set to the path (on the host) which will be the default directory for semihosting functions.
ESP_ONLYCPU	For multi-core targets, can be set to 1 to only enable single core debugging.

How Debugger Resets ESP32-P4? The board can be reset by entering `mon reset` or `mon reset halt` into GDB.

Can JTAG Pins Be Used for Other Purposes? Operation of JTAG may be disturbed, if some other hardware is connected to JTAG pins besides ESP32-P4 module and JTAG adapter. ESP32-P4 JTAG is using the following pins:

Table 6: ESP32-P4 pins and JTAG signals

ESP32-P4 Pin	JTAG Signal
MTDO / GPIO5	TDO
MTDI / GPIO3	TDI
MTCK / GPIO2	TCK
MTMS / GPIO4	TMS

JTAG communication will likely fail, if configuration of JTAG pins is changed by a user application. If OpenOCD initializes correctly (detects all the CPU cores in the SOC), but loses sync and spews out a lot of DTR/DIR errors when the program is running, it is likely that the application reconfigures the JTAG pins to something else, or the user forgot to connect Vtar to a JTAG adapter that requires it.

JTAG with Flash Encryption or Secure Boot By default, enabling Flash Encryption and/or Secure Boot will disable JTAG debugging. On first boot, the bootloader will burn an eFuse bit to permanently disable JTAG at the same time it enables the other features.

The project configuration option `CONFIG_SECURE_BOOT_ALLOW_JTAG` will keep JTAG enabled at this time, removing all physical security but allowing debugging. (Although the name suggests Secure Boot, this option can be applied even when only Flash Encryption is enabled).

However, OpenOCD may attempt to automatically read and write the flash in order to set *software breakpoints*. This has two problems:

- Software breakpoints are incompatible with Flash Encryption, OpenOCD currently has no support for encrypting or decrypting flash contents.
- If Secure Boot is enabled, setting a software breakpoint will change the digest of a signed app and make the signature invalid. This means if a software breakpoint is set and then a reset occurs, the signature verification will fail on boot.

To disable software breakpoints while using JTAG, add an extra argument `-c 'set ESP_FLASH_SIZE 0'` to the start of the OpenOCD command line, see *OpenOCD Configuration Variables*.

Note: For the same reason, the ESP-IDF app may fail bootloader verification of app signatures, when this option is enabled and a software breakpoint is set.

Reporting Issues with OpenOCD/GDB In case you encounter a problem with OpenOCD or GDB programs itself and do not find a solution searching available resources on the web, open an issue in the OpenOCD issue tracker under <https://github.com/espresif/openocd-esp32/issues>.

1. In issue report provide details of your configuration:
 - a. JTAG adapter type, and the chip/module being debugged.
 - b. Release of ESP-IDF used to compile and load application that is being debugged.
 - c. Details of OS used for debugging.
 - d. Is OS running natively on a PC or on a virtual machine?
2. Create a simple example that is representative to observed issue. Describe steps how to reproduce it. In such an example debugging should not be affected by non-deterministic behaviour introduced by the Wi-Fi stack, so problems will likely be easier to reproduce, if encountered once.
3. Prepare logs from debugging session by adding additional parameters to start up commands.
OpenOCD:

```
openocd -l openocd_log.txt -d3 -f board/esp32p4-builtin.cfg
```

Logging to a file this way will prevent information displayed on the terminal. This may be a good thing taken amount of information provided, when increased debug level `-d3` is set. If you still like to see the log on the screen, then use another command instead:

```
openocd -d3 -f board/esp32p4-builtin.cfg 2>&1 | tee openocd.log
```

Debugger:

```
riscv32-esp-elf-gdb -ex "set remotelogfile gdb_log.txt" <all other options>
```

Optionally add command `remotelogfile gdb_log.txt` to the `gdbinit` file.

4. Attach both `openocd_log.txt` and `gdb_log.txt` files to your issue report.

4.13.10 Related Documents

Using Debugger

This section covers the steps to configure and run a debugger using various methods, including:

- *Eclipse*
- *Command Line*
- *Idf.py Debug Targets*

For how to run a debugger from VS Code, see [Configuration for Visual Studio Code Debug](#).

Eclipse

Note: It is recommended to first check if debugger works using [Idf.py Debug Targets](#) or from [Command Line](#) and then move to using Eclipse.

Eclipse is an integrated development environment (IDE) that provides a powerful set of tools for developing and debugging software applications. For ESP-IDF applications, [IDF Eclipse plugin](#) provides two ways of debugging:

1. [ESP-IDF GDB OpenOCD Debugging](#)
2. GDB Hardware Debugging

By default, Eclipse supports OpenOCD Debugging via the GDB Hardware Debugging plugin, which requires starting the OpenOCD server from the command line and configuring the GDB client from Eclipse to start with the debugging. This approach can be time-consuming and error-prone.

To make the debugging process easier, the IDF Eclipse plugin has a customized ESP-IDF GDB OpenOCD Debugging functionality. This functionality supports configuring the OpenOCD server and GDB client from within Eclipse. All the required configuration parameters will be pre-filled by the plugin, and you can start debugging with just a click of a button.

Therefore, it is recommended to use the [ESP-IDF GDB OpenOCD Debugging](#) via the IDF Eclipse plugin.

GDB Hardware Debugging

Note: This approach is recommended only if you are unable to debug using [ESP-IDF GDB OpenOCD Debugging](#) for some reason.

To install the GDB Hardware Debugging plugin, open Eclipse and select **Help > Install New Software**.

After installation is complete, follow these steps to configure the debugging session. Please note that some configuration parameters are generic, while others are project-specific. This will be shown below by configuring debugging for "blink" example project. If not done already, add this project to Eclipse workspace following [Eclipse Plugin](#). The source of [get-started/blink](#) application is available in [examples](#) directory of ESP-IDF repository.

1. In Eclipse, go to **Run > Debug Configuration**. A new window will open. In the left pane of the window, double-click **GDB Hardware Debugging** (or select **GDB Hardware Debugging** and press the **New** button) to create a new configuration.
2. In a form that will show up on the right, enter the **Name** : of this configuration, e.g., "Blink checking".
3. On the **Main** tab below, under **Project :**, press the **Browse** button and select the **blink** project.
4. In the next line under **C/C++ Application :**, press the **Browse** button and select the **blink.elf** file. If **blink.elf** is not there, it is likely that this project has not been built yet. Refer to the [Eclipse Plugin](#) for instructions.
5. Finally, under **Build** (if required) before launching click **Disable auto build**.
A sample window with settings entered in points 1 - 5 is shown below.
6. Click the **Debugger** tab. In field **GDB Command**, enter **riscv32-esp-elf-gdb** to invoke the debugger.
7. Change the default configuration of the **Remote host** by entering **3333** under the **Port number**.
Configuration entered in points 6 and 7 is shown on the following picture.
8. The last tab that requires changing the default configuration is **Startup**. Under **Initialization Commands** uncheck **Reset** and **Delay (seconds)** and **Halt**. Then, in the entry field below, enter the following lines:

```
mon reset halt
maintenance flush register-cache
set remote hardware-watchpoint-limit 2
```

Note: To automatically update the image in the flash before starting a new debug session, add the following command lines to the beginning of the **Initialization Commands** textbox:




Fig. 11: Configuration of GDB Hardware Debugging - Main tab




Fig. 12: Configuration of GDB Hardware Debugging - Debugger tab

```
mon reset halt
mon program_esp ${workspace_loc:blink/build/blink.bin} 0x10000 verify
```

For description of `program_esp` command, see [Upload Application for Debugging](#).

9. Uncheck the Load image option under Load Image and Symbols.
10. Further down on the same tab, establish an initial breakpoint to halt CPUs after they are reset by debugger. The plugin will set this breakpoint at the beginning of the function entered under Set break point at:. Checkout this option and enter `app_main` in provided field.
11. Checkout Resume option. This will make the program to resume after `mon reset halt` is invoked per point 8. The program will then stop at breakpoint inserted at `app_main`. Configuration described in points 8 - 11 is shown below.




Fig. 13: Configuration of GDB Hardware Debugging - Startup tab

If the Startup sequence looks convoluted and respective Initialization Commands are unclear, check [What Is the Meaning of Debugger's Startup Commands?](#) for additional explanation.

12. If you have completed the [Configuring ESP32-P4 Target](#) steps described above, so the target is running and ready to talk to debugger, go right to debugging by pressing Debug button. Otherwise press Apply to save changes, go back to [Configuring ESP32-P4 Target](#) and return here to start debugging. Once all configuration steps 1-12 are satisfied, the new Eclipse perspective called "Debug" will open, as shown in the example picture below. If you are not quite sure how to use GDB, check [Eclipse](#) example debugging session in section [Debugging Examples](#).

Command Line




Fig. 14: Debug Perspective in Eclipse

1. Begin by completing the steps described under *Configuring ESP32-P4 Target*. This is prerequisite to start a debugging session.
2. Open a new terminal session and go to the directory that contains the project for debugging, e.g.,

```
cd ~/esp/blink
```

3. When launching a debugger, you will need to provide a couple of configuration parameters and commands. The build system generates several `.gdbinit` files to facilitate efficient debugging. Paths to these files can be found in the `build/project_description.json`, under the `gdbinit_files` section. The paths to these files are defined as follows:

```
"gdbinit_files": {
  "01_symbols": "application_path/build/gdbinit/symbols",
  "02_prefix_map": "application_path/build/gdbinit/prefix_map",
  "03_py_extensions": "application_path/build/gdbinit/py_extensions",
  "04_connect": "application_path/build/gdbinit/connect"
}
```

The `XX_` prefix in the JSON keys is included to have ability to sort them. Sorted fields indicate the recommended order in which to provide the data to GDB.

Descriptions of the generated `.gdbinit` files:

- `symbols` - contains symbol sources for debugging.
- `prefix_map` - configures the prefix map to modify source paths in GDB. For more details, see *Reproducible Builds and Debugging*.
- `py_extensions` - initializes Python extensions in GDB. This requires Python built with `libpython` and a version supported by GDB. To verify compatibility, run `riscv32-esp-elf-gdb --batch-silent --ex "python import os"`, which should complete without errors.
- `connect` - contains commands necessary for establishing a connection to the target device.

To enhance your debugging experience, you can also create custom `.gdbinit` files, used either alongside or in place of the generated configurations.

4. Now you are ready to launch GDB. Use the following example command to load symbols and connect to the target (`-q` option added to minimize startup output):

```
riscv32-esp-elf-gdb -q -x build/gdbinit/symbols -x build/gdbinit/prefix_map -x build/gdbinit/connect build/blink.elf
```

5. If the previous steps have been done correctly, you will see a similar log concluded with the (gdb) prompt:

```
riscv32-esp-elf-gdb -q -x build/gdbinit/symbols -x build/gdbinit/prefix_map -x build/gdbinit/connect build/blink.elf

user-name@computer-name:~/esp-idf/examples/get-started/blink$ riscv32-esp-elf-gdb -q -x build/gdbinit/symbols -x build/gdbinit/connect build/blink.elf
Reading symbols from build/blink.elf...
add symbol table from file "/home/user-name/esp-idf/examples/get-started/blink/build/bootloader/bootloader.elf"
[Switching to Thread 1070141764]
app_main () at /home/user-name/esp-idf/examples/get-started/blink/main/blink_example_main.c:95
95      configure_led();
add symbol table from file "/home/alex/.espressif/tools/esp-rom-elfs/20241011/esp32p4_rev0_rom.elf"
JTAG tap: esp32p4.tap0 tap/device found: 0x00005c25 (mfg: 0x612 (Espressif Systems), part: 0x0005, ver: 0x0)
[esp32p4] Reset cause (3) - (Software core reset)
Hardware assisted breakpoint 1 at 0x42009436: file /home/user-name/esp-idf/examples/get-started/blink/main/blink_example_main.c, line 92.
[Switching to Thread 1070139884]

Thread 2 "main" hit Temporary breakpoint 1, app_main () at /home/user-name/esp-idf/examples/get-started/blink/main/blink_example_main.c:92
```

(continues on next page)

(continued from previous page)

```
92      {
(gdb)
```

Note that the third-to-last line indicates the debugger has halted at a breakpoint established in `build/gdbinit/connect` file at function `app_main()`. Since the processor is halted, the LED should not be blinking. If this matches what you observe, you are ready to start debugging.

If you are not sure how to use GDB, check [Command Line](#) example debugging session in section [Debugging Examples](#).

Idf.py Debug Targets It is also possible to execute the described debugging tools conveniently from `idf.py`. These commands are supported:

1. `idf.py openocd`
Runs OpenOCD in a console with configuration defined in the environment or via command line. It uses default script directory defined as `OPENOCD_SCRIPTS` environmental variable, which is automatically added from an Export script (`export.sh` or `export.bat`). It is possible to override the script location using command line argument `--openocd-scripts`.
To configure the JTAG configuration for the current board, please use the environmental variable `OPENOCD_COMMANDS` or `--openocd-commands` command line argument. If none of the above is defined, OpenOCD is started with `-f board/esp32p4-builtin.cfg` board definition.
2. `idf.py gdb`
Starts GDB in the same way as the [Command Line](#), using generated GDB scripts referring to the current project's ELF file. For more details, see [Command Line](#).
3. `idf.py gdbtui`
The same as 2, but starts the gdb with `tui` argument, allowing for a simple source code view.
4. `idf.py gdbgui`
Starts [gdbgui](#) debugger frontend enabling out-of-the-box debugging in a browser window. To enable this option, run the install script with the `"--enable-gdbgui"` argument, e.g., `install.sh --enable-gdbgui`.
You can combine these debugging actions on a single command line, allowing for convenient setup of blocking and non-blocking actions in one step. `idf.py` implements a simple logic to move the background actions (such as `openocd`) to the beginning and the interactive ones (such as `gdb`, `monitor`) to the end of the action list. An example of a very useful combination is:

```
idf.py openocd gdbgui monitor
```

The above command runs OpenOCD in the background, starts [gdbgui](#) to open a browser window with active debugger frontend and opens a serial monitor in the active console.

Debugging Examples

This section describes debugging with GDB from [Eclipse](#) as well as from [Command Line](#).

Eclipse Verify if your target is ready and loaded with [get-started/blink](#) example. Configure and start debugger following steps in section [Eclipse](#). Pick up where target was left by debugger, i.e., having the application halted at breakpoint established at `app_main()`.

Examples in This Section

1. [Navigating Through the Code, Call Stack and Threads](#)
2. [Setting and Clearing Breakpoints](#)
3. [Halting the Target Manually](#)
4. [Stepping Through the Code](#)
5. [Checking and Setting Memory](#)
6. [Watching and Setting Program Variables](#)
7. [Setting Conditional Breakpoints](#)




Fig. 15: Debug Perspective in Eclipse

Navigating Through the Code, Call Stack and Threads When the target is halted, debugger shows the list of threads in "Debug" window. The line of code where program halted is highlighted in another window below, as shown on the following picture. The LED stops blinking.




Fig. 16: Target halted during debugging

Specific thread where the program halted is expanded showing the call stack. It represents function calls that lead up to the highlighted line of code, where the target halted. The first line of call stack under Thread #1 contains the last called function `app_main()`, that in turn was called from function `main_task()` shown in a line below. Each line of the stack also contains the file name and line number where the function was called. By clicking/highlighting the stack entries, in window below, you will see contents of this file.

By expanding threads you can navigate throughout the application. Expand Thread #5 that contains much longer call stack. You will see there, besides function calls, numbers like `0x4000000c`. They represent addresses of binary code not provided in source form.

In another window on right, you can see the disassembled machine code no matter if your project provides it in source or only the binary form.

Go back to the `app_main()` in Thread #1 to familiar code of `blink.c` file that will be examined in more details in the following examples. Debugger makes it easy to navigate through the code of entire application. This comes handy when stepping through the code and working with breakpoints and will be discussed below.

Setting and Clearing Breakpoints When debugging, we would like to be able to stop the application at critical lines of code and then examine the state of specific variables, memory and registers/peripherals. To do so we are using breakpoints. They provide a convenient way to quickly get to and halt the application at specific line.

Let's establish two breakpoints when the state of LED changes. Basing on code listing above, this happens at lines 33 and 36. To do so, hold the "Control" on the keyboard and double click on number 33 in file `blink.c` file. A dialog will open where you can confirm your selection by pressing "OK" button. If you do not like to see the dialog just double click the line number. Set another breakpoint in line 36.




Fig. 17: Navigate through the call stack




Fig. 18: Setting a breakpoint

Information how many breakpoints are set and where is shown in window "Breakpoints" on top right. Click "Show Breakpoints Supported by Selected Target" to refresh this list. Besides the two just set breakpoints the list may contain temporary breakpoint at function `app_main()` established at debugger start. As maximum two breakpoints are allowed (see [Breakpoints and Watchpoints Available](#)), you need to delete it, or debugging will fail.




Fig. 19: Three breakpoints are set / maximum two are allowed

If you now click "Resume" (click `blink_task()` under "Tread #8", if "Resume" button is grayed out), the processor will run and halt at a breakpoint. Clicking "Resume" another time will make it run again, halt on second breakpoint, and so on.

You will be also able to see that LED is changing the state after each click to "Resume" program execution.

Read more about breakpoints under [Breakpoints and Watchpoints Available](#) and [What Else Should I Know About Breakpoints?](#)

Halting the Target Manually When debugging, you may resume application and enter code waiting for some event or staying in infinite loop without any break points defined. In such case, to go back to debugging mode, you can break program execution manually by pressing "Suspend" button.

To check it, delete all breakpoints and click "Resume". Then click "Suspend". Application will be halted at some random point and LED will stop blinking. Debugger will expand tread and highlight the line of code where application halted.

In particular case above, the application has been halted in line 52 of code in file `freertos_hooks.c`. Now you can resume it again by pressing "Resume" button or do some debugging as discussed below.

Stepping Through the Code It is also possible to step through the code using "Step Into (F5)" and "Step Over (F6)" commands. The difference is that "Step Into (F5)" is entering inside subroutines calls, while "Step Over (F6)" steps over the call, treating it as a single source line.




Fig. 20: Target halted manually

Before being able to demonstrate this functionality, using information discussed in previous paragraph, make sure that you have only one breakpoint defined at line 36 of `blink.c`.

Resume program by entering pressing F8 and let it halt. Now press "Step Over (F6)", one by one couple of times, to see how debugger is stepping one program line at a time.




Fig. 21: Stepping through the code with "Step Over (F6)"

If you press "Step Into (F5)" instead, then debugger will step inside subroutine calls.

In this particular case debugger stepped inside `gpio_set_level(BLINK_GPIO, 0)` and effectively moved to `gpio.c` driver code.

See [Why Stepping with "next" Does Not Bypass Subroutine Calls?](#) for potential limitation of using `next` command.

Checking and Setting Memory To display or set contents of memory use "Memory" tab at the bottom of "Debug" perspective.

With the "Memory" tab, we will read from and write to the memory location `0x3FF44004` labeled as `GPIO_OUT_REG` used to set and clear individual GPIO's.

For more information, see *ESP32-P4 Technical Reference Manual > IO MUX and GPIO Matrix (GPIO, IO_MUX)* [PDF].

Being in the same `blink.c` project as before, set two breakpoints right after `gpio_set_level` instruction. Click "Memory" tab and then "Add Memory Monitor" button. Enter `0x3FF44004` in provided dialog.

Now resume program by pressing F8 and observe "Monitor" tab.

You should see one bit being flipped over at memory location `0x3FF44004` (and LED changing the state) each time F8 is pressed.

To set memory use the same "Monitor" tab and the same memory location. Type in alternate bit pattern as previously observed. Immediately after pressing enter you will see LED changing the state.




Fig. 22: Stepping through the code with "Step Into (F5)"




Fig. 23: Observing memory location 0x3FF44004 changing one bit to "ON"




Fig. 24: Observing memory location 0x3FF44004 changing one bit to "OFF"

Watching and Setting Program Variables A common debugging task is checking the value of a program variable as the program runs. To be able to demonstrate this functionality, update file `blink.c` by adding a declaration of a global variable `int i` above definition of function `blink_task`. Then add `i++` inside `while(1)` of this function to get `i` incremented on each blink.

Exit debugger, so it is not confused with new code, build and flash the code to the ESP and restart debugger. There is no need to restart OpenOCD.

Once application is halted, enter a breakpoint in the line where you put `i++`.

In next step, in the window with "Breakpoints", click the "Expressions" tab. If this tab is not visible, then add it by going to the top menu `Window > Show View > Expressions`. Then click "Add new expression" and enter `i`.

Resume program execution by pressing `F8`. Each time the program is halted you will see `i` value being incremented.

To modify `i` enter a new number in "Value" column. After pressing "Resume (`F8`)" the program will keep incrementing `i` starting from the new entered number.

Setting Conditional Breakpoints Here comes more interesting part. You may set a breakpoint to halt the program execution, if certain condition is satisfied. Right click on the breakpoint to open a context menu and select "Breakpoint Properties". Change the selection under "Type:" to "Hardware" and enter a "Condition:" like `i == 2`.

If current value of `i` is less than 2 (change it if required) and program is resumed, it will blink LED in a loop until condition `i == 2` gets true and then finally halt.

Command Line Verify if your target is ready and loaded with [get-started/blink](#) example. Configure and start debugger following steps in section [Command Line](#). Pick up where target was left by debugger, i.e. having the application halted at breakpoint established at `app_main()`:

```
Temporary breakpoint 1, app_main () at /home/user-name/esp/blink/main/./blink.c:43
43      xTaskCreate(&blink_task, "blink_task", configMINIMAL_STACK_SIZE, NULL, 5, &
      ↪NULL);
(gdb)
```

Examples in This Section

1. [Navigating Through the Code, Call Stack and Threads](#)
2. [Setting and Clearing Breakpoints](#)
3. [Halting and Resuming the Application](#)
4. [Stepping Through the Code](#)
5. [Checking and Setting Memory](#)




Fig. 25: Watching program variable "i"




Fig. 26: Setting a conditional breakpoint

6. *Watching and Setting Program Variables*
7. *Setting Conditional Breakpoints*
8. *Debugging FreeRTOS Objects*

Navigating Through the Code, Call Stack and Threads When you see the (gdb) prompt, the application is halted. LED should not be blinking.

To find out where exactly the code is halted, enter `l` or `list`, and debugger will show couple of lines of code around the halt point (line 43 of code in file `blink.c`)

```
(gdb) l
38     }
39 }
40
41 void app_main()
42 {
43     xTaskCreate(&blink_task, "blink_task", configMINIMAL_STACK_SIZE, NULL, 5, ↵
↵NULL);
44 }
(gdb)
```

Check how code listing works by entering, e.g., `l 30, 40` to see particular range of lines of code.

You can use `bt` or `backtrace` to see what function calls lead up to this code:

```
(gdb) bt
#0  app_main () at /home/user-name/esp/blink/main/./blink.c:43
#1  0x400d057e in main_task (args=0x0) at /home/user-name/esp/esp-idf/components/
↵esp32p4/./cpu_start.c:339
(gdb)
```

Line #0 of output provides the last function call before the application halted, i.e., `app_main ()` we have listed previously. The `app_main ()` was in turn called by function `main_task` from line 339 of code located in file `cpu_start.c`.

To get to the context of `main_task` in file `cpu_start.c`, enter `frame N`, where `N = 1`, because the `main_task` is listed under #1):

```
(gdb) frame 1
#1  0x400d057e in main_task (args=0x0) at /home/user-name/esp/esp-idf/components/
↵esp32p4/./cpu_start.c:339
339     app_main();
(gdb)
```

Enter `l` and this will reveal the piece of code that called `app_main ()` (in line 339):

```
(gdb) l
334     ;
335 }
336 #endif
337 //Enable allocation in region where the startup stacks were located.
338 heap_caps_enable_nonos_stack_heaps();
339 app_main();
340 vTaskDelete(NULL);
341 }
342
(gdb)
```

By listing some lines before, you will see the function name `main_task` we have been looking for:

```
(gdb) l 326, 341
326 static void main_task(void* args)
```

(continues on next page)

(continued from previous page)

```

327 {
328     // Now that the application is about to start, disable boot watchdogs
329     REG_CLR_BIT(TIMG_WDTCONFIG0_REG(0), TIMG_WDT_FLASHBOOT_MOD_EN_S);
330     REG_CLR_BIT(RTC_CNTL_WDTCONFIG0_REG, RTC_CNTL_WDT_FLASHBOOT_MOD_EN);
331 #if !CONFIG_FREERTOS_UNICORE
332     // Wait for FreeRTOS initialization to finish on APP CPU, before replacing
↳ its startup stack
333     while (port_xSchedulerRunning[1] == 0) {
334         ;
335     }
336 #endif
337     //Enable allocation in region where the startup stacks were located.
338     heap_caps_enable_nonos_stack_heaps();
339     app_main();
340     vTaskDelete(NULL);
341 }
(gdb)

```

To see the other code, enter `i threads`. This will show the list of threads running on target:

```

(gdb) i threads
  Id   Target Id           Frame
   8   Thread 1073411336 (dport) 0x400d0848 in dport_access_init_core (arg=
↳ <optimized out>)
      at /home/user-name/esp/esp-idf/components/esp32p4/./dport_access.c:170
   7   Thread 1073408744 (ipc0) xQueueGenericReceive (xQueue=0x3ffae694,
↳ pvBuffer=0x0, xTicksToWait=1644638200,
      xJustPeeking=0) at /home/user-name/esp/esp-idf/components/freertos/./queue.
↳ c:1452
   6   Thread 1073431096 (Tmr Svc) prvTimerTask (pvParameters=0x0)
      at /home/user-name/esp/esp-idf/components/freertos/./timers.c:445
   5   Thread 1073410208 (ipc1 : Running) 0x4000bfea in ?? ()
   4   Thread 1073432224 (dport) dport_access_init_core (arg=0x0)
      at /home/user-name/esp/esp-idf/components/esp32p4/./dport_access.c:150
   3   Thread 1073413156 (IDLE) prvIdleTask (pvParameters=0x0)
      at /home/user-name/esp/esp-idf/components/freertos/./tasks.c:3282
   2   Thread 1073413512 (IDLE) prvIdleTask (pvParameters=0x0)
      at /home/user-name/esp/esp-idf/components/freertos/./tasks.c:3282
*  1   Thread 1073411772 (main : Running) app_main () at /home/user-name/esp/blink/
↳ main/./blink.c:43
(gdb)

```

The thread list shows the last function calls per each thread together with the name of C source file if available.

You can navigate to specific thread by entering `thread N`, where `N` is the thread Id. To see how it works go to thread thread 5:

```

(gdb) thread 5
[Switching to thread 5 (Thread 1073410208)]
#0  0x4000bfea in ?? ()
(gdb)

```

Then check the backtrace:

```

(gdb) bt
#0  0x4000bfea in ?? ()
#1  0x40083a85 in vPortCPUReleaseMutex (mux=<optimized out>) at /home/user-name/
↳ esp/esp-idf/components/freertos/./port.c:415
#2  0x40083fc8 in vTaskSwitchContext () at /home/user-name/esp/esp-idf/components/
↳ freertos/./tasks.c:2846
#3  0x4008532b in _frxt_dispatch ()

```

(continues on next page)

(continued from previous page)

```
#4 0x4008395c in xPortStartScheduler () at /home/user-name/esp/esp-idf/components/
↳freertos/./port.c:222
#5 0x4000000c in ?? ()
#6 0x4000000c in ?? ()
#7 0x4000000c in ?? ()
#8 0x4000000c in ?? ()
(gdb)
```

As you see, the backtrace may contain several entries. This will let you check what exact sequence of function calls lead to the code where the target halted. Question marks ?? instead of a function name indicate that application is available only in binary format, without any source file in C language. The value like 0x4000bfea is the memory address of the function call.

Using `bt`, `i` threads, `thread N` and `list` commands we are now able to navigate through the code of entire application. This comes handy when stepping through the code and working with breakpoints and will be discussed below.

Setting and Clearing Breakpoints When debugging, we would like to be able to stop the application at critical lines of code and then examine the state of specific variables, memory and registers/peripherals. To do so we are using breakpoints. They provide a convenient way to quickly get to and halt the application at specific line.

Let's establish two breakpoints when the state of LED changes. Basing on code listing above this happens at lines 33 and 36. Breakpoints may be established using command `break M` where `M` is the code line number:

```
(gdb) break 33
Breakpoint 2 at 0x400db6f6: file /home/user-name/esp/blink/main/./blink.c, line 33.
(gdb) break 36
Breakpoint 3 at 0x400db704: file /home/user-name/esp/blink/main/./blink.c, line 36.
```

If you now enter `c`, the processor will run and halt at a breakpoint. Entering `c` another time will make it run again, halt on second breakpoint, and so on:

```
(gdb) c
Continuing.
Target halted. PRO_CPU: PC=0x400DB6F6 (active) APP_CPU: PC=0x400D10D8

Breakpoint 2, blink_task (pvParameter=0x0) at /home/user-name/esp/blink/main/./
↳blink.c:33
33      gpio_set_level(BLINK_GPIO, 0);
(gdb) c
Continuing.
Target halted. PRO_CPU: PC=0x400DB6F8 (active) APP_CPU: PC=0x400D10D8
Target halted. PRO_CPU: PC=0x400DB704 (active) APP_CPU: PC=0x400D10D8

Breakpoint 3, blink_task (pvParameter=0x0) at /home/user-name/esp/blink/main/./
↳blink.c:36
36      gpio_set_level(BLINK_GPIO, 1);
(gdb)
```

You will be also able to see that LED is changing the state only if you resume program execution by entering `c`.

To examine how many breakpoints are set and where, use command `info break`:

```
(gdb) info break
Num      Type      Disp Enb Address      What
2        breakpoint keep y  0x400db6f6 in blink_task at /home/user-name/esp/
↳blink/main/./blink.c:33
          breakpoint already hit 1 time
3        breakpoint keep y  0x400db704 in blink_task at /home/user-name/esp/
↳blink/main/./blink.c:36
```

(continues on next page)

(continued from previous page)

```
breakpoint already hit 1 time
(gdb)
```

Please note that breakpoint numbers (listed under Num) start with 2. This is because first breakpoint has been already established at function `app_main()` by running command `thb app_main` on debugger launch. As it was a temporary breakpoint, it has been automatically deleted and now is not listed anymore.

To remove breakpoints enter `delete N` command (in short `d N`), where `N` is the breakpoint number:

```
(gdb) delete 1
No breakpoint number 1.
(gdb) delete 2
(gdb)
```

Read more about breakpoints under [Breakpoints and Watchpoints Available](#) and [What Else Should I Know About Breakpoints?](#)

Halting and Resuming the Application When debugging, you may resume application and enter code waiting for some event or staying in infinite loop without any break points defined. In such case, to go back to debugging mode, you can break program execution manually by entering `Ctrl+C`.

To check it delete all breakpoints and enter `c` to resume application. Then enter `Ctrl+C`. Application will be halted at some random point and LED will stop blinking. Debugger will print the following:

```
(gdb) c
Continuing.
^CTarget halted. PRO_CPU: PC=0x400D0C00          APP_CPU: PC=0x400D0C00 (active)
[New Thread 1073433352]

Program received signal SIGINT, Interrupt.
[Switching to Thread 1073413512]
0x400d0c00 in esp_vApplicationIdleHook () at /home/user-name/esp/esp-idf/
↳components/esp32p4/./freertos_hooks.c:52
52          asm("waiti 0");
(gdb)
```

In particular case above, the application has been halted in line 52 of code in file `freertos_hooks.c`. Now you can resume it again by enter `c` or do some debugging as discussed below.

Stepping Through the Code It is also possible to step through the code using `step` and `next` commands (in short `s` and `n`). The difference is that `step` is entering inside subroutines calls, while `next` steps over the call, treating it as a single source line.

To demonstrate this functionality, using command `break` and `delete` discussed in previous paragraph, make sure that you have only one breakpoint defined at line 36 of `blink.c`:

```
(gdb) info break
Num      Type           Disp Enb Address      What
3        breakpoint      keep y   0x400db704 in blink_task at /home/user-name/esp/
↳blink/main/./blink.c:36
breakpoint already hit 1 time
(gdb)
```

Resume program by entering `c` and let it halt:

```
(gdb) c
Continuing.
Target halted. PRO_CPU: PC=0x400DB754 (active)   APP_CPU: PC=0x400D1128
```

(continues on next page)

(continued from previous page)

```
Breakpoint 3, blink_task (pvParameter=0x0) at /home/user-name/esp/blink/main/./
↳ blink.c:36
36         gpio_set_level(BLINK_GPIO, 1);
(gdb)
```

Then enter `n` couple of times to see how debugger is stepping one program line at a time:

```
(gdb) n
Target halted. PRO_CPU: PC=0x400DB756 (active)    APP_CPU: PC=0x400D1128
Target halted. PRO_CPU: PC=0x400DB758 (active)    APP_CPU: PC=0x400D1128
Target halted. PRO_CPU: PC=0x400DC04C (active)    APP_CPU: PC=0x400D1128
Target halted. PRO_CPU: PC=0x400DB75B (active)    APP_CPU: PC=0x400D1128
37         vTaskDelay(1000 / portTICK_PERIOD_MS);
(gdb) n
Target halted. PRO_CPU: PC=0x400DB75E (active)    APP_CPU: PC=0x400D1128
Target halted. PRO_CPU: PC=0x400846FC (active)    APP_CPU: PC=0x400D1128
Target halted. PRO_CPU: PC=0x400DB761 (active)    APP_CPU: PC=0x400D1128
Target halted. PRO_CPU: PC=0x400DB746 (active)    APP_CPU: PC=0x400D1128
33         gpio_set_level(BLINK_GPIO, 0);
(gdb)
```

If you enter `s` instead, then debugger will step inside subroutine calls:

```
(gdb) s
Target halted. PRO_CPU: PC=0x400DB748 (active)    APP_CPU: PC=0x400D1128
Target halted. PRO_CPU: PC=0x400DB74B (active)    APP_CPU: PC=0x400D1128
Target halted. PRO_CPU: PC=0x400DC04C (active)    APP_CPU: PC=0x400D1128
Target halted. PRO_CPU: PC=0x400DC04F (active)    APP_CPU: PC=0x400D1128
gpio_set_level (gpio_num=GPIO_NUM_4, level=0) at /home/user-name/esp/esp-idf/
↳ components/driver/gpio/gpio.c:183
183     GPIO_CHECK(GPIO_IS_VALID_OUTPUT_GPIO(gpio_num), "GPIO output gpio_num error
↳ ", ESP_ERR_INVALID_ARG);
(gdb)
```

In this particular case debugger stepped inside `gpio_set_level(BLINK_GPIO, 0)` and effectively moved to `gpio.c` driver code.

See [Why Stepping with "next" Does Not Bypass Subroutine Calls?](#) for potential limitation of using `next` command.

Checking and Setting Memory Displaying the contents of memory is done with command `x`. With additional parameters you may vary the format and count of memory locations displayed. Run `help x` to see more details. Companion command to `x` is `set` that let you write values to the memory.

We will demonstrate how `x` and `set` work by reading from and writing to the memory location `0x3FF44004` labeled as `GPIO_OUT_REG` used to set and clear individual GPIO's.

For more information, see [ESP32-P4 Technical Reference Manual > IO MUX and GPIO Matrix \(GPIO, IO_MUX\)](#) [PDF].

Being in the same `blink.c` project as before, set two breakpoints right after `gpio_set_level` instruction. Enter two times `c` to get to the break point followed by `x /1wx 0x3FF44004` to display contents of `GPIO_OUT_REG` memory location:

```
(gdb) c
Continuing.
Target halted. PRO_CPU: PC=0x400DB75E (active)    APP_CPU: PC=0x400D1128
Target halted. PRO_CPU: PC=0x400DB74E (active)    APP_CPU: PC=0x400D1128

Breakpoint 2, blink_task (pvParameter=0x0) at /home/user-name/esp/blink/main/./
↳ blink.c:34
34         vTaskDelay(1000 / portTICK_PERIOD_MS);
```

(continues on next page)

(continued from previous page)

```
(gdb) x /1wx 0x3FF44004
0x3ff44004: 0x00000000
(gdb) c
Continuing.
Target halted. PRO_CPU: PC=0x400DB751 (active)    APP_CPU: PC=0x400D1128
Target halted. PRO_CPU: PC=0x400DB75B (active)    APP_CPU: PC=0x400D1128

Breakpoint 3, blink_task (pvParameter=0x0) at /home/user-name/esp/blink/main/./
↳ blink.c:37
37      vTaskDelay(1000 / portTICK_PERIOD_MS);
(gdb) x /1wx 0x3FF44004
0x3ff44004: 0x00000010
(gdb)
```

If you are blinking LED connected to GPIO4, then you should see fourth bit being flipped each time the LED changes the state:

```
0x3ff44004: 0x00000000
...
0x3ff44004: 0x00000010
```

Now, when the LED is off, that corresponds to 0x3ff44004: 0x00000000 being displayed, try using set command to set this bit by writing 0x00000010 to the same memory location:

```
(gdb) x /1wx 0x3FF44004
0x3ff44004: 0x00000000
(gdb) set {unsigned int}0x3FF44004=0x000010
```

You should see the LED to turn on immediately after entering set {unsigned int}0x3FF44004=0x000010 command.

Watching and Setting Program Variables A common debugging task is checking the value of a program variable as the program runs. To be able to demonstrate this functionality, update file `blink.c` by adding a declaration of a global variable `int i` above definition of function `blink_task`. Then add `i++` inside `while(1)` of this function to get `i` incremented on each blink.

Exit debugger, so it is not confused with new code, build and flash the code to the ESP and restart debugger. There is no need to restart OpenOCD.

Once application is halted, enter the command `watch i`:

```
(gdb) watch i
Hardware watchpoint 2: i
(gdb)
```

This will insert so called "watchpoint" in each place of code where variable `i` is being modified. Now enter `continue` to resume the application and observe it being halted:

```
(gdb) c
Continuing.
Target halted. PRO_CPU: PC=0x400DB751 (active)    APP_CPU: PC=0x400D0811
[New Thread 1073432196]

Program received signal SIGTRAP, Trace/breakpoint trap.
[Switching to Thread 1073432196]
0x400db751 in blink_task (pvParameter=0x0) at /home/user-name/esp/blink/main/./
↳ blink.c:33
33      i++;
(gdb)
```


Resume application couple more times so `i` gets incremented. Now you can enter `print i` (in short `p i`) to check the current value of `i`:

```
(gdb) p i
$1 = 3
(gdb)
```

To modify the value of `i` use `set` command as below (you can then print it out to check if it has been indeed changed):

```
(gdb) set var i = 0
(gdb) p i
$3 = 0
(gdb)
```

You may have up to two watchpoints, see [Breakpoints and Watchpoints Available](#).

Setting Conditional Breakpoints Here comes more interesting part. You may set a breakpoint to halt the program execution, if certain condition is satisfied. Delete existing breakpoints and try this:

```
(gdb) break blink.c:34 if (i == 2)
Breakpoint 3 at 0x400db753: file /home/user-name/esp/blink/main/./blink.c, line 34.
(gdb)
```

Above command sets conditional breakpoint to halt program execution in line 34 of `blink.c` if `i == 2`.

If current value of `i` is less than 2 and program is resumed, it will blink LED in a loop until condition `i == 2` gets true and then finally halt:

```
(gdb) set var i = 0
(gdb) c
Continuing.
Target halted. PRO_CPU: PC=0x400DB755 (active)    APP_CPU: PC=0x400D112C
Target halted. PRO_CPU: PC=0x400DB753 (active)    APP_CPU: PC=0x400D112C
Target halted. PRO_CPU: PC=0x400DB755 (active)    APP_CPU: PC=0x400D112C
Target halted. PRO_CPU: PC=0x400DB753 (active)    APP_CPU: PC=0x400D112C

Breakpoint 3, blink_task (pvParameter=0x0) at /home/user-name/esp/blink/main/./
↪blink.c:34
34         gpio_set_level(BLINK_GPIO, 0);
(gdb)
```

Debugging FreeRTOS Objects This part might be interesting when you are debugging FreeRTOS tasks interactions.

Users that need to use the FreeRTOS task interactions can use the GDB `freertos` command. The `freertos` command is not native to GDB and comes from the [freertos-gdb](#) Python extension module. The `freertos` command contains a series of sub-commands as demonstrated in the code snippet:

```
(gdb) freertos
"freertos" must be followed by the name of a subcommand.
List of freertos subcommands:

freertos queue -- Generate a print out of the current queues info.
freertos semaphore -- Generate a print out of the current semaphores info.
freertos task -- Generate a print out of the current tasks and their states.
freertos timer -- Generate a print out of the current timers info.
```

For a more detailed description of this extension, please refer to <https://pypi.org/project/freertos-gdb>.

Note: The freertos-gdb Python module is included as a Python package requirement by ESP-IDF, thus should be automatically installed (see [Step 3. Set up the Tools](#) for more details).

The FreeRTOS extension automatically loads in case GDB is executed with command via `idf.py gdb`. Otherwise, the module could be enabled via the `python import freertos_gdb` command inside GDB.

Users only need to have Python 3.6 (or above) that contains a Python shared library.

Obtaining Help on Commands Commands presented so far should provide a very basis and intended to let you quickly get started with JTAG debugging. Check help what are the other commands at your disposal. To obtain help on syntax and functionality of particular command, being at (gdb) prompt type `help` and command name:

```
(gdb) help next
Step program, proceeding through subroutine calls.
Usage: next [N]
Unlike "step", if the current source line calls a subroutine,
this command does not enter the subroutine, but instead steps over
the call, in effect treating it as a single source line.
(gdb)
```

By typing just `help`, you will get top level list of command classes, to aid you drilling down to more details. Optionally refer to available GDB cheat sheets, for instance <https://darkdust.net/files/GDB%20Cheat%20Sheet.pdf>. Good to have as a reference (even if not all commands are applicable in an embedded environment).

Ending Debugger Session To quit debugger enter `q`:

```
(gdb) q
A debugging session is active.

    Inferior 1 [Remote target] will be detached.

Quit anyway? (y or n) y
Detaching from program: /home/user-name/esp/blink/build/blink.elf, Remote target
Ending remote debugging.
user-name@computer-name:~/esp/blink$
```

- [Using Debugger](#)
- [Debugging Examples](#)
- [Tips and Quirks](#)
- [Application Level Tracing Library](#)
- [Introduction to ESP-Prog Board](#)

4.14 Linker Script Generation

4.14.1 Overview

There are several *memory regions* where code and data can be placed. Code and read-only data are placed by default in flash, writable data in RAM, etc. However, it is sometimes necessary to change these default placements.

For example, it may be necessary to place:

- critical code in RAM for performance reasons.

- executable code in IRAM so that it can be ran while cache is disabled.
- code in RTC memory for use in a wake stub.

With the linker script generation mechanism, it is possible to specify these placements at the component level within ESP-IDF. The component presents information on how it would like to place its symbols, objects or the entire archive. During build, the information presented by the components are collected, parsed and processed; and the placement rules generated is used to link the app.

4.14.2 Quick Start

This section presents a guide for quickly placing code/data to RAM and RTC memory - placements ESP-IDF provides out-of-the-box.

For this guide, suppose we have the following:

```
components
└─ my_component
    ├── CMakeLists.txt
    ├── Kconfig
    ├── src/
    │   ├── my_src1.c
    │   ├── my_src2.c
    │   └── my_src3.c
    └─ my_linker_fragment_file.lf
```

- a component named `my_component` that is archived as library `libmy_component.a` during build
- three source files archived under the library, `my_src1.c`, `my_src2.c` and `my_src3.c` which are compiled as `my_src1.o`, `my_src2.o` and `my_src3.o`, respectively
- under `my_src1.o`, the function `my_function1` is defined; under `my_src2.o`, the function `my_function2` is defined
- there is bool-type config `PERFORMANCE_MODE` (y/n) and int type config `PERFORMANCE_LEVEL` (with range 0-3) in `my_component`'s `Kconfig`

Creating and Specifying a Linker Fragment File

Before anything else, a linker fragment file needs to be created. A linker fragment file is simply a text file with a `.lf` extension upon which the desired placements will be written. After creating the file, it is then necessary to present it to the build system. The instructions for the build systems supported by ESP-IDF are as follows:

In the component's `CMakeLists.txt` file, specify argument `LDFRAGMENTS` in the `idf_component_register` call. The value of `LDFRAGMENTS` can either be an absolute path or a relative path from the component directory to the created linker fragment file.

```
# file paths relative to CMakeLists.txt
idf_component_register(...
    LDFRAGMENTS "path/to/linker_fragment_file.lf" "path/to/
↪another_linker_fragment_file.lf"
    ...
)
```

Specifying Placements

It is possible to specify placements at the following levels of granularity:

- object file (`.obj` or `.o` files)
- symbol (function/variable)
- archive (`.a` files)

Placing Object Files Suppose the entirety of `my_src1.o` is performance-critical, so it is desirable to place it in RAM. On the other hand, the entirety of `my_src2.o` contains symbols needed coming out of deep sleep, so it needs to be put under RTC memory.

In the linker fragment file, we can write:

```
[mapping:my_component]
archive: libmy_component.a
entries:
    my_src1 (noflash)      # places all my_src1 code/read-only data under IRAM/DRAM
    my_src2 (rtc)          # places all my_src2 code/ data and read-only data under_
↪RTC fast memory/RTC slow memory
```

What happens to `my_src3.o`? Since it is not specified, default placements are used for `my_src3.o`. More on default placements [here](#).

Placing Symbols Continuing our example, suppose that among functions defined under `object1.o`, only `my_function1` is performance-critical; and under `object2.o`, only `my_function2` needs to execute after the chip comes out of deep sleep. This could be accomplished by writing:

```
[mapping:my_component]
archive: libmy_component.a
entries:
    my_src1:my_function1 (noflash)
    my_src2:my_function2 (rtc)
```

The default placements are used for the rest of the functions in `my_src1.o` and `my_src2.o` and the entire `object3.o`. Something similar can be achieved for placing data by writing the variable name instead of the function name, like so:

```
my_src1:my_variable (noflash)
```

Warning: There are [limitations](#) in placing code/data at symbol granularity. In order to ensure proper placements, an alternative would be to group relevant code and data into source files, and [use object-granularity placements](#).

Placing Entire Archive In this example, suppose that the entire component archive needs to be placed in RAM. This can be written as:

```
[mapping:my_component]
archive: libmy_component.a
entries:
    * (noflash)
```

Similarly, this places the entire component in RTC memory:

```
[mapping:my_component]
archive: libmy_component.a
entries:
    * (rtc)
```

Configuration-Dependent Placements Suppose that the entire component library should only have special placement when a certain condition is true; for example, when `CONFIG_PERFORMANCE_MODE == y`. This could be written as:

```
[mapping:my_component]
archive: libmy_component.a
```

(continues on next page)

(continued from previous page)

```
entries:
    if PERFORMANCE_MODE = y:
        * (noflash)
    else:
        * (default)
```

For a more complex config-dependent placement, suppose the following requirements: when `CONFIG_PERFORMANCE_LEVEL == 1`, only `object1.o` is put in RAM; when `CONFIG_PERFORMANCE_LEVEL == 2`, `object1.o` and `object2.o`; and when `CONFIG_PERFORMANCE_LEVEL == 3` all object files under the archive are to be put into RAM. When these three are false however, put entire library in RTC memory. This scenario is a bit contrived, but, it can be written as:

```
[mapping:my_component]
archive: libmy_component.a
entries:
    if PERFORMANCE_LEVEL = 1:
        my_src1 (noflash)
    elif PERFORMANCE_LEVEL = 2:
        my_src1 (noflash)
        my_src2 (noflash)
    elif PERFORMANCE_LEVEL = 3:
        my_src1 (noflash)
        my_src2 (noflash)
        my_src3 (noflash)
    else:
        * (rtc)
```

Nesting condition-checking is also possible. The following is equivalent to the snippet above:

```
[mapping:my_component]
archive: libmy_component.a
entries:
    if PERFORMANCE_LEVEL <= 3 && PERFORMANCE_LEVEL > 0:
        if PERFORMANCE_LEVEL >= 1:
            object1 (noflash)
            if PERFORMANCE_LEVEL >= 2:
                object2 (noflash)
                if PERFORMANCE_LEVEL >= 3:
                    object2 (noflash)
    else:
        * (rtc)
```

The 'default' Placements

Up until this point, the term 'default placements' has been mentioned as fallback placements when the placement rules `rtc` and `noflash` are not specified. It is important to note that the tokens `noflash` or `rtc` are not merely keywords, but are actually entities called fragments, specifically [schemes](#).

In the same manner as `rtc` and `noflash` are schemes, there exists a default scheme which defines what the default placement rules should be. As the name suggests, it is where code and data are usually placed, i.e., code/constants is placed in flash, variables placed in RAM, etc. More on the default scheme [here](#).

Note: For an example of an ESP-IDF component using the linker script generation mechanism, see [freertos/CMakeLists.txt](#). `freertos` uses this to place its object files to the instruction RAM for performance reasons.

This marks the end of the quick start guide. The following text discusses the internals of the mechanism in a little bit more detail. The following sections should be helpful in creating custom placements or modifying default behavior.

4.14.3 Linker Script Generation Internals

Linking is the last step in the process of turning C/C++ source files into an executable. It is performed by the toolchain's linker, and accepts linker scripts which specify code/data placements, among other things. With the linker script generation mechanism, this process is no different, except that the linker script passed to the linker is dynamically generated from: (1) the collected *linker fragment files* and (2) *linker script template*.

Note: The tool that implements the linker script generation mechanism lives under [tools/ldgen](#).

Linker Fragment Files

As mentioned in the quick start guide, fragment files are simple text files with the `.lf` extension containing the desired placements. This is a simplified description of what fragment files contain, however. What fragment files actually contain are 'fragments'. Fragments are entities which contain pieces of information which, when put together, form placement rules that tell where to place sections of object files in the output binary. There are three types of fragments: *sections*, *scheme* and *mapping*.

Grammar The three fragment types share a common grammar:

```
[type:name]
key: value
key:
    value
    value
    value
    ...
```

- **type:** Corresponds to the fragment type, can either be *sections*, *scheme* or *mapping*.
- **name:** The name of the fragment, should be unique for the specified fragment type.
- **key, value:** Contents of the fragment; each fragment type may support different keys and different grammars for the key values.
 - For *sections* and *scheme*, the only supported key is *entries*
 - For *mappings*, both *archive* and *entries* are supported.

Note: In cases where multiple fragments of the same type and name are encountered, an exception is thrown.

Note: The only valid characters for fragment names and keys are alphanumeric characters and underscore.

Condition Checking

Condition checking enable the linker script generation to be configuration-aware. Depending on whether expressions involving configuration values are true or not, a particular set of values for a key can be used. The evaluation uses `eval_string` from `kconfiglib` package and adheres to its required syntax and limitations. Supported operators are as follows:

- **comparison**
 - `LessThan <`
 - `LessThanOrEqualTo <=`
 - `MoreThan >`
 - `MoreThanOrEqualTo >=`
 - `Equal =`
 - `NotEqual !=`
- **logical**
 - `Or ||`

- And &&
 - Negation !
- **grouping**
 - Parenthesis ()

Condition checking behaves as you would expect an `if...elseif/elif...else` block in other languages. Condition-checking is possible for both key values and entire fragments. The two sample fragments below are equivalent:

```
# Value for keys is dependent on config
[type:name]
key_1:
    if CONDITION = y:
        value_1
    else:
        value_2
key_2:
    if CONDITION = y:
        value_a
    else:
        value_b
```

```
# Entire fragment definition is dependent on config
if CONDITION = y:
    [type:name]
    key_1:
        value_1
    key_2:
        value_a
else:
    [type:name]
    key_1:
        value_2
    key_2:
        value_b
```

Comments

Comment in linker fragment files begin with `#`. Like in other languages, comment are used to provide helpful descriptions and documentation and are ignored during processing.

Types Sections

Sections fragments defines a list of object file sections that the GCC compiler emits. It may be a default section (e.g., `.text`, `.data`) or it may be user defined section through the `__attribute__` keyword.

The use of an optional `'+'` indicates the inclusion of the section in the list, as well as sections that start with it. This is the preferred method over listing both explicitly.

```
[sections:name]
entries:
    .section+
    .section
    ...
```

Example:

```
# Non-preferred
[sections:text]
entries:
    .text
    .text.*
```

(continues on next page)

(continued from previous page)

```
.literal
.literal.*

# Preferred, equivalent to the one above
[sections:text]
entries:
    .text+           # means .text and .text.*
    .literal+        # means .literal and .literal.*
```

Scheme

Scheme fragments define what target a sections fragment is assigned to.

```
[scheme:name]
entries:
    sections -> target
    sections -> target
    ...
```

Example:

```
[scheme:noflash]
entries:
    text -> iram0_text           # the entries under the sections fragment named_
↪text will go to iram0_text
    rodata -> dram0_data        # the entries under the sections fragment named_
↪rodata will go to dram0_data
```

The default scheme

There exists a special scheme with the name `default`. This scheme is special because catch-all placement rules are generated from its entries. This means that, if one of its entries is `text -> flash_text`, the placement rule will be generated for the target `flash_text`.

```
*(.literal .literal.* .text .text.*)
```

These catch-all rules then effectively serve as fallback rules for those whose mappings were not specified.

The `default` scheme is defined in [esp_system/app.lf](#). The `noflash` and `rtc` scheme fragments which are built-in schemes referenced in the quick start guide are also defined in this file.

Mapping

Mapping fragments define what scheme fragment to use for mappable entities, i.e., object files, function names, variable names, archives.

```
[mapping:name]
archive: archive           # output archive file name, as built (i.e., libxxx.
↪a)
entries:
    object:symbol (scheme) # symbol granularity
    object (scheme)        # object granularity
    * (scheme)             # archive granularity
```

There are three levels of placement granularity:

- **symbol:** The object file name and symbol name are specified. The symbol name can be a function name or a variable name.
- **object:** Only the object file name is specified.
- **archive:** `*` is specified, which is a short-hand for all the object files under the archive.

To know what an entry means, let us expand a sample object-granularity placement:


```
object (scheme)
```

Then expanding the scheme fragment from its entries definitions, we have:

```
object (sections -> target,  
        sections -> target,  
        ...)
```

Expanding the sections fragment with its entries definition:

```
object (.section,      # given this object file  
        .section,      # put its sections listed here at this  
        ... -> target, # target  
  
        .section,  
        .section,      # same should be done for these sections  
        ... -> target,  
  
        ...)           # and so on
```

Example:

```
[mapping:map]  
archive: libfreertos.a  
entries:  
    * (noflash)
```

Aside from the entity and scheme, flags can also be specified in an entry. The following flags are supported (note: <> = argument name, [] = optional):

1. ALIGN(<alignment>[, pre, post])
Align the placement by the amount specified in alignment. Generates
2. SORT([<sort_by_first>, <sort_by_second>])
Emits SORT_BY_NAME, SORT_BY_ALIGNMENT, SORT_BY_INIT_PRIORITY or SORT in the input section description.
Possible values for sort_by_first and sort_by_second are: name, alignment, init_priority.
If both sort_by_first and sort_by_second are not specified, the input sections are sorted by name. If both are specified, then the nested sorting follows the same rules discussed in <https://sourceware.org/binutils/docs/ld/Input-Section-Wildcards.html>.
3. KEEP()
Prevent the linker from discarding the placement by surrounding the input section description with KEEP command. See <https://sourceware.org/binutils/docs/ld/Input-Section-Keep.html> for more details.

4.SURROUND(<name>)

Generate symbols before and after the placement. The generated symbols follow the naming _<name>_start and _<name>_end. For example, if name == sym1,

When adding flags, the specific section -> target in the scheme needs to be specified. For multiple section -> target, use a comma as a separator. For example,

```
# Notes:  
# A. semicolon after entity-scheme  
# B. comma before section2 -> target2  
# C. section1 -> target1 and section2 -> target2 should be defined in entries of ↪  
↪scheme1  
entity1 (scheme1);  
    section1 -> target1 KEEP() ALIGN(4, pre, post),  
    section2 -> target2 SURROUND(sym) ALIGN(4, post) SORT()
```

Putting it all together, the following mapping fragment, for example,

```
[mapping:name]
archive: lib1.a
entries:
    obj1 (noflash);
        rodata -> dram0_data KEEP() SORT() ALIGN(8) SURROUND(my_sym)
```

generates an output on the linker script:

```
. = ALIGN(8)
_my_sym_start = ABSOLUTE(.)
KEEP(lib1.a:obj1.*( SORT(.rodata) SORT(.rodata.*) ))
_my_sym_end = ABSOLUTE(.)
```

Note that **ALIGN** and **SURROUND**, as mentioned in the flag descriptions, are order sensitive. Therefore, if for the same mapping fragment these two are switched, the following is generated instead:

```
_my_sym_start = ABSOLUTE(.)
. = ALIGN(8)
KEEP(lib1.a:obj1.*( SORT(.rodata) SORT(.rodata.*) ))
_my_sym_end = ABSOLUTE(.)
```

On Symbol-Granularity Placements Symbol granularity placements is possible due to compiler flags `-ffunction-sections` and `-ffdata-sections`. ESP-IDF compiles with these flags by default. If the user opts to remove these flags, then the symbol-granularity placements will not work. Furthermore, even with the presence of these flags, there are still other limitations to keep in mind due to the dependence on the compiler's emitted output sections.

For example, with `-ffunction-sections`, separate sections are emitted for each function; with section names predictably constructed i.e., `.text.{func_name}` and `.literal.{func_name}`. This is not the case for string literals within the function, as they go to pooled or generated section names.

With `-ffdata-sections`, for global scope data the compiler predictably emits either `.data.{var_name}`, `.rodata.{var_name}` or `.bss.{var_name}`; and so Type I mapping entry works for these. However, this is not the case for static data declared in function scope, as the generated section name is a result of mangling the variable name with some other information.

Linker Script Template

The linker script template is the skeleton in which the generated placement rules are put into. It is an otherwise ordinary linker script, with a specific marker syntax that indicates where the generated placement rules are placed.

To reference the placement rules collected under a `target` token, the following syntax is used:

```
mapping[target]
```

Example:

The example below is an excerpt from a possible linker script template. It defines an output section `.iram0.text`, and inside is a marker referencing the target `iram0_text`.

```
.iram0.text :
{
    /* Code marked as running out of IRAM */
    _iram_text_start = ABSOLUTE(.);

    /* Marker referencing iram0_text */
    mapping[iram0_text]
```

(continues on next page)

(continued from previous page)

```

_iram_text_end = ABSOLUTE(.);
} > iram0_0_seg

```

Suppose the generator collected the fragment definitions below:

```

[sections:text]
    .text+
    .literal+

[sections:iram]
    .iram1+

[scheme:default]
entries:
    text -> flash_text
    iram -> iram0_text

[scheme:noflash]
entries:
    text -> iram0_text

[mapping:freertos]
archive: libfreertos.a
entries:
    * (noflash)

```

Then the corresponding excerpt from the generated linker script will be as follows:

```

.iram0.text :
{
    /* Code marked as running out of IRAM */
    _iram_text_start = ABSOLUTE(.);

    /* Placement rules generated from the processed fragments, placed where the
    ↪marker was in the template */
    *(.iram1 .iram1.*)
    *libfreertos.a:(.literal .text .literal.* .text.*)

    _iram_text_end = ABSOLUTE(.);
} > iram0_0_seg

```

```
*libfreertos.a:(.literal .text .literal.* .text.*)
```

Rule generated from the entry `* (noflash)` of the `freertos` mapping fragment. All text sections of all object files under the archive `libfreertos.a` will be collected under the target `iram0_text` (as per the `noflash` scheme) and placed wherever in the template `iram0_text` is referenced by a marker.

```
*(.iram1 .iram1.*)
```

Rule generated from the default scheme entry `iram -> iram0_text`. Since the default scheme specifies an `iram -> iram0_text` entry, it too is placed wherever `iram0_text` is referenced by a marker. Since it is a rule generated from the default scheme, it comes first among all other rules collected under the same target name.

The linker script template currently used is [esp_system/ld/esp32p4/sections.ld.in](#); the generated output script `sections.ld` is put under its build directory.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.15 lwIP

ESP-IDF uses the open source [lwIP lightweight TCP/IP stack](#). The ESP-IDF version of lwIP ([esp-lwip](#)) has some modifications and additions compared to the upstream project.

4.15.1 Supported APIs

ESP-IDF supports the following lwIP TCP/IP stack functions:

- [BSD Sockets API](#)
- [Netconn API](#) is enabled but not officially supported for ESP-IDF applications

Adapted APIs

Warning: When using any lwIP API other than the [BSD Sockets API](#), please make sure that the API is thread-safe. To check if a given API call is thread-safe, enable the [CONFIG_LWIP_CHECK_THREAD_SAFETY](#) configuration option and run the application. This enables lwIP to assert the correct access of the TCP/IP core functionality. If the API is not accessed or locked properly from the appropriate [lwIP FreeRTOS Task](#), the execution will be aborted. The general recommendation is to use the [ESP-NETIF](#) component to interact with lwIP.

Some common lwIP app APIs are supported indirectly by ESP-IDF:

- Dynamic Host Configuration Protocol (DHCP) Server & Client are supported indirectly via the [ESP-NETIF](#) functionality.
- Domain Name System (DNS) is supported in lwIP; DNS servers could be assigned automatically when acquiring a DHCP address, or manually configured using the [ESP-NETIF](#) API.

Note: DNS server configuration in lwIP is global, not interface-specific. If you are using multiple network interfaces with distinct DNS servers, exercise caution to prevent inadvertent overwrites of one interface's DNS settings when acquiring a DHCP lease from another interface.

- Simple Network Time Protocol (SNTP) is also supported via the [ESP-NETIF](#), or directly via the [lwip/include/apps/esp_sntp.h](#) functions, which also provide thread-safe API to [lwip/lwip/src/include/lwip/apps/sntp.h](#) functions, see also [SNTP Time Synchronization](#).
- ICMP Ping is supported using a variation on the lwIP ping API, see [ICMP Echo](#).
- ICMPv6 Ping, supported by lwIP's ICMPv6 Echo API, is used to test IPv6 network connectivity. For more information, see [protocols/sockets/icmpv6_ping](#).
- NetBIOS lookup is available using the standard lwIP API. [protocols/http_server/restful_server](#) has the option to demonstrate using NetBIOS to look up a host on the LAN.
- mDNS uses a different implementation to the lwIP default mDNS, see [mDNS Service](#). But lwIP can look up mDNS hosts using standard APIs such as `gethostbyname()` and the convention `hostname.local`, provided the [CONFIG_LWIP_DNS_SUPPORT_MDNS_QUERIES](#) setting is enabled.
- The PPP implementation in lwIP can be used to create PPPoS (PPP over serial) interface in ESP-IDF. Please refer to the documentation of the [ESP-NETIF](#) component to create and configure a PPP network interface, by means of the `ESP_NETIF_DEFAULT_PPP()` macro defined in [esp_netif/include/esp_netif_defaults.h](#). Additional runtime settings are provided via [esp_netif/include/esp_netif_ppp.h](#). PPPoS interfaces are typically used to interact with NBIoT/GSM/LTE modems. More application-level friendly API is supported by the [esp_modem](#) library, which uses this PPP lwIP module behind the scenes.

4.15.2 BSD Sockets API

The BSD Sockets API is a common cross-platform TCP/IP sockets API that originated in the Berkeley Standard Distribution of UNIX but is now standardized in a section of the POSIX specification. BSD Sockets are sometimes called POSIX Sockets or Berkeley Sockets.

As implemented in ESP-IDF, lwIP supports all of the common usages of the BSD Sockets API.

References

A wide range of BSD Sockets reference materials are available, including:

- [Single UNIX Specification - BSD Sockets page](#)
- [Berkeley Sockets - Wikipedia page](#)

Examples

A number of ESP-IDF examples show how to use the BSD Sockets APIs:

- [protocols/sockets/tcp_server](#)
- [protocols/sockets/tcp_client](#)
- [protocols/sockets/udp_server](#)
- [protocols/sockets/udp_client](#)
- [protocols/sockets/udp_multicast](#)
- [protocols/http_request](#): this simplified example uses a TCP socket to send an HTTP request, but [ESP HTTP Client](#) is a much better option for sending HTTP requests

Supported Functions

The following BSD socket API functions are supported. For full details, see [lwip/lwip/src/include/lwip/sockets.h](#).

- `socket()`
- `bind()`
- `accept()`
- `shutdown()`
- `getpeername()`
- `getsockopt()` & `setsockopt()`: see [Socket Options](#)
- `close()`: via [Virtual Filesystem Component](#)
- `read()`, `readv()`, `write()`, `writew()`: via [Virtual Filesystem Component](#)
- `recv()`, `recvmsg()`, `recvfrom()`
- `send()`, `sendmsg()`, `sendto()`
- `select()`: via [Virtual Filesystem Component](#)
- `poll()`: on ESP-IDF, `poll()` is implemented by calling `select()` internally, so using `select()` directly is recommended, if a choice of methods is available
- `fcntl()`: see [fcntl\(\)](#)

Non-standard functions:

- `ioctl()`: see [ioctl\(\)](#)

Note: Some lwIP application sample code uses prefixed versions of BSD APIs, e.g., `lwip_socket()`, instead of the standard `socket()`. Both forms can be used with ESP-IDF, but using standard names is recommended.

Socket Error Handling

BSD Socket error handling code is very important for robust socket applications. Normally, socket error handling involves the following aspects:

- Detecting the error
- Getting the error reason code
- Handling the error according to the reason code

In lwIP, we have two different scenarios for handling socket errors:

- Socket API returns an error. For more information, see [Socket API Errors](#).
- `select(int maxfdp1, fd_set *readset, fd_set *writeset, fd_set *exceptset, struct timeval *timeout)` has an exception descriptor indicating that the socket has an error. For more information, see [select\(\) Errors](#).

Socket API Errors Error detection

- We can know that the socket API fails according to its return value.

Get the error reason code

- When socket API fails, the return value does not contain the failure reason and the application can get the error reason code by accessing `errno`. Different values indicate different meanings. For more information, see [Socket Error Reason Code](#).

Example:

```
int err;
int sockfd;

if (sockfd = socket(AF_INET, SOCK_STREAM, 0) < 0) {
    // the error code is obtained from errno
    err = errno;
    return err;
}
```

select() Errors Error detection

- Socket error when `select()` has exception descriptor.

Get the error reason code

- If the `select()` indicates that the socket fails, we can not get the error reason code by accessing `errno`, instead we should call `getsockopt()` to get the failure reason code. Since `select()` has exception descriptor, the error code is not given to `errno`.

Note: The `getsockopt()` function has the following prototype: `int getsockopt(int s, int level, int optname, void *optval, socklen_t *optlen)`. Its purpose is to get the current value of the option of any type, any state socket, and store the result in `optval`. For example, when you get the error code on a socket, you can get it by `getsockopt(sockfd, SOL_SOCKET, SO_ERROR, &err, &optlen)`.

Example:

```
int err;

if (select(sockfd + 1, NULL, NULL, &exfds, &tval) <= 0) {
    err = errno;
    return err;
} else {
    if (FD_ISSET(sockfd, &exfds)) {
        // select() exception set using getsockopt()
        int optlen = sizeof(int);
        getsockopt(sockfd, SOL_SOCKET, SO_ERROR, &err, &optlen);
        return err;
    }
}
```

(continues on next page)

(continued from previous page)

```

    }
}

```

Socket Error Reason Code Below is a list of common error codes. For a more detailed list of standard POSIX/C error codes, please see [newlib errno.h](#) and the platform-specific extensions [newlib/platform_include/errno.h](#).

Error code	Description
ECONNREFUSED	Connection refused
EADDRINUSE	Address already in use
ECONNABORTED	Software caused connection abort
ENETUNREACH	Network is unreachable
ENETDOWN	Network interface is not configured
ETIMEDOUT	Connection timed out
EHOSTDOWN	Host is down
EHOSTUNREACH	Host is unreachable
EINPROGRESS	Connection already in progress
EALREADY	Socket already connected
EDESTADDRREQ	Destination address required
EPROTONOSUPPORT	Unknown protocol

Socket Options

The `getsockopt()` and `setsockopt()` functions allow getting and setting per-socket options respectively.

Not all standard socket options are supported by lwIP in ESP-IDF. The following socket options are supported:

Common Options Used with level argument `SOL_SOCKET`.

- `SO_REUSEADDR`: available if [CONFIG_LWIP_SO_REUSE](#) is set, whose behavior can be customized by setting [CONFIG_LWIP_SO_REUSE_RXTOALL](#)
- `SO_KEEPALIVE`
- `SO_BROADCAST`
- `SO_ACCEPTCONN`
- `SO_RCVBUF`: available if [CONFIG_LWIP_SO_RCVBUF](#) is set
- `SO_SNDTIMEO` / `SO_RCVTIMEO`
- `SO_ERROR`: only used with `select()`, see [Socket Error Handling](#)
- `SO_TYPE`
- `SO_NO_CHECK`: for UDP sockets only

IP Options Used with level argument `IPPROTO_IP`.

- `IP_TOS`
- `IP_TTL`
- `IP_PKTINFO`: available if [CONFIG_LWIP_NETBUF_RECVINFO](#) is set

For multicast UDP sockets:

- `IP_MULTICAST_IF`
- `IP_MULTICAST_LOOP`
- `IP_MULTICAST_TTL`
- `IP_ADD_MEMBERSHIP`
- `IP_DROP_MEMBERSHIP`

TCP Options TCP sockets only. Used with level argument `IPPROTO_TCP`.

- `TCP_NODELAY`

Options relating to TCP keepalive probes:

- `TCP_KEEPAIVE`: int value, TCP keepalive period in milliseconds
- `TCP_KEEPIIDLE`: same as `TCP_KEEPAIVE`, but the value is in seconds
- `TCP_KEEPIINTVL`: int value, the interval between keepalive probes in seconds
- `TCP_KEEPCNT`: int value, number of keepalive probes before timing out

IPv6 Options IPv6 sockets only. Used with level argument `IPPROTO_IPV6`.

- `IPV6_CHECKSUM`
- `IPV6_V6ONLY`

For multicast IPv6 UDP sockets:

- `IPV6_JOIN_GROUP` / `IPV6_ADD_MEMBERSHIP`
- `IPV6_LEAVE_GROUP` / `IPV6_DROP_MEMBERSHIP`
- `IPV6_MULTICAST_IF`
- `IPV6_MULTICAST_HOPS`
- `IPV6_MULTICAST_LOOP`

fcntl()

The `fcntl()` function is a standard API for manipulating options related to a file descriptor. In ESP-IDF, the *Virtual Filesystem Component* layer is used to implement this function.

When the file descriptor is a socket, only the following `fcntl()` values are supported:

- `O_NONBLOCK` to set or clear non-blocking I/O mode. Also supports `O_NDELAY`, which is identical to `O_NONBLOCK`.
- `O_RDONLY`, `O_WRONLY`, `O_RDWR` flags for different read or write modes. These flags can only be read using `F_GETFL`, and cannot be set using `F_SETFL`. A TCP socket returns a different mode depending on whether the connection has been closed at either end or is still open at both ends. UDP sockets always return `O_RDWR`.

ioctl()

The `ioctl()` function provides a semi-standard way to access some internal features of the TCP/IP stack. In ESP-IDF, the *Virtual Filesystem Component* layer is used to implement this function.

When the file descriptor is a socket, only the following `ioctl()` values are supported:

- `FIONREAD` returns the number of bytes of the pending data already received in the socket's network buffer.
- `FIONBIO` is an alternative way to set/clear non-blocking I/O status for a socket, equivalent to `fcntl(fd, F_SETFL, O_NONBLOCK, ...)`.

4.15.3 Netconn API

lwIP supports two lower-level APIs as well as the BSD Sockets API: the Netconn API and the Raw API.

The lwIP Raw API is designed for single-threaded devices and is not supported in ESP-IDF.

The Netconn API is used to implement the BSD Sockets API inside lwIP, and it can also be called directly from ESP-IDF apps. This API has lower resource usage than the BSD Sockets API. In particular, it can send and receive data without firstly copying it into internal lwIP buffers.

Important: Espressif does not test the Netconn API in ESP-IDF. As such, this functionality is **enabled but not supported**. Some functionality may only work correctly when used from the BSD Sockets API.

For more information about the Netconn API, consult [lwip/lwip/src/include/lwip/api.h](#) and [part of the ****unofficial**** lwIP Application Developers Manual](#).

4.15.4 lwIP FreeRTOS Task

lwIP creates a dedicated TCP/IP FreeRTOS task to handle socket API requests from other tasks.

A number of configuration items are available to modify the task and the queues (mailboxes) used to send data to/from the TCP/IP task:

- [CONFIG_LWIP_TCPIP_RECVMBOX_SIZE](#)
- [CONFIG_LWIP_TCPIP_TASK_STACK_SIZE](#)
- [CONFIG_LWIP_TCPIP_TASK_AFFINITY](#)

4.15.5 IPv6 Support

Both IPv4 and IPv6 are supported in a dual-stack configuration and are enabled by default. Both IPv6 and IPv4 may be disabled separately if they are not needed, see [Minimum RAM Usage](#).

IPv6 support is limited to **Stateless Autoconfiguration** only. **Stateful configuration** is not supported in ESP-IDF, nor in upstream lwIP.

IPv6 Address configuration is defined by means of these protocols or services:

- **SLAAC** IPv6 Stateless Address Autoconfiguration (RFC-2462)
- **DHCPv6** Dynamic Host Configuration Protocol for IPv6 (RFC-8415)

None of these two types of address configuration is enabled by default, so the device uses only Link Local addresses or statically-defined addresses.

Stateless Autoconfiguration Process

To enable address autoconfiguration using the Router Advertisement protocol, please enable:

- [CONFIG_LWIP_IPV6_AUTOCONFIG](#)

This configuration option enables IPv6 autoconfiguration for all network interfaces, which differs from the upstream lwIP behavior, where the autoconfiguration needs to be explicitly enabled for each `netif` with `netif->ip6_autoconfig_enabled=1`.

DHCPv6

DHCPv6 in lwIP is very simple and supports only stateless configuration. It could be enabled using:

- [CONFIG_LWIP_IPV6_DHCP6](#)

Since the DHCPv6 works only in its stateless configuration, the [Stateless Autoconfiguration Process](#) has to be enabled as well via [CONFIG_LWIP_IPV6_AUTOCONFIG](#).

Moreover, the DHCPv6 needs to be explicitly enabled from the application code using:

```
dhcp6_enable_stateless(netif);
```

DNS Servers in IPv6 Autoconfiguration

In order to autoconfigure DNS server(s), especially in IPv6-only networks, we have these two options:

- Recursive Domain Name System (DNS): this belongs to the Neighbor Discovery Protocol (NDP) and uses *Stateless Autoconfiguration Process*. The number of servers must be set `CONFIG_LWIP_IPV6_RDNSS_MAX_DNS_SERVERS`, this option is disabled by default, i.e., set to 0.
- DHCPv6 stateless configuration, uses *DHCPv6* to configure DNS servers. Note that this configuration assumes IPv6 Router Advertisement Flags (RFC-5175) to be set to
 - Managed Address Configuration Flag = 0
 - Other Configuration Flag = 1

4.15.6 ESP-lwIP Custom Modifications

Additions

The following code is added, which is not present in the upstream lwIP release:

Thread-Safe Sockets It is possible to `close()` a socket from a different thread than the one that created it. The `close()` call blocks, until any function calls currently using that socket from other tasks have returned.

It is, however, not possible to delete a task while it is actively waiting on `select()` or `poll()` APIs. It is always necessary that these APIs exit before destroying the task, as this might corrupt internal structures and cause subsequent crashes of the lwIP. These APIs allocate globally referenced callback pointers on the stack so that when the task gets destroyed before unrolling the stack, the lwIP could still hold pointers to the deleted stack.

On-Demand Timers lwIP IGMP and MLD6 feature both initialize a timer in order to trigger timeout events at certain times.

The default lwIP implementation is to have these timers enabled all the time, even if no timeout events are active. This increases CPU usage and power consumption when using automatic Light-sleep mode. ESP-lwIP default behavior is to set each timer on demand, so it is only enabled when an event is pending.

To return to the default lwIP behavior, which is always-on timers, disable `CONFIG_LWIP_TIMERS_ONDEMAND`.

lwIP Timers API When not using Wi-Fi, the lwIP timer can be turned off via the API to reduce power consumption.

The following API functions are supported. For full details, see `lwip/lwip/src/include/lwip/timeouts.h`.

- `sys_timeouts_init()`
- `sys_timeouts_deinit()`

Additional Socket Options

- Some standard IPV4 and IPV6 multicast socket options are implemented, see *Socket Options*.
- Possible to set IPV6-only UDP and TCP sockets with `IPV6_V6ONLY` socket option, while normal lwIP is TCP-only.

IP Layer Features

- IPV4-source-based routing implementation is different
- IPV4-mapped IPV6 addresses are supported

Customized lwIP Hooks The original lwIP supports implementing custom compile-time modifications via `LWIP_HOOK_FILENAME`. This file is already used by the ESP-IDF port layer, but ESP-IDF users could still include and implement any custom additions via a header file defined by the macro `ESP_IDF_LWIP_HOOK_FILENAME`. Here is an example of adding a custom hook file to the build process, and the hook is called `my_hook.h`, located in the project's `main` folder:

```
idf_component_get_property(lwip lwip COMPONENT_LIB)
target_compile_options(${lwip} PRIVATE "-I${PROJECT_DIR}/main")
target_compile_definitions(${lwip} PRIVATE "-DESP_IDF_LWIP_HOOK_FILENAME=\"my_hook.
↪h\"")
```

Customized lwIP Options From ESP-IDF Build System The most common lwIP options are configurable through the component configuration menu. However, certain definitions need to be injected from the command line. The CMake function `target_compile_definitions()` can be employed to define macros, as illustrated below:

```
idf_component_get_property(lwip lwip COMPONENT_LIB)
target_compile_definitions(${lwip} PRIVATE "-DETHARP_SUPPORT_VLAN=1")
```

This approach may not work for function-like macros, as there is no guarantee that the definition will be accepted by all compilers, although it is supported in GCC. To address this limitation, the `add_definitions()` function can be utilized to define the macro for the entire project, for example: `add_definitions("-DFALLBACK_DNS_SERVER_ADDRESS(addr)=\"IP_ADDR4((addr), 8, 8, 8, 8)\")`.

Alternatively, you can define your function-like macro in a header file which will be pre-included as an lwIP hook file, see [Customized lwIP Hooks](#).

Limitations

ESP-IDF additions to lwIP still suffer from the global DNS limitation, described in [Adapted APIs](#). To address this limitation from application code, the `FALLBACK_DNS_SERVER_ADDRESS()` macro can be utilized to define a global DNS fallback server accessible from all interfaces. Alternatively, you have the option to maintain per-interface DNS servers and reconfigure them whenever the default interface changes.

The number of IP addresses returned by network database APIs such as `getaddrinfo()` and `gethostbyname()` is restricted by the macro `DNS_MAX_HOST_IP`. By default, the value of this macro is set to 1.

In the implementation of `getaddrinfo()`, the canonical name is not available. Therefore, the `ai_canonname` field of the first returned `addrinfo` structure will always refer to the `nodename` argument or a string with the same contents.

Calling `send()` or `sendto()` repeatedly on a UDP socket may eventually fail with `errno` equal to `ENOMEM`. This failure occurs due to the limitations of buffer sizes in the lower-layer network interface drivers. If all driver transmit buffers are full, the UDP transmission will fail. For applications that transmit a high volume of UDP datagrams and aim to avoid any dropped datagrams by the sender, it is advisable to implement error code checking and employ a retransmission mechanism with a short delay.

4.15.7 Performance Optimization

TCP/IP performance is a complex subject, and performance can be optimized toward multiple goals. The default settings of ESP-IDF are tuned for a compromise between throughput, latency, and moderate memory usage.

Maximum Throughput

Espressif tests ESP-IDF TCP/IP throughput using the [wifi/iperf](#) example in an RF-sealed enclosure.

The [wifi/iperf/sdkconfig.defaults](#) file for the iperf example contains settings known to maximize TCP/IP throughput, usually at the expense of higher RAM usage. To get maximum TCP/IP throughput in an application at the expense of other factors, it is suggested to apply settings from this file into the project `sdkconfig`.

Important: Suggest applying changes a few at a time and checking the performance each time with a particular application workload.

- If a lot of tasks are competing for CPU time on the system, consider that the lwIP task has configurable CPU affinity ([CONFIG_LWIP_TCPIP_TASK_AFFINITY](#)) and runs at fixed priority (18, `ESP_TASK_TCPIP_PRIO`). To optimize CPU utilization, consider assigning competing tasks to different cores or adjusting their priorities to lower values. For additional details on built-in task priorities, please refer to [Built-in Task Priorities](#).
- If using `select()` function with socket arguments only, disabling [CONFIG_VFS_SUPPORT_SELECT](#) will make `select()` calls faster.
- If there is enough free IRAM, select [CONFIG_LWIP_IRAM_OPTIMIZATION](#) and [CONFIG_LWIP_EXTRA_IRAM_OPTIMIZATION](#) to improve TX/RX throughput.

Minimum Latency

Except for increasing buffer sizes, most changes that increase throughput also decrease latency by reducing the amount of CPU time spent in lwIP functions.

- For TCP sockets, lwIP supports setting the standard `TCP_NODELAY` flag to disable Nagle's algorithm.

Minimum RAM Usage

Most lwIP RAM usage is on-demand, as RAM is allocated from the heap as needed. Therefore, changing lwIP settings to reduce RAM usage may not change RAM usage at idle, but can change it at peak.

- Reducing [CONFIG_LWIP_MAX_SOCKETS](#) reduces the maximum number of sockets in the system. This also causes TCP sockets in the `WAIT_CLOSE` state to be closed and recycled more rapidly when needed to open a new socket, further reducing peak RAM usage.
- Reducing [CONFIG_LWIP_TCPIP_RECVMBOX_SIZE](#), [CONFIG_LWIP_TCP_RECVMBOX_SIZE](#) and [CONFIG_LWIP_UDP_RECVMBOX_SIZE](#) reduce RAM usage at the expense of throughput, depending on usage.
- Reducing [CONFIG_LWIP_TCP_MSL](#) and [CONFIG_LWIP_TCP_FIN_WAIT_TIMEOUT](#) reduces the maximum segment lifetime in the system. This also causes TCP sockets in the `TIME_WAIT` and `FIN_WAIT_2` states to be closed and recycled more rapidly.
- Disabling [CONFIG_LWIP_IPV6](#) can save about 39 KB for firmware size and 2 KB RAM when the system is powered up and 7 KB RAM when the TCP/IP stack is running. If there is no requirement for supporting IPv6, it can be disabled to save flash and RAM footprint.
- Disabling [CONFIG_LWIP_IPV4](#) can save about 26 KB of firmware size and 600 B RAM on power up and 6 KB RAM when the TCP/IP stack is running. If the local network supports IPv6-only configuration, IPv4 can be disabled to save flash and RAM footprint.

Peak Buffer Usage The peak heap memory that lwIP consumes is the **theoretically-maximum memory** that the lwIP driver consumes. Generally, the peak heap memory that lwIP consumes depends on:

- the memory required to create a UDP connection: `lwip_udp_conn`
- the memory required to create a TCP connection: `lwip_tcp_conn`
- the number of UDP connections that the application has: `lwip_udp_con_num`
- the number of TCP connections that the application has: `lwip_tcp_con_num`
- the TCP TX window size: `lwip_tcp_tx_win_size`
- the TCP RX window size: `lwip_tcp_rx_win_size`

So, the peak heap memory that the lwIP consumes can be calculated with the following formula:

$$\text{lwip\_dynamic\_peek\_memory} = (\text{lwip\_udp\_con\_num} * \text{lwip\_udp\_conn}) + (\text{lwip\_tcp\_con\_num} * (\text{lwip\_tcp\_tx\_win\_size} + \text{lwip\_tcp\_rx\_win\_size} + \text{lwip\_tcp\_conn}))$$

Some TCP-based applications need only one TCP connection. However, they may choose to close this TCP connection and create a new one when an error occurs (e.g., a sending failure). This may result in multiple TCP connections

existing in the system simultaneously, because it may take a long time for a TCP connection to close, according to the TCP state machine, refer to RFC793.

4.16 Memory Types

ESP32-P4 chip has multiple memory types and flexible memory mapping features. This section describes how ESP-IDF uses these features by default.

ESP-IDF distinguishes between instruction memory bus (IRAM, IROM, RTC FAST memory) and data memory bus (DRAM, DROM). Instruction memory is executable, and can only be read or written via 4-byte aligned words. Data memory is not executable and can be accessed via individual byte operations. For more information about the different memory buses consult the *ESP32-P4 Technical Reference Manual > System and Memory* [PDF].

4.16.1 DRAM (Data RAM)

Non-constant static data (.data) and zero-initialized data (.bss) is placed by the linker into Internal SRAM as data memory. The remaining space in this region is used for the runtime heap.

By applying the `EXT_RAM_BSS_ATTR` macro, zero-initialized data can also be placed into external RAM. To use this macro, the `CONFIG_SPIRAM_ALLOW_BSS_SEG_EXTERNAL_MEMORY` needs to be enabled. See *Allow .bss Segment to Be Placed in External Memory*.

Note: The maximum statically allocated DRAM size is reduced by the *IRAM (Instruction RAM)* size of the compiled application. The available heap memory at runtime is reduced by the total static IRAM and DRAM usage of the application.

Constant data may also be placed into DRAM, for example if it is used in an non-flash-safe ISR (see explanation under *How to Place Code in IRAM*).

"noinit" DRAM

The macro `__NOINIT_ATTR` can be used as attribute to place data into `.noinit` section. The values placed into this section will not be initialized at startup and should keep its value after software restart.

By applying the `EXT_RAM_NOINIT_ATTR` macro, non-initialized value could also be placed in external RAM. To do this, the `CONFIG_SPIRAM_ALLOW_NOINIT_SEG_EXTERNAL_MEMORY` needs to be enabled. See *Allow .noinit Segment to Be Placed in External Memory*. If the `CONFIG_SPIRAM_ALLOW_NOINIT_SEG_EXTERNAL_MEMORY` is not enabled, `EXT_RAM_NOINIT_ATTR` will behave just as `__NOINIT_ATTR`, it will make data to be placed into `.noinit` segment in internal RAM.

Example:

```
__NOINIT_ATTR uint32_t noinit_data;
```

4.16.2 IRAM (Instruction RAM)

Note: Any internal SRAM which is not used for Instruction RAM will be made available as *DRAM (Data RAM)* for static data and dynamic allocation (heap).

When to Place Code in IRAM

Cases when parts of the application should be placed into IRAM:

- Interrupt handlers must be placed into IRAM if `ESP_INTR_FLAG_IRAM` is used when registering the interrupt handler. For more information, see [IRAM-Safe Interrupt Handlers](#).
- Some timing critical code may be placed into IRAM to reduce the penalty associated with loading the code from flash. ESP32-P4 reads code and data from flash via the MMU cache. In some cases, placing a function into IRAM may reduce delays caused by a cache miss and significantly improve that function's performance.

How to Place Code in IRAM

Some code is automatically placed into the IRAM region using the linker script.

If some specific application code needs to be placed into IRAM, it can be done by using the [Linker Script Generation](#) feature and adding a linker script fragment file to your component that targets at the entire source files or functions with the `noflash` placement. See the [Linker Script Generation](#) docs for more information.

Alternatively, it is possible to specify IRAM placement in the source code using the `IRAM_ATTR` macro:

```
#include "esp_attr.h"

void IRAM_ATTR gpio_isr_handler(void* arg)
{
    // ...
}
```

There are some possible issues with placement in IRAM, that may cause problems with IRAM-safe interrupt handlers:

- Strings or constants inside an `IRAM_ATTR` function may not be placed in RAM automatically. It is possible to use `DRAM_ATTR` attributes to mark these, or using the linker script method will cause these to be automatically placed correctly.

```
void IRAM_ATTR gpio_isr_handler(void* arg)
{
    const static DRAM_ATTR uint8_t INDEX_DATA[] = { 45, 33, 12, 0 };
    const static char *MSG = DRAM_STR("I am a string stored in RAM");
}
```

Note that knowing which data should be marked with `DRAM_ATTR` can be hard, the compiler will sometimes recognize that a variable or expression is constant (even if it is not marked `const`) and optimize it into flash, unless it is marked with `DRAM_ATTR`.

- GCC optimizations that automatically generate jump tables or switch/case lookup tables place these tables in flash. IDF by default builds all files with `-fno-jump-tables -fno-tree-switch-conversion` flags to avoid this.

Jump table optimizations can be re-enabled for individual source files that do not need to be placed in IRAM. For instructions on how to add the `-fno-jump-tables -fno-tree-switch-conversion` options when compiling individual source files, see [Controlling Component Compilation](#).

4.16.3 IROM (Code Executed from flash)

If a function is not explicitly placed into [IRAM \(Instruction RAM\)](#) or RTC memory, it is placed into flash. As IRAM is limited, most of an application's binary code must be placed into IROM instead.

During [Application Startup Flow](#), the bootloader (which runs from IRAM) configures the MMU flash cache to map the app's instruction code region to the instruction space. Flash accessed via the MMU is cached using some internal SRAM and accessing cached flash data is as fast as accessing other types of internal memory.

4.16.4 DROM (Data Stored in flash)

By default, constant data is placed by the linker into a region mapped to the MMU flash cache. This is the same as the *IROM (Code Executed from flash)* section, but is for read-only data not executable code.

The only constant data not placed into this memory type by default are literal constants which are embedded by the compiler into application code. These are placed as the surrounding function's executable instructions.

The `DRAM_ATTR` attribute can be used to force constants from DROM into the *DRAM (Data RAM)* section (see above).

4.16.5 RTC FAST Memory

The same region of RTC FAST memory can be accessed as both instruction and data memory. Code which has to run after wake-up from deep sleep mode has to be placed into RTC memory. Please check detailed description in *deep sleep* documentation.

Remaining RTC FAST memory is added to the heap unless the option *CONFIG_ESP_SYSTEM_ALLOW_RTC_FAST_MEM_AS_HEAP* is disabled. This memory can be used interchangeably with *DRAM (Data RAM)*, but is slightly slower to access.

4.16.6 TCM (Tightly-Coupled Memory)

TCM is memory placed near the CPU, accessible at CPU frequency without passing through a cache. Even though on average, it may not surpass the efficiency or speed of cached memory, it does provide predictable and consistent access times. TCM can be useful for time-critical routines where having a deterministic access speed is important.

4.16.7 DMA-Capable Requirement

Most peripheral DMA controllers (e.g., SPI, sdmmc, etc.) have requirements that sending/receiving buffers should be placed in DRAM and word-aligned. We suggest to place DMA buffers in static variables rather than in the stack. Use macro `DMA_ATTR` to declare global/local static variables like:

```
DMA_ATTR uint8_t buffer[]="I want to send something";

void app_main()
{
    // initialization code...
    spi_transaction_t temp = {
        .tx_buffer = buffer,
        .length = 8 * sizeof(buffer),
    };
    spi_device_transmit(spi, &temp);
    // other stuff
}
```

Or:

```
void app_main()
{
    DMA_ATTR static uint8_t buffer[] = "I want to send something";
    // initialization code...
    spi_transaction_t temp = {
        .tx_buffer = buffer,
        .length = 8 * sizeof(buffer),
    };
    spi_device_transmit(spi, &temp);
    // other stuff
}
```


It is also possible to allocate DMA-capable memory buffers dynamically by using the [MALLOC_CAP_DMA](#) capabilities flag.

4.16.8 DMA Buffer in the Stack

Placing DMA buffers in the stack is possible but discouraged. If doing so, pay attention to the following:

- Placing DRAM buffers on the stack is not recommended if the stack may be in PSRAM. If the stack of a task is placed in the PSRAM, several steps have to be taken as described in [Support for External RAM](#).
- Use macro `WORD_ALIGNED_ATTR` in functions before variables to place them in proper positions like:

```
void app_main()
{
    uint8_t stuff;
    WORD_ALIGNED_ATTR uint8_t buffer[] = "I want to send something";    //or_
    ↪the buffer will be placed right after stuff.
    // initialization code...
    spi_transaction_t temp = {
        .tx_buffer = buffer,
        .length = 8 * sizeof(buffer),
    };
    spi_device_transmit(spi, &temp);
    // other stuff
}
```

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.17 OpenThread

OpenThread is an IP stack running on the 802.15.4 MAC layer which features mesh network and low power consumption.

4.17.1 Modes of the OpenThread Stack

OpenThread can run under the following modes on Espressif chips:

Standalone Node

The full OpenThread stack and the application layer run on the same chip. This mode is available on chips with 15.4 radio such as ESP32-H2 and ESP32-C6.

Radio Co-Processor (RCP)

The chip is connected to another host running the OpenThread IP stack. It sends and receives 15.4 packets on behalf of the host. This mode is available on chips with 15.4 radio such as ESP32-H2 and ESP32-C6. The underlying transport between the chip and the host can be SPI or UART. For the sake of latency, we recommend using SPI as the underlying transport.

OpenThread Host

For chips without a 15.4 radio, it can be connected to an RCP and run OpenThread under host mode. This mode enables OpenThread on Wi-Fi chips such as ESP32, ESP32-S2, ESP32-S3, and ESP32-C3. The following diagram shows how devices work under different modes:




Fig. 27: OpenThread device modes

4.17.2 How to Write an OpenThread Application

The OpenThread `openthread/ot_cli` example is a good place to start at. It demonstrates basic OpenThread initialization and simple socket-based server and client.

Before OpenThread Initialization

- s1.1: The main task calls `esp_vfs_eventfd_register()` to initialize the eventfd virtual file system. The eventfd file system is used for task notification in the OpenThread driver.
- s1.2: The main task calls `nvs_flash_init()` to initialize the NVS where the Thread network data is stored.
- s1.3: **Optional.** The main task calls `esp_netif_init()` only when it wants to create the network interface for Thread.
- s1.4: The main task calls `esp_event_loop_create()` to create the system Event task and initialize an application event's callback function.

OpenThread Stack Initialization

- s2.1: Call `esp_openthread_init()` to initialize the OpenThread stack.

OpenThread Network Interface Initialization

The whole stage is **optional** and only required if the application wants to create the network interface for Thread.

- s3.1: Call `esp_netif_new()` with `ESP_NETIF_DEFAULT_OPENTHREAD` to create the interface.
- s3.2: Call `esp_openthread_netif_glue_init()` to create the OpenThread interface handlers.
- s3.3: Call `esp_netif_attach()` to attach the handlers to the interface.

The OpenThread Main Loop

- s4.3: Call `esp_openthread_launch_mainloop()` to launch the OpenThread main loop. Note that this is a busy loop and does not return until the OpenThread stack is terminated.

Calling OpenThread APIs

The OpenThread APIs are not thread-safe. When calling OpenThread APIs from other tasks, make sure to hold the lock with `esp_openthread_lock_acquire()` and release the lock with `esp_openthread_lock_release()` afterwards.

Deinitialization

The following steps are required to deinitialize the OpenThread stack:

- Call `esp_netif_destroy()` and `esp_openthread_netif_glue_deinit()` to deinitialize the OpenThread network interface if you have created one.
- Call `esp_openthread_deinit()` to deinitialize the OpenThread stack.

4.17.3 OpenThread Macro Definitions

In the OpenThread protocol stack, defining macros to enable features and configure parameters is a common practice. Users can define macro values to enable or disable specific features and adjust parameters. ESP provides the following methods for defining OpenThread macros:

- Using configuration menu (menuconfig): Some macros are mapped to Kconfig files and can be configured through `idf.py menuconfig` → Component config → OpenThread. This allows enabling or disabling features and setting related parameters.
- Using user-defined header files: Users can create a custom header file and enable it via `idf.py menuconfig` → Component config → OpenThread → Thread Extended Features → Use a header file defined by customer. The priority of the custom header file is second only to the `menuconfig`.
- Using `openthread-core-esp32x-xxx-config.h` for configuration: Some macros have default values set in the OpenThread private header files. These cannot currently be modified through the `menuconfig`, but can be modified via user-defined header files.
- Using OpenThread stack default configurations: Other macros are assigned default values when defined in the OpenThread stack.

Note: The priority of the above configuration methods, from highest to lowest, is as follows: Configuration Menu → User-defined Header File → `openthread-core-esp32x-xxx-config.h` → OpenThread Stack Default Configuration

4.17.4 The OpenThread Border Router

The OpenThread border router connects the Thread network with other IP networks. It provides IPv6 connectivity, service registration, and commission functionality.

To launch an OpenThread border router on an ESP chip, you need to connect an RCP to a Wi-Fi capable chip such as ESP32.

Calling `esp_openthread_border_router_init()` during the initialization launches all the border routing functionalities.

You may refer to the `openthread/ot_br` example and the README for further border router details.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.18 Partition Tables

4.18.1 Overview

A single ESP32-P4's flash can contain multiple apps, as well as many different kinds of data (calibration data, filesystems, parameter storage, etc). For this reason a partition table is flashed to (*default offset*) 0x8000 in the flash.

The partition table length is 0xC00 bytes, as we allow a maximum of 95 entries. An MD5 checksum, used for checking the integrity of the partition table at runtime, is appended after the table data. Thus, the partition table occupies an entire flash sector, which size is 0x1000 (4 KB). As a result, any partition following it must be at least located at (*default offset*) + 0x1000.

Each entry in the partition table has a name (label), type (app, data, or something else), subtype and the offset in flash where the partition is loaded.

The simplest way to use the partition table is to open the project configuration menu (`idf.py menuconfig`) and choose one of the simple predefined partition tables under `CONFIG_PARTITION_TABLE_TYPE`:

- "Single factory app, no OTA"
- "Factory app, two OTA definitions"

In both cases the factory app is flashed at offset 0x10000. If you execute `idf.py partition-table` then it will print a summary of the partition table.

4.18.2 Built-in Partition Tables

Here is the summary printed for the "Single factory app, no OTA" configuration:

```
# ESP-IDF Partition Table
# Name,   Type, SubType, Offset,   Size, Flags
nvs,      data, nvs,      0x9000,  0x6000,
phy_init, data, phy,      0xf000,  0x1000,
factory,  app,  factory, 0x10000, 1M,
```

- At a 0x10000 (64 KB) offset in the flash is the app labelled "factory". The bootloader will run this app by default.
- There are also two data regions defined in the partition table for storing NVS library partition and PHY init data.

Here is the summary printed for the "Factory app, two OTA definitions" configuration:

```
# ESP-IDF Partition Table
# Name,   Type, SubType, Offset,   Size, Flags
nvs,      data, nvs,      0x9000,  0x4000,
otadata,  data, ota,       0xd000,  0x2000,
phy_init, data, phy,      0xf000,  0x1000,
factory,  app,  factory, 0x10000, 1M,
ota_0,    app,  ota_0,    0x110000, 1M,
ota_1,    app,  ota_1,    0x210000, 1M,
```

- There are now three app partition definitions. The type of the factory app (at 0x10000) and the next two "OTA" apps are all set to "app", but their subtypes are different.

- There is also a new "otadata" slot, which holds the data for OTA updates. The bootloader consults this data in order to know which app to execute. If "ota data" is empty, it will execute the factory app.

4.18.3 Creating Custom Tables

If you choose "Custom partition table CSV" in menuconfig then you can also enter the name of a CSV file (in the project directory) to use for your partition table. The CSV file can describe any number of definitions for the table you need.

The CSV format is the same format as printed in the summaries shown above. However, not all fields are required in the CSV. For example, here is the "input" CSV for the OTA partition table:

#	Name,	Type,	SubType,	Offset,	Size,	Flags
	nvs,	data,	nvs,	0x9000,	0x4000	
	otadata,	data,	ota,	0xd000,	0x2000	
	phy_init,	data,	phy,	0xf000,	0x1000	
	factory,	app,	factory,	0x10000,	1M	
	ota_0,	app,	ota_0,	,	1M	
	ota_1,	app,	ota_1,	,	1M	
	nvs_key,	data,	nvs_keys,	,	0x1000	

- Whitespace between fields is ignored, and so is any line starting with # (comments).
- Each non-comment line in the CSV file is a partition definition.
- The "Offset" field for each partition is empty. The `gen_esp32part.py` tool fills in each blank offset, starting after the partition table and making sure each partition is aligned correctly.

Name Field

Name field can be any meaningful name. It is not significant to the ESP32-P4. The maximum length of names is 16 bytes, including one null terminator. Names longer than the maximum length will be truncated.

Type Field

Partition type field can be specified as `app` (0x00) or `data` (0x01). Or it can be a number 0-254 (or as hex 0x00-0xFE). Types 0x00-0x3F are reserved for ESP-IDF core functions.

If your app needs to store data in a format not already supported by ESP-IDF, then please add a custom partition type value in the range 0x40-0xFE.

See [`esp\_partition\_type\_t`](#) for the enum definitions for `app` and `data` partitions.

If writing in C++ then specifying a application-defined partition type requires casting an integer to [`esp\_partition\_type\_t`](#) in order to use it with the [partition API](#). For example:

```
static const esp_partition_type_t APP_PARTITION_TYPE_A = (esp_partition_type_t) 0x40;
```

The ESP-IDF bootloader ignores any partition types other than `app` (0x00) and `data` (0x01).

SubType

The 8-bit SubType field is specific to a given partition type. ESP-IDF currently only specifies the meaning of the subtype field for `app` and `data` partition types.

See enum [`esp\_partition\_subtype\_t`](#) for the full list of subtypes defined by ESP-IDF, including the following:

- When type is `app`, the SubType field can be specified as `factory` (0x00), `ota_0` (0x10) ... `ota_15` (0x1F) or `test` (0x20).

- `factory` (0x00) is the default app partition. The bootloader will execute the factory app unless there it sees a partition of type `data/ota`, in which case it reads this partition to determine which OTA image to boot.
 - * OTA never updates the factory partition.
 - * If you want to conserve flash usage in an OTA project, you can remove the factory partition and use `ota_0` instead.
- `ota_0` (0x10) ... `ota_15` (0x1F) are the OTA app slots. When [OTA](#) is in use, the OTA data partition configures which app slot the bootloader should boot. When using OTA, an application should have at least two OTA application slots (`ota_0` & `ota_1`). Refer to the [OTA documentation](#) for more details.
- `test` (0x20) is a reserved subtype for factory test procedures. It will be used as the fallback boot partition if no other valid app partition is found. It is also possible to configure the bootloader to read a GPIO input during each boot, and boot this partition if the GPIO is held low, see [Boot from Test Firmware](#).
- When type is `data`, the subtype field can be specified as `ota` (0x00), `phy` (0x01), `nvs` (0x02), `nvs_keys` (0x04), or a range of other component-specific subtypes (see [subtype enum](#)).
 - `ota` (0) is the [OTA data partition](#) which stores information about the currently selected OTA app slot. This partition should be 0x2000 bytes in size. Refer to the [OTA documentation](#) for more details.
 - `phy` (1) is for storing PHY initialisation data. This allows PHY to be configured per-device, instead of in firmware.
 - * In the default configuration, the phy partition is not used and PHY initialisation data is compiled into the app itself. As such, this partition can be removed from the partition table to save space.
 - * To load PHY data from this partition, open the project configuration menu (`idf.py menuconfig`) and enable NOT UPDATED YET option. You will also need to flash your devices with phy init data as the esp-idf build system does not do this automatically.
 - `nvs` (2) is for the [Non-Volatile Storage \(NVS\) API](#).
 - * NVS is used to store per-device PHY calibration data (different to initialisation data).
 - * The NVS API can also be used for other application data.
 - * It is strongly recommended that you include an NVS partition of at least 0x3000 bytes in your project.
 - * If using NVS API to store a lot of data, increase the NVS partition size from the default 0x6000 bytes.
 - `nvs_keys` (4) is for the NVS key partition. See [Non-Volatile Storage \(NVS\) API](#) for more details.
 - * It is used to store NVS encryption keys when [NVS Encryption](#) feature is enabled.
 - * The size of this partition should be 4096 bytes (minimum partition size).
 - There are other predefined data subtypes for data storage supported by ESP-IDF. These include:
 - * `coredump` (0x03) is for storing core dumps while using a custom partition table CSV file. See [Core Dump](#) for more details.
 - * `efuse` (0x05) is for emulating eFuse bits using [Virtual eFuses](#).
 - * `undefined` (0x06) is implicitly used for data partitions with unspecified (empty) subtype, but it is possible to explicitly mark them as undefined as well.
 - * `fat` (0x81) is for [FAT Filesystem Support](#).
 - * `spiffs` (0x82) is for [SPIFFS Filesystem](#).
 - * `littlefs` (0x83) is for [LittleFS filesystem](#). See [storage/littlefs](#) example for more details.
- If the partition type is any application-defined value (range 0x40-0xFE), then `subtype` field can be any value chosen by the application (range 0x00-0xFE).

Note that when writing in C++, an application-defined subtype value requires casting to type [esp_partition_subtype_t](#) in order to use it with the [partition API](#).

Extra Partition SubTypes

A component can define a new partition subtype by setting the `EXTRA_PARTITION_SUBTYPES` property. This property is a CMake list, each entry of which is a comma separated string with `<type>`, `<subtype>`, `<value>` format. The build system uses this property to add extra subtypes and creates fields named `ESP_PARTITION_SUBTYPE_<type>_<subtype>` in [esp_partition_subtype_t](#). The project can use this subtype to define partitions in the partitions table CSV file and use the new fields in [esp_partition_subtype_t](#).

Offset & Size

The offset represents the partition address in the SPI flash, which sector size is 0x1000 (4 KB). Thus, the offset must be a multiple of 4 KB.

Partitions with blank offsets in the CSV file will start after the previous partition, or after the partition table in the case of the first partition.

Partitions of type `app` have to be placed at offsets aligned to 0x10000 (64 K). If you leave the offset field blank, `gen_esp32part.py` will automatically align the partition. If you specify an unaligned offset for an `app` partition, the tool will return an error.

Sizes and offsets can be specified as decimal numbers, hex numbers with the prefix 0x, or size multipliers K or M (1024 and 1024*1024 bytes).

If you want the partitions in the partition table to work relative to any placement ([CONFIG_PARTITION_TABLE_OFFSET](#)) of the table itself, leave the offset field (in CSV file) for all partitions blank. Similarly, if changing the partition table offset then be aware that all blank partition offsets may change to match, and that any fixed offsets may now collide with the partition table (causing an error).

Flags

Two flags are currently supported, `encrypted` and `readonly`:

- If `encrypted` flag is set, the partition will be encrypted if [Flash Encryption](#) is enabled.

Note: `app` type partitions will always be encrypted, regardless of whether this flag is set or not.

- If `readonly` flag is set, the partition will be read-only. This flag is only supported for data type partitions except `ota`` and `coredump`` subtypes. This flag can help to protect against accidental writes to a partition that contains critical device-specific configuration data, e.g., factory data partition.

Note: Using C file I/O API to open a file (`fopen``) in any write mode (`w`, `w+`, `a`, `a+`, `r+`) will fail and return `NULL`. Using `open` with any other flag than `O_RDONLY` will fail and return `-1` while `errno` global variable will be set to `EROFS`. This is also true for any other POSIX syscall function performing write or erase operations. Opening a handle in read-write mode for NVS on a read-only partition will fail and return [ESP_ERR_NOT_ALLOWED](#) error code. Using a lower level API like `esp_partition`, `spi_flash`, etc. to write to a read-only partition will result in [ESP_ERR_NOT_ALLOWED](#) error code.

You can specify multiple flags by separating them with a colon. For example, `encrypted:readonly`.

4.18.4 Generating Binary Partition Table

The partition table which is flashed to the ESP32-P4 is in a binary format, not CSV. The tool [partition_table/gen_esp32part.py](#) is used to convert between CSV and binary formats.

If you configure the partition table CSV name in the project configuration (`idf.py menuconfig`) and then build the project or run `idf.py partition-table`, this conversion is done as part of the build process.

To convert CSV to Binary manually:

```
python gen_esp32part.py input_partitions.csv binary_partitions.bin
```

To convert binary format back to CSV manually:

```
python gen_esp32part.py binary_partitions.bin input_partitions.csv
```

To display the contents of a binary partition table on stdout (this is how the summaries displayed when running `idf.py partition-table` are generated:

```
python gen_esp32part.py binary_partitions.bin
```

4.18.5 Partition Size Checks

The ESP-IDF build system will automatically check if generated binaries fit in the available partition space, and will fail with an error if a binary is too large.

Currently these checks are performed for the following binaries:

- Bootloader binary must fit in space before partition table (see [Bootloader Size](#)).
- App binary should fit in at least one partition of type "app". If the app binary does not fit in any app partition, the build will fail. If it only fits in some of the app partitions, a warning is printed about this.

Note: Although the build process will fail if the size check returns an error, the binary files are still generated and can be flashed (although they may not work if they are too large for the available space.)

MD5 Checksum

The binary format of the partition table contains an MD5 checksum computed based on the partition table. This checksum is used for checking the integrity of the partition table during the boot.

The MD5 checksum generation can be disabled by the `--disable-md5sum` option of `gen_esp32part.py` or by the [CONFIG_PARTITION_TABLE_MD5](#) option.

4.18.6 Flashing the Partition Table

- `idf.py partition-table-flash`: will flash the partition table with `esptool.py`.
- `idf.py flash`: Will flash everything including the partition table.

A manual flashing command is also printed as part of `idf.py partition-table` output.

Note: Note that updating the partition table does not erase data that may have been stored according to the old partition table. You can use `idf.py erase-flash` (or `esptool.py erase_flash`) to erase the entire flash contents.

4.18.7 Partition Tool (`parttool.py`)

The component `partition_table` provides a tool [parttool.py](#) for performing partition-related operations on a target device. The following operations can be performed using the tool:

- reading a partition and saving the contents to a file (`read_partition`)
- writing the contents of a file to a partition (`write_partition`)
- erasing a partition (`erase_partition`)
- retrieving info such as name, offset, size and flag ("encrypted") of a given partition (`get_partition_info`)

The tool can either be imported and used from another Python script or invoked from shell script for users wanting to perform operation programmatically. This is facilitated by the tool's Python API and command-line interface, respectively.

Python API

Before anything else, make sure that the *parttool* module is imported.

```
import sys
import os

idf_path = os.environ["IDF_PATH"] # get value of IDF_PATH from environment
parttool_dir = os.path.join(idf_path, "components", "partition_table") # parttool.
↳py lives in $IDF_PATH/components/partition_table

sys.path.append(parttool_dir) # this enables Python to find parttool module
from parttool import * # import all names inside parttool module
```

The starting point for using the tool's Python API to do is create a *ParttoolTarget* object:

```
# Create a parttool.py target device connected on serial port /dev/ttyUSB1
target = ParttoolTarget("/dev/ttyUSB1")
```

The created object can now be used to perform operations on the target device:

```
# Erase partition with name 'storage'
target.erase_partition(PartitionName("storage"))

# Read partition with type 'data' and subtype 'spiffs' and save to file 'spiffs.bin'
↳'
target.read_partition(PartitionType("data", "spiffs"), "spiffs.bin")

# Write to partition 'factory' the contents of a file named 'factory.bin'
target.write_partition(PartitionName("factory"), "factory.bin")

# Print the size of default boot partition
storage = target.get_partition_info(PARTITION_BOOT_DEFAULT)
print(storage.size)
```

The partition to operate on is specified using *PartitionName* or *PartitionType* or *PARTITION_BOOT_DEFAULT*. As the name implies, these can be used to refer to partitions of a particular name, type-subtype combination, or the default boot partition.

More information on the Python API is available in the docstrings for the tool.

Command-line Interface

The command-line interface of *parttool.py* has the following structure:

```
parttool.py [command-args] [subcommand] [subcommand-args]

- command-args - These are arguments that are needed for executing the main_
↳command (parttool.py), mostly pertaining to the target device
- subcommand - This is the operation to be performed
- subcommand-args - These are arguments that are specific to the chosen operation
```

```
# Erase partition with name 'storage'
parttool.py --port "/dev/ttyUSB1" erase_partition --partition-name=storage

# Read partition with type 'data' and subtype 'spiffs' and save to file 'spiffs.bin'
↳'
parttool.py --port "/dev/ttyUSB1" read_partition --partition-type=data --partition-
↳subtype=spiffs --output "spiffs.bin"

# Write to partition 'factory' the contents of a file named 'factory.bin'
```

(continues on next page)

(continued from previous page)

```
parttool.py --port "/dev/ttyUSB1" write_partition --partition-name=factory --input
↪ "factory.bin"

# Print the size of default boot partition
parttool.py --port "/dev/ttyUSB1" get_partition_info --partition-boot-default --
↪ info size
```

More information can be obtained by specifying `--help` as argument:

```
# Display possible subcommands and show main command argument descriptions
parttool.py --help

# Show descriptions for specific subcommand arguments
parttool.py [subcommand] --help
```

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.19 Performance

ESP-IDF ships with default settings that are designed for a trade-off between performance, resource usage, and available functionality.

These guides describe how to optimize a firmware application for a particular aspect of performance. Usually this involves some trade-off in terms of limiting available functions, or swapping one aspect of performance (such as execution speed) for another (such as RAM usage).

4.19.1 How to Optimize Performance

1. Decide the performance-critical aspects of your application, such as achieving a particular response time for a certain network operation, meeting a particular startup time limit, or maintaining a certain level of peripheral data throughput.
2. Find a way to measure this performance (some methods are outlined in the guides below).
3. Modify the code and project configuration and compare the new measurement to the old measurement.
4. Repeat step 3 until the performance meets the requirements set out in step 1.

4.19.2 Guides

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Speed Optimization

Overview Optimizing execution speed is a key element of software performance. Code that executes faster can also have other positive effects, e.g., reducing overall power consumption. However, improving execution speed may have trade-offs with other aspects of performance such as [Minimizing Binary Size](#).

Choose What to Optimize If a function in the application firmware is executed once per week in the background, it may not matter if that function takes 10 ms or 100 ms to execute. If a function is executed constantly at 10 Hz, it matters greatly if it takes 10 ms or 100 ms to execute.

Most kinds of application firmware only have a small set of functions that require optimal performance. Perhaps those functions are executed very often, or have to meet some application requirements for latency or throughput. Optimization efforts should be targeted at these particular functions.

Measuring Performance The first step to improving something is to measure it.

Basic Performance Measurements You may be able to measure directly the performance relative to an external interaction with the world, e.g., see the examples [wifi/iperf](#) and [ethernet/iperf](#) for measuring general network performance. Or you can use an oscilloscope or logic analyzer to measure the timing of an interaction with a device peripheral.

Otherwise, one way to measure performance is to augment the code to take timing measurements:

```
#include "esp_timer.h"

void measure_important_function(void) {
    const unsigned MEASUREMENTS = 5000;
    uint64_t start = esp_timer_get_time();

    for (int retries = 0; retries < MEASUREMENTS; retries++) {
        important_function(); // This is the thing you need to measure
    }

    uint64_t end = esp_timer_get_time();

    printf("%u iterations took %llu milliseconds (%llu microseconds per_
↪invocation)\n",
           MEASUREMENTS, (end - start)/1000, (end - start)/MEASUREMENTS);
}
```

Executing the target multiple times can help average out factors, e.g., RTOS context switches, overhead of measurements, etc.

- Using `esp_timer_get_time()` generates "wall clock" timestamps with microsecond precision, but has moderate overhead each time the timing functions are called.
- It is also possible to use the standard Unix `gettimeofday()` and `utime()` functions, although the overhead is slightly higher.
- Otherwise, including `hal/cpu_hal.h` and calling the HAL function `cpu_hal_get_cycle_count()` returns the number of CPU cycles executed. This function has lower overhead than the others, which is good for measuring very short execution times with high precision. The CPU cycles are counted per-core, so only use this method from an interrupt handler, or a task that is pinned to a single core.
- While performing "microbenchmarks" (i.e., benchmarking only a very small routine of code that runs in less than 1-2 milliseconds), the flash cache performance can sometimes cause big variations in timing measurements depending on the binary. This happens because binary layout can cause different patterns of cache misses in a particular sequence of execution. If the test code is larger, then this effect usually averages out. Executing a small function multiple times when benchmarking can help reduce the impact of flash cache misses. Alternatively, move this code to IRAM (see [Targeted Optimizations](#)).

External Tracing The [Application Level Tracing Library](#) allows measuring code execution with minimal impact on the code itself.

Tasks If the option `CONFIG_FREERTOS_GENERATE_RUN_TIME_STATS` is enabled, then the FreeRTOS API `vTaskGetRunTimeStats()` can be used to retrieve runtime information about the processor time used by each FreeRTOS task.

SEGGER SystemView is an excellent tool for visualizing task execution and looking for performance issues or improvements in the system as a whole.

Improving Overall Speed The following optimizations improve the execution of nearly all code, including boot times, throughput, latency, etc:

- Set `CONFIG_ESPTOOLPY_FLASHMODE` to QIO or QOUT mode (Quad I/O). Both almost double the speed at which code is loaded or executed from flash compared to the default DIO mode. QIO is slightly faster than QOUT if both are supported. Note that both the flash chip model, and the electrical connections between the ESP32-P4 and the flash chip must support quad I/O modes or the SoC will not work correctly.
- Set `CONFIG_COMPILER_OPTIMIZATION` to `Optimize for performance (-O2)`. This may slightly increase binary size compared to the default setting, but almost certainly increases the performance of some code. Note that if your code contains C or C++ Undefined Behavior, then increasing the compiler optimization level may expose bugs that otherwise are not seen.
- Avoid using floating point arithmetic `float`. Even though ESP32-P4 has a single precision hardware floating point unit, floating point calculations are always slower than integer calculations. If possible then use fixed point representations, a different method of integer representation, or convert part of the calculation to be integer only before switching to floating point.
- Avoid using double precision floating point arithmetic `double`. These calculations are emulated in software and are very slow. If possible then use an integer-based representation, or single-precision floating point.

Reduce Logging Overhead Although standard output is buffered, it is possible for an application to be limited by the rate at which it can print data to log output once buffers are full. This is particularly relevant for startup time if a lot of output is logged, but such problem can happen at other times as well. There are multiple ways to solve this problem:

- Reduce the volume of log output by lowering the app `CONFIG_LOG_DEFAULT_LEVEL` (the equivalent boot-loader setting is `CONFIG_BOOTLOADER_LOG_LEVEL`). This also reduces the binary size, and saves some CPU time spent on string formatting.
- Increase the speed of logging output by increasing the `CONFIG_ESP_CONSOLE_UART_BAUDRATE`.

Not Recommended The following options also increase execution speed, but are not recommended as they also reduce the debuggability of the firmware application and may increase the severity of any bugs.

- Set `CONFIG_COMPILER_OPTIMIZATION_ASSERTION_LEVEL` to disabled. This also reduces firmware binary size by a small amount. However, it may increase the severity of bugs in the firmware including security-related bugs. If it is necessary to do this to optimize a particular function, consider adding `#define NDEBUG` at the top of that single source file instead.

Targeted Optimizations The following changes increase the speed of a chosen part of the firmware application:

- Move frequently executed code to IRAM. By default, all code in the app is executed from flash cache. This means that it is possible for the CPU to have to wait on a "cache miss" while the next instructions are loaded from flash. Functions which are copied into IRAM are loaded once at boot time, and then always execute at full speed.
IRAM is a limited resource, and using more IRAM may reduce available DRAM, so a strategic approach is needed when moving code to IRAM. See *IRAM (Instruction RAM)* for more information.
- Jump table optimizations can be re-enabled for individual source files that do not need to be placed in IRAM. For hot paths in large `switch` cases, this improves performance. For instructions on how to add the `-fjump-tables` and `-ftree-switch-conversion` options when compiling individual source files, see *Controlling Component Compilation*

Improving Startup Time In addition to the overall performance improvements shown above, the following options can be tweaked to specifically reduce startup time:

- Minimizing the `CONFIG_LOG_DEFAULT_LEVEL` and `CONFIG_BOOTLOADER_LOG_LEVEL` has a large impact on startup time. To enable more logging after the app starts up, set the `CONFIG_LOG_MAXIMUM_LEVEL` as well, and then call `esp_log_level_set()` to restore higher level logs. The `system/startup_time` main function shows how to do this.
- If using Deep-sleep mode, setting `CONFIG_BOOTLOADER_SKIP_VALIDATE_IN_DEEP_SLEEP` allows a faster wake from sleep. Note that if using Secure Boot, this represents a security compromise, as Secure Boot validation are not performed on wake.
- Setting `CONFIG_BOOTLOADER_SKIP_VALIDATE_ON_POWER_ON` skips verifying the binary on every boot from the power-on reset. How much time this saves depends on the binary size and the flash settings. Note that this setting carries some risk if the flash becomes corrupt unexpectedly. Read the help text of the [config item](#) for an explanation and recommendations if using this option.
- It is possible to save a small amount of time during boot by disabling RTC slow clock calibration. To do so, set `CONFIG_RTC_CLK_CAL_CYCLES` to 0. Any part of the firmware that uses RTC slow clock as a timing source will be less accurate as a result.

The example project `system/startup_time` is pre-configured to optimize startup time. The file `system/startup_time/sdkconfig.defaults` contain all of these settings. You can append these to the end of your project's own `sdkconfig` file to merge the settings, but please read the documentation for each setting first.

Task Priorities As ESP-IDF FreeRTOS is a real-time operating system, it is necessary to ensure that high-throughput or low-latency tasks are granted a high priority in order to run immediately. Priority is set when calling `xTaskCreate()` or `xTaskCreatePinnedToCore()` and can be changed at runtime by calling `vTaskPrioritySet()`.

It is also necessary to ensure that tasks yield CPU (by calling `vTaskDelay()`, `sleep()`, or by blocking on semaphores, queues, task notifications, etc) in order to not starve lower-priority tasks and cause problems for the overall system. The [Task Watchdog Timer \(TWDT\)](#) provides a mechanism to automatically detect if task starvation happens. However, note that a TWDT timeout does not always indicate a problem, because sometimes the correct operation of the firmware requires some long-running computation. In these cases, tweaking the TWDT timeout or even disabling the TWDT may be necessary.

Built-in Task Priorities ESP-IDF starts a number of system tasks at fixed priority levels. Some are automatically started during the boot process, while some are started only if the application firmware initializes a particular feature. To optimize performance, structure the task priorities of your application properly to ensure the tasks are not delayed by the system tasks, while also not starving system tasks and impacting other functions of the system.

This may require splitting up a particular task. For example, perform a time-critical operation in a high-priority task or an interrupt handler and do the non-time-critical part in a lower-priority task.

Header `components/esp_system/include/esp_task.h` contains macros for the priority levels used for built-in ESP-IDF tasks system. See [Background Tasks](#) for more details about the system tasks.

Common priorities are:

- [Running the Main Task](#) that executes `app_main` function has minimum priority (1). This task is pinned to Core 0 by default ([configurable](#)).
- [High Resolution Timer \(ESP Timer\)](#) system task to manage high precision timer events and execute callbacks has high priority (22, `ESP_TASK_TIMER_PRIO`). This task is pinned to Core 0.
- FreeRTOS Timer Task to handle FreeRTOS timer callbacks is created when the scheduler initializes and has minimum task priority (1, [configurable](#)). This task is pinned to Core 0.
- [Event Loop Library](#) system task to manage the default system event loop and execute callbacks has high priority (20, `ESP_TASK_EVENT_PRIO`) and it is pinned to Core 0. This configuration is only used if the application calls `esp_event_loop_create_default()`, it is possible to call `esp_event_loop_create()` with a custom task configuration instead.

- **lwIP TCP/IP task has high priority (18, `ESP_TASK_TCPIP_PRIO`) and is not pinned to any core (*configurable*).**
 - Stack event callback task ("BTC") has high priority (19).
 - Stack BTU layer task has high priority (20).
 - Host HCI host task has high priority (22).
- All Bluedroid Tasks are pinned to the same core, which is Core 0 by default (*configurable*).
- The Ethernet driver creates a task for the MAC to receive Ethernet frames. If using the default config `ETH_MAC_DEFAULT_CONFIG` then the priority is medium-high (15) and the task is not pinned to any core. These settings can be changed by passing a custom `eth_mac_config_t` struct when initializing the Ethernet MAC.
 - If using the *ESP-MQTT* component, it creates a task with default priority 5 (*configurable*, depending on *CONFIG_MQTT_USE_CUSTOM_CONFIG*) and not pinned to any core (*configurable*).
 - To see what is the task priority for mDNS service, please check [Performance Optimization](#).

Choosing Task Priorities of the Application With a few exceptions, most importantly the lwIP TCP/IP task, in the default configuration most built-in tasks are pinned to Core 0. This makes it quite easy for the application to place high priority tasks on Core 1. Using priority 19 or higher guarantees that an application task can run on Core 1 without being preempted by any built-in task. To further isolate the tasks running on each CPU, configure the *lwIP task* to only run on Core 0 instead of either core, which may reduce total TCP/IP throughput depending on what other tasks are running.

In general, it is not recommended to set task priorities on Core 0 higher than the built-in Wi-Fi/Bluetooth operations as starving them of CPU may make the system unstable. Choosing priority 19 and Core 0 allows lower-layer Wi-Fi/Bluetooth functionality to run without delays, but still pre-empt the lwIP TCP/IP stack and other less time-critical internal functionality. This is an option for time-critical tasks that do not perform network operations. Any task that does TCP/IP network operations should run at lower priority than the lwIP TCP/IP task (18) to avoid priority-inversion issues.

Note: Setting a task to always run in preference to built-in ESP-IDF tasks does not require pinning the task to Core 1. Instead, the task can be left unpinned and assigned a priority of 17 or lower. This allows the task to optionally run on Core 0 if there are no higher-priority built-in tasks running on that core. Using unpinned tasks can improve the overall CPU utilization, however it makes reasoning about task scheduling more complex.

Note: Task execution is always completely suspended when writing to the built-in SPI flash chip. Only *IRAM-Safe Interrupt Handlers* continues executing.

Improving Interrupt Performance ESP-IDF supports dynamic *Interrupt Allocation* with interrupt preemption. Each interrupt in the system has a priority, and higher-priority interrupts preempts lower priority ones.

Interrupt handlers execute in preference to any task, provided the task is not inside a critical section. For this reason, it is important to minimize the amount of time spent in executing an interrupt handler.

To obtain the best performance for a particular interrupt handler:

- Assign more important interrupts a higher priority using a flag such as `ESP_INTR_FLAG_LEVEL2` or `ESP_INTR_FLAG_LEVEL3` when calling `esp_intr_alloc()`.
- Assign the interrupt on a CPU where built-in Wi-Fi/Bluetooth tasks are not configured to run, which means assigning the interrupt on Core 1 by default, see *Built-in Task Priorities*. Interrupts are assigned on the same CPU where the `esp_intr_alloc()` function call is made.
- If you are sure the entire interrupt handler can run from IRAM (see *IRAM-Safe Interrupt Handlers*) then set the `ESP_INTR_FLAG_IRAM` flag when calling `esp_intr_alloc()` to assign the interrupt. This prevents it being temporarily disabled if the application firmware writes to the internal SPI flash.
- Even if the interrupt handler is not IRAM-safe, if it is going to be executed frequently then consider moving the handler function to IRAM anyhow. This minimizes the chance of a flash cache miss when the interrupt code is

executed (see [Targeted Optimizations](#)). It is possible to do this without adding the `ESP_INTR_FLAG_IRAM` flag to mark the interrupt as IRAM-safe, if only part of the handler is guaranteed to be in IRAM.

Improving Network Speed

- For lwIP TCP/IP (Wi-Fi and Ethernet), see [Performance Optimization](#)

Improving I/O Performance Using standard C library functions like `fread` and `fwrite` instead of platform specific unbuffered syscalls such as `read` and `write` can be slow. These functions are designed to be portable, so they are not necessarily optimized for speed, have a certain overhead and are buffered.

[FAT Filesystem Support](#) specific information and tips:

- Maximum size of the R/W request = FatFS cluster size (allocation unit size).
- Use `read` and `write` instead of `fread` and `fwrite`.
- To increase speed of buffered reading functions like `fread` and `fgets`, you can increase a size of the file buffer (Newlib's default is 128 bytes) to a higher number like 4096, 8192 or 16384. This can be done locally via the `setvbuf` function used on a certain file pointer or globally applied to all files via modifying [CONFIG_FATFS_VFS_FSTAT_BLKSIZE](#).

Note: Setting a bigger buffer size also increases the heap memory usage.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Minimizing Binary Size

The ESP-IDF build system compiles all source files in the project and ESP-IDF, but only functions and variables that are actually referenced by the program are linked into the final binary. In some cases, it is necessary to reduce the total size of the firmware binary, e.g., in order to fit it into the available flash partition size.

The first step to reducing the total firmware binary size is measuring what is causing the size to increase.

Measuring Static Sizes To optimize both the firmware binary size and the memory usage, it is necessary to measure statically-allocated RAM (`data`, `bss`), code (`text`), and read-only data (`rodata`) in your project.

Using the [idf.py](#) sub-commands `size`, `size-components`, and `size-files` provides a summary of memory used by the project:

Size Summary `idf.py size`

```
$ idf.py size
[...]
Total sizes:
Used stat D/IRAM:  53743 bytes ( 122385 remain, 30.5% used)
    .data size:    6504 bytes
    .bss  size:    1984 bytes
    .text size:    44228 bytes
    .vectors size: 1027 bytes
Used Flash size : 118879 bytes
    .text:        83467 bytes
```

(continues on next page)

(continued from previous page)

```

        .rodata:    35156 bytes
Total image size: 170638 bytes (.bin may be padded larger)

```

This output breaks down the size of all static memory regions in the firmware binary:

```

$ idf.py size
[...]
Total sizes:
Used stat D/IRAM: 53743 bytes ( 122385 remain, 30.5% used)
    .data size: 6504 bytes
    .bss size: 1984 bytes
    .text size: 44228 bytes
    .vectors size: 1027 bytes
Used Flash size : 118879 bytes
    .text: 83467 bytes
    .rodata: 35156 bytes
Total image size: 170638 bytes (.bin may be padded larger)

```

- **Used stat D/IRAM:** Total amount of D/IRAM used at compile time. `remain` indicates the amount of D/IRAM left to be used as heap memory at runtime. Note that due to meta data overhead, implementation constraints, and startup heap allocations, the actual size of the DRAM heap is smaller.
 - **.data size:** Amount of D/IRAM allocated at compile time for the `.data` (i.e., all statically allocated variables that are initialized to non-zero values). `.data` also consumes space in the binary image to store the non-zero initialization values.
 - **.bss size:** Amount of D/IRAM allocated at compile time for `.bss` (i.e., all statically allocated variables that are initialized to zero). `.bss` does not consume extra space in flash.
 - **.text size:** Amount of D/IRAM used for `.text` (i.e., all code that is executed from internal RAM). `.text` also consumes space in the binary image as the code is initially stored there and is then copied over to D/IRAM on startup.
- **Used Flash size:** Total amount of flash used (excluding usage by D/IRAM)
 - **.text:** Amount of flash used for `.text` (i.e., all code that is executed via the flash cache, see [IROM](#)).
 - **.rodata:** Amount of flash used for `.rodata` (i.e., read-only data that is loaded via the flash cache, see [DROM](#)).
- **Total image size** is the estimated total size of the binary file.

Component Usage Summary `idf.py size-components` The summary output provided by `idf.py size` does not give enough details to find the main contributor to excessive binary size. To analyze in detail, use `idf.py size-components`.

```

$ idf.py size-components
[...]
Total sizes:
DRAM .data size: 14956 bytes
DRAM .bss size: 15808 bytes
Used static DRAM: 30764 bytes ( 149972 available, 17.0% used)
Used static IRAM: 83918 bytes ( 47154 available, 64.0% used)
Flash code: 559943 bytes
Flash rodata: 176736 bytes
Total image size:~ 835553 bytes (.bin may be padded larger)
Per-archive contributions to ELF file:
    Archive File DRAM .data & .bss & other    IRAM    D/IRAM Flash code &
↪rodata    Total
    libnet80211.a      1267    6044      0    5490      0    107445
↪18484    138730
    liblwip.a          21    3838      0      0      0    97465
↪16116    117440
    libmbedtls.a       60     524      0      0      0    27655
↪69907    98146
    libmbedcrypto.a   64      81      0     30      0    76645
↪11661    88481

```

(continues on next page)

(continued from previous page)

		libpp.a	2427	1292	0	20851	0	37208	└
↪4708	66486								
		libc.a	4	0	0	0	0	57056	└
↪6455	63515								
		libphy.a	1439	715	0	7798	0	33074	└
↪ 0	43026								
		libwpa_suppliment.a	12	848	0	0	0	35505	└
↪1446	37811								
		libfreertos.a	3104	740	0	15711	0	367	└
↪4228	24150								
		libnvs_flash.a	0	24	0	0	0	14347	└
↪2924	17295								
		libspi_flash.a	1562	294	0	8851	0	1840	└
↪1913	14460								
		libesp_system.a	245	206	0	3078	0	5990	└
↪3817	13336								
		libesp-tls.a	0	4	0	0	0	5637	└
↪3524	9165								
[... removed some lines here ...]									
		libesp_rom.a	0	0	0	112	0	0	└
↪ 0	112								
		libcxx.a	0	0	0	0	0	47	└
↪ 0	47								
		(exe)	0	0	0	3	0	3	└
↪ 12	18								
		libesp_pm.a	0	0	0	0	0	8	└
↪ 0	8								
		libesp_eth.a	0	0	0	0	0	0	└
↪ 0	0								
		libmesh.a	0	0	0	0	0	0	└
↪ 0	0								

The first lines of the output from `idf.py size-components` are the same as that from `idf.py size`. After this, a table is printed as Per-archive contributions to ELF file. This means how much each static library archive has contributed to the final binary size.

Generally, one static library archive is built per component, although some are binary libraries included by a particular component, for example, `libnet80211.a` is included by `esp_wifi` component. There are also toolchain libraries such as `libc.a` and `libgcc.a` listed here, these provide Standard C/C++ Library and toolchain built-in functionality.

If your project is simple and only has a `main` component, then all of the project's code will be shown under `lib-main.a`. If your project includes its own components (see [Build System](#)), then they will each be shown on a separate line.

The table is sorted in descending order of the total contribution of the static archive to the binary size.

The columns are as follows:

- `DRAM .data & .bss & other - .data` and `.bss` are the same as for the totals shown above. Both are static variables and reduce the total available RAM at runtime, but `.bss` does not contribute to the binary file size. `other` is a column for any custom section types that also contribute to RAM size. Usually, the value is 0.
- `IRAM` - is the same as for the totals shown above. It refers to code linked to execute from IRAM, which uses space in the binary file and also reduces DRAM available as heap at runtime.
- `Flash code & rodata` - these are the same as the totals above, IROM and DROM space accessed from the flash cache that contribute to the binary size.

Source File Usage Summary `idf.py size-files` For even more details, run `idf.py size-files` to get a summary of the contribution each object file has made to the final binary size. Each object file corresponds to

a single source file.

```
$ idf.py size-files
[...]
Total sizes:
  DRAM .data size:  14956 bytes
  DRAM .bss size:  15808 bytes
Used static DRAM:  30764 bytes ( 149972 available, 17.0% used)
Used static IRAM:  83918 bytes ( 47154 available, 64.0% used)
  Flash code: 559943 bytes
  Flash rodata: 176736 bytes
Total image size:~ 835553 bytes (.bin may be padded larger)
Per-file contributions to ELF file:
```

Object File	DRAM .data	& .bss	& other	IRAM	D/IRAM	Flash code	& rodata
→rodata Total							
x509_crt_bundle.S.o	0	0	0	0	0	0	└
→64212 64212							
wl_cnx.o	2	3183	0	221	0	13119	└
→3286 19811							
phy_chip_v7.o	721	614	0	1642	0	16820	└
→ 0 19797							
ieee80211_ioctl.o	740	96	0	437	0	15325	└
→2627 19225							
pp.o	1142	45	0	8871	0	5030	└
→537 15625							
ieee80211_output.o	2	20	0	2118	0	11617	└
→914 14671							
ieee80211_sta.o	1	41	0	1498	0	10858	└
→2218 14616							
lib_a-vfprintf.o	0	0	0	0	0	13829	└
→752 14581							
lib_a-svfprintf.o	0	0	0	0	0	13251	└
→752 14003							
ssl_tls.c.o	60	0	0	0	0	12769	└
→463 13292							
sockets.c.o	0	648	0	0	0	11096	└
→1030 12774							
nd6.c.o	8	932	0	0	0	11515	└
→314 12769							
phy_chip_v7_cal.o	477	53	0	3499	0	8561	└
→ 0 12590							
pm.o	32	364	0	2673	0	7788	└
→782 11639							
ieee80211_scan.o	18	288	0	0	0	8889	└
→1921 11116							
lib_a-svfprintf.o	0	0	0	0	0	9654	└
→1206 10860							
lib_a-vfprintf.o	0	0	0	0	0	10069	└
→734 10803							
ieee80211_ht.o	0	4	0	1186	0	8628	└
→898 10716							
phy_chip_v7_ana.o	241	48	0	2657	0	7677	└
→ 0 10623							
bignum.c.o	0	4	0	0	0	9652	└
→752 10408							
tcp_in.c.o	0	52	0	0	0	8750	└
→1282 10084							
trc.o	664	88	0	1726	0	6245	└
→1108 9831							
tasks.c.o	8	704	0	7594	0	0	└
→1475 9781							
ecp_curves.c.o	28	0	0	0	0	7384	└
→2325 9737							

(continues on next page)

(continued from previous page)

↪286	9214	ecp.c.o	0	64	0	0	0	8864	↪
↪585	9205	ieee80211_hostap.o	1	41	0	0	0	8578	↪
↪580	9009	wdev.o	121	125	0	4499	0	3684	↪
↪2161	7847	tcp_out.c.o	0	0	0	0	0	5686	↪
↪1617	7806	tcp.c.o	2	26	0	0	0	6161	↪
↪973	7770	ieee80211_input.o	0	0	0	0	0	6797	↪
↪55	7539	wpa.c.o	0	656	0	0	0	6828	↪
[... additional lines removed ...]									

After the summary of total sizes, a table of Per-file contributions to ELF file is printed.

The columns are the same as shown above for `idf.py` size-components, but this time the granularity is the contribution of each individual object file to the binary size.

For example, we can see that the file `x509_crt_bundle.S.o` contributed 64,212 bytes to the total firmware size, all as `.rodata` in flash. Therefore we can guess that this application is using the [ESP x509 Certificate Bundle](#) feature and not using this feature would save at last this many bytes from the firmware size.

Some of the object files are linked from binary libraries and therefore you will not find a corresponding source file. To locate which component a source file belongs to, it is generally possible to search in the ESP-IDF source tree or look in the [Linker Map File](#) for the full path.

Comparing Two Binaries If making some changes that affect binary size, it is possible to use an ESP-IDF tool to break down the exact differences in size.

This operation is not part of `idf.py`, it is necessary to run the `esp_idf_size` Python tool directly.

To do so, first, locate the linker map file with the name `PROJECTNAME.map` in the build directory. The `esp_idf_size` tool performs its analysis based on the output of the linker map file.

To compare with another binary, you also need its corresponding `.map` file saved from the build directory.

For example, to compare two builds, one of which with the default [CONFIG_COMPILER_OPTIMIZATION](#) setting Debug (`-Og`) configuration while another with Optimize for size (`-Os`):

```
$ python -m esp_idf_size --diff build_Og/https_request.map build_Os/https_request.
↪map
<CURRENT> MAP file: build_Os/https_request.map
<REFERENCE> MAP file: build_Og/https_request.map
Difference is counted as <CURRENT> - <REFERENCE>, i.e. a positive number means
↪that <CURRENT> is larger.
Total sizes of <CURRENT>:
↪<REFERENCE>      Difference
DRAM .data size:   14516 bytes
↪14956             -440
DRAM .bss size:    15792 bytes
↪15808             -16
Used static DRAM:  30308 bytes ( 150428 available, 16.8% used)
↪30764             -456 ( +456 available, +0 total)
Used static IRAM:  78498 bytes ( 52574 available, 59.9% used)
↪83918             -5420 ( +5420 available, +0 total)
Flash code:        509183 bytes
↪559943            -50760
Flash rodata:      170592 bytes
↪176736            -6144
```

(continues on next page)

(continued from previous page)

```
Total image size:~ 772789 bytes (.bin may be padded larger)
↪835553          -62764
```

We can see from the `Difference` column that changing this one setting caused the whole binary to be over 60 KB smaller and over 5 KB more RAM is available.

It is also possible to use the `diff` mode to output a table of component-level (static library archive) differences:

Note: To get the output in JSON or CSV format using `esp_idf_size`, it is possible to use the `--format` option.

```
python -m esp_idf_size --archives --diff build_Og/https_request.map build_Oshttps_
↪request.map
```

Also at the individual source file level:

```
python -m esp_idf_size --files --diff build_Og/https_request.map build_Oshttps_
↪request.map
```

Other options, like writing the output to a file, are available, pass `--help` to see the full list.

Showing Size When Linker Fails If too much static memory is allocated, the linker will fail with an error such as DRAM segment data does not fit, region ``iram0_0_seg'` overflowed by 44 bytes, or similar.

In these cases, `idf.py size` will not succeed either. However, it is possible to run `esp_idf_size` manually to view the **partial static memory usage**. The memory usage will miss the variables that could not be linked, so there still appears to be some free space.

The map file argument is `<projectname>.map` in the build directory.

```
python -m esp_idf_size build/project_name.map
```

It is also possible to view the equivalent of `size-components` or `size-files` output:

```
python -m esp_idf_size --archives build/project_name.map
python -m esp_idf_size --files build/project_name.map
```

Linker Map File

Note: This is an advanced analysis method, but it can be very useful. Feel free to skip ahead to [Reducing Overall Size](#) and possibly come back to this later.

The `idf.py size` analysis tools all work by parsing the GNU binutils linker map file, which is a summary of everything the linker did when it created (i.e., linked) the final firmware binary file.

Linker map files themselves are plain text files, so it is possible to read them and find out exactly what the linker did. However, they are also very complex and long, often exceeding 100,000 lines.

The map file itself is broken into parts and each part has a heading. The parts are:

- Archive member included to satisfy reference by file (symbol)
 - This shows you: for each object file included in the link, what symbol (function or variable) was the linker searching for when it included that object file.
 - If you are wondering why some object file in particular was included in the binary, this part may give a clue. This part can be used in conjunction with the Cross Reference Table at the end of the file.

Note: Not every object file shown in this list ends up included in the final binary, some end up in the `Discarded input sections` list instead.

- `Allocating common symbols`
 - This is a list of some global variables along with their sizes. Common symbols have a particular meaning in ELF binary files, but ESP-IDF does not make much use of them.
- `Discarded input sections`
 - These sections were read by the linker as part of an object file to be linked into the final binary, but then nothing else referred to them, so they were discarded from the final binary.
 - For ESP-IDF, this list can be very long, as we compile each function and static variable to a unique section in order to minimize the final binary size. Specifically, ESP-IDF uses compiler options `-ffunction-sections` `-fdata-sections` and linker option `--gc-sections`.
 - Items mentioned in this list **do not** contribute to the final binary.
- `Memory Configuration, Linker script and memory map`
 - These two parts go together. Some of the output comes directly from the linker command line and the Linker Script, both provided by *Build System*. The linker script is partially generated from the ESP-IDF project using the *Linker Script Generation* feature.
 - As the output of the `Linker script` and `memory map` part of the map unfolds, you can see each symbol (function or static variable) linked into the final binary along with its address (as a 16 digit hex number), its length (also in hex), and the library and object file it was linked from (which can be used to determine the component and the source file).
 - Following all of the output sections that take up space in the final `.bin` file, the `memory map` also includes some sections in the ELF file that are only used for debugging, e.g., ELF sections `.debug_*`, etc. These do not contribute to the final binary size. You can notice the address of these symbols is a very small number, starting from `0x0000000000000000` and counting up.
- `Cross Reference Table`
 - This table shows the symbol (function or static variable) that the list of object file(s) refers to. If you are wondering why a particular thing is included in the binary, this will help determine what included it.

Note: Unfortunately, the `Cross Reference Table` does not only include symbols that made it into the final binary. It also includes symbols in discarded sections. Therefore, just because something is shown here does not mean that it was included in the final binary - this needs to be checked separately.

Note: Linker map files are generated by the GNU binutils linker `ld`, not ESP-IDF. You can find additional information online about the linker map file format. This quick summary is written from the perspective of ESP-IDF build system in particular.

Reducing Overall Size The following configuration options reduces the final binary size of almost any ESP-IDF project:

- Set `CONFIG_COMPILER_OPTIMIZATION` to `Optimize for size (-Os)`. In some cases, `Optimize for performance (-O2)` will also reduce the binary size compared to the default. Note that if your code contains C or C++ Undefined Behavior then increasing the compiler optimization level may expose bugs that otherwise do not happen.
- Reduce the compiled-in log output by lowering the app `CONFIG_LOG_DEFAULT_LEVEL`. If the `CONFIG_LOG_MAXIMUM_LEVEL` is changed from the default then this setting controls the binary size instead. Reducing compiled-in logging reduces the number of strings in the binary, and also the code size of the calls to logging functions.
- Set the `CONFIG_COMPILER_OPTIMIZATION_ASSERTION_LEVEL` to `Silent`. This avoids compiling in a dedicated assertion string and source file name for each `assert` that may fail. It is still possible to find the failed `assert` in the code by looking at the memory address where the assertion failed.

- Besides the [CONFIG_COMPILER_OPTIMIZATION_ASSERTION_LEVEL](#), you can disable or silent the assertion for the HAL component separately by setting [CONFIG_HAL_DEFAULT_ASSERTION_LEVEL](#). It is to notice that ESP-IDF lowers the HAL assertion level in bootloader to be silent even if [CONFIG_HAL_DEFAULT_ASSERTION_LEVEL](#) is set to full-assertion level. This is to reduce the bootloader size.
- Setting [CONFIG_COMPILER_OPTIMIZATION_CHECKS_SILENT](#) removes specific error messages for particular internal ESP-IDF error check macros. This may make it harder to debug some error conditions by reading the log output.
- Do not enable [CONFIG_COMPILER_CXX_EXCEPTIONS](#), [CONFIG_COMPILER_CXX_RTTI](#), or set the [CONFIG_COMPILER_STACK_CHECK_MODE](#) to Overall. All of these options are already disabled by default, but they have a large impact on binary size.
- Disabling [CONFIG_ESP_ERR_TO_NAME_LOOKUP](#) removes the lookup table to translate user-friendly names for error values (see [Error Handling](#)) in error logs, etc. This saves some binary size, but error values will be printed as integers only.
- Setting [CONFIG_ESP_SYSTEM_PANIC](#) to `Silent` `reboot` saves a small amount of binary size, however this is **only** recommended if no one will use UART output to debug the device.
- Setting [CONFIG_COMPILER_SAVE_RESTORE_LIBCALLS](#) reduces binary size by replacing inlined prologues/epilogues with library calls.
- If the application binary uses only one of the security versions of the protocomm component, then the support for others can be disabled to save some code size. The support can be disabled through [CONFIG_ESP_PROTOCOMM_SUPPORT_SECURITY_VERSION_0](#), [CONFIG_ESP_PROTOCOMM_SUPPORT_SECURITY_VERSION_1](#) or [CONFIG_ESP_PROTOCOMM_SUPPORT_SECURITY_VERSION_2](#) respectively.

Note: In addition to the many configuration items shown here, there are a number of configuration options where changing the option from the default increases binary size. These are not noted here. Where the increase is significant is usually noted in the configuration item help text.

Targeted Optimizations The following binary size optimizations apply to a particular component or a function:

lwIP IPv6

- Setting [CONFIG_LWIP_IPV6](#) to `false` will reduce the size of the lwIP TCP/IP stack, at the cost of only supporting IPv4.

Note: IPv6 is required by some components such as `coap` and [ASIO Port](#). These components will not be available if IPV6 is disabled.

lwIP IPv4

- If IPv4 connectivity is not required, setting [CONFIG_LWIP_IPV4](#) to `false` will reduce the size of the lwIP, supporting IPv6-only TCP/IP stack.

Note: Before disabling IPv4 support, please note that IPv6 only network environments are not ubiquitous and must be supported in the local network, e.g., by your internet service provider or using constrained local network settings.

Newlib Nano Formatting By default, ESP-IDF uses Newlib "full" formatting for I/O functions (`printf()`, `scanf()`, etc.)

Enabling "Nano" formatting reduces the stack usage of each function that calls `printf()` or another string formatting function, see [Reducing Stack Sizes](#).

"Nano" formatting does not support 64-bit integers, or C99 formatting features. For a full list of restrictions, search for `--enable-newlib-nano-formatted-io` in the [Newlib documentation](#).

MbedTLS Features Under **Component Config > mbedTLS**, there are multiple mbedTLS features enabled default, some of which can be disabled if not needed to save code size.

These include:

- [*CONFIG_MBEDTLS_HAVE_TIME*](#)
- [*CONFIG_MBEDTLS_ECDSA_DETERMINISTIC*](#)
- [*CONFIG_MBEDTLS_SHA512_C*](#)
- [*CONFIG_MBEDTLS_SHA3_C*](#)
- [*CONFIG_MBEDTLS_CLIENT_SSL_SESSION_TICKETS*](#)
- [*CONFIG_MBEDTLS_SERVER_SSL_SESSION_TICKETS*](#)
- [*CONFIG_MBEDTLS_SSL_CONTEXT_SERIALIZATION*](#)
- [*CONFIG_MBEDTLS_SSL_ALPN*](#)
- [*CONFIG_MBEDTLS_SSL_RENEGOTIATION*](#)
- [*CONFIG_MBEDTLS_CCM_C*](#)
- [*CONFIG_MBEDTLS_GCM_C*](#)
- [*CONFIG_MBEDTLS_ECP_C*](#) (Alternatively: Leave this option enabled but disable some of the elliptic curves listed in the sub-menu.)
- [*CONFIG_MBEDTLS_ECP_NIST_OPTIM*](#)
- [*CONFIG_MBEDTLS_ECP_FIXED_POINT_OPTIM*](#)
- Change [*CONFIG_MBEDTLS_TLS_MODE*](#) if both server & client functionalities are not needed.
- Consider disabling some cipher suites listed in the TLS Key Exchange Methods sub-menu (i.e., [*CONFIG_MBEDTLS_KEY_EXCHANGE_RSA*](#)).
- Consider disabling [*CONFIG_MBEDTLS_ERROR_STRINGS*](#) if the application is already pulling in mbedTLS error strings through using `MBEDTLS_strerror()`.

The help text for each option has some more information for reference.

Important: It is **strongly not recommended to disable all these mbedTLS options**. Only disable options of which you understand the functionality and are certain that it is not needed in the application. In particular:

- Ensure that any TLS server(s) the device connects to can still be used. If the server is controlled by a third party or a cloud service, it is recommended to ensure that the firmware supports at least two of the supported cipher suites in case one is disabled in a future update.
- Ensure that any TLS client(s) that connect to the device can still connect with supported/recommended cipher suites. Note that future versions of client operating systems may remove support for some features, so it is recommended to enable multiple supported cipher suites, or algorithms for redundancy.

If depending on third party clients or servers, always pay attention to announcements about future changes to supported TLS features. If not, the ESP32-P4 device may become inaccessible if support changes.

Note: Not every combination of mbedTLS compile-time config is tested in ESP-IDF. If you find a combination that fails to compile or function as expected, please report the details on [GitHub](#).

VFS *Virtual Filesystem Component* feature in ESP-IDF allows multiple filesystem drivers and file-like peripheral drivers to be accessed using standard I/O functions (`open`, `read`, `write`, etc.) and C library functions (`fopen`, `fread`, `fwrite`, etc.). When filesystem or file-like peripheral driver functionality is not used in the application, this feature can be fully or partially disabled. VFS component provides the following configuration options:

- [*CONFIG_VFS_SUPPORT_TERMIOS*](#) — can be disabled if the application does not use `termios` family of functions. Currently, these functions are implemented only for UART VFS driver. Most applications can disable this option. Disabling this option reduces the code size by about 1.8 KB.
- [*CONFIG_VFS_SUPPORT_SELECT*](#) — can be disabled if the application does not use the `select` function with file descriptors. Currently, only the UART and eventfd VFS drivers implement `select` support. Note that when this option is disabled, `select` can still be used for socket file descriptors. Disabling this option reduces the code size by about 2.7 KB.

- [CONFIG_VFS_SUPPORT_DIR](#) —can be disabled if the application does not use directory-related functions, such as `readdir` (see the description of this option for the complete list). Applications that only open, read and write specific files and do not need to enumerate or create directories can disable this option, reducing the code size by 0.5 KB or more, depending on the filesystem drivers in use.
- [CONFIG_VFS_SUPPORT_IO](#) —can be disabled if the application does not use filesystems or file-like peripheral drivers. This disables all VFS functionality, including the three options mentioned above. When this option is disabled, [Console](#) can not be used. Note that the application can still use standard I/O functions with socket file descriptors when this option is disabled. Compared to the default configuration, disabling this option reduces code size by about 9.4 KB.

HAL

- Enabling [CONFIG_HAL_SYSTIMER_USE_ROM_IMPL](#) can reduce the IRAM usage and binary size by linking in the systimer HAL driver of ROM implementation.
- Enabling [CONFIG_HAL_WDT_USE_ROM_IMPL](#) can reduce the IRAM usage and binary size by linking in the watchdog HAL driver of ROM implementation.

Heap

- Enabling [CONFIG_HEAP_PLACE_FUNCTION_INTO_FLASH](#) can reduce the IRAM usage and binary size by placing the entirety of the heap functionalities in flash memory.

Bootloader Size This document deals with the size of an ESP-IDF app binary only, and not the ESP-IDF [Second Stage Bootloader](#).

For a discussion of ESP-IDF bootloader binary size, see [Bootloader Size](#).

IRAM Binary Size If the IRAM section of a binary is too large, this issue can be resolved by reducing IRAM memory usage. See [Optimizing IRAM Usage](#).

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

Minimizing RAM Usage

In some cases, a firmware application's available RAM may run low or run out entirely. In these cases, it is necessary to tune the memory usage of the firmware application.

In general, firmware should aim to leave some headroom of free internal RAM to deal with extraordinary situations or changes in RAM usage in future updates.

Background Before optimizing ESP-IDF RAM usage, it is necessary to understand the basics of ESP32-P4 memory types, the difference between static and dynamic memory usage in C, and the way ESP-IDF uses stack and heap. This information can all be found in [Heap Memory Allocation](#).

Measuring Static Memory Usage The `idf.py` tool can be used to generate reports about the static memory usage of an application, see [Measuring Static Sizes](#).

Measuring Dynamic Memory Usage ESP-IDF contains a range of heap APIs for measuring free heap at runtime, see [Heap Memory Debugging](#).

Note: In embedded systems, heap fragmentation can be a significant issue alongside total RAM usage. The heap measurement APIs provide ways to measure the largest free block. Monitoring this value along with the total number of free bytes can give a quick indication of whether heap fragmentation is becoming an issue.

Reducing Static Memory Usage

- Reducing the static memory usage of the application increases the amount of RAM available for heap at runtime, and vice versa.
- Generally speaking, minimizing static memory usage requires monitoring the `.data` and `.bss` sizes. For tools to do this, see [Measuring Static Sizes](#).
- Internal ESP-IDF functions do not make heavy use of static RAM in C. In many instances (such as Wi-Fi library, Bluetooth controller), static buffers are still allocated from the heap. However, the allocation is performed only once during feature initialization and will be freed if the feature is deinitialized. This approach is adopted to optimize the availability of free memory at various stages of the application's life cycle.

To minimize static memory use:

- Constant data can be stored in flash memory instead of RAM, thus it is recommended to declare structures, buffers, or other variables as `const`. This approach may require modifying firmware functions to accept `const *` arguments instead of mutable pointer arguments. These changes can also help reduce the stack usage of certain functions.
- If using OpenThread, enabling the option [CONFIG_OPENTHREAD_PLATFORM_MSGPOOL_MANAGEMENT](#) will cause OpenThread to allocate message pool buffers from PSRAM, which will reduce static memory use.

Reducing Stack Sizes In FreeRTOS, task stacks are usually allocated from the heap. The stack size for each task is fixed and passed as an argument to `xTaskCreate()`. Each task can use up to its allocated stack size, but using more than this will cause an otherwise valid program to crash, with a stack overflow or heap corruption.

Therefore, determining the optimum sizes of each task stack, minimizing the required size of each task stack, and minimizing the number of task stacks as whole, can all substantially reduce RAM usage.

To determine the optimum size for a particular task stack, users can consider the following methods:

- At runtime, call the function `uxTaskGetStackHighWaterMark()` with the handle of any task where you think there is unused stack memory. This function returns the minimum lifetime free stack memory in bytes.
 - The easiest time to call `uxTaskGetStackHighWaterMark()` is from the task itself: call `uxTaskGetStackHighWaterMark(NULL)` to get the current task's high water mark after the time that the task has achieved its peak stack usage, i.e., if there is a main loop, execute the main loop a number of times with all possible states, and then call `uxTaskGetStackHighWaterMark()`.
 - Often, it is possible to subtract almost the entire value returned here from the total stack size of a task, but allow some safety margin to account for unexpected small increases in stack usage at runtime.
- Call `uxTaskGetSystemState()` at runtime to get a summary of all tasks in the system. This includes their individual stack high watermark values.
- When debugger watchpoints are not being used, users can set the [CONFIG_FREERTOS_WATCHPOINT_END_OF_STACK](#) option. This will cause one of the watchpoints to watch the last word of the task's stack. If that word is overwritten (such as in a stack overflow), a panic is triggered immediately. This is slightly more reliable than the default [CONFIG_FREERTOS_CHECK_STACKOVERFLOW](#) option of Check using canary bytes, because the panic happens immediately, rather than on the next RTOS context switch. Neither option is perfect. In some cases, it is possible that the stack pointer skips the watchpoint or canary bytes and corrupts another region of RAM instead.

To reduce the required size of a particular task stack, users can consider the following methods:

- Avoid stack heavy functions. String formatting functions (like `printf()`) are particularly heavy users of the stack, so any task which does not ever call these can usually have its stack size reduced.
 - Enabling *Newlib Nano Formatting* reduces the stack usage of any task that calls `printf()` or other C string formatting functions.
- Avoid allocating large variables on the stack. In C, any large structures or arrays allocated as an automatic variable (i.e., default scope of a C declaration) uses space on the stack. To minimize the sizes of these, allocate them statically and/or see if you can save memory by dynamically allocating them from the heap only when they are needed.
- Avoid deep recursive function calls. Individual recursive function calls do not always add a lot of stack usage each time they are called, but if each function includes large stack-based variables then the overhead can get quite high.

To reduce the total number of tasks, users can consider the following method:

- Combine tasks. If a particular task is never created, the task's stack is never allocated, thus reducing RAM usage significantly. Unnecessary tasks can typically be removed if those tasks can be combined with another task. In an application, tasks can typically be combined or removed if:
 - The work done by the tasks can be structured into multiple functions that are called sequentially.
 - The work done by the tasks can be structured into smaller jobs that are serialized (via a FreeRTOS queue or similar) for execution by a worker task.

Internal Task Stack Sizes ESP-IDF allocates a number of internal tasks for housekeeping purposes or operating system functions. Some are created during the startup process, and some are created at runtime when particular features are initialized.

The default stack sizes for these tasks are usually set conservatively high to allow all common usage patterns. Many of the stack sizes are configurable, and it may be possible to reduce them to match the real runtime stack usage of the task.

Important: If internal task stack sizes are set too small, ESP-IDF will crash unpredictably. Even if the root cause is task stack overflow, this is not always clear when debugging. It is recommended that internal stack sizes are only reduced carefully (if at all), with close attention to high water mark free space under load. If reporting an issue that occurs when internal task stack sizes have been reduced, please always include the following information and the specific configuration that is being used.

- *Running the Main Task* has stack size `CONFIG_ESP_MAIN_TASK_STACK_SIZE`.
- *High Resolution Timer (ESP Timer)* system task which executes callbacks has stack size `CONFIG_ESP_TIMER_TASK_STACK_SIZE`.
- FreeRTOS Timer Task to handle FreeRTOS timer callbacks has stack size `CONFIG_FREERTOS_TIMER_TASK_STACK_DEPTH`.
- *Event Loop Library* system task to execute callbacks for the default system event loop has stack size `CONFIG_ESP_SYSTEM_EVENT_TASK_STACK_SIZE`.
- *lwIP* TCP/IP task has stack size `CONFIG_LWIP_TCPIP_TASK_STACK_SIZE`.
- The Ethernet driver creates a task for the MAC to receive Ethernet frames. If using the default config `ETH_MAC_DEFAULT_CONFIG` then the task stack size is 4 KB. This setting can be changed by passing a custom `eth_mac_config_t` struct when initializing the Ethernet MAC.
- FreeRTOS idle task stack size is configured by `CONFIG_FREERTOS_IDLE_TASK_STACKSIZE`.
- If using the *ESP-MQTT* component, it creates a task with stack size configured by `CONFIG_MQTT_TASK_STACK_SIZE`. MQTT stack size can also be configured using `task_stack` field of `esp_mqtt_client_config_t`.
- To see how to optimize RAM usage when using mDNS, please check [Minimizing RAM Usage](#).

Note: Aside from built-in system features such as ESP-timer, if an ESP-IDF feature is not initialized by the firmware, then no associated task is created. In those cases, the stack usage is zero, and the stack-size configuration for the task is not relevant.

Reducing Heap Usage For functions that assist in analyzing heap usage at runtime, see [Heap Memory Debugging](#).

Normally, optimizing heap usage consists of analyzing the usage and removing calls to `malloc()` that are not being used, reducing the corresponding sizes, or freeing previously allocated buffers earlier.

There are some ESP-IDF configuration options that can reduce heap usage at runtime:

- lwIP documentation has a section to configure [Minimum RAM Usage](#).
- Several Mbed TLS configuration options can be used to reduce heap memory usage. See the [Reducing Heap Usage](#) docs for details.

Note: There are other configuration options that increases heap usage at runtime if changed from the defaults. These options are not listed above, but the help text for the configuration item will mention if there is some memory impact.

Optimizing IRAM Usage The available DRAM at runtime for heap usage is also reduced by the static IRAM usage. Therefore, one way to increase available DRAM is to reduce IRAM usage.

If the app allocates more static IRAM than available, then the app will fail to build, and linker errors such as `section '.iram0.text' will not fit in region 'iram0_0_seg', IRAM0 segment data does not fit, and region 'iram0_0_seg' overflowed by 84-bytes` will be seen. If this happens, it is necessary to find ways to reduce static IRAM usage in order to link the application.

To analyze the IRAM usage in the firmware binary, use [Measuring Static Sizes](#). If the firmware failed to link, steps to analyze are shown at [Showing Size When Linker Fails](#).

The following options will reduce IRAM usage of some ESP-IDF features:

- Enable [CONFIG_FREERTOS_PLACE_FUNCTIONS_INTO_FLASH](#). Provided these functions are not incorrectly used from ISRs, this option is safe to enable in all configurations.
- Enable [CONFIG_RINGBUF_PLACE_FUNCTIONS_INTO_FLASH](#). Provided these functions are not incorrectly used from ISRs, this option is safe to enable in all configurations.
- Enable [CONFIG_RINGBUF_PLACE_ISR_FUNCTIONS_INTO_FLASH](#). This option is not safe to use if the ISR ringbuf functions are used from an IRAM interrupt context, e.g., if [CONFIG_UART_ISR_IN_IRAM](#) is enabled. For the ESP-IDF drivers where this is the case, you can get an error at run-time when installing the driver in question.
- Disabling [CONFIG_ESP_EVENT_POST_FROM_IRAM_ISR](#) prevents posting `esp_event` events from [IRAM-Safe Interrupt Handlers](#) but saves some IRAM.
- Disabling [CONFIG_SPI_MASTER_ISR_IN_IRAM](#) prevents `spi_master` interrupts from being serviced while writing to flash, and may otherwise reduce `spi_master` performance, but saves some IRAM.
- Disabling [CONFIG_SPI_SLAVE_ISR_IN_IRAM](#) prevents `spi_slave` interrupts from being serviced while writing to flash, which saves some IRAM.
- Setting [CONFIG_HAL_DEFAULT_ASSERTION_LEVEL](#) to disable assertion for HAL component saves some IRAM, especially for HAL code who calls `HAL_ASSERT` a lot and resides in IRAM.
- Refer to the `sdkconfig` menu `Auto-detect Flash chips`, and you can disable flash drivers which you do not need to save some IRAM.
- Enable [CONFIG_HEAP_PLACE_FUNCTION_INTO_FLASH](#). Provided that [CONFIG_SPI_MASTER_ISR_IN_IRAM](#) is not enabled and the heap functions are not incorrectly used from ISRs, this option is safe to enable in all configurations.

Note: Moving frequently-called functions from IRAM to flash may increase their execution time.

Note: Other configuration options exist that will increase IRAM usage by moving some functionality into IRAM, usually for performance, but the default option is not to do this. These are not listed here. The IRAM size impact of

enabling these options is usually noted in the configuration item help text.

4.20 Reproducible Builds

4.20.1 Introduction

ESP-IDF build system has support for [reproducible builds](#).

When reproducible builds are enabled, the application built with ESP-IDF does not depend on the build environment. Both the .elf file and .bin files of the application remains exactly the same, even if the following variables change:

- Directory where the project is located
- Directory where ESP-IDF is located (`IDF_PATH`)
- Build time

4.20.2 Reasons for Non-Reproducible Builds

There are several reasons why an application may depend on the build environment, even when the same source code and tools versions are used.

- In C code, `__FILE__` preprocessor macro is expanded to the full path of the source file.
- `__DATE__` and `__TIME__` preprocessor macros are expanded to compilation date and time.
- When the compiler generates object files, it adds sections with debug information. These sections help debuggers, like GDB, to locate the source code which corresponds to a particular location in the machine code. These sections typically contain paths of relevant source files. These paths may be absolute, and will include the path to ESP-IDF or to the project.

There are also other possible reasons, such as unstable order of inputs and non-determinism in the build system.

4.20.3 Enabling Reproducible Builds in ESP-IDF

Reproducible builds can be enabled in ESP-IDF using `CONFIG_APP_REPRODUCIBLE_BUILD` option.

This option is disabled by default. It can be enabled in `menuconfig`.

The option may also be added into `sdkconfig.defaults`. If adding the option into `sdkconfig.defaults`, delete the `sdkconfig` file and run the build again. See [Custom Sdkconfig Defaults](#) for more information.

4.20.4 How Reproducible Builds Are Achieved

ESP-IDF achieves reproducible builds using the following measures:

- In ESP-IDF source code, `__DATE__` and `__TIME__` macros are not used when reproducible builds are enabled. Note, if the application source code uses these macros, the build will not be reproducible.
- ESP-IDF build system passes a set of `-fmacro-prefix-map` and `-fdebug-prefix-map` flags to replace base paths with placeholders:
 - Path to ESP-IDF is replaced with `/IDF`
 - Path to the project is replaced with `/IDF_PROJECT`
 - Path to the build directory is replaced with `/IDF_BUILD`
 - Paths to components are replaced with `/COMPONENT_NAME_DIR` (where NAME is the name of the component)
- Build date and time are not included into the *application metadata structure* and *bootloader metadata structure* if `CONFIG_APP_REPRODUCIBLE_BUILD` is enabled.
- ESP-IDF build system ensures that source file lists, component lists and other sequences are sorted before passing them to CMake. Various other parts of the build system, such as the linker script generator also perform sorting to ensure that same output is produced regardless of the environment.

4.20.5 Reproducible Builds and Debugging

When reproducible builds are enabled, file names included in debug information sections are altered as shown in the previous section. Due to this fact, the debugger (GDB) is not able to locate the source files for the given code location.

This issue can be solved using GDB `set substitute-path` command. For example, by adding the following command to GDB init script, the altered paths can be reverted to the original ones:

```
set substitute-path /COMPONENT_FREERTOS_DIR /home/user/esp/esp-idf/components/  
↪freertos
```

ESP-IDF build system generates a file with the list of such `set substitute-path` commands automatically during the build process. The file is called `prefix_map` and is located in the project `build/gdbinit` directory.

When `idf.py gdb` is used to start debugging, this additional `gdbinit` file is automatically passed to GDB. When launching GDB manually or from an IDE, please pass this additional `gdbinit` script to GDB using `-x build/prefix_map_gdbinit` argument.

4.20.6 Factors Which Still Affect Reproducible Builds

Note that the built application still depends on:

- ESP-IDF version
- Versions of the build tools (CMake, Ninja) and the cross-compiler

IDF Docker Image can be used to ensure that these factors do not affect the build.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.21 Thread Local Storage

4.21.1 Overview

Thread-local storage (TLS) is a mechanism by which variables are allocated such that there is one instance of the variable per extant thread. ESP-IDF provides three ways to make use of such variables:

- *FreeRTOS Native APIs*: ESP-IDF FreeRTOS native APIs.
- *Pthread APIs*: ESP-IDF pthread APIs.
- *C11 Standard*: C11 standard introduces special keywords to declare variables as thread local.

4.21.2 FreeRTOS Native APIs

The ESP-IDF FreeRTOS provides the following APIs to manage thread local variables:

- `vTaskSetThreadLocalStoragePointer()`
- `pvTaskGetThreadLocalStoragePointer()`
- `vTaskSetThreadLocalStoragePointerAndDelCallback()`

In this case, the maximum number of variables that can be allocated is limited by `CONFIG_FREERTOS_THREAD_LOCAL_STORAGE_POINTERS`. Variables are kept in the task control block (TCB) and accessed by their index. Note that index 0 is reserved for ESP-IDF internal uses.

Using the APIs above, you can allocate thread local variables of an arbitrary size, and assign them to any number of tasks. Different tasks can have different sets of TLS variables.

If size of the variable is more than 4 bytes, then you need to allocate/deallocate memory for it. Variable's deallocation is initiated by FreeRTOS when task is deleted, but user must provide callback function to do proper cleanup.

4.21.3 Pthread APIs

The ESP-IDF provides the following *POSIX Threads Support* to manage thread local variables:

- `pthread_key_create()`
- `pthread_key_delete()`
- `pthread_getspecific()`
- `pthread_setspecific()`

These APIs have all benefits of the ones above, but eliminates some their limits. The number of variables is limited only by size of available memory on the heap. Due to the dynamic nature, this APIs introduce additional performance overhead compared to the native one.

4.21.4 C11 Standard

The ESP-IDF FreeRTOS supports thread local variables according to C11 standard, ones specified with `__thread` keyword. For details on this feature, please refer to the [GCC documentation](#).

Storage for that kind of variables is allocated on the task stack. Note that area for all such variables in the program is allocated on the stack of every task in the system even if that task does not use such variables at all. For example, ESP-IDF system tasks (e.g., `ipc`, `timer` tasks etc.) will also have that extra stack space allocated. Thus feature should be used with care.

Using C11 thread local variables comes at a trade-off. One one hand, they are quite handy to use in programming and can be accessed using minimal CPU instructions. However, this benefit comes at the cost of additional stack usage for all tasks in the system. Due to static nature of variables allocation, all tasks in the system have the same sets of C11 thread local variables.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.22 Tools

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.22.1 IDF Frontend - `idf.py`

The `idf.py` command-line tool provides a front-end for easily managing your project builds, deployment and debugging, and more. It manages several tools, for example:

- [CMake](#), which configures the project to be built.
- [Ninja](#), which builds the project.
- [esptool.py](#), which flashes the target.

The [Step 5. First Steps on ESP-IDF](#) contains a brief introduction on how to set up `idf.py` to configure, build, and flash projects.

Important: `idf.py` should be run in an ESP-IDF project directory, i.e., a directory containing a `CMakeLists.txt` file. Older style projects that contain a `Makefile` will not work with `idf.py`.

Commands

Start a New Project: `create-project`

```
idf.py create-project <project name>
```

This command creates a new ESP-IDF project. Additionally, the folder where the project will be created in can be specified by the `--path` option.

Create a New Component: `create-component`

```
idf.py create-component <component name>
```

This command creates a new component, which will have a minimum set of files necessary for building. The `-C` option can be used to specify the directory the component will be created in. For more information about components see the [Component CMakeLists Files](#).

Select the Target Chip: `set-target` ESP-IDF supports multiple targets (chips). A full list of supported targets in your version of ESP-IDF can be seen by running `idf.py --list-targets`.

```
idf.py set-target <target>
```

This command sets the current project target.

Important: `idf.py set-target` will clear the build directory and re-generate the `sdkconfig` file from scratch. The old `sdkconfig` file will be saved as `sdkconfig.old`.

Note: The behavior of the `idf.py set-target` command is equivalent to:

1. clearing the build directory (`idf.py fullclean`)
 2. removing the `sdkconfig` file (`mv sdkconfig sdkconfig.old`)
 3. configuring the project with the new target (`idf.py -DIDF_TARGET=esp32 reconfigure`)
-

It is also possible to pass the desired `IDF_TARGET` as an environment variable (e.g., `export IDF_TARGET=esp32s2`) or as a CMake variable (e.g., `-DIDF_TARGET=esp32s2` argument to CMake or `idf.py`). Setting the environment variable is a convenient method if you mostly work with one type of the chip.

To specify the default value of `IDF_TARGET` for a given project, please add the `CONFIG_IDF_TARGET` option to the project's `sdkconfig.defaults` file, e.g., `CONFIG_IDF_TARGET="esp32s2"`. This value of the option will be used if `IDF_TARGET` is not specified by other methods, such as using an environment variable, a CMake variable, or the `idf.py set-target` command.

If the target has not been set by any of these methods, the build system will default to `esp32` target.

Start the Graphical Configuration Tool: `menuconfig`

```
idf.py menuconfig
```

Build the Project: `build`

```
idf.py build
```

This command builds the project found in the current directory. This can involve multiple steps:

- Create the build directory if needed. The sub-directory "build" is used to hold build output, although this can be changed with the `-B` option.
- Run [CMake](#) as necessary to configure the project and generate build files for the main build tool.
- Run the main build tool ([Ninja](#) or *GNU Make*). By default, the build tool is automatically detected but it can be explicitly set by passing the `-G` option to `idf.py`.

Building is incremental, so if no source files or configuration has changed since the last build, nothing will be done.

Additionally, the command can be run with `app`, `bootloader` and `partition-table` arguments to build only the app, bootloader or partition table as applicable.

Remove the Build Output: `clean`

```
idf.py clean
```

This command removes the project build output files from the build directory, and the project will be fully rebuilt on next build. Using this command does not remove the CMake configuration output inside the build folder.

Delete the Entire Build Contents: `fullclean`

```
idf.py fullclean
```

This command deletes the entire "build" directory contents, which includes all CMake configuration output. The next time the project is built, CMake will configure it from scratch. Note that this option recursively deletes **all** files in the build directory, so use with care. Project configuration is not deleted.

Flash the Project: `flash`

```
idf.py flash
```

This command automatically builds the project if necessary, and then flash it to the target. You can use `-p` and `-b` options to set serial port name and flasher baud rate, respectively.

Note: The environment variables `ESPPORT` and `ESPBAUD` can be used to set default values for the `-p` and `-b` options, respectively. Providing these options on the command line overrides the default.

`idf.py` uses the `write_flash` command of `esptool.py` under the hood to flash the target. You can pass additional arguments to configure the flash writing process using the `--extra-args` option. For example, to [write to an external SPI flash chip](#), use the following command: `idf.py flash --extra-args="--spi-connection <CLK>, <Q>, <D>, <HD>, <CS>"`. To see the full list of available arguments, run `esptool.py write_flash --help` or see the [esptool.py documentation](#).

Similarly to the `build` command, the command can be run with `app`, `bootloader` and `partition-table` arguments to flash only the app, bootloader or partition table as applicable.

Hints on How to Resolve Errors

`idf.py` will try to suggest hints on how to resolve errors. It works with a database of hints stored in [tools/idf_py_actions/hints.yml](#) and the hints will be printed if a match is found for the given error. The menuconfig target is not supported at the moment by automatic hints on resolving errors.

The `--no-hints` argument of `idf.py` can be used to turn the hints off in case they are not desired.

Important Notes

Multiple `idf.py` commands can be combined into one. For example, `idf.py -p COM4 clean flash monitor` will clean the source tree, then build the project and flash it to the target before running the serial monitor.

The order of multiple `idf.py` commands on the same invocation is not important, as they will automatically be executed in the correct order for everything to take effect (e.g., building before flashing, erasing before flashing).

For commands that are not known to `idf.py`, an attempt to execute them as a build system target will be made.

The command `idf.py` supports [shell autocompletion](#) for bash, zsh and fish shells.

In order to make [shell autocompletion](#) supported, please make sure you have at least Python 3.5 and [click 7.1](#) or newer ([Software](#)).

To enable autocompletion for `idf.py`, use the `export` command ([Step 4. Set up the environment variables](#)). Autocompletion is initiated by pressing the TAB key. Type `idf.py -` and press the TAB key to autocomplete options.

The autocomplete support for PowerShell is planned in the future.

Advanced Commands

Open the Documentation: `docs`

```
idf.py docs
```

This command opens the documentation for the projects target and ESP-IDF version in the browser.

Show Size: `size`

```
idf.py size
```

This command prints app size information including the occupied RAM and flash and section (i.e., .bss) sizes.

```
idf.py size-components
```

Similarly, this command prints the same information for each component used in the project.

```
idf.py size-files
```

This command prints size information per source file in the project.

Options

- `--format` specifies the output format with available options: `text`, `csv`, `json`, default being `text`.
- `--output-file` optionally specifies the name of the file to print the command output to instead of the standard output.

Reconfigure the Project: `reconfigure`

```
idf.py reconfigure
```

This command forces [CMake](#) to be rerun regardless of whether it is necessary. It is unnecessary during normal usage, but can be useful after adding/removing files from the source tree, or when modifying CMake cache variables. For example, `idf.py -DNAME='VALUE' reconfigure` can be used to set variable `NAME` in CMake cache to value `VALUE`.

Clean the Python Byte Code: `python-clean`


```
idf.py python-clean
```

This command deletes generated python byte code from the ESP-IDF directory. The byte code may cause issues when switching between ESP-IDF and Python versions. It is advised to run this target after switching versions of Python.

Generate a UF2 Binary: `uf2`

```
idf.py uf2
```

This command generates a UF2 ([USB Flashing Format](#)) binary `uf2.bin` in the build directory. This file includes all the necessary binaries (bootloader, app, and partition table) for flashing the target.

This UF2 file can be copied to a USB mass storage device exposed by another ESP running the [ESP USB Bridge](#) project. The bridge MCU will use it to flash the target MCU. This is as simple as copying (or "drag-and-dropping") the file to the exposed disk accessed by a file explorer in your machine.

To generate a UF2 binary for the application only (not including the bootloader and partition table), use the `uf2-app` command.

```
idf.py uf2-app
```

Global Options

To list all available root level options, run `idf.py --help`. To list options that are specific for a subcommand, run `idf.py <command> --help`, e.g., `idf.py monitor --help`. Here is a list of some useful options:

- `-C <dir>` allows overriding the project directory from the default current working directory.
- `-B <dir>` allows overriding the build directory from the default "build" subdirectory of the project directory.
- `--ccache` enables [CCache](#) when compiling source files if the [CCache](#) tool is installed. This can dramatically reduce the build time.

Important: Note that some older versions of [CCache](#) may exhibit bugs on some platforms, so if files are not rebuilt as expected, try disabling [CCache](#) and rebuilding the project. To enable [CCache](#) by default, set the `IDF_CCACHE_ENABLE` environment variable to a non-zero value.

- `-v` flag causes both `idf.py` and the build system to produce verbose build output. This can be useful for debugging build problems.
- `--cmake-warn-uninitialized` (or `-w`) causes CMake to print uninitialized variable warnings found in the project directory only. This only controls CMake variable warnings inside CMake itself, not other types of build warnings. This option can also be set permanently by setting the `IDF_CMAKE_WARN_UNINITIALIZED` environment variable to a non-zero value.
- `--no-hints` flag disables hints on resolving errors and disable capturing output.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.22.2 IDF Docker Image

IDF Docker image (`espressif/idf`) is intended for building applications and libraries with specific versions of ESP-IDF when doing automated builds.

The image contains:

- Common utilities such as `git`, `wget`, `curl`, and `zip`.
- Python 3.8 or newer.
- A copy of a specific version of ESP-IDF. See below for information about versions. `IDF_PATH` environment variable is set and points to the ESP-IDF location in the container.
- All the build tools required for the specific version of ESP-IDF: CMake, Ninja, cross-compiler toolchains, etc.
- All Python packages required by ESP-IDF are installed in a virtual environment.

The image `ENTRYPOINT` sets up the `PATH` environment variable to point to the correct version of tools, and activates the Python virtual environment. As a result, the environment is ready to use the ESP-IDF build system.

The image can also be used as a base for custom images, if additional utilities are required.

Tags

Multiple tags of this image are maintained:

- `latest`: tracks master branch of ESP-IDF
- `vX.Y`: corresponds to ESP-IDF release `vX.Y`
- `release-vX.Y`: tracks `release/vX.Y` branch of ESP-IDF

Note: Versions of ESP-IDF released before this feature was introduced do not have corresponding Docker image versions. You can check the up-to-date list of available tags at <https://hub.docker.com/r/espressif/idf/tags>.

Usage

Setting up Docker Before using the `espressif/idf` Docker image locally, make sure you have Docker installed. Follow the instructions at <https://docs.docker.com/install/>, if it is not installed yet.

If using the image in a CI environment, consult the documentation of your CI service on how to specify the image used for the build process.

Building a Project with CMake In the project directory, run:

```
docker run --rm -v $PWD:/project -w /project -u $UID -e HOME=/tmp espressif/idf_
↳idf.py build
```

The above command explained:

- `docker run`: runs a Docker image. It is a shorter form of the command `docker container run`.
- `--rm`: removes the container when the build is finished.
- `-v $PWD:/project`: mounts the current directory on the host (`$PWD`) as `/project` directory in the container.
- `-w /project`: makes `/project` the working directory for the command.
- `-u $UID`: makes the command run with your user ID so that files are created as you (instead of root).
- `-e HOME=/tmp`: gives the user a home directory for storing temporary files created by `idf.py` in `~/ .cache`.
- `espressif/idf`: uses Docker image `espressif/idf` with tag `latest`. The `latest` tag is implicitly added by Docker when no tag is specified.
- `idf.py build`: runs this command inside the container.

Note: When the mounted directory, `/project`, contains a git repository owned by a different user (UID) than the one running the Docker container, git commands executed within `/project` might fail, displaying an error message `fatal: detected dubious ownership in repository at '/project'`. To resolve this issue, you can designate the `/project` directory as safe by setting the `IDF_GIT_SAFE_DIR` environment variable

during the Docker container startup. For instance, you can achieve this by including `-e IDF_GIT_SAFE_DIR=' /project '` as a parameter. Additionally, multiple directories can be specified by using a `:` separator. To entirely disable this git security check, `*` can be used.

To build with a specific Docker image tag, specify it as `espressif/idf:TAG`, for example:

```
docker run --rm -v $PWD:/project -w /project -u $UID -e HOME=/tmp espressif/
↳idf:release-v4.4 idf.py build
```

You can check the up-to-date list of available tags at <https://hub.docker.com/r/espressif/idf/tags>.

Using the Image Interactively It is also possible to do builds interactively, to debug build issues or test the automated build scripts. Start the container with `-i -t` flags:

```
docker run --rm -v $PWD:/project -w /project -u $UID -e HOME=/tmp -it espressif/idf
```

Then inside the container, use `idf.py` as usual:

```
idf.py menuconfig
idf.py build
```

Note: Commands which communicate with the development board, such as `idf.py flash` and `idf.py monitor` does not work in the container, unless the serial port is passed through into the container. This can be done with Docker for Linux with the [device option](#). However, currently, this is not possible with Docker for Windows (<https://github.com/docker/for-win/issues/1018>) and Docker for Mac (<https://github.com/docker/for-mac/issues/900>). This limitation may be overcome by using [remote serial ports](#). An example of how to do this can be found in the following [using remote serial port](#) section.

Using Remote Serial Port The [RFC2217](#) (Telnet) protocol can be used to remotely connect to a serial port. For more information please see the [remote serial ports](#) documentation in the ESP tool project. This method can also be used to access the serial port inside a Docker container if it cannot be accessed directly. Following is an example of how to use the Flash command from within a Docker container.

On host install and start `esp_rfc2217_server`:

- On Windows, the package is available as a one-file bundled executable created by `pyinstaller` and it can be downloaded from the [esptool releases](#) page in a ZIP archive along with other ESP tool utilities:

```
esp_rfc2217_server -v -p 4000 COM3
```

- On Linux or macOS, the package is available as part of `esptool`, which can be found in the ESP-IDF environment or by installing using `pip`:

```
pip install esptool
```

And then starting the server by executing

```
esp_rfc2217_server.py -v -p 4000 /dev/ttyUSB0
```

Now the device attached to the host can be flashed from inside a Docker container by using:

```
docker run --rm -v <host_path>:/<container_path> -w /<container_path> espressif/
↳idf idf.py --port 'rfc2217://host.docker.internal:4000?ign_set_control' flash
```

Please make sure that `<host_path>` is properly set to your project path on the host, and `<container_path>` is set as a working directory inside the container with the `-w` option. The `host.docker.internal` is a special Docker DNS name to access the host. This can be replaced with a host IP if necessary.

Building Custom Images

The Docker file in ESP-IDF repository provides several build arguments which can be used to customize the Docker image:

- `IDF_CLONE_URL`: URL of the repository to clone ESP-IDF from. Can be set to a custom URL when working with a fork of ESP-IDF. The default is `https://github.com/espressif/esp-idf.git`.
- `IDF_CLONE_BRANCH_OR_TAG`: Name of a git branch or tag used when cloning ESP-IDF. This value is passed to the `git clone` command using the `--branch` argument. The default is `master`.
- `IDF_CHECKOUT_REF`: If this argument is set to a non-empty value, `git checkout $IDF_CHECKOUT_REF` command performs after cloning. This argument can be set to the SHA of the specific commit to check out, for example, if some specific commit on a release branch is desired.
- `IDF_CLONE_SHALLOW`: If this argument is set to a non-empty value, `--depth=1 --shallow-submodules` arguments are used when performing `git clone`. Depth can be customized using `IDF_CLONE_SHALLOW_DEPTH`. Doing a shallow clone significantly reduces the amount of data downloaded and the size of the resulting Docker image. However, if switching to a different branch in such a "shallow" repository is necessary, an additional `git fetch origin <branch>` command must be executed first.
- `IDF_CLONE_SHALLOW_DEPTH`: This argument specifies the depth value to use when doing a shallow clone. If not set, `--depth=1` will be used. This argument has effect only if `IDF_CLONE_SHALLOW` is used. Use this argument if you are building a Docker image for a branch, and the image has to contain the latest tag on that branch. To determine the required depth, run `git describe` for the given branch and note the offset number. Increment it by 1, then use it as the value of this argument. The resulting image will contain the latest tag on the branch, and consequently `git describe` command inside the Docker image will work as expected.
- `IDF_INSTALL_TARGETS`: Comma-separated list of ESP-IDF targets to install toolchains for, or `all` to install toolchains for all targets. Selecting specific targets reduces the amount of data downloaded and the size of the resulting Docker image. The default is `all`.
- `IDF_GITHUB_ASSETS`: Hostname used to download ESP-IDF tools from a mirror server instead of directly from `github.com`. This is useful when `github.com` is not accessible from your network. The default value is empty.

To use these arguments, pass them via the `--build-arg` command line option. For example, the following command builds a Docker image with a shallow clone of ESP-IDF v4.4.1 and tools for ESP32-C3 only:

```
docker build -t idf-custom:v4.4.1-esp32c3 \
  --build-arg IDF_CLONE_BRANCH_OR_TAG=v4.4.1 \
  --build-arg IDF_CLONE_SHALLOW=1 \
  --build-arg IDF_INSTALL_TARGETS=esp32c3 \
  tools/docker
```

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.22.3 IDF Windows Installer

Command-Line Parameters

Windows Installer `esp-idf-tools-setup` provides the following command-line parameters:

- `/CONFIG=[PATH]` - Path to `ini` configuration file to override default configuration of the installer. Default: `config.ini`.
- `/GITCLEAN=[yes|no]` - Perform `git clean` and remove untracked directories in offline-mode installation. Default: `yes`.

- /GITRECURSIVE=[yes|no] - Clone recursively all Git repository submodules. Default: yes.
- /GITREPO=[URL|PATH] - URL of repository to clone ESP-IDF. Default: <https://github.com/espressif/esp-idf.git>.
- /GITRESET=[yes|no] - Enable/Disable git reset of repository during installation. Default: yes.
- /HELP - Display command line options provided by Inno Setup installer.
- /IDFDIR=[PATH] - Path to directory where it is installed. Default: {userdesktop}\esp-idf.
- /IDFVERSION=[v4.3|v4.1|master] - Use specific ESP-IDF version. E.g., v4.1, v4.2, master. Default: empty, pick the first version in the list.
- /IDFVERSIONSURL=[URL] - Use URL to download list of ESP-IDF versions. Default: https://dl.espressif.com/dl/esp-idf/idf_versions.txt.
- /LOG=[PATH] - Store installation log file in specific directory. Default: empty.
- /OFFLINE=[yes|no] - Execute installation of Python packages by pip in offline mode. The same result can be achieved by setting the environment variable PIP_NO_INDEX. Default: no.
- /USEEMBEDDEDPYTHON=[yes|no] - Use Embedded Python version for the installation. Set to no to allow the Python selection screen in the installer. Default: yes.
- /PYTHONNOUSERSITE=[yes|no] - Set PYTHONNOUSERSITE variable before launching any Python command to avoid loading Python packages from AppDataRoaming. Default: yes.
- /PYTHONWHEELSURL=[URL] - Specify URLs to PyPi repositories for resolving binary Python Wheel dependencies. The same result can be achieved by setting the environment variable PIP_EXTRA_INDEX_URL. Default: <https://dl.espressif.com/pypi>.
- /SKIPSYSTEMCHECK=[yes|no] - Skip System Check page. Default: no.
- /VERYSILENT /SUPPRESSMSGBOXES /SP- /NOCANCEL - Perform silent installation.

Unattended Installation

The unattended installation of ESP-IDF can be achieved by following command-line parameters:

```
esp-idf-tools-setup-x.x.exe /VERYSILENT /SUPPRESSMSGBOXES /SP- /NOCANCEL
```

When running the installer from the command line, it detaches its process from the command line and starts a separate process in the background to perform the installation without blocking the use of the command line. The following PowerShell script allows you to wait for the installer to complete:

```
esp-idf-tools-setup-x.x.exe /VERYSILENT /SUPPRESSMSGBOXES /SP- /NOCANCEL
$InstallerProcess = Get-Process esp-idf-tools-setup
Wait-Process -Id $InstallerProcess.id
```

Custom Python and Custom Location of Python Wheels

The IDF installer is using by default embedded Python with reference to the Python Wheel mirror.

The following parameters allow to select custom Python and custom location of Python wheels:

```
esp-idf-tools-setup-x.x.exe /USEEMBEDDEDPYTHON=no /PYTHONWHEELSURL=https://pypi.
↳org/simple/
```

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.22.4 IDF Component Manager

The IDF Component Manager is a tool that downloads dependencies for any ESP-IDF CMake project. The download happens automatically during a run of CMake. It can source components either from the [component registry](#) or from a Git repository.

A list of components can be found on <https://components.espressif.com/>.

For detailed information about the IDF Component Manager, see the [IDF Component Manager and ESP Component Registry Documentation](#).

Using with a Project

Dependencies for each component in the project are defined in a separate manifest file named `idf_component.yml` placed in the root of the component. The manifest file template can be created by running `idf.py create-manifest`. By default, a manifest file is created for the main component. You can explicitly either specify the directory where the manifest should be created using the `--path` option or specify the component in your `components` folder using `--component=my_component`. The `create-manifest` command can be run in the following ways:

- `idf.py create-manifest` creates a manifest file for the main component
- `idf.py create-manifest --component=my_component` creates a manifest file for the component **my_component** in the `components` directory
- `idf.py create-manifest --path="../../my_component"` creates a manifest file for the component **my_component** in the `my_component` directory

When a new manifest is added to one of the components in the project, it is necessary to reconfigure the project manually by running `idf.py reconfigure`. The build will then track changes in `idf_component.yml` manifests and automatically trigger CMake when necessary.

To add a dependency to a component (e.g., `my_component`) in your ESP-IDF project, you can run the command `idf.py add-dependency DEPENDENCY`. The `DEPENDENCY` argument represents an additional component managed by the IDF Component Manager that `my_component` depends on. It is defined in the format `namespace/name=1.0.0`, where `namespace/name` is the name of the component and `=1.0.0` is a version range of the component, see the [Versioning Documentation](#). By default, dependencies are added to the main component. You can either explicitly specify a directory where the manifest is located using the `--path` option, or specify the component in your `components` folder using `--component=my_component`. The `add-dependency` command can be run in the following ways:

- `idf.py add-dependency example/cmp` adds a dependency on the most recent version of `example/cmp` to the main component
- `idf.py add-dependency --component=my_component example/cmp<=3.3.3` adds a dependency on the version `<=3.3.3` of `example/cmp` to the component `my_component` in the `components` directory
- `idf.py add-dependency --path="../../my_component" example/cmp^3.3.3` adds a dependency on the version `^3.3.3` of `example/cmp` to the component `my_component` in the `my_component` directory

Note: The command `add-dependency` adds dependencies to your project explicitly from the [Espressif Component Registry](#).

To update dependencies of the ESP-IDF project, you can run the command `idf.py update-dependencies`. You can also specify the path to the project directory using `--project-dir PATH`.

There is an example application [build_system/cmake/component_manager](#) that uses components installed by the component manager.

It is not necessary to have a manifest for components that do not need any managed dependencies.

When CMake configures the project (e.g., `idf.py reconfigure`) component manager does a few things:

- Processes `idf_component.yml` manifests for every component in the project and recursively solves dependencies.

- Creates a `dependencies.lock` file in the root of the project with a full list of dependencies.
- Downloads all dependencies to the `managed_components` directory.

The lock file `dependencies.lock` and the content of the `managed_components` directory are not supposed to be modified by a user. When the component manager runs, it always makes sure they are up to date. If these files were accidentally modified, it is possible to re-run the component manager by triggering CMake with `idf.py reconfigure`.

You may set the build property `DEPENDENCIES_LOCK` to specify the lock-file path in the top-level CMakeLists.txt. For example, adding `idf_build_set_property(DEPENDENCIES_LOCK dependencies.lock.${IDF_TARGET})` before `project(PROJECT_NAME)` could help generate different lock files for different targets.

Creating a Project From an Example

Some components on the registry contain example projects. To create a new project from an example you can run the command `idf.py create-project-from-example EXAMPLE`. The `EXAMPLE` argument should be in the format `namespace/name=1.0.0:example` where `namespace/name` is the name of the component, `=1.0.0` is a version range of the component (see the [Versioning Documentation](#)) and `example` is the example's name. You can find the list of examples for every component and the command to start a project for it in the [Espressif Component Registry](#).

Defining Dependencies in the Manifest

You can easily define dependencies in the manifest file `idf_component.yml` by editing it directly in the text editor. Below are some basic examples that demonstrate how to define dependencies.

You can define a dependency from the registry by specifying the component name and the version range:

```
dependencies:
  # Define a dependency from the registry (https://components.espressif.com/
  ↪ component/example/cmp)
  example/cmp: ">=1.0.0"
```

To define a dependency from a Git repository, provide the path to the component within the repository and the repository's URL:

```
dependencies:
  # Define a dependency from a Git repository
  test_component:
    path: test_component
    git: ssh://git@gitlab.com/user/components.git
```

During the development of components, you can use components from a local directory by specifying either a relative or an absolute path:

```
dependencies:
  # Define local dependency with relative path
  some_local_component:
    path: ../../projects/component
```

For detailed information about the manifest file format, see [Manifest File Format Documentation](#).

Disabling the Component Manager

The component manager can be explicitly disabled by setting the `IDF_COMPONENT_MANAGER` environment variable to 0.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.22.5 IDF Clang-Tidy

The IDF Clang Tidy is a tool that uses [clang-tidy](#) to run static analysis on your current app.

Warning: This functionality and the toolchain it relies on are still under development. There may be breaking changes before a final release.

Warning: This tool does not support RISC-V based chips yet. For now, we do not provide clang based toolchain for RISC-V.

Prerequisites

If you have never run this tool before, take the following steps to get this tool prepared.

1. Run `idf_tools.py install esp-clang` to install the clang-tidy required binaries

Note: This toolchain is still under development. After the final release, you do not have to install them manually.

2. Run the export scripts (`export.sh / export.bat / ...`) again to refresh the environment variables.

Extra Commands

clang-check Run `idf.py clang-check` to re-generate the compilation database and run `clang-tidy` under your current project folder. The output would be written to `<project_dir>/warnings.txt`.

Run `idf.py clang-check --help` to see the full documentation.

clang-html-report

1. Run `pip install codereport` to install the additional dependency.
2. Run `idf.py clang-html-report` to generate an HTML report in folder `<project_dir>/html_report` according to the `warnings.txt`. Please open the `<project_dir>/html_report/index.html` in your browser to check the report.

Bug Report

This tool is hosted in [espressif/clang-tidy-runner](#). If you were to face any bugs or have any feature request, please report them via [Github issues](#)

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.22.6 Downloadable IDF Tools

The ESP-IDF build process relies on a number of tools: cross-compiler toolchains, CMake build system, and others.

Installing the tools using an OS-specific package manager (e.g., apt, yum, brew, etc.) is the preferred method, when the required version of the tool is available. This recommendation is reflected in the [Get Started](#). For example, on Linux and macOS, it is recommended to install CMake using an OS package manager.

However, some of the tools are specific to ESP-IDF and are not available in OS package repositories. Furthermore, different ESP-IDF versions require different tool versions for proper operation. To solve these two problems, ESP-IDF provides a set of scripts that can download and install the correct tool versions and set up the environment accordingly.

The rest of the document refers to these downloadable tools simply as "tools". Other kinds of tools used in ESP-IDF are:

- Python scripts bundled with ESP-IDF such as `idf.py`
- Python packages installed from PyPI

The following sections explain the installation method and provide the list of tools installed on each platform.

Note: This document is provided for advanced users who need to customize their installation, users who wish to understand the installation process, and ESP-IDF developers.

If you are looking for instructions on how to install the tools, see [Get Started](#).

Tools Metadata File

The list of tools and tool versions required for each platform is located in [tools/tools.json](#). The schema of this file is defined by [tools/tools_schema.json](#).

This file is used by the [tools/idf_tools.py](#) script when installing the tools or setting up the environment variables.

Tools Installation Directory

The `IDF_TOOLS_PATH` environment variable specifies the location where the tools are to be downloaded and installed. If not set, the default location will be `HOME/.espressif` on Linux and macOS, and `%USER_PROFILE%\espressif` on Windows.

Inside the `IDF_TOOLS_PATH` directory, the tools installation scripts create the following directories and files:

- `dist` —where the archives of the tools are downloaded.
- `tools` —where the tools are extracted. The tools are extracted into subdirectories: `tools/TOOL_NAME/VERSION/`. This arrangement allows different versions of tools to be installed side by side.
- `idf-env.json` —user install options, such as targets and features, are stored in this file. Targets are selected chip targets for which tools are installed and kept up-to-date. Features determine the Python package set which should be installed. These options will be discussed later.
- `python_env` —not related to the tools; virtual Python environments are installed in the sub-directories. Note that the Python environment directory can be placed elsewhere by setting the `IDF_PYTHON_ENV_PATH` environment variable.
- `espidf.constraints.*.txt` —one constraint file for each ESP-IDF release containing Python package version requirements.

GitHub Assets Mirror

Most of the tools downloaded by the tools installer are GitHub Release Assets, which are files attached to a software release on GitHub.

If GitHub downloads are inaccessible or slow to access, a GitHub assets mirror can be configured.

To use Espressif's download server, set the environment variable `IDF_GITHUB_ASSETS` to `dl.espressif.com/github_assets`, or `dl.espressif.cn/github_assets` for faster download in China. When the install process is downloading a tool from `github.com`, the URL will be rewritten to use this server instead.

Any mirror server can be used provided the URL matches the `github.com` download URL format. For any GitHub asset URL that the install process downloads, it will replace `https://github.com` with `https://${IDF_GITHUB_ASSETS}`.

Note: The Espressif download server currently does not mirror everything from GitHub, but only files attached as Assets to some releases, as well as source archives for some releases.

`idf_tools.py` Script

The `tools/idf_tools.py` script bundled with ESP-IDF performs several functions:

- `install`: Download the tool into the `${IDF_TOOLS_PATH}/dist` directory and extract it into `${IDF_TOOLS_PATH}/tools/TOOL_NAME/VERSION`.
The `install` command accepts the list of tools to install in the `TOOL_NAME` or `TOOL_NAME@VERSION` format. If `all` is given, all the tools, including required and optional ones, are installed. If no argument or `required` is given, only the required tools are installed.
- `download`: Similar to `install` but doesn't extract the tools. An optional `--platform` argument may be used to download the tools for the specific platform.
- `export`: Lists the environment variables that need to be set to use the installed tools. For most of the tools, setting the `PATH` environment variable is sufficient, but some tools require extra environment variables. The environment variables can be listed in either `shell` or `key-value` formats, which can be set using the `--format` parameter:
 - `export` optional parameters:
 - * `--unset`: Creates a statement that unsets specific global variables and restores the environment to its state before calling `export.{sh/fish}`.
 - * `--add_paths_extras`: Adds extra ESP-IDF-related paths of `$PATH` to `${IDF_TOOLS_PATH}/esp-idf.json`, which is used to remove global variables when the active ESP-IDF environment is deactivated. For example, while processing the `export.{sh/fish}` script, if new paths are added to the global variable `$PATH`, this option saves these new paths to the `${IDF_TOOLS_PATH}/esp-idf.json` file.
 - `shell`: Produces output suitable for evaluation in the shell. For example, produce the following output on Linux and macOS:

```
export PATH="/home/user/.espressif/tools/tool/v1.0.0/bin:$PATH"
```

Produce the following output on Windows:

```
set "PATH=C:\Users\user\.espressif\tools\v1.0.0\bin;%PATH%"
```

Note: Exporting environment variables in Powershell format is not supported at the moment. `key-value` format may be used instead.

The output of this command may be used to update the environment variables if the shell supports it. For example

```
eval $(${IDF_PATH}/tools/idf_tools.py export)
```

- `key-value`: Produces output in the `VARIABLE=VALUE` format that is suitable for parsing by other scripts

```
PATH=/home/user/.espressif/tools/tool/v1.0.0:$PATH
```

Note that the script consuming this output has to perform expansion of `$VAR` or `%VAR%` patterns found in the output.

- `list`: Lists the known versions of the tools, and indicates which ones are installed.
The following option is available to customize the output.
 - `--outdated`: Lists only outdated versions of tools installed in `IDF_TOOLS_PATH`.
- `check`: For each tool, checks whether the tool is available in the system path and in `IDF_TOOLS_PATH`.
- `install-python-env`: Creates a Python virtual environment in the `${IDF_TOOLS_PATH}/python_env` directory or directly in the directory set by the `IDF_PYTHON_ENV_PATH` environment variable, and install the required Python packages there.
 - An optional `--features` argument allows one to specify a comma-separated list of features to be added or removed.
 1. A feature that begins with `-` will be removed, and features with `+` or without any sign will be added. Example syntax for removing feature `XY` is `--features=-XY`, and for adding feature `XY` is `--features=+XY` or `--features=XY`. If both removing and adding options are provided with the same feature, no operation is performed.
 2. For each feature, a requirements file must exist. For example, feature `XY` is a valid feature if `${IDF_PATH}/tools/requirements/requirements.XY.txt` is an existing file with a list of Python packages to be installed.
 3. There is one mandatory `core` feature ensuring the core functionality of ESP-IDF, e.g., `build`, `flash`, `monitor`, `debug in console`. There can be an arbitrary number of optional features.
 4. The selected list of features is stored in `idf-env.json`.
 5. The requirement files contain a list of the desired Python packages to be installed and the `espidf.constraints.*.txt` file downloaded from <https://dl.espressif.com> and stored in `${IDF_TOOLS_PATH}`, which contains the package version requirements for a given ESP-IDF version.

Note: Although it is **not recommended**, the download and use of constraint files can be disabled with the `--no-constraints` argument or setting the `IDF_PYTHON_CHECK_CONSTRAINTS` environment variable to `no`.

- `check-python-dependencies`: Checks if all required Python packages are installed. Packages from `${IDF_PATH}/tools/requirements/requirements.*.txt` files selected by the feature list of `idf-env.json` are checked with the package versions specified in the `espidf.constraints.*.txt` file.

Note: The constraint file is downloaded with the `install-python-env` command. Similar to the `install-python-env` command, the use of constraint files can be disabled with the `--no-constraints` argument or setting the `IDF_PYTHON_CHECK_CONSTRAINTS` environment variable to `no`.

- `uninstall`: Prints and removes tools that are currently not used by the active ESP-IDF version.
 - `--dry-run`: Prints installed unused tools.
 - `--remove-archives`: Additionally removes all older versions of previously downloaded installation packages.

Install Scripts

Shell-specific user-facing installation scripts are provided in the root directory of ESP-IDF repository to facilitate tools installation. These are:

- `install.bat` for Windows Command Prompt
- `install.ps1` for Powershell
- `install.sh` for Bash
- `install.fish` for Fish

Apart from downloading and installing the tools in `IDF_TOOLS_PATH`, these scripts prepare a Python virtual environment, and install the required packages into that environment.

These scripts accept optionally a comma-separated list of chip targets and `--enable-*` arguments for enabling features. These arguments are passed to the `idf_tools.py` script which stores them in `idf-env.json`. Therefore, chip targets and features can be enabled incrementally.

To install tools for all chip targets, run the scripts without any optional arguments using `idf_tools.py install --targets=all`. Similarly, to install Python packages for core ESP-IDF functionality, run `idf_tools.py install-python-env --features=core`.

It is also possible to install tools for specific chip targets. For example, `install.sh esp32` installs tools only for ESP32. See [Step 3. Set up the Tools](#) for more examples.

`install.sh --enable-XY` enables feature XY (by running `idf_tools.py install-python-env --features=core,XY`).

Export Scripts

Since the installed tools are not permanently added to the user or system `PATH` environment variable, an extra step is required to use them in the command line. The following scripts modify the environment variables in the current shell to make the correct versions of the tools available:

- `export.bat` for Windows Command Prompt
- `export.ps1` for Powershell
- `export.sh` for Bash
- `export.fish` for Fish

Note: To modify the shell environment in Bash, `export.sh` must be "sourced" by using the `./export.sh` command. Please ensure to include the leading dot and space.

`export.sh` may be used with shells other than Bash (such as `zsh`). However, in this case, it is required to set the `IDF_PATH` environment variable before running the script. When used in Bash, the script guesses the `IDF_PATH` value from its own location.

In addition to calling `idf_tools.py`, these scripts list the directories that have been added to the `PATH`.

Other Installation Methods

Depending on the environment, more user-friendly wrappers for `idf_tools.py` are provided:

- [ESP-IDF Tools Installer](#) can download and install the tools. Internally the installer uses `idf_tools.py`.
- [ESP-IDF Eclipse Plugin](#) includes a menu item to set up the tools. Internally the plugin calls `idf_tools.py`.
- [VSCode ESP-IDF Extension](#) includes an onboarding flow. This flow helps set up the tools. Although the extension does not rely on `idf_tools.py`, the same installation method is used.

Custom Installation

Although the methods above are recommended for ESP-IDF users, they are not a must for building ESP-IDF applications. ESP-IDF build system expects that all the necessary tools are installed somewhere, and made available in the `PATH`.

Uninstall ESP-IDF

Uninstalling ESP-IDF requires removing both the tools and the environment variables that have been configured during the installation.

- Windows users using the [Windows ESP-IDF Tools Installer](#) can simply run the uninstall wizard to remove ESP-IDF.
- To remove an installation performed by running the supported [install scripts](#), simply delete the [tools installation directory](#) including the downloaded and installed tools. Any environment variables set by the [export scripts](#) are not permanent and will not be present after opening a new environment.

- When dealing with a custom installation, in addition to deleting the tools as mentioned above, you may also need to manually revert any changes to environment variables or system paths that were made to accommodate the ESP-IDF tools (e.g., `IDF_PYTHON_ENV_PATH` or `IDF_TOOLS_PATH`). If you manually copied any tools, you would need to track and delete those files manually.
- If you installed any plugins like the [ESP-IDF Eclipse Plugin](#) or [VSCode ESP-IDF Extension](#), you should follow the specific uninstallation instructions described in the documentation of those components.

Note: Uninstalling the ESP-IDF tools does not remove any project files or your code. Be mindful of what you are deleting to avoid losing any work. If you are unsure about a step, refer back to the installation instructions.

These instructions assume that the tools were installed following the procedures in this provided document. If you've used a custom installation method, you might need to adapt these instructions accordingly.

List of ESP-IDF Tools

xtensa-esp-elf-gdb GDB for Xtensa

License: [GPL-3.0-or-later](#)

More info: <https://github.com/espressif/binutils-gdb>

Platform	Required	Download
linux-amd64	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/xtensa-esp-elf-gdb-16.3_20250913-x86_64-linux-gnu.tar.gz SHA256: 16d05c9104ff84529ac3799abb04d5666c193131ab461f153040721728b48730
linux-arm64	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/xtensa-esp-elf-gdb-16.3_20250913-aarch64-linux-gnu.tar.gz SHA256: ecbd53ba28cf24301be8260249bfcfb60567f938f4402797617c8a0fc170dc7d
linux-armel	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/xtensa-esp-elf-gdb-16.3_20250913-arm-linux-gnueabi.tar.gz SHA256: c0895e88797089fd6b16e1cb986c5c85a880e0e8dc03bde1016c7771bc10ddba
linux-armhf	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/xtensa-esp-elf-gdb-16.3_20250913-arm-linux-gnueabi-hf.tar.gz SHA256: d939e273150f78c79e62c15564cb880f253b8f23670968905ea00687f9a7af2f
linux-i686	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/xtensa-esp-elf-gdb-16.3_20250913-i586-linux-gnu.tar.gz SHA256: 64ffefb7625edae77a03a13fd9bd07db088dec9d145eb1124de66f11510f7558
macos	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/xtensa-esp-elf-gdb-16.3_20250913-x86_64-apple-darwin24.5.tar.gz SHA256: 8341493abc87e6ae468f4eda16c768b2ddb20c98336e1c491a3801ad823680ae
macos-arm64	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/xtensa-esp-elf-gdb-16.3_20250913-aarch64-apple-darwin24.5.tar.gz SHA256: 251e3be9c9436d9ab7fee6c05519fd816a05e63bd47495e24ea4e354881a851c
win32	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/xtensa-esp-elf-gdb-16.3_20250913-i686-w64-mingw32.zip SHA256: 8fc9fa6a934523b6ad6e787cf1664d48496bae456fd85ea7589e3684ce3bbbe5
win64	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/xtensa-esp-elf-gdb-16.3_20250913-x86_64-w64-mingw32.zip SHA256: 99a2243b9a75bbac95a672cc3ab4b36013429ab5b4583e7a28339e3015a3fdfa

riscv32-esp-elf-gdb GDB for RISC-V

License: [GPL-3.0-or-later](#)

More info: <https://github.com/espressif/binutils-gdb>

Platform	Required	Download
linux-amd64	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/riscv32-esp-elf-gdb-16.3_20250913-x86_64-linux-gnu.tar.gz SHA256: 4e3cf8b7d11c7a2d1b50f40b1c50c0671dfe7eb13782c27c8a8cfdc8548bcdd4
linux-arm64	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/riscv32-esp-elf-gdb-16.3_20250913-aarch64-linux-gnu.tar.gz SHA256: 8f1f4f24fa534c76ed9d71effbf728cc30169e911742d7bd67dd0fdcf5f3ae3
linux-armel	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/riscv32-esp-elf-gdb-16.3_20250913-arm-linux-gnueabi.tar.gz SHA256: ac4fc85e3daf190b21598ec468933dc2659033580715560f45827da25e15b285
linux-armhf	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/riscv32-esp-elf-gdb-16.3_20250913-arm-linux-gnueabi.tar.gz SHA256: 628d0375f80f63834f68be1b2c4d71690d32ac71b4a9a32f514829a9eb90439f
linux-i686	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/riscv32-esp-elf-gdb-16.3_20250913-i586-linux-gnu.tar.gz SHA256: 905dcd78558d7d559a95dc1eacc4572ea908be4ae6b1c937ea63a98df4482ca9
macos	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/riscv32-esp-elf-gdb-16.3_20250913-x86_64-apple-darwin24.5.tar.gz SHA256: 2d5e5efead0b189e13cfe2670ca9d6d5965378ef3632d0b163a14f2f0536c274
macos-arm64	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/riscv32-esp-elf-gdb-16.3_20250913-aarch64-apple-darwin24.5.tar.gz SHA256: 92771492084746fd22521c7c5b52bf1ed6dd86ef3cafe60e771bbdb4f0943f5a
win32	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/riscv32-esp-elf-gdb-16.3_20250913-i686-w64-mingw32.zip SHA256: c6a36c469d3b76e2b442be207814f7c3f71f21faf6faab4dd33fdedd56d89c01
win64	required	https://github.com/espressif/binutils-gdb/releases/download/esp-gdb-v16.3_20250913/riscv32-esp-elf-gdb-16.3_20250913-x86_64-w64-mingw32.zip SHA256: 32e79cb43b40f3b256193139b1fef2782e3eaf82ee317b757ec8ba18b35159d

xtensa-esp-elf Toolchain for 32-bit Xtensa based on GCC

License: [GPL-3.0-with-GCC-exception](#)

More info: <https://github.com/espressif/crosstool-NG>

Platform	Required	Download
linux-amd64	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/xtensa-esp-elf-13.2.0_20250707-x86_64-linux-gnu.tar.xz SHA256: d1f1378394baa3788bf78062e5d041cb14d4858bbb57f86d59e5b83fc9d5525f
linux-arm64	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/xtensa-esp-elf-13.2.0_20250707-aarch64-linux-gnu.tar.xz SHA256: 166c0a6268254acb6d7a72fa8953f3f569a63b1f2dc9adff752a52debb12a364
linux-armel	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/xtensa-esp-elf-13.2.0_20250707-arm-linux-gnueabi.tar.xz SHA256: 93d510b3dee9a5d59390213e617bf23c1b69c2b179e7ef415c9a803d54f9e1b1
linux-armhf	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/xtensa-esp-elf-13.2.0_20250707-arm-linux-gnueabi-hf.tar.xz SHA256: fdcbaa8cd7584e45deb77f12ea51fc25b408b9a45ab9335ff0613145b119091
linux-i686	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/xtensa-esp-elf-13.2.0_20250707-i586-linux-gnu.tar.xz SHA256: 75a2abf8137e0f90927b5fc57a098bccd031d489fcc82c93bec244950d9044d2
macos	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/xtensa-esp-elf-13.2.0_20250707-x86_64-apple-darwin.tar.xz SHA256: 8152d897b559b0833d81fb4e891b42dbf225bd1166a4b5a5d91a1d94475444e1
macos-arm64	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/xtensa-esp-elf-13.2.0_20250707-aarch64-apple-darwin.tar.xz SHA256: 28e2a33b04397dcfeb8bf1beaea7f4d76eb1edf163626ca377913a1859cfb7a3
win32	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/xtensa-esp-elf-13.2.0_20250707-i686-w64-mingw32.zip SHA256: 6450e261299c601b5f5e4d25f5a6be35cf6ec4fc8d36323f686f116a2e6cd1b4
win64	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/xtensa-esp-elf-13.2.0_20250707-x86_64-w64-mingw32.zip SHA256: e0eae3964d603d15fca309d2c02af40c25429ced443460b9cb07980b261e5069

esp-clang Toolchain for all Espressif chips based on clang

License: [Apache-2.0](#)

More info: <https://github.com/espressif/llvm-project>

Platform	Required	Download
linux-amd64	optional	https://github.com/espressif/llvm-project/releases/download/esp-16.0.0-20230516/llvm-esp-16.0.0-20230516-linux-amd64.tar.xz SHA256: 3dbd8dd290913a93e8941da8a451ecd49f9798cc2d74bb9b63ef5cf5c4fee37f
linux-arm64	optional	https://github.com/espressif/llvm-project/releases/download/esp-16.0.0-20230516/llvm-esp-16.0.0-20230516-linux-arm64.tar.xz SHA256: 4b115af6ddd04a9bffc1908fc05837998ee71d450891d741c446186f2aa9b961
linux-armhf	optional	https://github.com/espressif/llvm-project/releases/download/esp-16.0.0-20230516/llvm-esp-16.0.0-20230516-linux-armhf.tar.xz SHA256: 935082bb0704420c5ca42b35038bba8702135348a50cac454ae2fb55af0b4c32
macos	optional	https://github.com/espressif/llvm-project/releases/download/esp-16.0.0-20230516/llvm-esp-16.0.0-20230516-macos.tar.xz SHA256: d9824acafd3e7b1d17ace084243b82a95bbdcb149a26b085bba487ab3d3716d7
macos-arm64	optional	https://github.com/espressif/llvm-project/releases/download/esp-16.0.0-20230516/llvm-esp-16.0.0-20230516-macos-arm64.tar.xz SHA256: ed5621396dc3e48413e14e8b6caed8e2993e7f2ab5fca1410081f40c940a1060
win64	optional	https://github.com/espressif/llvm-project/releases/download/esp-16.0.0-20230516/llvm-esp-16.0.0-20230516-win64.tar.xz SHA256: 598c8241c8bf10fd1be8bd21845307cfc404e127041b4ba4e828350a88692883

riscv32-esp-elf Toolchain for 32-bit RISC-V based on GCC

License: [GPL-3.0-with-GCC-exception](#)

More info: <https://github.com/espressif/crosstool-NG>

Platform	Required	Download
linux-amd64	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/riscv32-esp-elf-13.2.0_20250707-x86_64-linux-gnu.tar.xz SHA256: bd1ad6cdcee2d5de288c0c05b70f56f959dd270e2df254159594d0b05ad5dea1
linux-arm64	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/riscv32-esp-elf-13.2.0_20250707-aarch64-linux-gnu.tar.xz SHA256: c97ef1a0d7cf675ec1ba61c6733f4a821912d0af60575dbb9e90a80a0c6e598f
linux-armel	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/riscv32-esp-elf-13.2.0_20250707-arm-linux-gnueabi.tar.xz SHA256: bcf72e08e601e82dd848ced4b87e712bb88ed76ac6aa6b208c087f5f615a0e8f
linux-armhf	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/riscv32-esp-elf-13.2.0_20250707-arm-linux-gnueabi.tar.xz SHA256: 249ecf34be5a7743798d86a53db9a5a3df618721afedd3bacf64e02046a21225
linux-i686	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/riscv32-esp-elf-13.2.0_20250707-i586-linux-gnu.tar.xz SHA256: 50277ae3a0adba0363762325e10c509c21dd42c4d2263d8bc445270934467d99
macos	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/riscv32-esp-elf-13.2.0_20250707-x86_64-apple-darwin.tar.xz SHA256: bf433caf46f2c124921d0a1e317d777f6f9697ac84d0b45c8aaed91fa3bc954
macos-arm64	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/riscv32-esp-elf-13.2.0_20250707-aarch64-apple-darwin.tar.xz SHA256: 984819225ecd1ccfd03fbc1417061b34c33e04ba76510f95a932d5cbc01be45c
win32	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/riscv32-esp-elf-13.2.0_20250707-i686-w64-mingw32.zip SHA256: 474be67f9c0c7bdf901265f8bd63141a1fec6e94e2fcd887693edd1d6dc106ba
win64	required	https://github.com/espressif/crosstool-NG/releases/download/esp-13.2.0_20250707/riscv32-esp-elf-13.2.0_20250707-x86_64-w64-mingw32.zip SHA256: 79a65a85be948c2d64972314d785f13f4984acca556166422dbc23814c591c3b

esp32ulp-elf Toolchain for ESP32 ULP coprocessor

License: [GPL-3.0-or-later](#)

More info: <https://github.com/espressif/binutils-gdb>

Platform	Required	Download
linux-amd64	required	https://github.com/espressif/binutils-gdb/releases/download/esp32ulp-elf-v2.35_20220830/esp32ulp-elf-2.35_20220830-linux-amd64.tar.gz SHA256: b1f7801c3a16162e72393ebb772c0cbfe4d22d907be7c2c2dac168736e9195fd
linux-arm64	required	https://github.com/espressif/binutils-gdb/releases/download/esp32ulp-elf-v2.35_20220830/esp32ulp-elf-2.35_20220830-linux-arm64.tar.gz SHA256: d6671b31bab31b9b13aea25bb7d60f15484cb8bf961ddbf67a62867e5563eae5
linux-armel	required	https://github.com/espressif/binutils-gdb/releases/download/esp32ulp-elf-v2.35_20220830/esp32ulp-elf-2.35_20220830-linux-armel.tar.gz SHA256: e107e7a9cd50d630b034f435a16a52db5a57388dc639a99c4c393c5e429711e9
linux-armhf	required	https://github.com/espressif/binutils-gdb/releases/download/esp32ulp-elf-v2.35_20220830/esp32ulp-elf-2.35_20220830-linux-armhf.tar.gz SHA256: 6c6dd25477b2e758d4669da3774bf664d1f012442c880f17dfdf0339e9c3dae9
linux-i686	required	https://github.com/espressif/binutils-gdb/releases/download/esp32ulp-elf-v2.35_20220830/esp32ulp-elf-2.35_20220830-linux-i686.tar.gz SHA256: beb9b6737c975369b6959007739c88f44eb5afbb220f40737071540b2c1a9064
macos	required	https://github.com/espressif/binutils-gdb/releases/download/esp32ulp-elf-v2.35_20220830/esp32ulp-elf-2.35_20220830-macos.tar.gz SHA256: 5a952087b621ced16af1e375feac1371a61cb51ab7e7b44cbefb5afda2d573de
macos-arm64	required	https://github.com/espressif/binutils-gdb/releases/download/esp32ulp-elf-v2.35_20220830/esp32ulp-elf-2.35_20220830-macos-arm64.tar.gz SHA256: 73bda8476ef92d4f4abee96519abbba40e5ee32f368427469447b83cc7bb9b42
win32	required	https://github.com/espressif/binutils-gdb/releases/download/esp32ulp-elf-v2.35_20220830/esp32ulp-elf-2.35_20220830-win32.zip SHA256: 77344715ea7d7a7a9fd0b27653f880efaf3bcc1ac843f61492d8a0365d91f731
win64	required	https://github.com/espressif/binutils-gdb/releases/download/esp32ulp-elf-v2.35_20220830/esp32ulp-elf-2.35_20220830-win64.zip SHA256: 525e5b4c8299869a3fdddb51baad76612c5c104bd96952ae6460ad7e5b5a4e21

cmake CMake build system

On Linux and macOS, it is recommended to install CMake using the OS-specific package manager (like apt, yum, brew, etc.). However, for convenience it is possible to install CMake using idf_tools.py along with the other tools.

License: [BSD-3-Clause](#)

More info: <https://github.com/Kitware/CMake>

Platform	Required	Download
linux-amd64	optional	https://github.com/Kitware/CMake/releases/download/v3.30.2/cmake-3.30.2-linux-x86_64.tar.gz SHA256: cdd7fb352605cee3ae53b0e18b5929b642900e33d6b0173e19f6d4f2067ebf16
linux-arm64	optional	https://github.com/Kitware/CMake/releases/download/v3.30.2/cmake-3.30.2-linux-aarch64.tar.gz SHA256: d18f50f01b001303d21f53c6c16ff12ee3aa45df5da1899c2fe95be7426aa026
linux-armel	optional	https://dl.espressif.com/dl/cmake/cmake-3.30.2-Linux-armv7l.tar.gz SHA256: 446650c69ea74817a770f96446c162bb7ad24ffecaacb35fcd4845ec7d3c9099
linux-armhf	optional	https://dl.espressif.com/dl/cmake/cmake-3.30.2-Linux-armv7l.tar.gz SHA256: 446650c69ea74817a770f96446c162bb7ad24ffecaacb35fcd4845ec7d3c9099
macos	optional	https://github.com/Kitware/CMake/releases/download/v3.30.2/cmake-3.30.2-macos-universal.tar.gz SHA256: c6fdda745f9ce69bca048e91955c7d043ba905d6388a62e0ff52b681ac17183c
macos-arm64	optional	https://github.com/Kitware/CMake/releases/download/v3.30.2/cmake-3.30.2-macos-universal.tar.gz SHA256: c6fdda745f9ce69bca048e91955c7d043ba905d6388a62e0ff52b681ac17183c
win32	required	https://github.com/Kitware/CMake/releases/download/v3.30.2/cmake-3.30.2-windows-x86_64.zip SHA256: 48bf4b3dc2d668c578e0884cac7878e146b036ca6b5ce4f8b5572f861b004c25
win64	required	https://github.com/Kitware/CMake/releases/download/v3.30.2/cmake-3.30.2-windows-x86_64.zip SHA256: 48bf4b3dc2d668c578e0884cac7878e146b036ca6b5ce4f8b5572f861b004c25

openocd-esp32 OpenOCD for ESP32License: [GPL-2.0-only](#)More info: <https://github.com/espressif/openocd-esp32>

Platform	Required	Download
linux-amd64	required	https://github.com/espressif/openocd-esp32/releases/download/v0.12.0-esp32-20260304/openocd-esp32-linux-amd64-0.12.0-esp32-20260304.tar.gz SHA256: dbd7ecf751431c70628176fbf1ce404c3ff28027e91b66bda7f834a2d5ff5b81
linux-arm64	required	https://github.com/espressif/openocd-esp32/releases/download/v0.12.0-esp32-20260304/openocd-esp32-linux-arm64-0.12.0-esp32-20260304.tar.gz SHA256: 7f8e82e36f8e34a7a3118045fd7888754afbfe4c60cfaee0ac70663fd5965f63
linux-armel	required	https://github.com/espressif/openocd-esp32/releases/download/v0.12.0-esp32-20260304/openocd-esp32-linux-armel-0.12.0-esp32-20260304.tar.gz SHA256: c717a6ff87b07be729850fd7662fda3f1d4d7125d44dd15b0694e3021bed2bfb
linux-armhf	required	https://github.com/espressif/openocd-esp32/releases/download/v0.12.0-esp32-20260304/openocd-esp32-linux-armhf-0.12.0-esp32-20260304.tar.gz SHA256: 847df6f58308fddbb00d0db71ad971d9ab6346d091bb060bd98c053a0d4e4322
macos	required	https://github.com/espressif/openocd-esp32/releases/download/v0.12.0-esp32-20260304/openocd-esp32-macos-0.12.0-esp32-20260304.tar.gz SHA256: be6951d9766f88fad11060314f6c3469c56715a60f2715aeb7d806afc935c0d
macos-arm64	required	https://github.com/espressif/openocd-esp32/releases/download/v0.12.0-esp32-20260304/openocd-esp32-macos-arm64-0.12.0-esp32-20260304.tar.gz SHA256: a36099d3a47241e816693d9bd719198e4667ad67f0a027404d90584d44b6842d
win32	required	https://github.com/espressif/openocd-esp32/releases/download/v0.12.0-esp32-20260304/openocd-esp32-win32-0.12.0-esp32-20260304.zip SHA256: a9db16887fb0df26d1c3e495203c9edcd86d9262b2be7b7d929f8017194add31
win64	required	https://github.com/espressif/openocd-esp32/releases/download/v0.12.0-esp32-20260304/openocd-esp32-win64-0.12.0-esp32-20260304.zip SHA256: ad29bd55f2b7ad39669fbeecc32012954359dcfc0ecfa5a03068589b4d0e8613

ninja Ninja build system

On Linux and macOS, it is recommended to install ninja using the OS-specific package manager (like apt, yum, brew, etc.). However, for convenience it is possible to install ninja using idf_tools.py along with the other tools.

License: [Apache-2.0](#)

More info: <https://github.com/ninja-build/ninja>

Platform	Required	Download
linux-amd64	optional	https://github.com/ninja-build/ninja/releases/download/v1.12.1/ninja-linux.zip SHA256: 6f98805688d19672bd699fbbfa2c2cf0fc054ac3df1f0e6a47664d963d530255
macos	optional	https://github.com/ninja-build/ninja/releases/download/v1.12.1/ninja-mac.zip SHA256: 89a287444b5b3e98f88a945afa50ce937b8ffd1dcc59c555ad9b1baf855298c9
macos-arm64	optional	https://github.com/ninja-build/ninja/releases/download/v1.12.1/ninja-mac.zip SHA256: 89a287444b5b3e98f88a945afa50ce937b8ffd1dcc59c555ad9b1baf855298c9
win64	required	https://github.com/ninja-build/ninja/releases/download/v1.12.1/ninja-win.zip SHA256: f550fec705b6d6ff58f2db3c374c2277a37691678d6aba463adcbb129108467a

idf-exe IDF wrapper tool for Windows

License: [Apache-2.0](#)

More info: https://github.com/espressif/idf_py_exe_tool

Platform	Required	Download
win32	required	https://github.com/espressif/idf_py_exe_tool/releases/download/v1.0.3/idf-exe-v1.0.3.zip SHA256: 7c81ef534c562354a5402ab6b90a6eb1cc8473a9f4a7b7a7f93ebbd23b4a2755
win64	required	https://github.com/espressif/idf_py_exe_tool/releases/download/v1.0.3/idf-exe-v1.0.3.zip SHA256: 7c81ef534c562354a5402ab6b90a6eb1cc8473a9f4a7b7a7f93ebbd23b4a2755

ccache Ccache (compiler cache)

License: [GPL-3.0-or-later](#)

More info: <https://github.com/ccache/ccache>

Platform	Required	Download
win64	required	https://github.com/ccache/ccache/releases/download/v4.12.1/ccache-4.12.1-windows-x86_64.zip SHA256: 98aea520d66905b8ba7a8e648a4cc0ca941d5e119d441f1e879a4a9045bf18f6

dfu-util dfu-util (Device Firmware Upgrade Utilities)

License: [GPL-2.0-only](#)

More info: <http://dfu-util.sourceforge.net/>

Platform	Required	Download
win64	required	https://dl.espressif.com/dl/dfu-util-0.11-win64.zip SHA256: 652eb94cb1c074c6dbead9e47adb628922aeb198a4d440a346ab32e7a0e9bf64

esp-rom-elfs ESP ROM ELFsLicense: [Apache-2.0](#)More info: <https://github.com/espressif/esp-rom-elfs>

Platform	Required	Download
any	required	https://github.com/espressif/esp-rom-elfs/releases/download/20240305/esp-rom-elfs-20240305.tar.gz SHA256: a26609b415710f0163d785850c769752717004059c129c472e9a0cbd54e0422c

qemu-xtensa QEMU for XtensaSome ESP-specific instructions for running QEMU for Xtensa chips are here: <https://github.com/espressif/esp-toolchain-docs/blob/main/qemu/esp32/README.md>License: [GPL-2.0-only](#)More info: <https://github.com/espressif/qemu>

Platform	Required	Download
linux-amd64	optional	https://github.com/espressif/qemu/releases/download/esp-develop-9.0.0-20240606/qemu-xtensa-softmmu-esp_develop_9.0.0_20240606-x86_64-linux-gnu.tar.xz SHA256: 071d117c44a6e9a1bc8664ab63b592d3e17ceb779119dcb46c59571a4a7a88c9
linux-arm64	optional	https://github.com/espressif/qemu/releases/download/esp-develop-9.0.0-20240606/qemu-xtensa-softmmu-esp_develop_9.0.0_20240606-aarch64-linux-gnu.tar.xz SHA256: 43552f32b303a6820d0d9551903e54fc221aca98ccbd04e5cbccbca881548008
macos	optional	https://github.com/espressif/qemu/releases/download/esp-develop-9.0.0-20240606/qemu-xtensa-softmmu-esp_develop_9.0.0_20240606-x86_64-apple-darwin.tar.xz SHA256: 0096734280ce04f558cd9bd72f35db39667f80d44309a35565f2f8c02d1f9cc3
macos-arm64	optional	https://github.com/espressif/qemu/releases/download/esp-develop-9.0.0-20240606/qemu-xtensa-softmmu-esp_develop_9.0.0_20240606-aarch64-apple-darwin.tar.xz SHA256: fb4ca6be7b1a4dbcf153879cf0582300f974371def0826c0c5b728f12812ad08
win64	optional	https://github.com/espressif/qemu/releases/download/esp-develop-9.0.0-20240606/qemu-xtensa-softmmu-esp_develop_9.0.0_20240606-x86_64-w64-mingw32.tar.xz SHA256: 281659f7a1d49761ac6f54d0aeb14366cb93c002f21948b847a0e15c0b8f5425

qemu-riscv32 QEMU for RISC-VSome ESP-specific instructions for running QEMU for RISC-V chips are here: <https://github.com/espressif/esp-toolchain-docs/blob/main/qemu/esp32c3/README.md>License: [GPL-2.0-only](#)More info: <https://github.com/espressif/qemu>

Platform	Required	Download
linux-amd64	optional	https://github.com/espressif/qemu/releases/download/esp-develop-9.0.0-20240606/qemu-riscv32-softmmu-esp_develop_9.0.0_20240606-x86_64-linux-gnu.tar.xz SHA256: 47120e826cfec7180db8cb611a7a4aed2e9b2191c2a739194f8ce085e63cdd8d
linux-arm64	optional	https://github.com/espressif/qemu/releases/download/esp-develop-9.0.0-20240606/qemu-riscv32-softmmu-esp_develop_9.0.0_20240606-aarch64-linux-gnu.tar.xz SHA256: 3b6221a8b1881d2c9b9fa0b0bf8d7065c84153d2a54e429307bde9feae235c27
macos	optional	https://github.com/espressif/qemu/releases/download/esp-develop-9.0.0-20240606/qemu-riscv32-softmmu-esp_develop_9.0.0_20240606-x86_64-apple-darwin.tar.xz SHA256: 3afa55d5abea52ccf18d0bc41fe819d568bd4ee1582989b1ee9b1ee4a609a31e
macos-arm64	optional	https://github.com/espressif/qemu/releases/download/esp-develop-9.0.0-20240606/qemu-riscv32-softmmu-esp_develop_9.0.0_20240606-aarch64-apple-darwin.tar.xz SHA256: 69ba5154594fb2922d5490a49ea6b4925c024c6c37f875b42f9885f513e0bcdd
win64	optional	https://github.com/espressif/qemu/releases/download/esp-develop-9.0.0-20240606/qemu-riscv32-softmmu-esp_develop_9.0.0_20240606-x86_64-w64-mingw32.tar.xz SHA256: f49bb5c8f4d6e2cfbf7ecec21eb8ef190a57307778705bc689536ac13bde511c

4.23 Unit Testing in ESP32-P4

ESP-IDF provides the following methods to test software.

- Target based tests using a central unit test application which runs on the esp32p4. These tests use the [Unity](#) unit test framework. They can be integrated into an ESP-IDF component by placing them in the component's `test` subdirectory. This document mainly introduces this target based tests.
- Linux-host based unit tests in which part of the hardware can be abstracted via mocks. Currently, Linux-host based tests are still under development and only a small fraction of IDF components support them. More information on running IDF applications on the host can be found here: [Running Applications on the Host Machine](#).

4.23.1 Normal Test Cases

Unit tests are located in the `test` subdirectory of a component. Tests are written in C, and a single C source file can contain multiple test cases. Test files start with the word "test".

Each test file should include the `unity.h` header and the header for the C module to be tested.

Tests are added in a function in the C file as follows:

```
TEST_CASE("test name", "[module name]")
{
    // Add test here
}
```

- The first argument is a descriptive name for the test.
- The second argument is an identifier in square brackets. Identifiers are used to group related test, or tests with specific properties.

Note: There is no need to add a main function with `UNITY_BEGIN()` and `UNITY_END()` in each test case. `unity_platform.c` will run `UNITY_BEGIN()` autonomously, and run the test cases, then call `UNITY_END()`.

The `test` subdirectory should contain a [component CMakeLists.txt](#), since they are themselves components (i.e., a test component). ESP-IDF uses the Unity test framework located in the `unity` component. Thus, each test component

should specify the `unity` component as a component requirement using the `REQUIRES` argument. Normally, components *should list their sources manually*; for component tests however, this requirement is relaxed and the use of the `SRC_DIRS` argument in `idf_component_register` is advised.

Overall, the minimal test subdirectory `CMakeLists.txt` file should contain the following:

```
idf_component_register(SRC_DIRS "."
                      INCLUDE_DIRS "."
                      REQUIRES unity)
```

See <http://www.throwtheswitch.org/unity> for more information about writing tests in Unity.

4.23.2 Multi-device Test Cases

The normal test cases will be executed on one DUT (Device Under Test). However, components that require some form of communication (e.g., GPIO, SPI) require another device to communicate with, thus cannot be tested through normal test cases. Multi-device test cases involve writing multiple test functions, and running them on multiple DUTs.

The following is an example of a multi-device test case:

```
void gpio_master_test()
{
    gpio_config_t slave_config = {
        .pin_bit_mask = 1 << MASTER_GPIO_PIN,
        .mode = GPIO_MODE_INPUT,
    };
    gpio_config(&slave_config);
    unity_wait_for_signal("output high level");
    TEST_ASSERT(gpio_get_level(MASTER_GPIO_PIN) == 1);
}

void gpio_slave_test()
{
    gpio_config_t master_config = {
        .pin_bit_mask = 1 << SLAVE_GPIO_PIN,
        .mode = GPIO_MODE_OUTPUT,
    };
    gpio_config(&master_config);
    gpio_set_level(SLAVE_GPIO_PIN, 1);
    unity_send_signal("output high level");
}

TEST_CASE_MULTIPLE_DEVICES("gpio multiple devices test example", "[driver]", gpio_
↪master_test, gpio_slave_test);
```

The macro `TEST_CASE_MULTIPLE_DEVICES` is used to declare a multi-device test case.

- The first argument is test case name.
- The second argument is test case description.
- From the third argument, up to 5 test functions can be defined, each function will be the entry point of tests running on each DUT.

Running test cases from different DUTs could require synchronizing between DUTs. We provide `unity_wait_for_signal` and `unity_send_signal` to support synchronizing with UART. As the scenario in the above example, the slave should get GPIO level after master set level. DUT UART console will prompt and user interaction is required:

DUT1 (master) console:

```
Waiting for signal: [output high level]!
Please press "Enter" key to once any board send this signal.
```

DUT2 (slave) console:

```
Send signal: [output high level]!
```

Once the signal is sent from DUT2, you need to press "Enter" on DUT1, then DUT1 unblocks from `unity_wait_for_signal` and starts to change GPIO level.

4.23.3 Multi-stage Test Cases

The normal test cases are expected to finish without reset (or only need to check if reset happens). Sometimes we expect to run some specific tests after certain kinds of reset. For example, we want to test if the reset reason is correct after a wake up from deep sleep. We need to create a deep-sleep reset first and then check the reset reason. To support this, we can define multi-stage test cases, to group a set of test functions:

```
static void trigger_deepsleep(void)
{
    esp_sleep_enable_timer_wakeup(2000);
    esp_deep_sleep_start();
}

void check_deepsleep_reset_reason()
{
    soc_reset_reason_t reason = esp_rom_get_reset_reason(0);
    TEST_ASSERT(reason == RESET_REASON_CORE_DEEP_SLEEP);
}

TEST_CASE_MULTIPLE_STAGES("reset reason check for deepsleep", "[esp32p4]", trigger_
↪deepsleep, check_deepsleep_reset_reason);
```

Multi-stage test cases present a group of test functions to users. It needs user interactions (select cases and select different stages) to run the case.

4.23.4 Tests For Different Targets

Some tests (especially those related to hardware) cannot run on all targets. Below is a guide how to make your unit tests run on only specified targets.

1. Wrap your test code by `!(TEMPORARY_)DISABLED_FOR_TARGETS()` macros and place them either in the original test file, or separate the code into files grouped by functions, but make sure all these files will be processed by the compiler. E.g.:

```
#if !TEMPORARY_DISABLED_FOR_TARGETS(ESP32, ESP8266)
TEST_CASE("a test that is not ready for esp32 and esp8266 yet", "[]")
{
}
#endif //!TEMPORARY_DISABLED_FOR_TARGETS(ESP32, ESP8266)
```

Once you need one of the tests to be compiled on a specified target, just modify the targets in the disabled list. It's more encouraged to use some general conception that can be described in `soc_caps.h` to control the disabling of tests. If this is done but some of the tests are not ready yet, use both of them (and remove `!(TEMPORARY_)DISABLED_FOR_TARGETS()` later). E.g.:

```
#if SOC_SDIO_SLAVE_SUPPORTED
#if !TEMPORARY_DISABLED_FOR_TARGETS(ESP64)
TEST_CASE("a sdio slave tests that is not ready for esp64 yet", "[sdio_slave]")
{
    //available for esp32 now, and will be available for esp64 in the future
}
#endif //!TEMPORARY_DISABLED_FOR_TARGETS(ESP64)
#endif //SOC_SDIO_SLAVE_SUPPORTED
```


2. For test code that you are 100% for sure that will not be supported (e.g., no peripheral at all), use `DISABLED_FOR_TARGETS`; for test code that should be disabled temporarily, or due to lack of runners, etc., use `TEMPORARY_DISABLED_FOR_TARGETS`.

Some old ways of disabling unit tests for targets, that have obvious disadvantages, are deprecated:

- DON'T put the test code under `test/target` folder and use `CMakeLists.txt` to choose one of the target folder. This is prevented because test code is more likely to be reused than the implementations. If you put something into `test/esp32` just to avoid building it on `esp32s2`, it's hard to make the code tidy if you want to enable the test again on `esp32s3`.
- DON'T use `CONFIG_IDF_TARGET_XXX` macros to disable the test items any more. This makes it harder to track disabled tests and enable them again. Also, a black-list style `#if !disabled` is preferred to white-list style `#if CONFIG_IDF_TARGET_XXX`, since you will not silently disable cases when new targets are added in the future. But for test implementations, it's allowed to use `#if CONFIG_IDF_TARGET_XXX` to pick one of the implementation code.
 - Test item: some items that will be performed on some targets, but skipped on other targets. E.g. There are three test items SD 1-bit, SD 4-bit and SDSPI. For ESP32-S2, which doesn't have SD host, among the tests only SDSPI is enabled on ESP32-S2.
 - Test implementation: some code will always happen, but in different ways. E.g. There is no SDIO PKT_LEN register on ESP8266. If you want to get the length from the slave as a step in the test process, you can have different implementation code protected by `#if CONFIG_IDF_TARGET_` reading in different ways. But please avoid using `#else` macro. When new target is added, the test case will fail at building stage, so that the maintainer will be aware of this, and choose one of the implementations explicitly.

4.23.5 Building Unit Test App

Follow the setup instructions in the top-level esp-idf README. Make sure that `IDF_PATH` environment variable is set to point to the path of esp-idf top-level directory.

Change into `tools/unit-test-app` directory to configure and build it:

- `idf.py menuconfig` - configure unit test app.
- `idf.py -T all build` - build unit test app with tests for each component having tests in the `test` subdirectory.
- `idf.py -T "xxx yyy" build` - build unit test app with tests for some space-separated specific components (For instance: `idf.py -T heap build` - build unit tests only for `heap` component directory).
- `idf.py -T all -E "xxx yyy" build` - build unit test app with all unit tests, except for unit tests of some components (For instance: `idf.py -T all -E "ulp mbedtls" build` - build all unit tests excludes `ulp` and `mbedtls` components).

Note: Due to inherent limitations of Windows command prompt, following syntax has to be used in order to build `unit-test-app` with multiple components: `idf.py -T xxx -T yyy build` or with escaped quotes: `idf.py -T \"xxx yyy\" build` in PowerShell or `idf.py -T ^"ssd1306 hts221\" build` in Windows command prompt.

When the build finishes, it will print instructions for flashing the chip. You can simply run `idf.py flash` to flash all build output.

You can also run `idf.py -T all flash` or `idf.py -T xxx flash` to build and flash. Everything needed will be rebuilt automatically before flashing.

Use `menuconfig` to set the serial port for flashing. For more information, see [tools/unit-test-app/README.md](https://github.com/espressif/esp-idf/blob/master/tools/unit-test-app/README.md).

4.23.6 Running Unit Tests

Note: We also provide the pytest-based framework [pytest-embedded](#) to help make running unit-tests more convenient and efficient. If you need to run tests in CI or run multiple tests in a row we recommend checking out this project. For more information see [Pytest-embedded Docs](#) and [pytest in ESP-IDF](#).

After flashing reset the ESP32-P4 and it will boot the unit test app.

When unit test app is idle, press "Enter" will make it print test menu with all available tests:

```
Here's the test menu, pick your combo:
(1)    "esp_ota_begin() verifies arguments" [ota]
(2)    "esp_ota_get_next_update_partition logic" [ota]
(3)    "Verify bootloader image in flash" [bootloader_support]
(4)    "Verify unit test app image" [bootloader_support]
(5)    "can use new and delete" [cxx]
(6)    "can call virtual functions" [cxx]
(7)    "can use static initializers for non-POD types" [cxx]
(8)    "can use std::vector" [cxx]
(9)    "static initialization guards work as expected" [cxx]
(10)   "global initializers run in the correct order" [cxx]
(11)   "before scheduler has started, static initializers work correctly" [cxx]
(12)   "adc2 work with wifi" [adc]
(13)   "gpio master/slave test example" [ignore][misc][test_env=UT_T2_1][multi_
↪device]
      (1)    "gpio_master_test"
      (2)    "gpio_slave_test"
(14)   "SPI Master clockdiv calculation routines" [spi]
(15)   "SPI Master test" [spi][ignore]
(16)   "SPI Master test, interaction of multiple devs" [spi][ignore]
(17)   "SPI Master no response when switch from host1 (SPI2) to host2 (SPI3)" ↪
↪[spi]
(18)   "SPI Master DMA test, TX and RX in different regions" [spi]
(19)   "SPI Master DMA test: length, start, not aligned" [spi]
(20)   "reset reason check for deepsleep" [esp32p4][test_env=UT_T2_1][multi_stage]
      (1)    "trigger_deepsleep"
      (2)    "check_deepsleep_reset_reason"
```

The normal case will print the case name and description. Master-slave cases will also print the sub-menu (the registered test function names).

Test cases can be run by inputting one of the following:

- Test case name in quotation marks to run a single test case
- Test case index to run a single test case
- Module name in square brackets to run all test cases for a specific module
- An asterisk to run all test cases

[multi_device] and [multi_stage] tags tell the test runner whether a test case is a multiple devices or multiple stages of test case. These tags are automatically added by `TEST_CASE_MULTIPLE_STAGES` and `TEST_CASE_MULTIPLE_DEVICES` macros.

After you select a multi-device test case, it will print sub-menu:

```
Running gpio master/slave test example...
gpio master/slave test example
      (1)    "gpio_master_test"
      (2)    "gpio_slave_test"
```

You need to input a number to select the test running on the DUT.

Similar to multi-device test cases, multi-stage test cases will also print sub-menu:

```
Running reset reason check for deepsleep...
reset reason check for deepsleep
(1)      "trigger_deepsleep"
(2)      "check_deepsleep_reset_reason"
```

First time you execute this case, input 1 to run first stage (trigger deepsleep). After DUT is rebooted and able to run test cases, select this case again and input 2 to run the second stage. The case only passes if the last stage passes and all previous stages trigger reset.

4.23.7 Timing Code with Cache Compensated Timer

Instructions and data stored in external memory (e.g., SPI Flash and SPI RAM) are accessed through the CPU's unified instruction and data cache. When code or data is in cache, access is very fast (i.e., a cache hit).

However, if the instruction or data is not in cache, it needs to be fetched from external memory (i.e., a cache miss). Access to external memory is significantly slower, as the CPU must execute stall cycles whilst waiting for the instruction or data to be retrieved from external memory. This can cause the overall code execution speed to vary depending on the number of cache hits or misses.

Code and data placements can vary between builds, and some arrangements may be more favorable with regards to cache access (i.e., minimizing cache misses). This can technically affect execution speed, however these factors are usually irrelevant as their effect 'average out' over the device's operation.

The effect of the cache on execution speed, however, can be relevant in benchmarking scenarios (especially micro benchmarks). There might be some variability in measured time between runs and between different builds. A technique for eliminating for some of the variability is to place code and data in instruction or data RAM (IRAM/DRAM), respectively. The CPU can access IRAM and DRAM directly, eliminating the cache out of the equation. However, this might not always be viable as the size of IRAM and DRAM is limited.

The cache compensated timer is an alternative to placing the code/data to be benchmarked in IRAM/DRAM. This timer uses the processor's internal event counters in order to determine the amount of time spent on waiting for code/data in case of a cache miss, then subtract that from the recorded wall time.

```
// Start the timer
ccomp_timer_start();

// Function to time
func_code_to_time();

// Stop the timer, and return the elapsed time in microseconds relative to
// ccomp_timer_start
int64_t t = ccomp_timer_stop();
```

One limitation of the cache compensated timer is that the task that benchmarked functions should be pinned to a core. This is due to each core having its own event counters that are independent of each other. For example, if `ccomp_timer_start` gets called on one core, put to sleep by the scheduler, wakes up, and gets rescheduled on the other core, then the corresponding `ccomp_timer_stop` will be invalid.

4.23.8 Mocks

Note: Currently, mocking is only possible with some selected components when running on the Linux host. In the future, we plan to make essential components in IDF mock-able. This will also include mocking when running on the ESP32-P4.

One of the biggest problems regarding unit testing on embedded systems are the strong hardware dependencies. Running unit tests directly on the ESP32-P4 can be especially difficult for higher layer components for the following reasons:

- Decreased test reliability due to lower layer components and/or hardware setup.
- Increased difficulty in testing edge cases due to limitations of lower layer components and/or hardware setup
- Increased difficulty in identifying the root cause due to the large number of dependencies influencing the behavior

When testing a particular component, (i.e., the component under test), mocking allows the dependencies of the component under test to be substituted (i.e., mocked) entirely in software. Through mocking, hardware details are emulated and specified at run time, but only if necessary. To allow mocking, ESP-IDF integrates the [CMock](#) mocking framework as a component. With the addition of some CMake functions in the ESP-IDF build system, it is possible to conveniently mock the entirety (or a part) of an IDF component.

Ideally, all components that the component under test is dependent on should be mocked, thus allowing the test environment complete control over all interactions with the component under test. However, if mocking all dependent components becomes too complex or too tedious (e.g., because you need to mock too many function calls) you have the following options:

- Include more "real" IDF code in the tests. This may work but increases the dependency on the "real" code's behavior. Furthermore, once a test fails, you may not know if the failure is in your actual code under test or the "real" IDF code.
- Re-evaluate the design of the code under test and attempt to reduce its dependencies by dividing the code under test into more manageable components. This may seem burdensome but it is quite common that unit tests expose software design weaknesses. Fixing design weaknesses will not only help with unit testing in the short term, but will help future code maintenance as well.

Refer to [cmock/CMock/docs/CMock_Summary.md](#) for more details on how CMock works and how to create and use mocks.

Requirements

Mocking with CMock requires `Ruby` on the host machine. Furthermore, since mocking currently only works on the Linux target, the requirements of the latter also need to be fulfilled:

- Installed ESP-IDF including all ESP-IDF requirements
- System package requirements (`libbsd`, `libbsd-dev`)
- A recent enough Linux or macOS version and GCC compiler
- All components the application depends on must be either supported on the Linux target (Linux/POSIX simulator) or mock-able

An application that runs on the Linux target has to set the `COMPONENTS` variable to `main` in the `CMakeLists.txt` of the application's root directory:

```
set(COMPONENTS main)
```

This prevents the automatic inclusion of all components from ESP-IDF to the build process which is otherwise done for convenience.

Mock a Component

If a mocked component, called a *component mock*, is already available in ESP-IDF, then it can be used right away as long as it satisfies the required functionality. Refer to [Component Linux/Mock Support Overview](#) to see which components are mocked already. Then refer to [Adjustments in Unit Test](#) in order to use the component mock.

It is necessary to create component mocks if they are not yet provided in ESP-IDF. To create a component mock, the component needs to be overwritten in a particular way. Overriding a component entails creating a component with the exact same name as the original component, then letting the build system discover it later than the original component (see [Multiple components with the same name](#) for more details).

In the component mock, the following parts are specified:

- The headers providing the functions to generate mocks for
- Include paths of the aforementioned headers
- Dependencies of the mock component (this is necessary e.g. if the headers include files from other components)

All these parts have to be specified using the IDF build system function `idf_component_mock`. You can use the IDF build system function `idf_component_get_property` with the tag `COMPONENT_OVERRIDEN_DIR` to access the component directory of the original component and then register the mock component parts using `idf_component_mock`:

```
idf_component_get_property(original_component_dir <original-component-name>_
↪COMPONENT_OVERRIDEN_DIR)
...
idf_component_mock(INCLUDE_DIRS "${original_component_dir}/include"
    REQUIRES freertos
    MOCK_HEADER_FILES ${original_component_dir}/include/header_containing_
↪functions_to_mock.h)
```

The component mock also requires a separate mock directory containing a `mock_config.yaml` file that configures CMock. A simple `mock_config.yaml` could look like this:

```
:cmock:
:plugins:
- expect
- expect_any_args
```

For more details about the CMock configuration yaml file, have a look at [cmock/CMock/docs/CMock_Summary.md](#).

Note that the component mock does not have to mock the original component in its entirety. As long as the test project's dependencies and dependencies of other code to the original components are satisfied by the component mock, partial mocking is adequate. In fact, most of the component mocks in IDF in `tools/mocks` are only partially mocking the original component.

Examples of component mocks can be found under [tools/mocks](#) in the IDF directory. General information on how to *override an IDF component* can be found in [Multiple components with the same name](#). There are several examples for testing code while mocking dependencies with CMock (non-exhaustive list):

- [unit test for the NVS Page class](#) .
- [unit test for esp_event](#) .
- [unit test for mqtt](#) .

Adjustments in Unit Test

The unit test needs to inform the cmake build system to mock dependent components (i.e., it needs to override the original component with the mock component). This is done by either placing the component mock into the project's `components` directory or adding the mock component's directory using the following line in the project's root `CMakeLists.txt`:

```
list(APPEND EXTRA_COMPONENT_DIRS "<mock_component_dir>")
```

Both methods will override existing components in ESP-IDF with the component mock. The latter is particularly convenient if you use component mocks that are already supplied by IDF.

Users can refer to the `esp_event` host-based unit test and its [esp_event/host_test/esp_event_unit_test/CMakeLists.txt](#) as an example of a component mock.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

4.24 Running ESP-IDF Applications on Host

Note: Running ESP-IDF applications on host is currently still an experimental feature, thus there is no guarantee for API stability. However, user feedback via the [ESP-IDF GitHub repository](#) or the [ESP32 forum](#) is highly welcome, and may help influence the future of design of the ESP-IDF host-based applications.

This document provides an overview of the methods to run ESP-IDF applications on Linux, and what type of ESP-IDF applications can typically be run on Linux.

4.24.1 Introduction

Typically, an ESP-IDF application is built (cross-compiled) on a host machine, uploaded (i.e., flashed) to an ESP chip for execution, and monitored by the host machine via a UART/USB port. However, execution of an ESP-IDF application on an ESP chip can be limiting in various development/usage/testing scenarios.

Therefore, it is possible for an ESP-IDF application to be built and executed entirely within the same Linux host machine (henceforth referred to as "running on host"). Running ESP-IDF applications on host has several advantages:

- No need to upload to a target.
- Faster execution on a host machine, compared to running on an ESP chip.
- No requirements for any specific hardware, except the host machine itself.
- Easier automation and setup for software testing.
- Large number of tools for code and runtime analysis, e.g., Valgrind.

A large number of ESP-IDF components depend on chip-specific hardware. These hardware dependencies must be mocked or simulated when running on host. ESP-IDF currently supports the following mocking and simulation approaches:

1. Using the [FreeRTOS POSIX/Linux simulator](#) that simulates FreeRTOS scheduling. On top of this simulation, other APIs are also simulated or implemented when running on host.
2. Using [CMock](#) to mock all dependencies and run the code in complete isolation.

In principle, it is possible to mix both approaches (POSIX/Linux simulator and mocking using CMock), but this has not been done yet in ESP-IDF. Note that despite the name, the FreeRTOS POSIX/Linux simulator currently also works on macOS. Running ESP-IDF applications on host machines is often used for testing. However, simulating the environment and mocking dependencies does not fully represent the target device. Thus, testing on the target device is still necessary, though with a different focus that usually puts more weight on integration and system testing.

Note: Another possibility to run applications on the host is to use the QEMU simulator. However, QEMU development for ESP-IDF applications is still a work in progress and has not been documented yet.

CMock-Based Approach

This approach uses the [CMock](#) framework to solve the problem of missing hardware and software dependencies. CMock-based applications running on the host machine have the added advantage that they usually only compile the necessary code, i.e., the (mostly mocked) dependencies instead of the entire system. For a general introduction to Mocks and how to configure and use them in ESP-IDF, please refer to [Mocks](#).

POSIX/Linux Simulator Approach

The [FreeRTOS POSIX/Linux simulator](#) is available on ESP-IDF as a preview target already. This simulator allows ESP-IDF components to be implemented on the host, making them accessible to ESP-IDF applications when running on host. Currently, only a limited number of components are ready to be built on Linux. Furthermore, the

functionality of each component ported to Linux may also be limited or different compared to the functionality when building that component for a chip target. For more information about whether the desired components are supported on Linux, please refer to [Component Linux/Mock Support Overview](#).

4.24.2 Requirements for Using Mocks

- Installed ESP-IDF including all ESP-IDF requirements
- System package requirements (`libbsd`, `libbsd-dev`)
- A recent enough Linux or macOS version and GCC compiler
- All components the application depends on must be either supported on the Linux target (Linux/POSIX simulator) or mock-able

An application that runs on the Linux target has to set the `COMPONENTS` variable to `main` in the `CMakeLists.txt` of the application's root directory:

```
set(COMPONENTS main)
```

This prevents the automatic inclusion of all components from ESP-IDF to the build process which is otherwise done for convenience.

If any mocks are used, then Ruby is required, too.

4.24.3 Build and Run

To build the application on Linux, the target has to be set to `linux` and then it can be built and run:

```
idf.py --preview set-target linux
idf.py build
idf.py monitor
```

4.24.4 Component Linux/Mock Support Overview

Note that any "Yes" here does not necessarily mean a full implementation or mocking. It can also mean a partial implementation or mocking of functionality. Usually, the implementation or mocking is done to a point where enough functionality is provided to build and run a test application.

Component	Mock	Simulation
cmock	No	Yes
driver	Yes	No
esp_common	No	Yes
esp_event	Yes	Yes
esp_http_client	No	Yes
esp_http_server	No	Yes
esp_https_server	No	Yes
esp_hw_support	Yes	Yes
esp_netif	Yes	Yes
esp_netif_stack	No	Yes
esp_partition	Yes	Yes
esp_rom	No	Yes
esp_system	No	Yes
esp_timer	Yes	No
esp_tls	Yes	Yes
fatfs	No	Yes
freertos	Yes	Yes
hal	No	Yes

continues on next page

Table 7 – continued from previous page

Component	Mock	Simulation
heap	No	Yes
http_parser	Yes	Yes
json	No	Yes
linux	No	Yes
log	No	Yes
lwip	Yes	Yes
mbedtls	No	Yes
mqtt	No	Yes
nvs_flash	No	Yes
partition_table	No	Yes
protobuf-c	No	Yes
pthread	No	Yes
soc	No	Yes
spiffs	No	Yes
spi_flash	Yes	No
tcp_transport	Yes	No
unity	No	Yes

4.25 Low Power Mode User Guide

The document has not been translated into English yet. In the meantime, please refer to the Chinese version.

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.

This warning was automatically inserted due to the source file being in the *add_warnings_pages* list.

Chapter 5

Security Guides

5.1 Overview

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

5.1.1 Security

This guide provides an overview of the overall security features available in various Espressif solutions. It is highly recommended to consider this guide while designing the products with the Espressif platform and the ESP-IDF software stack from the **security** perspective.

Goals

High level security goals are as follows:

1. Preventing untrustworthy code from being executed
2. Protecting the identity and integrity of the code stored in the off-chip flash memory
3. Securing device identity
4. Secure storage for confidential data
5. Authenticated and encrypted communication from the device

Platform Security

Secure Boot The Secure Boot feature ensures that only authenticated software can execute on the device. The Secure Boot process forms a chain of trust by verifying all **mutable** software entities involved in the [Application Startup Flow](#). Signature verification happens during both boot-up as well as in OTA updates.

Please refer to [Secure Boot V2](#) for detailed documentation about this feature.

Important: It is highly recommended that Secure Boot be enabled on all production devices.

Secure Boot Best Practices

- Generate the signing key on a system with a quality source of entropy.
- Always keep the signing key private. A leak of this key will compromise the Secure Boot system.
- Do not allow any third party to observe any aspects of the key generation or signing process using `espsecure.py`. Both processes are vulnerable to timing or other side-channel attacks.
- Ensure that all security eFuses have been correctly programmed, including disabling of the debug interfaces, non-required boot mediums (e.g., UART DL mode), etc.

Flash Encryption The Flash Encryption feature helps to encrypt the contents on the off-chip flash memory and thus provides the **confidentiality** aspect to the software or data stored in the flash memory.

Please refer to [Flash Encryption](#) for detailed information about this feature.

If ESP32-P4 is connected to an external SPI RAM, the contents written to or read from the SPI RAM will also be encrypted and decrypted respectively (via the MMU's flash cache, provided that Flash Encryption is enabled). This provides an additional safety layer for the data stored in SPI RAM, hence configurations like `CONFIG_MBEDTLS_EXTERNAL_MEM_ALLOC` can be safely enabled in this case.

Flash Encryption Best Practices

- It is recommended to use flash Encryption release mode for the production use-cases.
- It is recommended to have a unique flash encryption key per device.
- Enable [Secure Boot](#) as an extra layer of protection, and to prevent an attacker from selectively corrupting any part of the flash before boot.

Device Identity The Digital Signature peripheral in ESP32-P4 produces hardware-accelerated RSA digital signatures with the assistance of HMAC, without the RSA private key being accessible by software. This allows the private key to be kept secured on the device without anyone other than the device hardware being able to access it.

ESP32-P4 also supports ECDSA peripheral for generating hardware-accelerated ECDSA digital signatures. ECDSA private key can be directly programmed in an eFuse block and marked as read protected from the software.

DS or ECDSA peripheral can help to establish the **Secure Device Identity** to the remote endpoint, e.g., in the case of TLS mutual authentication based on the RSA or ECDSA cipher scheme.

Please refer to the [Elliptic Curve Digital Signature Algorithm \(ECDSA\)](#) and [Digital Signature \(DS\)](#) guides for detailed documentation.

Memory Protection ESP32-P4 supports the **Memory Protection** scheme, either through architecture or special peripheral like PMS, which provides an ability to enforce and monitor permission attributes to memory and, in some cases, peripherals. ESP-IDF application startup code configures the permissions attributes like Read/Write access on data memories and Read/Execute access on instruction memories using the relevant peripheral. If there is any attempt made that breaks these permission attributes, e.g., a write operation to instruction memory region, then a violation interrupt is raised, and it results in system panic.

This feature depends on the config option `CONFIG_ESP_SYSTEM_MEMPROT_FEATURE` and it is kept enabled by default. Please note that the API for this feature is **private** and used exclusively by ESP-IDF code only.

Note: This feature can help to prevent the possibility of remote code injection due to the existing vulnerabilities in the software.

Debug Interfaces

JTAG

- JTAG interface stays disabled if any of the security features are enabled. Please refer to [JTAG with Flash Encryption or Secure Boot](#) for more information.
- JTAG interface can also be disabled in the absence of any other security features using [eFuse API](#).
- ESP32-P4 supports soft disabling the JTAG interface and it can be re-enabled by programming a secret key through HMAC. ([HMAC for Enabling JTAG](#))

UART Download Mode In ESP32-P4, Secure UART Download mode gets activated if any of the security features are enabled.

- Secure UART Download mode can also be enabled by calling `esp_efuse_enable_rom_secure_download_mode()`.
- This mode does not allow any arbitrary code to execute if downloaded through the UART download mode.
- It also limits the available commands in Download mode to update SPI config, e.g., changing baud rate, basic flash write, and the command to return a summary of currently enabled security features (`get_security_info`).
- To disable Download Mode entirely, select the [CONFIG_SECURE_UART_ROM_DL_MODE](#) to the recommended option Permanently disable ROM Download Mode or call `esp_efuse_disable_rom_download_mode()` at runtime.

Important: In Secure UART Download mode, `esptool.py` can only work with the argument `--no-stub`.

Product Security

Secure Provisioning Secure Provisioning refers to a process of secure on-boarding of the ESP device on to the Wi-Fi network. This mechanism also allows provision of additional custom configuration data during the initial provisioning phase from the provisioning entity, e.g., Smartphone.

ESP-IDF provides various security schemes to establish a secure session between ESP and the provisioning entity, they are highlighted at `provisioning_security_schemes`.

Please refer to the `../api-reference/provisioning/wifi_provisioning` documentation for details and the example code for this feature.

Note: Espressif provides Android and iOS Phone Apps along with their sources, so that it could be easy to further customize them as per the product requirement.

Secure OTA (Over-the-air) Updates

- OTA Updates must happen over secure transport, e.g., HTTPS.
- ESP-IDF provides a simplified abstraction layer [ESP HTTPS OTA](#) for this.
- If [Secure Boot](#) is enabled, then the server should host the signed application image.
- If [Flash Encryption](#) is enabled, then no additional steps are required on the server side, encryption shall be taken care on the device itself during flash write.
- OTA update [Rollback Process](#) can help to switch the application as `active` only after its functionality has been verified.

Anti-Rollback Protection Anti-rollback protection feature ensures that device only executes the application that meets the security version criteria as stored in its eFuse. So even though the application is trusted and signed by legitimate key, it may contain some revoked security feature or credential. Hence, device must reject any such application.

ESP-IDF allows this feature for the application only and it is managed through 2nd stage bootloader. The security version is stored in the device eFuse and it is compared against the application image header during both bootup and over-the-air updates.

Please see more information to enable this feature in the [Anti-rollback](#) guide.

Encrypted Firmware Distribution Encrypted firmware distribution during over-the-air updates ensures that the application stays encrypted **in transit** from the server to the device. This can act as an additional layer of protection on top of the TLS communication during OTA updates and protect the identity of the application.

Please see working example for this documented in [OTA Upgrades with Pre-Encrypted Firmware](#) section.

Secure Storage Secure storage refers to the application-specific data that can be stored in a secure manner on the device, i.e., off-chip flash memory. This is typically a read-write flash partition and holds device specific configuration data, e.g., Wi-Fi credentials.

ESP-IDF provides the **NVS (Non-volatile Storage)** management component which allows encrypted data partitions. This feature is tied with the platform [Flash Encryption](#) feature described earlier.

Please refer to the [NVS Encryption](#) for detailed documentation on the working and instructions to enable this feature.

Important: By default, ESP-IDF components writes the device specific data into the default NVS partition, including Wi-Fi credentials too, and it is recommended to protect this data using **NVS Encryption** feature.

Secure Device Control ESP-IDF provides capability to control an ESP device over Wi-Fi + HTTP or BLE in a secure manner using ESP Local Control component.

Please refer to the [ESP Local Control](#) for detailed documentation about this feature.

Security Policy

The ESP-IDF GitHub repository has attached [Security Policy Brief](#).

Advisories

- Espressif publishes critical [Security Advisories](#), which includes security advisories regarding both hardware and software.
- The specific advisories of the ESP-IDF software components are published through the [GitHub repository](#).

Software Updates Critical security issues in the ESP-IDF components, and third-party libraries are fixed as and when we find them or when they are reported to us. Gradually, we make the fixes available in all applicable release branches in ESP-IDF.

Applicable security issues and CVEs for the ESP-IDF components, third-party libraries are mentioned in the ESP-IDF release notes.

Important: We recommend periodically updating to the latest bugfix version of the ESP-IDF release to have all critical security fixes available.

5.2 Features

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

5.2.1 Flash Encryption

This is a quick start guide to ESP32-P4's flash encryption feature. Using application code as an example, it demonstrates how to test and verify flash encryption operations during development and production.

Introduction

Flash encryption is intended for encrypting the contents of the ESP32-P4's off-chip flash memory. Once this feature is enabled, firmware is flashed as plaintext, and then the data is encrypted in place on the first boot. As a result, physical readout of flash will not be sufficient to recover most flash contents.

Important: For production use, flash encryption should be enabled in the "Release" mode only.

Important: Enabling flash encryption limits the options for further updates of ESP32-P4. Before using this feature, read the document and make sure to understand the implications.

Encrypted Partitions

With flash encryption enabled, the following types of data are encrypted by default:

- [Second Stage Bootloader](#) (Firmware Bootloader)
- Partition Table
- [NVS Key Partition](#)
- Otadata
- All app type partitions

Other types of data can be encrypted conditionally:

- Any partition marked with the `encrypted` flag in the partition table. For details, see [Encrypted Partition Flag](#).
- Secure Boot bootloader digest if Secure Boot is enabled (see below).

Relevant eFuses

The flash encryption operation is controlled by various eFuses available on ESP32-P4. The list of eFuses and their descriptions is given in the table below. The names in eFuse column are also used by `espfuse.py` tool. For usage in the eFuse API, modify the name by adding `ESP_EFUSE_`, for example: `esp_efuse_read_field_bit(ESP_EFUSE_DISABLE_DL_ENCRYPT)`.

Table 1: eFuses Used in Flash Encryption

eFuse	Description	Bit Depth
BLOCK_KEYN	AES key storage. N is between 0 and 5.	256 bit key block
KEY_PURPOSE_N	Control the purpose of eFuse block BLOCK_KEYN, where N is between 0 and 5. For flash encryption, the only valid value is 4 for XTS_AES_128_KEY.	4
DIS_DOWNLOAD_MANUAL_ENCRYPT	If set, disable flash encryption when in download boot-modes.	1
SPI_BOOT_CRYPT_CNT	Enable encryption and decryption, when an SPI boot mode is set. Feature is enabled if 1 or 3 bits are set in the eFuse, disabled otherwise.	3

Note:

- R/W access control is available for all the eFuse bits listed in the table above.
 - The default value of these bits is 0 after manufacturing.
-

Read and write access to eFuse bits is controlled by appropriate fields in the registers `WR_DIS` and `RD_DIS`. For more information on ESP32-P4 eFuses, see [eFuse manager](#). To change protection bits of eFuse field using `espefuse.py`, use these two commands: `read_protect_efuse` and `write_protect_efuse`. Example `espefuse.py write_protect_efuse DISABLE_DL_ENCRYPT`.

Flash Encryption Process

Assuming that the eFuse values are in their default states and the firmware bootloader is compiled to support flash encryption, the flash encryption process executes as shown below:

1. On the first power-on reset, all data in flash is un-encrypted (plaintext). The ROM bootloader loads the firmware bootloader.
2. Firmware bootloader reads the `SPI_BOOT_CRYPT_CNT` eFuse value (0b000). Since the value is 0 (even number of bits set), it configures and enables the flash encryption block. For more information on the flash encryption block, see [ESP32-P4 Technical Reference Manual](#).
3. Firmware bootloader uses RNG (random) module to generate an 256 bit key and then writes it into `BLOCK_KEYN` eFuse. The software also updates the `KEY_PURPOSE_N` for the block where the key is stored. The key cannot be accessed via software as the write and read protection bits for `BLOCK_KEYN` eFuse are set. `KEY_PURPOSE_N` field is write-protected as well. The flash encryption is completely conducted by hardware, and the key cannot be accessed via software. If a valid key is already present in the eFuse (e.g., burned using `espefuse` tool) then the process of key generation is skipped and the same key is used for flash encryption process.
4. Flash encryption block encrypts the flash contents - the firmware bootloader, applications and partitions marked as encrypted. Encrypting in-place can take time, up to a minute for large partitions.
5. Firmware bootloader sets the first available bit in `SPI_BOOT_CRYPT_CNT` (0b001) to mark the flash contents as encrypted. Odd number of bits is set.
6. For [Development Mode](#), the firmware bootloader allows the UART bootloader to re-flash encrypted binaries. Also, the `SPI_BOOT_CRYPT_CNT` eFuse bits are NOT write-protected. In addition, the firmware bootloader by default sets the eFuse bits `DIS_DOWNLOAD_ICACHE`, `DIS_PAD_JTAG`, `DIS_USB_JTAG` and `DIS_LEGACY_SPI_BOOT`.
7. For [Release Mode](#), the firmware bootloader sets all the eFuse bits set under development mode as well as `DIS_DOWNLOAD_MANUAL_ENCRYPT`. It also write-protects the `SPI_BOOT_CRYPT_CNT` eFuse bits. To modify this behavior, see [Enabling UART Bootloader Encryption/Decryption](#).
8. The device is then rebooted to start executing the encrypted image. The firmware bootloader calls the flash decryption block to decrypt the flash contents and then loads the decrypted contents into IRAM.

During the development stage, there is a frequent need to program different plaintext flash images and test the flash encryption process. This requires that Firmware Download mode is able to load new plaintext images as many times as it might be needed. However, during manufacturing or production stages, Firmware Download mode should not be allowed to access flash contents for security reasons.

Hence, two different flash encryption configurations were created: for development and for production. For details on these configurations, see Section [Flash Encryption Configuration](#).

Flash Encryption Configuration

The following flash encryption modes are available:

- [Development Mode](#) - recommended for use only during development. In this mode, it is still possible to flash new plaintext firmware to the device, and the bootloader will transparently encrypt this firmware using the key stored in hardware. This allows, indirectly, to read out the plaintext of the firmware in flash.
- [Release Mode](#) - recommended for manufacturing and production. In this mode, flashing plaintext firmware to the device without knowing the encryption key is no longer possible.

This section provides information on the mentioned flash encryption modes and step by step instructions on how to use them.

Development Mode During development, you can encrypt flash using either an ESP32-P4 generated key or external host-generated key.

Using ESP32-P4 Generated Key Development mode allows you to download multiple plaintext images using Firmware Download mode.

To test flash encryption process, take the following steps:

1. Ensure that you have an ESP32-P4 device with default flash encryption eFuse settings as shown in [Relevant eFuses](#).

See how to check [ESP32-P4 Flash Encryption Status](#).

2. In [Project Configuration Menu](#), do the following:

- [Enable flash encryption on boot](#).
- [Select encryption mode](#) (**Development mode** by default).
- [Select UART ROM download mode](#) (**enabled** by default).
- [Select the appropriate bootloader log verbosity](#).
- Save the configuration and exit.

Enabling flash encryption will increase the size of bootloader, which might require updating partition table offset. See [Bootloader Size](#).

3. Run the command given below to build and flash the complete images.

```
idf.py flash monitor
```

Note: This command does not include any user files which should be written to the partitions on the flash memory. Please write them manually before running this command otherwise the files should be encrypted separately before writing.

This command will write to flash memory unencrypted images: the firmware bootloader, the partition table and applications. Once the flashing is complete, ESP32-P4 will reset. On the next boot, the firmware bootloader encrypts: the firmware bootloader, application partitions and partitions marked as encrypted then resets. Encrypting in-place can take time, up to a minute for large partitions. After that, the application is decrypted at runtime and executed.

A sample output of the first ESP32-P4 boot after enabling flash encryption is given below:

```
TO BE UPDATED TODO IDF-7747
```

A sample output of subsequent ESP32-P4 boots just mentions that flash encryption is already enabled:

```
TO BE UPDATED TODO IDF-7747
```

At this stage, if you need to update and re-flash binaries, see [Re-flashing Updated Partitions](#).

Using Host Generated Key It is possible to pre-generate a flash encryption key on the host computer and burn it into the eFuse. This allows you to pre-encrypt data on the host and flash already encrypted data without needing a plaintext flash update. This feature can be used in both [Development Mode](#) and [Release Mode](#). Without a pre-generated key, data is flashed in plaintext and then ESP32-P4 encrypts the data in-place.

Note: This option is not recommended for production, unless a separate key is generated for each individual device.

To use a host generated key, take the following steps:

1. Ensure that you have an ESP32-P4 device with default flash encryption eFuse settings as shown in [Relevant eFuses](#).

See how to check [ESP32-P4 Flash Encryption Status](#).

2. Generate a random key by running:

```
espsecure.py generate_flash_encryption_key my_flash_encryption_key.bin
```

3. **Before the first encrypted boot**, burn the key into your device's eFuse using the command below. This action can be done **only once**.

```
espefuse.py --port PORT burn_key BLOCK my_flash_encryption_key.bin XTS_
→AES_128_KEY
```

where BLOCK is a free keyblock between BLOCK_KEY0 and BLOCK_KEY5.

If the key is not burned and the device is started after enabling flash encryption, the ESP32-P4 will generate a random key that software cannot access or modify.

4. In [Project Configuration Menu](#), do the following:
 - [Enable flash encryption on boot](#)
 - [Select encryption mode](#) (**Development mode** by default)
 - [Select the appropriate bootloader log verbosity](#)
 - Save the configuration and exit.

Enabling flash encryption will increase the size of bootloader, which might require updating partition table offset. See [Bootloader Size](#).

5. Run the command given below to build and flash the complete images.

```
idf.py flash monitor
```

Note: This command does not include any user files which should be written to the partitions on the flash memory. Please write them manually before running this command otherwise the files should be encrypted separately before writing.

This command will write to flash memory unencrypted images: the firmware bootloader, the partition table and applications. Once the flashing is complete, ESP32-P4 will reset. On the next boot, the firmware bootloader encrypts: the firmware bootloader, application partitions and partitions marked as `encrypted` then resets. Encrypting in-place can take time, up to a minute for large partitions. After that, the application is decrypted at runtime and executed.

If using Development Mode, then the easiest way to update and re-flash binaries is [Re-flashing Updated Partitions](#).

If using Release Mode, then it is possible to pre-encrypt the binaries on the host and then flash them as ciphertext. See [Manually Encrypting Files](#).

Re-flashing Updated Partitions If you update your application code (done in plaintext) and want to re-flash it, you will need to encrypt it before flashing. To encrypt the application and flash it in one step, run:

```
idf.py encrypted-app-flash monitor
```

If all partitions needs to be updated in encrypted format, run:

```
idf.py encrypted-flash monitor
```


Release Mode In Release mode, UART bootloader cannot perform flash encryption operations. New plaintext images can ONLY be downloaded using the over-the-air (OTA) scheme which will encrypt the plaintext image before writing to flash.

To use this mode, take the following steps:

1. Ensure that you have an ESP32-P4 device with default flash encryption eFuse settings as shown in [Relevant eFuses](#).

See how to check [ESP32-P4 Flash Encryption Status](#).

2. In [Project Configuration Menu](#), do the following:

- [Enable flash encryption on boot](#).
- [Select Release mode](#). (Note that once Release mode is selected, the EFUSE_DIS_DOWNLOAD_MANUAL_ENCRYPT eFuse bit will be burned to disable flash encryption hardware in ROM Download Mode.)
- [Select UART ROM download mode \(Permanently switch to Secure mode \(recommended\)\)](#). This is the default option, and is recommended. It is also possible to change this configuration setting to permanently disable UART ROM download mode, if this mode is not needed.
- [Select the appropriate bootloader log verbosity](#).
- Save the configuration and exit.

Enabling flash encryption will increase the size of bootloader, which might require updating partition table offset. See [Bootloader Size](#).

3. Run the command given below to build and flash the complete images.

```
idf.py flash monitor
```

Note: This command does not include any user files which should be written to the partitions on the flash memory. Please write them manually before running this command otherwise the files should be encrypted separately before writing.

This command will write to flash memory unencrypted images: the firmware bootloader, the partition table and applications. Once the flashing is complete, ESP32-P4 will reset. On the next boot, the firmware bootloader encrypts: the firmware bootloader, application partitions and partitions marked as `encrypted` then resets. Encrypting in-place can take time, up to a minute for large partitions. After that, the application is decrypted at runtime and executed.

Once the flash encryption is enabled in Release mode, the bootloader will write-protect the SPI_BOOT_CRYPT_CNT eFuse.

For subsequent plaintext field updates, use [OTA scheme](#).

Note: If you have pre-generated the flash encryption key and stored a copy, and the UART download mode is not permanently disabled via [CONFIG_SECURE_UART_ROM_DL_MODE](#), then it is possible to update the flash locally by pre-encrypting the files and then flashing the ciphertext. See [Manually Encrypting Files](#).

Best Practices When using Flash Encryption in production:

- Do not reuse the same flash encryption key between multiple devices. This means that an attacker who copies encrypted data from one device cannot transfer it to a second device.
- The UART ROM Download Mode should be disabled entirely if it is not needed, or permanently set to "Secure Download Mode" otherwise. Secure Download Mode permanently limits the available commands to updating SPI config, changing baud rate, basic flash write, and returning a summary of the currently enabled security features with the `get_security_info` command. The default behaviour is to set Se-

cure Download Mode on first boot in Release mode. To disable Download Mode entirely, select [CONFIG_SECURE_UART_ROM_DL_MODE](#) to "Permanently disable ROM Download Mode (recommended)" or call `esp_efuse_disable_rom_download_mode()` at runtime.

- Enable [Secure Boot](#) as an extra layer of protection, and to prevent an attacker from selectively corrupting any part of the flash before boot.

Enable Flash Encryption Externally

In the process mentioned above, flash encryption related eFuses which ultimately enable flash encryption are programmed through the firmware bootloader. Alternatively, all the eFuses can be programmed with the help of `es-pefuse` tool. Please refer [Enable Flash Encryption Externally](#) for more details.

Possible Failures

Once flash encryption is enabled, the `SPI_BOOT_CRYPT_CNT` eFuse value will have an odd number of bits set. It means that all the partitions marked with the encryption flag are expected to contain encrypted ciphertext. Below are the three typical failure cases if the ESP32-P4 is erroneously loaded with plaintext data:

1. If the bootloader partition is re-flashed with a **plaintext firmware bootloader image**, the ROM bootloader will fail to load the firmware bootloader resulting in the following failure:

```
rst:0x3 (SW_RESET),boot:0x13 (SPI_FAST_FLASH_BOOT)
invalid header: 0xb414f76b
invalid header: 0xb414f76b
invalid header: 0xb414f76b
invalid header: 0xb414f76b
invalid header: 0xb414f76b
invalid header: 0xb414f76b
invalid header: 0xb414f76b
```

Note: The value of invalid header will be different for every application.

Note: This error also appears if the flash contents are erased or corrupted.

2. If the firmware bootloader is encrypted, but the partition table is re-flashed with a **plaintext partition table image**, the bootloader will fail to read the partition table resulting in the following failure:

```
rst:0x3 (SW_RESET),boot:0x13 (SPI_FAST_FLASH_BOOT)
configsip: 0, SPIWP:0xee
clk_drv:0x00,q_drv:0x00,d_drv:0x00,cs0_drv:0x00,hd_drv:0x00,wp_drv:0x00
mode:DIO, clock div:2
load:0x3fff0018,len:4
load:0x3fff001c,len:10464
ho 0 tail 12 room 4
load:0x40078000,len:19168
load:0x40080400,len:6664
entry 0x40080764
I (60) boot: ESP-IDF v4.0-dev-763-g2c55fae6c-dirty 2nd stage bootloader
I (60) boot: compile time 19:15:54
I (62) boot: Enabling RNG early entropy source...
I (67) boot: SPI Speed      : 40MHz
I (72) boot: SPI Mode      : DIO
I (76) boot: SPI Flash Size : 4MB
E (80) flash_parts: partition 0 invalid magic number 0x94f6
E (86) boot: Failed to verify partition table
E (91) boot: load partition table error!
```

3. If the bootloader and partition table are encrypted, but the application is re-flashed with a **plaintext application image**, the bootloader will fail to load the application resulting in the following failure:

```
rst:0x3 (SW_RESET),boot:0x13 (SPI_FAST_FLASH_BOOT)
configsip: 0, SPIWP:0xee
clk_drv:0x00,q_drv:0x00,d_drv:0x00,cs0_drv:0x00,hd_drv:0x00,wp_drv:0x00
mode:DIO, clock div:2
load:0x3fff0018,len:4
load:0x3fff001c,len:8452
load:0x40078000,len:13616
load:0x40080400,len:6664
entry 0x40080764
I (56) boot: ESP-IDF v4.0-dev-850-gc4447462d-dirty 2nd stage bootloader
I (56) boot: compile time 15:37:14
I (58) boot: Enabling RNG early entropy source...
I (64) boot: SPI Speed      : 40MHz
I (68) boot: SPI Mode      : DIO
I (72) boot: SPI Flash Size : 4MB
I (76) boot: Partition Table:
I (79) boot: ## Label                Usage            Type ST Offset   Length
I (87) boot:  0 nvs                  WiFi data        01 02 0000a000 00006000
I (94) boot:  1 phy_init             RF data          01 01 00010000 00001000
I (102) boot:  2 factory              factory app       00 00 00020000 00100000
I (109) boot: End of partition table
E (113) esp_image: image at 0x20000 has invalid magic byte
W (120) esp_image: image at 0x20000 has invalid SPI mode 108
W (126) esp_image: image at 0x20000 has invalid SPI size 11
E (132) boot: Factory app partition is not bootable
E (138) boot: No bootable app partitions in the partition table
```

ESP32-P4 Flash Encryption Status

1. Ensure that you have an ESP32-P4 device with default flash encryption eFuse settings as shown in [Relevant eFuses](#).

To check if flash encryption on your ESP32-P4 device is enabled, do one of the following:

- flash the application example [security/flash_encryption](#) onto your device. This application prints the SPI_BOOT_CRYPT_CNT eFuse value and if flash encryption is enabled or disabled.
- [Find the serial port name](#) under which your ESP32-P4 device is connected, replace PORT with your port name in the following command, and run it:

```
espefuse.py -p PORT summary
```

Reading and Writing Data in Encrypted Flash

ESP32-P4 application code can check if flash encryption is currently enabled by calling `esp_flash_encryption_enabled()`. Also, a device can identify the flash encryption mode by calling `esp_get_flash_encryption_mode()`.

Once flash encryption is enabled, be more careful with accessing flash contents from code.

Scope of Flash Encryption Whenever the SPI_BOOT_CRYPT_CNT eFuse is set to a value with an odd number of bits, all flash content accessed via the MMU's flash cache is transparently decrypted. It includes:

- Executable application code in flash (IROM).
- All read-only data stored in flash (DROM).
- Any data accessed via `spi_flash_mmap()`.
- The firmware bootloader image when it is read by the ROM bootloader.

Important: The MMU flash cache unconditionally decrypts all existing data. Data which is stored unencrypted in flash memory will also be "transparently decrypted" via the flash cache and will appear to software as random garbage.

Reading from Encrypted Flash To read data without using a flash cache MMU mapping, you can use the partition read function `esp_partition_read()`. This function will only decrypt data when it is read from an encrypted partition. Data read from unencrypted partitions will not be decrypted. In this way, software can access encrypted and non-encrypted flash in the same way.

You can also use the following SPI flash API functions:

- `esp_flash_read()` to read raw (encrypted) data which will not be decrypted
- `esp_flash_read_encrypted()` to read and decrypt data

Data stored using the Non-Volatile Storage (NVS) API is always stored and read decrypted from the perspective of flash encryption. It is up to the library to provide encryption feature if required. Refer to [NVS Encryption](#) for more details.

Writing to Encrypted Flash It is recommended to use the partition write function `esp_partition_write()`. This function will only encrypt data when it is written to an encrypted partition. Data written to unencrypted partitions will not be encrypted. In this way, software can access encrypted and non-encrypted flash in the same way.

You can also pre-encrypt and write data using the function `esp_flash_write_encrypted()`

Also, the following ROM function exist but not supported in esp-idf applications:

- `esp_rom_spiflash_write_encrypted` pre-encrypts and writes data to flash
- `SPIWrite` writes unencrypted data to flash

Since data is encrypted in blocks, the minimum write size for encrypted data is 16 bytes and the alignment is also 16 bytes.

Updating Encrypted Flash

OTA Updates OTA updates to encrypted partitions will automatically write encrypted data if the function `esp_partition_write()` is used.

Before building the application image for OTA updating of an already encrypted device, enable the option [Enable flash encryption on boot](#) in project configuration menu.

For general information about ESP-IDF OTA updates, please refer to [OTA](#)

Updating Encrypted Flash via Serial Flashing an encrypted device via serial bootloader requires that the serial bootloader download interface has not been permanently disabled via eFuse.

In Development Mode, the recommended method is [Re-flashing Updated Partitions](#).

In Release Mode, if a copy of the same key stored in eFuse is available on the host then it is possible to pre-encrypt files on the host and then flash them. See [Manually Encrypting Files](#).

Disabling Flash Encryption

If flash encryption was enabled accidentally, flashing of plaintext data will soft-brick the ESP32-P4. The device will reboot continuously, printing the error `flash read err, 1000` or `invalid header: 0xxxxxxx`.

For flash encryption in Development mode, encryption can be disabled by burning the `SPI_BOOT_CRYPT_CNT` eFuse. It can only be done one time per chip by taking the following steps:

1. In *Project Configuration Menu*, disable *Enable flash encryption on boot*, then save and exit.
2. Open project configuration menu again and **double-check** that you have disabled this option! If this option is left enabled, the bootloader will immediately re-enable encryption when it boots.
3. With flash encryption disabled, build and flash the new bootloader and application by running `idf.py flash`.
4. Use `espefuse.py` (in `components/esptool_py/esptool`) to disable the `SPI_BOOT_CRYPT_CNT` by running:

```
espefuse.py burn_efuse SPI_BOOT_CRYPT_CNT
```

Reset the ESP32-P4. Flash encryption will be disabled, and the bootloader will boot as usual.

Key Points About Flash Encryption

- Flash memory contents is encrypted using XTS-AES-128. The flash encryption key is 256 bits and stored in one `BLOCK_KEYN` eFuse internal to the chip and, by default, is protected from software access.
- Flash access is transparent via the flash cache mapping feature of ESP32-P4 - any flash regions which are mapped to the address space will be transparently decrypted when read.
Some data partitions might need to remain unencrypted for ease of access or might require the use of flash-friendly update algorithms which are ineffective if the data is encrypted. NVS partitions for non-volatile storage cannot be encrypted since the NVS library is not directly compatible with flash encryption. For details, refer to *NVS Encryption*.
- If flash encryption might be used in future, the programmer must keep it in mind and take certain precautions when writing code that *uses encrypted flash*.
- If secure boot is enabled, re-flashing the bootloader of an encrypted device requires a "Re-flashable" secure boot digest (see *Flash Encryption and Secure Boot*).

Enabling flash encryption will increase the size of bootloader, which might require updating partition table offset. See *Bootloader Size*.

Important: Do not interrupt power to the ESP32-P4 while the first boot encryption pass is running. If power is interrupted, the flash contents will be corrupted and will require flashing with unencrypted data again. In this case, re-flashing will not count towards the flashing limit.

Limitations of Flash Encryption

Flash encryption protects firmware against unauthorised readout and modification. It is important to understand the limitations of the flash encryption feature:

- Flash encryption is only as strong as the key. For this reason, we recommend keys are generated on the device during first boot (default behaviour). If generating keys off-device, ensure proper procedure is followed and do not share the same key between all production devices.
- Not all data is stored encrypted. If storing data on flash, check if the method you are using (library, API, etc.) supports flash encryption.
- Flash encryption does not prevent an attacker from understanding the high-level layout of the flash. This is because the same AES key is used for every pair of adjacent 16 byte AES blocks. When these adjacent 16 byte blocks contain identical content (such as empty or padding areas), these blocks will encrypt to produce matching pairs of encrypted blocks. This may allow an attacker to make high-level comparisons between encrypted devices (i.e., to tell if two devices are probably running the same firmware version).
- Flash encryption alone may not prevent an attacker from modifying the firmware of the device. To prevent unauthorised firmware from running on the device, use flash encryption in combination with *Secure Boot*.

Flash Encryption and Secure Boot

It is recommended to use flash encryption in combination with Secure Boot. However, if Secure Boot is enabled, additional restrictions apply to device re-flashing:

- [OTA Updates](#) are not restricted, provided that the new app is signed correctly with the Secure Boot signing key.

Advanced Features

The following section covers advanced features of flash encryption.

Encrypted Partition Flag Some partitions are encrypted by default. Other partitions can be marked in the partition table description as requiring encryption by adding the flag `encrypted` to the partitions' flag field. As a result, data in these marked partitions will be treated as encrypted in the same manner as an app partition.

#	Name,	Type,	SubType,	Offset,	Size,	Flags
nvs,		data,	nvs,	0x9000,	0x6000	
phy_init,		data,	phy,	0xf000,	0x1000	
factory,		app,	factory,	0x10000,	1M	
secret_data,		0x40,	0x01,	0x20000,	256K,	encrypted

For details on partition table description, see [partition table](#).

Further information about encryption of partitions:

- Default partition tables do not include any encrypted data partitions.
- With flash encryption enabled, the `app` partition is always treated as encrypted and does not require marking.
- If flash encryption is not enabled, the flag "encrypted" has no effect.
- You can also consider protecting `phy_init` data from physical access, readout, or modification, by marking the optional `phy` partition with the flag `encrypted`.
- The `nvs` partition cannot be encrypted, because the NVS library is not directly compatible with flash encryption.

Enabling UART Bootloader Encryption/Decryption On the first boot, the flash encryption process burns by default the following eFuses:

- `DIS_DOWNLOAD_MANUAL_ENCRYPT` which disables flash encryption operation when running in UART bootloader boot mode.
- `DIS_DOWNLOAD_ICACHE` which disables the entire MMU flash cache when running in UART bootloader mode.
- `DIS_PAD_JTAG` and `DIS_USB_JTAG` which disables JTAG.
- `DIS_DIRECT_BOOT` (old name `DIS_LEGACY_SPI_BOOT`) which disables direct boot mode

However, before the first boot you can choose to keep any of these features enabled by burning only selected eFuses and write-protect the rest of eFuses with unset value 0. For example:

```
espefuse.py --port PORT burn_efuse DIS_DOWNLOAD_MANUAL_ENCRYPT
espefuse.py --port PORT write_protect_efuse DIS_DOWNLOAD_MANUAL_ENCRYPT
```

Note: Set all appropriate bits before write-protecting!

Write protection of all the three eFuses is controlled by one bit. It means that write-protecting one eFuse bit will inevitably write-protect all unset eFuse bits.

Write protecting these eFuses to keep them unset is not currently very useful, as `esptool.py` does not support reading encrypted flash.

JTAG Debugging By default, when Flash Encryption is enabled (in either Development or Release mode) then JTAG debugging is disabled via eFuse. The bootloader does this on first boot, at the same time it enables flash encryption.

See *JTAG with Flash Encryption or Secure Boot* for more information about using JTAG Debugging with Flash Encryption.

Manually Encrypting Files Manually encrypting or decrypting files requires the flash encryption key to be pre-burned in eFuse (see *Using Host Generated Key*) and a copy to be kept on the host. If the flash encryption is configured in Development Mode then it is not necessary to keep a copy of the key or follow these steps, the simpler *Re-flashing Updated Partitions* steps can be used.

The key file should be a single raw binary file (example: `key.bin`).

For example, these are the steps to encrypt the file `build/my-app.bin` to flash at offset `0x10000`. Run `espsecure.py` as follows:

```
espsecure.py encrypt_flash_data --aes_xts --keyfile /path/to/key.bin --address_
↪0x10000 --output my-app-ciphertext.bin build/my-app.bin
```

The file `my-app-ciphertext.bin` can then be flashed to offset `0x10000` using `esptool.py`. To see all of the command line options recommended for `esptool.py`, see the output printed when `idf.py build` succeeds.

Note: If the flashed ciphertext file is not recognized by the ESP32-P4 when it boots, check that the keys match and that the command line arguments match exactly, including the correct offset.

The command `espsecure.py decrypt_flash_data` can be used with the same options (and different input/output files), to decrypt ciphertext flash contents or a previously encrypted file.

External RAM

When Flash Encryption is enabled any data read from and written to external SPI RAM through the cache will also be encrypted/decrypted. This happens the same way and with the same key as for Flash Encryption. If Flash Encryption is enabled then encryption for external SPI RAM is also always enabled, it is not possible to separately control this functionality.

Technical Details

The following sections provide some reference information about the operation of flash encryption.

Flash Encryption Algorithm

- ESP32-P4 use the XTS-AES block cipher mode with 256 bit size for flash encryption.
- XTS-AES is a block cipher mode specifically designed for disc encryption and addresses the weaknesses other potential modes (e.g., AES-CTR) have for this use case. A detailed description of the XTS-AES algorithm can be found in [IEEE Std 1619-2007](#).
- The flash encryption key is stored in one `BLOCK_KEYN` eFuse and, by default, is protected from further writes or software readout.
- To see the full flash encryption algorithm implemented in Python, refer to the `_flash_encryption_operation()` function in the `espsecure.py` source code.

5.2.2 Secure Boot V2

Important: This document is about Secure Boot V2, supported on ESP32-P4

Secure Boot V2 uses RSA-PSS or ECDSA based app and bootloader ([Second Stage Bootloader](#)) verification. This document can also be used as a reference for signing apps using the RSA-PSS or ECDSA scheme without signing the bootloader.

Background

Secure Boot protects a device from running any unauthorized (i.e., unsigned) code by checking that each piece of software that is being booted is signed. On an ESP32-P4, these pieces of software include the second stage bootloader and each application binary. Note that the first stage bootloader does not require signing as it is ROM code thus cannot be changed.

ESP32-P4 has provision to choose between a RSA-PSS or ECDSA based secure boot verification scheme.

The Secure Boot process on the ESP32-P4 involves the following steps:

1. When the first stage bootloader loads the second stage bootloader, the second stage bootloader's RSA-PSS or ECDSA signature is verified. If the verification is successful, the second stage bootloader is executed.
2. When the second stage bootloader loads a particular application image, the application's RSA-PSS or ECDSA signature is verified. If the verification is successful, the application image is executed.

Advantages

- The RSA-PSS or ECDSA public key is stored on the device. The corresponding RSA-PSS or ECDSA private key is kept at a secret place and is never accessed by the device.
- Up to three public keys can be generated and stored in the chip during manufacturing.
- ESP32-P4 provides the facility to permanently revoke individual public keys. This can be configured conservatively or aggressively.
- Conservatively - The old key is revoked after the bootloader and application have successfully migrated to a new key. Aggressively - The key is revoked as soon as verification with this key fails.
- Same image format and signature verification method is applied for applications and software bootloader.
- No secrets are stored on the device. Therefore, it is immune to passive side-channel attacks (timing or power analysis, etc.)

Secure Boot V2 Process

This is an overview of the Secure Boot V2 Process. Instructions how to enable Secure Boot are supplied in section [How To Enable Secure Boot V2](#).

Secure Boot V2 verifies the bootloader image and application binary images using a dedicated *signature block*. Each image has a separately generated signature block which is appended to the end of the image.

Up to 3 signature blocks can be appended to the bootloader or application image in ESP32-P4.

Each signature block contains a signature of the preceding image as well as the corresponding RSA-3072, ECDSA-256, or ECDSA-192 public key. For more details about the format, refer to [Signature Block Format](#). A digest of the RSA-3072, ECDSA-256, or ECDSA-192 public key is stored in the eFuse.

The application image is not only verified on every boot but also on each over the air (OTA) update. If the currently selected OTA app image cannot be verified, the bootloader will fall back and look for another correctly signed application image.

The Secure Boot V2 process follows these steps:

1. On startup, the ROM code checks the Secure Boot V2 bit in the eFuse. If Secure Boot is disabled, a normal boot will be executed. If Secure Boot is enabled, the boot will proceed according to the following steps.
2. The ROM code verifies the bootloader's signature block ([Verifying a Signature Block](#)). If this fails, the boot process will be aborted.

3. The ROM code verifies the bootloader image using the raw image data, its corresponding signature block(s), and the eFuse (*Verifying an Image*). If this fails, the boot process will be aborted.
4. The ROM code executes the bootloader.
5. The bootloader verifies the application image's signature block (*Verifying a Signature Block*). If this fails, the boot process will be aborted.
6. The bootloader verifies the application image using the raw image data, its corresponding signature blocks and the eFuse (*Verifying an Image*). If this fails, the boot process will be aborted. If the verification fails but another application image is found, the bootloader will then try to verify that other image using steps 5 to 7. This repeats until a valid image is found or no other images are found.
7. The bootloader executes the verified application image.

Signature Block Format

The signature block starts on a 4 KB aligned boundary and has a flash sector of its own. The signature is calculated over all bytes in the image including the padding bytes (*Secure Padding*).

Note: ESP32-P4 has a provision to choose between RSA scheme and ECDSA scheme. Only one scheme can be used per device.

ECDSA provides similar security strength, compared to RSA, with shorter key lengths. Current estimates are that ECDSA with curve P-256 has an approximate equivalent strength to RSA with 3072-bit keys. However, ECDSA signature verification takes considerably more amount of time as compared to RSA signature verification.

RSA is recommended for use cases where fast bootup time is required whereas ECDSA is recommended for use cases where shorter key length is required.

The content of each signature block is shown in the following table:

Table 2: Content of a RSA Signature Block

Offset	Size (bytes)	Description
0	1	Magic byte
1	1	Version number byte (currently 0x02), 0x01 is for Secure Boot V1.
2	2	Padding bytes, Reserved. Should be zero.
4	32	SHA-256 hash of only the image content, not including the signature block.
36	384	RSA Public Modulus used for signature verification. (value 'n' in RFC8017).
420	4	RSA Public Exponent used for signature verification (value 'e' in RFC8017).
424	384	Pre-calculated R, derived from 'n'.
808	4	Pre-calculated M', derived from 'n'
812	384	RSA-PSS Signature result (section 8.1.1 of RFC8017) of image content, computed using following PSS parameters: SHA256 hash, MGF1 function, salt length 32 bytes, default trailer field (0xBC).
1196	4	CRC32 of the preceding 1196 bytes.
1200	16	Zero padding to length 1216 bytes.

Note: R and M' are used for hardware-assisted Montgomery Multiplication.

Table 3: Content of a ECDSA Signature Block

Offset	Size (bytes)	Description
0	1	Magic byte.
1	1	Version number byte (currently 0x03).
2	2	Padding bytes, Reserved. Should be zero.
4	32	SHA-256 hash of only the image content, not including the signature block.
36	1	Curve ID (1 for NIST192p curve. 2 for NIST256p curve).
37	64	ECDSA Public key: 32 byte X coordinate followed by 32 byte Y coordinate.
101	64	ECDSA Signature result (section 5.3.2 of RFC6090) of the image content: 32 byte R component followed by 32 byte S component.
165	1031	Reserved.
1196	4	CRC32 of the preceding 1196 bytes.
1200	16	Zero padding to length 1216 bytes.

The remainder of the signature sector is erased flash (0xFF) which allows writing other signature blocks after previous signature block.

Secure Padding

In Secure Boot V2 scheme, the application image is padded to the flash MMU page size boundary to ensure that only verified contents are mapped in the internal address space. This is known as secure padding. Signature of the image is calculated after padding and then signature block (4KB) gets appended to the image.

- Default flash MMU page size is 64KB
- Secure padding is applied through the option `--secure-pad-v2` in the `elf2image` conversion using `esptool.py`

Following table explains the Secure Boot V2 signed image with secure padding and signature block appended:

Table 4: Contents of a signed application

Offset	Size (KB)	Description
0	580	Unsigned application size (as an example)
580	60	Secure padding (aligned to next 64KB boundary)
640	4	Signature block

Note: Please note that the application image always starts on the next flash MMU page size boundary (default 64KB) and hence the space left over after the signature block shown above can be utilized to store any other data partitions (e.g., `nvs`).

Verifying a Signature Block

A signature block is "valid" if the first byte is 0xe7 and a valid CRC32 is stored at offset 1196. Otherwise it is invalid.

Verifying an Image

An image is "verified" if the public key stored in any signature block is valid for this device, and if the stored signature is valid for the image data read from flash.

1. Compare the SHA-256 hash digest of the public key embedded in the bootloader's signature block with the digest(s) saved in the eFuses. If public key's hash does not match any of the hashes from the eFuses, the verification fails.
2. Generate the application image digest and match it with the image digest in the signature block. If the digests do not match, the verification fails.
3. Use the public key to verify the signature of the bootloader image, using either RSA-PSS (section 8.1.2 of RFC8017) or ECDSA signature verification (section 5.3.3 of RFC6090) with the image digest calculated in step (2) for comparison.

Bootloader Size

Enabling Secure boot and/or flash encryption will increase the size of bootloader, which might require updating partition table offset. See [Bootloader Size](#).

In the case when [CONFIG_SECURE_BOOT_BUILD_SIGNED_BINARIES](#) is disabled, the bootloader is sector padded (4KB) using the `--pad-to-size` option in `elf2image` command of `esptool`.

eFuse Usage

- `SECURE_BOOT_EN` - Enables Secure Boot protection on boot.
- `KEY_PURPOSE_X` - Set the purpose of the key block on ESP32-P4 by programming `SECURE_BOOT_DIGESTX` ($X = 0, 1, 2$) into `KEY_PURPOSE_X` ($X = 0, 1, 2, 3, 4, 5$). Example: If `KEY_PURPOSE_2` is set to `SECURE_BOOT_DIGEST1`, then `BLOCK_KEY2` will have the Secure Boot V2 public key digest. The write-protection bit must be set (this field does not have a read-protection bit).
- `BLOCK_KEYX` - The block contains the data corresponding to its purpose programmed in `KEY_PURPOSE_X`. Stores the SHA-256 digest of the public key. SHA-256 hash of public key modulus, exponent, pre-calculated R & M' values (represented as 776 bytes –offsets 36 to 812 - as per the [Signature Block Format](#)) is written to an eFuse key block. The write-protection bit must be set, but the read-protection bit must not.
- `KEY_REVOKEX` - The revocation bits corresponding to each of the 3 key block. Ex. Setting `KEY_REVOKE2` revokes the key block whose key purpose is `SECURE_BOOT_DIGEST2`.
- `SECURE_BOOT_AGGRESSIVE_REVOKE` - Enables aggressive revocation of keys. The key is revoked as soon as verification with this key fails.

To ensure no trusted keys can be added later by an attacker, each unused key digest slot should be revoked (`KEY_REVOKEX`). It will be checked during app startup in `esp_secure_boot_init_checks()` and fixed unless [CONFIG_SECURE_BOOT_ALLOW_UNUSED_DIGEST_SLOTS](#) is enabled.

The key(s) must be readable in order to give software access to it. If the key(s) is read-protected then the software reads the key(s) as all zeros and the signature verification process will fail, and the boot process will be aborted.

How To Enable Secure Boot V2

1. Open the [Project Configuration Menu](#), in "Security features" set "Enable hardware Secure Boot in bootloader" to enable Secure Boot.
2. The "Secure Boot V2" option will be selected and the "App Signing Scheme" would be set to RSA by default. RSA is recommended because of faster verification time. You can choose between RSA and ECDSA scheme from the menu.
3. Specify the path to Secure Boot signing key, relative to the project directory.
4. Select the desired UART ROM download mode in "UART ROM download mode". By default, it is set to "Permanently switch to Secure mode" which is generally recommended. For production devices, the most secure option is to set it to "Permanently disabled".
5. Set other menuconfig options (as desired). Then exit menuconfig and save your configuration.
6. The first time you run `idf.py build`, if the signing key is not found then an error message will be printed with a command to generate a signing key via `espsecure.py generate_signing_key`.

Important: A signing key generated this way will use the best random number source available to the OS and its Python installation (`/dev/urandom` on OSX/Linux and `CryptGenRandom()` on Windows). If this random number source is weak, then the private key will be weak.

Important: For production environments, we recommend generating the key pair using openssl or another industry standard encryption program. See [Generating Secure Boot Signing Key](#) for more details.

7. Run `idf.py bootloader` to build a Secure Boot enabled bootloader. The build output will include a prompt for a flashing command, using `esptool.py write_flash`.
8. When you are ready to flash the bootloader, run the specified command (you have to enter it yourself, this step is not performed by the build system) and then wait for flashing to complete.
9. Run `idf.py flash` to build and flash the partition table and the just-built app image. The app image will be signed using the signing key you generated in step 6.

Note: `idf.py flash` does not flash the bootloader if Secure Boot is enabled.

10. Reset the ESP32-P4 and it will boot the software bootloader you flashed. The software bootloader will enable Secure Boot on the chip, and then it verifies the app image signature and boots the app. You should watch the serial console output from the ESP32-P4 to verify that Secure Boot is enabled and no errors have occurred due to the build configuration.

Note: Secure boot will not be enabled until after a valid partition table and app image have been flashed. This is to prevent accidents before the system is fully configured.

Note: If the ESP32-P4 is reset or powered down during the first boot, it will start the process again on the next boot.

11. On subsequent boots, the Secure Boot hardware will verify the software bootloader has not changed and the software bootloader will verify the signed app image (using the validated public key portion of its appended signature block).

Restrictions After Secure Boot Is Enabled

- Any updated bootloader or app will need to be signed with a key matching the digest already stored in eFuse.
- Please note that enabling Secure Boot or flash encryption disables the USB-OTG USB stack in the ROM, disallowing updates via the serial emulation or Device Firmware Update (DFU) on that port.

Burning read-protected keys After Secure Boot is enabled, no further eFuses can be read-protected, because this might allow an attacker to read-protect the efuse block holding the public key digest, causing an immediate denial of service and possibly allowing an additional fault injection attack to bypass the signature protection.

If [Flash Encryption](#) is enabled by the 2nd stage bootloader, it ensures that the flash encryption key generated on the first boot shall already be read-protected.

In case you need to read-protect a key after Secure Boot has been enabled on the device, for example,

- the flash encryption key
- HMAC keys
- ECDSA keys

you need to enable the config [CONFIG_SECURE_BOOT_V2_ALLOW_EFUSE_RD_DIS](#) at the same time when you enable Secure Boot to prevent disabling the ability to read-protect further efuses.

It is highly recommended that all such keys must be burned before enabling secure boot.

In case you need to enable [CONFIG_SECURE_BOOT_V2_ALLOW_EFUSE_RD_DIS](#), make sure that you burn the efuse ESP_EFUSE_WR_DIS_RD_DIS, using `esp_efuse_write_field_bit()` API from `esp_efuse.h`, once all the read-protected efuse keys have been programmed.

Generating Secure Boot Signing Key

The build system will prompt you with a command to generate a new signing key via `espsecure.py generate_signing_key`.

The `--version 2` parameter will generate the RSA 3072 private key for Secure Boot V2. Additionally `--scheme rsa3072` can be passed as well to generate RSA 3072 private key

Select the ECDSA scheme by passing `--version 2 --scheme ecdsa256` or `--version 2 --scheme ecdsa192` to generate corresponding ECDSA private key

The strength of the signing key is proportional to (a) the random number source of the system, and (b) the correctness of the algorithm used. For production devices, we recommend generating signing keys from a system with a quality entropy source, and using the best available RSA-PSS or ECDSA key generation utilities.

For example, to generate a signing key using the `openssl` command line:

For RSA 3072

```
` openssl genrsa -out my_secure_boot_signing_key.pem 3072 `
```

For ECC NIST192p curve

```
` openssl ecparam -name prime192v1 -genkey -noout -out  
my_secure_boot_signing_key.pem `
```

For ECC NIST256p curve

```
` openssl ecparam -name prime256v1 -genkey -noout -out  
my_secure_boot_signing_key.pem `
```

Remember that the strength of the Secure Boot system depends on keeping the signing key private.

Remote Signing of Images

Signing Using `espsecure.py` For production builds, it can be good practice to use a remote signing server rather than have the signing key on the build machine (which is the default `esp-idf` Secure Boot configuration). The `espsecure.py` command line program can be used to sign app images & partition table data for Secure Boot, on a remote system.

To use remote signing, disable the option [CONFIG_SECURE_BOOT_BUILD_SIGNED_BINARIES](#) and build the firmware. The private signing key does not need to be present on the build system.

After the app image and partition table are built, the build system will print signing steps using `espsecure.py`:

```
espsecure.py sign_data BINARY_FILE --version 2 --keyfile PRIVATE_SIGNING_KEY
```

The above command appends the image signature to the existing binary. You can use the `--output` argument to write the signed binary to a separate file:

```
espsecure.py sign_data --version 2 --keyfile PRIVATE_SIGNING_KEY --output SIGNED_  
↪BINARY_FILE BINARY_FILE
```

Signing Using Pre-calculated Signatures If you have valid pre-calculated signatures generated for an image and their corresponding public keys, you can use these signatures to generate a signature sector and append it to the image. Note that the pre-calculated signature should be calculated over all bytes in the image including the secure-padding bytes.

In such cases, the firmware image should be built by disabling the option `CONFIG_SECURE_BOOT_BUILD_SIGNED_BINARIES`. This image will be secure-padded and to generate a signed binary use the following command:

```
espssecure.py sign_data --version 2 --pub-key PUBLIC_SIGNING_KEY --signature_  
↪SIGNATURE_FILE --output SIGNED_BINARY_FILE BINARY_FILE
```

The above command verifies the signature, generates a signature block (refer to [Signature Block Format](#)) and appends it to the binary file.

Signing Using an External Hardware Security Module (HSM) For security reasons, you might also use an external Hardware Security Module (HSM) to store your private signing key, which cannot be accessed directly but has an interface to generate the signature of a binary file and its corresponding public key.

In such cases, disable the option `CONFIG_SECURE_BOOT_BUILD_SIGNED_BINARIES` and build the firmware. This secure-padded image then can be used to supply the external HSM for generating a signature. Refer to [Signing using an External HSM](#) to generate a signed image.

Note: For all the above three remote signing workflows, the signed binary is written to the filename provided to the `--output` argument and the option `--append_signatures` allows us to append multiple signatures (up to 3) the image.

Secure Boot Best Practices

- Generate the signing key on a system with a quality source of entropy.
- Keep the signing key private at all times. A leak of this key will compromise the Secure Boot system.
- Do not allow any third party to observe any aspects of the key generation or signing process using `espssecure.py`. Both processes are vulnerable to timing or other side-channel attacks.
- Enable all Secure Boot options in the Secure Boot Configuration. These include flash encryption, disabling of JTAG, disabling BASIC ROM interpreter, and disabling the UART bootloader encrypted flash access.
- Use Secure Boot in combination with [Flash Encryption](#) to prevent local readout of the flash contents.

Key Management

- Between 1 and 3 RSA-3072, ECDSA-256, or ECDSA-192 public key pairs (Keys #0, #1, #2) should be computed independently and stored separately.
- The KEY_DIGEST eFuses should be write protected after being programmed.
- The unused KEY_DIGEST slots must have their corresponding KEY_REVOKE eFuse burned to permanently disable them. This must happen before the device leaves the factory.
- The eFuses can either be written by the software bootloader during first boot after enabling "Secure Boot V2" from menuconfig or can be done using `espefuse.py` which communicates with the serial bootloader program in ROM.
- The KEY_DIGESTs should be numbered sequentially beginning at key digest #0. (i.e., if key digest #1 is used, key digest #0 should be used. If key digest #2 is used, key digest #0 & #1 must be used.)
- The software bootloader (non OTA upgradeable) is signed using at least one, possibly all three, private keys and flashed in the factory.
- Apps should only be signed with a single private key (the others being stored securely elsewhere), however they may be signed with multiple private keys if some are being revoked (see Key Revocation, below).

Multiple Keys

- The bootloader should be signed with all the private key(s) that are needed for the life of the device, before it is flashed.
- The build system can sign with at most one private key, user has to run manual commands to append more signatures if necessary.
- **You can use the append functionality of `espsecure.py`, this command would also printed at the end of the Secure B**
`espsecure.py sign_data -k secure_boot_signing_key2.pem -v 2 --append_signatures -o signed_bootloader.bin build/bootloader/bootloader.bin`
- While signing with multiple private keys, it is recommended that the private keys be signed independently, if possible on different servers and stored separately.
- **You can check the signatures attached to a binary using -** `espsecure.py signature_info_v2 datafile.bin`

Key Revocation

- Keys are processed in a linear order. (key #0, key #1, key #2).
- Applications should be signed with only one key at a time, to minimize the exposure of unused private keys.
- The bootloader can be signed with multiple keys from the factory.

Conservative Approach: Assuming a trusted private key (N-1) has been compromised, to update to new key pair (N).

1. Server sends an OTA update with an application signed with the new private key (#N).
 2. The new OTA update is written to an unused OTA app partition.
 3. The new application's signature block is validated. The public keys are checked against the digests programmed in the eFuse & the application is verified using the verified public key.
 4. The active partition is set to the new OTA application's partition.
 5. Device resets, loads the bootloader (verified with key #N-1 and #N) which then boots new app (verified with key #N).
 6. The new app verifies bootloader and application with key #N (as a final check) and then runs code to revoke key #N-1 (sets `KEY_REVOKE` eFuse bit).
 7. The API `esp_ota_revoke_secure_boot_public_key()` can be used to revoke the key #N-1.
- A similar approach can also be used to physically re-flash with a new key. For physical re-flashing, the bootloader content can also be changed at the same time.

Aggressive Approach: ROM code has an additional feature of revoking a public key digest if the signature verification fails.

To enable this feature, you need to burn `SECURE_BOOT_AGGRESSIVE_REVOKE` efuse or enable [*CONFIG_SECURE_BOOT_ENABLE_AGGRESSIVE_KEY_REVOKE*](#)

Key revocation is not applicable unless secure boot is successfully enabled. Also, a key is not revoked in case of invalid signature block or invalid image digest, it is only revoked in case the signature verification fails, i.e., revoke key only if failure in step 3 of [*Verifying an Image*](#)

Once a key is revoked, it can never be used for verifying a signature of an image. This feature provides strong resistance against physical attacks on the device. However, this could also brick the device permanently if all the keys are revoked because of signature verification failure.

Technical Details

The following sections contain low-level reference descriptions of various Secure Boot elements:

Manual Commands Secure boot is integrated into the esp-idf build system, so `idf.py build` will sign an app image and `idf.py bootloader` will produce a signed bootloader if secure signed binaries on build is enabled.

However, it is possible to use the `espsecure.py` tool to make standalone signatures and digests.

To sign a binary image:

```
espsecure.py sign_data --version 2 --keyfile ./my_signing_key.pem --output ./image_
↪signed.bin image-unsigned.bin
```

Keyfile is the PEM file containing an RSA-3072, ECDSA-256, or ECDSA-192 private signing key.

Secure Boot & Flash Encryption

If Secure Boot is used without [Flash Encryption](#), it is possible to launch "time-of-check to time-of-use" attack, where flash contents are swapped after the image is verified and running. Therefore, it is recommended to use both the features together.

Signed App Verification Without Hardware Secure Boot

The Secure Boot V2 signature of apps can be checked on OTA update, without enabling the hardware Secure Boot option. This option uses the same app signature scheme as Secure Boot V2, but unlike hardware Secure Boot it does not prevent an attacker who can write to flash from bypassing the signature protection.

This may be desirable in cases where the delay of Secure Boot verification on startup is unacceptable, and/or where the threat model does not include physical access or attackers writing to bootloader or app partitions in flash.

In this mode, the public key which is present in the signature block of the currently running app will be used to verify the signature of a newly updated app. (The signature on the running app is not verified during the update process, it is assumed to be valid.) In this way the system creates a chain of trust from the running app to the newly updated app.

For this reason, it is essential that the initial app flashed to the device is also signed. A check is run on app startup and the app will abort if no signatures are found. This is to try and prevent a situation where no update is possible. The app should have only one valid signature block in the first position. Note again that, unlike hardware Secure Boot V2, the signature of the running app is not verified on boot. The system only verifies a signature block in the first position and ignores any other appended signatures.

Although multiple trusted keys are supported when using hardware Secure Boot, only the first public key in the signature block is used to verify updates if signature checking without Secure Boot is configured. If multiple trusted public keys are required, it is necessary to enable the full Secure Boot feature instead.

Note: In general, it is recommended to use full hardware Secure Boot unless certain that this option is sufficient for application security needs.

How To Enable Signed App Verification

1. Open [Project Configuration Menu](#) -> Security features
2. Choose *App Signing Scheme*. Either *RSA* or *ECDSA (V2)*
3. Enable [CONFIG_SECURE_SIGNED_APPS_NO_SECURE_BOOT](#)
4. By default, "Sign binaries during build" will be enabled on selecting "Require signed app images" option, which will sign binary files as a part of build process. The file named in "Secure boot private signing key" will be used to sign the image.
5. If you disable "Sign binaries during build" option then all app binaries must be manually signed by following instructions in [Remote Signing of Images](#).

Warning: It is very important that all apps flashed have been signed, either during the build or after the build.

Advanced Features

JTAG Debugging By default, when Secure Boot is enabled then JTAG debugging is disabled via eFuse. The bootloader does this on first boot, at the same time it enables Secure Boot.

See [JTAG with Flash Encryption or Secure Boot](#) for more information about using JTAG Debugging with either Secure Boot or signed app verification enabled.

5.3 Workflows

Warning: This document is not updated for ESP32P4 yet, so some of the content may not be correct.
This warning was automatically inserted due to the source file being in the `add_warnings_pages` list.

5.3.1 Host-Based Security Workflows

Introduction

It is recommended to have an uninterrupted power supply while enabling security features on ESP32 SoCs. Power failures during the secure manufacturing process could cause issues that are hard to debug and, in some cases, may cause permanent boot-up failures.

This guide highlights an approach where security features are enabled with the assistance of an external host machine. Security workflows are broken down into various stages and key material is generated on the host machine; thus, allowing greater recovery chances in case of power or other failures. It also offers better timings for secure manufacturing, e.g., in the case of encryption of firmware on the host machine vs. on the device.

Goals

1. Simplify the traditional workflow with stepwise instructions.
2. Design a more flexible workflow as compared to the traditional firmware-based workflow.
3. Improve reliability by dividing the workflow into small operations.
4. Eliminate dependency on [Second Stage Bootloader](#) (firmware bootloader).

Pre-requisite

- `esptool`: Please make sure the `esptool` has been installed. It can be installed by running:

```
pip install esptool
```

Scope

- [Enable Flash Encryption and Secure Boot V2 Externally](#)
- [Enable Flash Encryption Externally](#)
- [Enable Secure Boot V2 Externally](#)

Security Workflows

Enable Flash Encryption and Secure Boot V2 Externally

Important: It is recommended to enable both Flash Encryption and Secure Boot V2 for a production use case.

When enabling the Flash Encryption and Secure Boot V2 externally we need to enable them in the following order:

1. Enable the Flash Encryption feature by following the steps listed in [Enable Flash Encryption Externally](#).
2. Enable the Secure Boot V2 feature by following the steps listed in [Enable Secure Boot V2 Externally](#).

The reason for this order is as follows:

To enable the Secure Boot (SB) V2, it is necessary to keep the SB V2 key readable. To protect the key's readability, the write protection for RD_DIS (ESP_EFUSE_WR_DIS_RD_DIS) is applied. However, this action poses a challenge when attempting to enable Flash Encryption, as the Flash Encryption (FE) key needs to remain unreadable. This conflict arises because the RD_DIS is already write-protected, making it impossible to read protect the FE key.

Enable Flash Encryption Externally In this case, all the eFuses related to flash encryption are written with help of the espfuse tool. More details about flash encryption can be found in the [Flash Encryption Guide](#)

1. Ensure that you have an ESP32-P4 device with default flash encryption eFuse settings as shown in [Relevant eFuses](#).

See how to check [ESP32-P4 Flash Encryption Status](#).

In this case, the flash on the chip must be erased and flash encryption must not be enabled. The chip can be erased by running:

```
esptool.py --port PORT erase_flash
```

2. Generate a flash encryption key.

A random flash encryption key can be generated by running:

```
espsecure.py generate_flash_encryption_key my_flash_encryption_key.bin
```

3. Burn the flash encryption key into eFuse.

This action **cannot be reverted**. It can be done by running:

```
espfuse.py --port PORT burn_key BLOCK my_flash_encryption_key.bin XTS_
↪AES_128_KEY
```

where BLOCK is a free keyblock between BLOCK_KEY0 and BLOCK_KEY5.

4. Burn the SPI_BOOT_CRYPT_CNT eFuse.

If you only want to enable flash encryption in **Development** mode and want to keep the ability to disable it in the future, Update the SPI_BOOT_CRYPT_CNT value in the below command from 7 to 0x1 (not recommended for production).

```
espfuse.py --port PORT --chip esp32p4 burn_efuse SPI_BOOT_CRYPT_CNT 7
```

Note: At this point, the flash encryption on the device has been enabled. You may test the flash encryption process as given in step 5. Please note that the security-related eFuses have not been burned at this point. It is recommended that they should be burned in production use cases as explained in step 6.

5. Encrypt and flash the binaries.

The bootloader and the application binaries for the project must be built with Flash Encryption Release mode with default configurations.

Flash encryption Release mode can be set in the menuconfig as follows:

- [Enable flash encryption on boot](#)
- [Select Release mode](#) (Note that once Release mode is selected, the EFUSE_DIS_DOWNLOAD_MANUAL_ENCRYPT eFuse bit will be burned to disable flash encryption hardware in ROM Download Mode.)
- [Select UART ROM download mode \(Permanently switch to Secure mode \(recommended\)\)](#). This is the default option and is recommended. It is also possible to change this configuration setting to permanently disable UART ROM download mode, if this mode is not needed.
- [Select the appropriate bootloader log verbosity](#)
- Save the configuration and exit.

The binaries can be encrypted on the host machine by running:

```
espsecure.py encrypt_flash_data --aes_xts --keyfile my_flash_encryption_
↪key.bin --address 0x0 --output bootloader-enc.bin build/my-app.bin

espsecure.py encrypt_flash_data --aes_xts --keyfile my_flash_encryption_
↪key.bin --address 0x8000 --output partition-table-enc.bin build/
↪partition_table/partition-table.bin

espsecure.py encrypt_flash_data --aes_xts --keyfile my_flash_encryption_
↪key.bin --address 0x10000 --output my-app-enc.bin build/my-app.bin
```

The above files can then be flashed to their respective offset using `esptool.py`. To see all of the command line options recommended for `esptool.py`, see the output printed when `idf.py build` succeeds. In the above command the offsets are used for a sample firmware, the actual offset for your firmware can be obtained by checking the partition table entry or by running `idf.py partition-table`. When the application contains the following partition: `otadata`, `nvs_encryption_keys` they need to be encrypted as well. Please refer to [Encrypted Partitions](#) for more details about encrypted partitions.

Note: If the flashed ciphertext file is not recognized by the ESP32-P4 when it boots, check that the keys match and that the command line arguments match exactly, including the correct offset. It is important to provide the correct offset as the ciphertext changes when the offset changes.

The command `espsecure.py decrypt_flash_data` can be used with the same options (and different input/output files), to decrypt ciphertext flash contents or a previously encrypted file.

6. Burn flash encryption-related security eFuses as listed below:

A) Burn security eFuses:

Important: For production use cases, it is highly recommended to burn all the eFuses listed below.

- `DIS_DOWNLOAD_ICACHE`: Disable UART cache
- `DIS_DIRECT_BOOT`: Disable direct boot (legacy SPI boot mode)
- `DIS_USB_JTAG`: Disable USB switch to JTAG
- `DIS_PAD_JTAG`: Disable JTAG permanently
- `DIS_DOWNLOAD_MANUAL_ENCRYPT`: Disable UART bootloader encryption access

The respective eFuses can be burned by running:

```
espfuse.py burn_efuse --port PORT EFUSE_NAME 0x1
```

Note: Please update the `EFUSE_NAME` with the eFuse that you need to burn. Multiple

eFuses can be burned at the same time by appending them to the above command (e.g., EFUSE_NAME VAL EFUSE_NAME2 VAL2). More documentation about *espefuse.py* can be found [here](#).

B) Write protect security eFuses:

After burning the respective eFuses we need to write_protect the security configurations

```
espefuse.py --port PORT write_protect_efuse DIS_ICACHE
```

Note: The write protection of above eFuse also write protects multiple other eFuses, Please refer to the ESP32-P4 eFuse table for more details.

C) Enable Security Download mode:

Warning: Please burn the following bit at the very end. After this bit is burned, the *espefuse* tool can no longer be used to burn additional eFuses.

- ENABLE_SECURITY_DOWNLOAD: Enable Secure ROM download mode

The eFuse can be burned by running:

```
espefuse.py --port PORT burn_efuse ENABLE_SECURITY_DOWNLOAD
```

Important:

7. Delete flash encryption key on host:

Once the flash encryption has been enabled for the device, the key **must be deleted immediately**. This ensures that the host cannot produce encrypted binaries for the same device going forward. This step is important to reduce the vulnerability of the flash encryption key.

Flash Encryption Guidelines

- It is recommended to generate a unique flash encryption key for each device for production use-cases.
- It is recommended to ensure that the RNG used by host machine to generate the flash encryption key has good entropy.
- See [Limitations of Flash Encryption](#) for more details.

Enable Secure Boot V2 Externally In this workflow, we shall use *espsecure* tool to generate signing keys and use the *espefuse* tool to burn the relevant eFuses. The details about the Secure Boot V2 process can be found at [Secure Boot V2 Guide](#)

1. Generate Secure Boot V2 Signing Private Key.

The Secure Boot V2 signing key for the RSA3072 scheme can be generated by running:

```
espsecure.py generate_signing_key --version 2 --scheme rsa3072 secure_  
↪boot_signing_key.pem
```

The Secure Boot V2 signing key for ECDSA scheme can be generated by running:

```
espsecure.py generate_signing_key --version 2 --scheme ecdsa256_  
↪secure_boot_signing_key.pem
```

The scheme in the above command can be changed to `ecdsa192` to generate `ecdsa192` private key.

A total of 3 keys can be used for Secure Boot V2 at once. These should be computed independently and stored separately. The same command with different key file names can be used to generate multiple Secure Boot V2 signing keys. It is recommended to use multiple keys in order to reduce dependency on a single key.

2. Generate Public Key Digest.

The public key digest for the private key generated in the previous step can be generated by running:

```
espsecure.py digest_sbv2_public_key --keyfile secure_boot_signing_key.pem  
→--output digest.bin
```

In case of multiple digests, each digest should be kept in a separate file.

3. Burn the key digest in eFuse.

The public key digest can be burned in the eFuse by running:

```
espefuse.py --port PORT --chip esp32p4 burn_key digest.bin BLOCK SECURE_  
→BOOT_DIGEST0
```

where `BLOCK` is a free keyblock between `BLOCK_KEY0` and `BLOCK_KEY5`.

In case of multiple digests, the other digests can be burned sequentially by changing the key purpose to `SECURE_BOOT_DIGEST1` and `SECURE_BOOT_DIGEST2` respectively.

4. Enable Secure Boot V2.

Secure Boot V2 eFuse can be enabled by running:

```
espefuse.py --port PORT --chip esp32p4 burn_efuse SECURE_BOOT_EN
```

Note: At this stage the secure boot V2 has been enabled for the ESP32-P4. The ROM bootloader shall now verify the authenticity of the *Second Stage Bootloader* on every bootup. You may test the secure boot process by executing Steps 5 & 6. Please note that security-related eFuses have not been burned at this point. For production use cases, all security-related eFuses must be burned as given in step 7.

5. Build the binaries.

By default, the ROM bootloader would only verify the *Second Stage Bootloader* (firmware bootloader). The firmware bootloader would verify the app partition only when the `CONFIG_SECURE_BOOT` option is enabled (and `CONFIG_SECURE_BOOT_VERSION` is set to `SECURE_BOOT_V2_ENABLED`) while building the bootloader.

- Open the *Project Configuration Menu*, in "Security features" set "Enable hardware Secure Boot in bootloader" to enable Secure Boot.

The "Secure Boot V2" option will be selected and the "App Signing Scheme" will be set to RSA by default.

- Disable the option `CONFIG_SECURE_BOOT_BUILD_SIGNED_BINARIES` for the project in the *Project Configuration Menu*. This shall make sure that all the generated binaries are secure padded and unsigned. This step is done to avoid generating signed binaries as we are going to manually sign the binaries using `espsecure` tool.

After the above configurations, the bootloader and application binaries can be built with `idf.py build` command.

6. Sign and Flash the binaries.

The Secure Boot V2 workflow only verifies the `bootloader` and `application` binaries, hence only those binaries need to be signed. The other binaries (e.g., `partition-table.bin`) can be flashed as they are generated in the build stage.

The `bootloader.bin` and `app.bin` binaries can be signed by running:

```
espsecure.py sign_data --version 2 --keyfile secure_boot_signing_key.pem -  
→-output bootloader-signed.bin build/bootloader/bootloader.bin  
  
espsecure.py sign_data --version 2 --keyfile secure_boot_signing_key.pem -  
→-output my-app-signed.bin build/my-app.bin
```

If multiple keys secure boot keys are to be used then the same signed binary can be appended with a signature block signed with the new key as follows:

```
espsecure.py sign_data --keyfile secure_boot_signing_key2.pem --version 2  
→--append_signatures -o bootloader-signed2.bin bootloader-signed.bin  
  
espsecure.py sign_data --keyfile secure_boot_signing_key2.pem --version 2  
→--append_signatures -o my-app-signed2.bin my-app-signed.bin
```

The same process can be repeated for the third key. Note that the names of the input and output files must not be the same.

The signatures attached to a binary can be checked by running:

```
espsecure.py signature_info_v2 bootloader-signed.bin
```

The above files along with other binaries (e.g., partition table) can then be flashed to their respective offset using `esptool.py`. To see all of the command line options recommended for `esptool.py`, see the output printed when `idf.py build` succeeds. The flash offset for your firmware can be obtained by checking the partition table entry or by running `idf.py partition-table`.

7. Burn relevant eFuses.

A) Burn security eFuses:

Important: For production use cases, it is highly recommended to burn all the eFuses listed below.

- `SOFT_DIS_JTAG`: Disable software access to JTAG peripheral
- `DIS_DIRECT_BOOT`: Disable direct boot (legacy SPI boot mode)
- `DIS_USB_JTAG`: Disable USB switch to JTAG
- `DIS_PAD_JTAG`: Disable JTAG permanently
- `SECURE_BOOT_AGGRESSIVE_REVOKE`: Aggressive revocation of key digests, see [Aggressive Approach](#) for more details.

The respective eFuses can be burned by running:

```
espefuse.py burn_efuse --port PORT EFUSE_NAME 0x1
```

Note: Please update the `EFUSE_NAME` with the eFuse that you need to burn. Multiple eFuses can be burned at the same time by appending them to the above command (e.g., `EFUSE_NAME VAL EFUSE_NAME2 VAL2`). More documentation about `espefuse.py` can be found [here](#)

B) Secure Boot V2-related eFuses:

i) Disable the ability for read protection:

The secure boot digest burned in the eFuse must be kept readable otherwise secure boot operation would result in a failure. To prevent the accidental enabling of read protection for this key block we need to burn the following eFuse:

```
espefuse.py -p $ESP_PORT write_protect_efuse RD_DIS
```

Important: After this eFuse has been burned, read protection cannot be enabled for any key. E.g., if flash encryption which requires read protection for its key is not enabled at this point then it cannot be enabled afterwards. Please ensure that no eFuse keys are going to need read protection after this.

ii) Revoke key digests:

The unused digest slots need to be revoked when we are burning the secure boot key. The respective slots can be revoked by running

```
espefuse.py --port PORT --chip esp32p4 burn_efuse EFUSE_REVOKE_BIT
```

The `EFUSE_REVOKE_BIT` in the above command can be `SECURE_BOOT_KEY_REVOKE0` or `SECURE_BOOT_KEY_REVOKE1` or `SECURE_BOOT_KEY_REVOKE2`. Please note that only the unused key digests must be revoked. Once revoked, the respective digest cannot be used again.

C) Enable Security Download mode:

Warning: Please burn the following bit at the very end. After this bit is burned, the `espefuse` tool can no longer be used to burn additional eFuses.

- `ENABLE_SECURITY_DOWNLOAD`: Enable Secure ROM download mode

The eFuse can be burned by running:

```
espefuse.py --port PORT burn_efuse ENABLE_SECURITY_DOWNLOAD
```

Secure Boot V2 Guidelines

- It is recommended to store the secure boot key in a highly secure place. A physical or a cloud HSM may be used for secure storage of the secure boot private key. Please take a look at [Remote Signing of Images](#) for more details.
- It is recommended to use all the available digest slots to reduce dependency on a single private key.

Chapter 6

Migration Guides

6.1 ESP-IDF 5.x Migration Guide

6.1.1 Migration from 4.4 to 5.0

Build System

Migrating from GNU Make Build System ESP-IDF v5.0 no longer supports GNU make-based projects. Please follow the [build system](#) guide for migration.

Update Fragment File Grammar The former grammar, supported in ESP-IDF v3.x, was dropped in ESP-IDF v5.0. Here are a few notes on how to migrate properly:

1. Indentation is now enforced: improperly indented fragment files generate a runtime parse exception. Although the former version did not enforce this, the previous documentation and examples demonstrated properly indented grammar.
2. Migrate the old condition entry to the `if...elif...else` structure for conditionals. You can refer to the [Configuration-Dependent Placements](#) for detailed grammar.
3. Mapping fragments now requires a name like other fragment types.

Specify Component Requirements Explicitly In previous versions of ESP-IDF, some components were always added as public requirements (dependencies) to every component in the build, in addition to the [common component requirements](#):

- driver
- efuse
- esp_timer
- lwip
- vfs
- esp_wifi
- esp_event
- esp_netif
- esp_eth
- esp_phy

This means that it was possible to include header files of those components without specifying them as requirements in `idf_component_register`. This behavior was caused by transitive dependencies of various common components.

In ESP-IDF v5.0, this behavior is fixed and these components are no longer added as public requirements by default.

Every component depending on one of the components which isn't part of common requirements has to declare this dependency explicitly. This can be done by adding `REQUIRES <component_name>` or `PRIV_REQUIRES <component_name>` in `idf_component_register` call inside component's `CMakeLists.txt`. See [Component Requirements](#) for more information on specifying requirements.

Setting `COMPONENT_DIRS` and `EXTRA_COMPONENT_DIRS` Variables ESP-IDF v5.0 includes a number of improvements to support building projects with space characters in their paths. To make that possible, there are some changes related to setting `COMPONENT_DIRS` and `EXTRA_COMPONENT_DIRS` variables in project `CMakeLists.txt` files.

Adding non-existent directories to `COMPONENT_DIRS` or `EXTRA_COMPONENT_DIRS` is no longer supported and will result in an error.

Using string concatenation to define `COMPONENT_DIRS` or `EXTRA_COMPONENT_DIRS` variables is now deprecated. These variables should be defined as CMake lists, instead. For example, use:

```
set(EXTRA_COMPONENT_DIRS path1 path2)
list(APPEND EXTRA_COMPONENT_DIRS path3)
```

instead of:

```
set(EXTRA_COMPONENT_DIRS "path1 path2")
set(EXTRA_COMPONENT_DIRS "${EXTRA_COMPONENT_DIRS} path3")
```

Defining these variables as CMake lists is compatible with previous ESP-IDF versions.

Update Usage of `target_link_libraries` with `project_elf` ESP-IDF v5.0 fixes CMake variable propagation issues for components. This issue caused compiler flags and definitions that were supposed to apply to one component to be applied to every component in the project.

As a side effect of this, user projects from ESP-IDF v5.0 onwards must use `target_link_libraries` with `project_elf` explicitly and custom CMake projects must specify `PRIVATE`, `PUBLIC`, or `INTERFACE` arguments. This is a breaking change and is not backward compatible with previous ESP-IDF versions.

For example:

```
target_link_libraries(${project_elf} PRIVATE "-Wl,--wrap=esp_panic_handler")
```

instead of:

```
target_link_libraries(${project_elf} "-Wl,--wrap=esp_panic_handler")
```

Update CMake Version In ESP-IDF v5.0 minimal CMake version was increased to 3.16 and versions lower than 3.16 are not supported anymore. Run `tools/idf_tools.py install cmake` to install a suitable version if your OS version doesn't have one.

This affects ESP-IDF users who use system-provided CMake and custom CMake.

Reorder the Applying of the Target-Specific Config Files ESP-IDF v5.0 reorders the applying order of target-specific config files and other files listed in `SDKCONFIG_DEFAULTS`. Now, target-specific files will be applied right after the file brings it in, before all latter files in `SDKCONFIG_DEFAULTS`.

For example:

```
If ``SDKCONFIG_DEFAULTS="sdkconfig.defaults;sdkconfig_devkit1"`` , and there is a
↳file ``sdkconfig.defaults.esp32`` in the same folder, then the files will be
↳applied in the following order: (1) sdkconfig.defaults (2) sdkconfig.defaults.
↳esp32 (3) sdkconfig_devkit1.
```

If you have a key with different values in the target-specific files of the former item (e.g., `sdkconfig.defaults.esp32` above) and the latter item (e.g., `sdkconfig_devkit1` above), please note the latter will override the target-specific file of the former.

If you do want to have some target-specific config values, please put it into the target-specific file of the latter item (e.g., `sdkconfig_devkit1.esp32`).

GCC

GCC Version The previous GCC version was GCC 8.4.0. This has now been upgraded to GCC 11.2.0 on all targets. Users that need to port their code from GCC 8.4.0 to 11.2.0 should refer to the series of official GCC porting guides listed below:

- [Porting to GCC 9](#)
- [Porting to GCC 10](#)
- [Porting to GCC 11](#)

Warnings The upgrade to GCC 11.2.0 has resulted in the addition of new warnings, or enhancements to existing warnings. The full details of all GCC warnings can be found in [GCC Warning Options](#). Users are advised to double-check their code, then fix the warnings if possible. Unfortunately, depending on the warning and the complexity of the user's code, some warnings will be false positives that require non-trivial fixes. In such cases, users can choose to suppress the warning in multiple ways. This section outlines some common warnings that users are likely to encounter, and ways to suppress them.

Warning: Users are advised to check that a warning is indeed a false positive before attempting to suppress them it.

-Wstringop-overflow, -Wstringop-overread, -Wstringop-truncation, and -Warray-bounds Users that use memory/string copy/compare functions will run into one of the `-Wstringop` warnings if the compiler cannot properly determine the size of the memory/string. The examples below demonstrate code that triggers these warnings and how to suppress them.

```
#pragma GCC diagnostic push
#pragma GCC diagnostic ignored "-Wstringop-overflow"
#pragma GCC diagnostic ignored "-Warray-bounds"
    memset(RTC_SLOW_MEM, 0, CONFIG_ULP_COPROC_RESERVE_MEM); // <-- This line
↳leads to warnings
#pragma GCC diagnostic pop
```

```
#pragma GCC diagnostic push
#if __GNUC__ >= 11
#pragma GCC diagnostic ignored "-Wstringop-overread" // <-- This key had been
↳introduced since GCC 11
#endif
#pragma GCC diagnostic ignored "-Warray-bounds"
    memcpy(backup_write_data, (void *)EFUSE_PGM_DATA0_REG, sizeof(backup_
↳write_data)); // <-- This line leads to warnings
#pragma GCC diagnostic pop
```

-Waddress-of-packed-member GCC will issue this warning when accessing an unaligned member of a packed struct due to the incurred penalty of unaligned memory access. However, all ESP chips (on both Xtensa and RISC-V architectures) allow for unaligned memory access and incur no extra penalty. Thus, this warning can be ignored in most cases.

```
components/bt/host/bluedroid/btc/profile/std/gatt/btc_gatt_util.c: In function
↳ 'btc_to_bta_gatt_id':
components/bt/host/bluedroid/btc/profile/std/gatt/btc_gatt_util.c:105:21: warning:↳
↳ taking address of packed member of 'struct <anonymous>' may result in an↳
↳ unaligned pointer value [-Waddress-of-packed-member]
105 |         btc_to_bta_uuid(&p_dest->uuid, &p_src->uuid);
    |                        ^~~~~~
```

If the warning occurs in multiple places across multiple source files, users can suppress the warning at the CMake level as demonstrated below.

```
set_source_files_properties(
    "host/bluedroid/bta/gatt/bta_gattc_act.c"
    "host/bluedroid/bta/gatt/bta_gattc_cache.c"
    "host/bluedroid/btc/profile/std/gatt/btc_gatt_util.c"
    "host/bluedroid/btc/profile/std/gatt/btc_gatts.c"
    PROPERTIES COMPILE_FLAGS -Wno-address-of-packed-member)
```

However, if there are only one or two instances, users can suppress the warning directly in the source code itself as demonstrated below.

```
#pragma GCC diagnostic push
#if __GNUC__ >= 9
#pragma GCC diagnostic ignored "-Waddress-of-packed-member" <-- This key had been↳
↳ introduced since GCC 9
#endif
uint32_t* reg_ptr = (uint32_t*)src;
#pragma GCC diagnostic pop
```

labs () for 64-bit Integers The function `abs ()` from `stdlib.h` takes `int` argument. Please use `labs ()` for types that are intended to be 64-bit. It is particularly important for `time_t`.

Espressif Toolchain Changes

int32_t and uint32_t for Xtensa Compiler The types `int32_t` and `uint32_t` have been changed from the previous `int` and `unsigned int` to `long` and `unsigned long` respectively for the Xtensa compiler. This change now matches upstream GCC which `long` integers for `int32_t` and `uint32_t` on Xtensa, RISC-V, and other architectures.

	2021r2 and older, GCC 8	2022r1, GCC 11
Xtensa	(unsigned) int	(unsigned) long
riscv32	(unsigned) long	(unsigned) long

The change mostly affects code that formats strings using types provided by `<inttypes.h>`. When using these fixed-width types (e.g., `uint32_t`), users will need to replace placeholders such as `%i` and `%x` with `PRIi32` and `PRIx32` respectively. Types *not* defined in `<inttypes.h>` (e.g., `int`) do *not* need this special formatting.

In other cases, it should be noted that enums have the `int` type.

In common, `int32_t` and `int`, as well as `uint32_t` and `unsigned int`, are different types.

If users do not make the aforementioned updates to format strings in their applications, the following error will be reported during compilation:

```

/Users/name/esp/esp-rainmaker/components/esp-insights/components/esp_diagnostics/
↳include/esp_diagnostics.h:238:29: error: format '%u' expects argument of type
↳'unsigned int', but argument 3 has type 'uint32_t' {aka 'long unsigned int'} [-
↳Werror=format=]
238 |         esp_diag_log_event(tag, "EV (%u) %s: " format, esp_log_timestamp(), tag,
↳##__VA_ARGS__); \
    |                                     ^~~~~~
    |                                     |
    |                                     uint32_t {aka long_
↳unsigned int}
                                     uint32_t {aka long unsigned int}

```

Removing **CONFIG_COMPILER_DISABLE_GCC8_WARNINGS** **Build Option** **CON-**
FIG_COMPILER_DISABLE_GCC8_WARNINGS option was introduced to allow building of legacy code
 dating from the rigid GCC 5 toolchain. However, enough time has passed to allow for the warnings to be fixed, thus
 this option has been removed.

For now in GCC 11, users are advised to review their code and fix the compiler warnings where possible.

Networking

Wi-Fi

Callback Function Type `esp_now_recv_cb_t` Previously, the first parameter of `esp_now_recv_cb_t`
 was of type `const uint8_t *mac_addr`, which only included the address of ESP-NOW peer device.

This now changes. The first parameter is of type `esp_now_recv_info_t`, which has members `src_addr`,
`des_addr` and `rx_ctrl`. Therefore, the following updates are required:

- Redefine ESP-NOW receive callback function.
- `src_addr` can be used to replace original `mac_addr`.
- `des_addr` is the destination MAC address of ESP-NOW packet, which can be unicast or broadcast address.
 With `des_addr`, the user can distinguish unicast and broadcast ESP-NOW packets where broadcast ESP-
 NOW packets can be non-encrypted even when encryption policy is configured for the ESP-NOW.
- `rx_ctrl` is Rx control info of the ESP-NOW packet, which provides more information about the packet.

Please refer to the ESP-NOW example: [wifi/espnow/main/espnow_example_main.c](#)

Ethernet

`esp_eth_ioctl()` API Previously, the `esp_eth_ioctl()` API had the following issues:

- The third parameter (which is of type `void *`) would accept an `int/bool` type arguments (i.e., not pointers)
 as input in some cases. However, these cases were not documented properly.
- To pass `int/bool` type argument as the third parameter, the argument had to be "unnaturally" casted to a
`void *` type, to prevent a compiler warning as demonstrated in the code snippet below. This casting could
 lead to misuse of the `esp_eth_ioctl()` function.

```
esp_eth_ioctl(eth_handle, ETH_CMD_S_FLOW_CTRL, (void *)true);
```

Therefore, the usage of `esp_eth_ioctl()` is now unified. Arguments to the third parameter must be passed as
 pointers to a specific data type to/from where data is stored/read by `esp_eth_ioctl()`. The code snippets below
 demonstrate the usage of `esp_eth_ioctl()`.

Usage example to set Ethernet configuration:

```
eth_duplex_t new_duplex_mode = ETH_DUPLEX_HALF;
esp_eth_ioctl(eth_handle, ETH_CMD_S_DUPLEX_MODE, &new_duplex_mode);
```

Usage example to get Ethernet configuration:

```
eth_duplex_t duplex_mode;
esp_eth_ioctl(eth_handle, ETH_CMD_G_DUPLEX_MODE, &duplex_mode);
```

KSZ8041/81 and LAN8720 Driver Update The KSZ8041/81 and LAN8720 drivers are updated to support more devices (i.e., generations) from their associated product families. The drivers can recognize particular chip numbers and their potential support by the driver.

As a result, the specific "chip number" functions calls are replaced by generic ones as follows:

- Removed `esp_eth_phy_new_ksz8041()` and `esp_eth_phy_new_ksz8081()`, and use `esp_eth_phy_new_ksz80xx()` instead
- Removed `esp_eth_phy_new_lan8720()`, and use `esp_eth_phy_new_lan87xx()` instead

ESP NETIF Glue Event Handlers `esp_eth_set_default_handlers()` and `esp_eth_clear_default_handlers()` functions are removed. Registration of the default IP layer handlers for Ethernet is now handled automatically. If you have already followed the suggestion to fully initialize the Ethernet driver and network interface before registering their Ethernet/IP event handlers, then no action is required (except for deleting the affected functions). Otherwise, you may start the Ethernet driver right after they register the user event handler.

PHY Address Auto-detect The Ethernet PHY address auto-detect function `esp_eth_detect_phy_addr()` is renamed to `esp_eth_phy_802_3_detect_phy_addr()` and its header declaration is moved to `esp_eth/include/esp_eth_phy_802_3.h`.

SPI-Ethernet Module Initialization The SPI-Ethernet Module initialization is now simplified. Previously, you had to manually allocate an SPI device using `spi_bus_add_device()` before instantiating the SPI-Ethernet MAC.

Now, you no longer need to call `spi_bus_add_device()` as SPI devices are allocated internally. As a result, the `eth_dm9051_config_t`, `eth_w5500_config_t`, and `eth_ksz8851snl_config_t` configuration structures are updated to include members for SPI device configuration (e.g., to allow fine tuning of SPI timing which may be dependent on PCB design). Likewise, the `ETH_DM9051_DEFAULT_CONFIG`, `ETH_W5500_DEFAULT_CONFIG`, and `ETH_KSZ8851SNL_DEFAULT_CONFIG` configuration initialization macros are updated to accept new input parameters. Refer to *Ethernet API Reference Guide* for an example of SPI-Ethernet Module initialization.

TCP/IP Adapter The TCP/IP Adapter was a network interface abstraction component used in ESP-IDF prior to v4.1. This section outlines migration from `tcpip_adapter` API to its successor *ESP-NETIF*.

Updating Network Connection Code

Network Stack Initialization

- You may simply replace `tcpip_adapter_init()` with `esp_netif_init()`. However, please should note that the `esp_netif_init()` function now returns standard error codes. See *ESP-NETIF* for more details.
- The `esp_netif_deinit()` function is provided to de-initialize the network stack.
- You should also replace `#include "tcpip_adapter.h"` with `#include "esp_netif.h"`.

Network Interface Creation Previously, the TCP/IP Adapter defined the following network interfaces statically:

- WiFi Station
- WiFi Access Point
- Ethernet

This now changes. Network interface instance should be explicitly constructed, so that the *ESP-NETIF* can connect to the TCP/IP stack. For example, after the TCP/IP stack and the event loop are initialized, the initialization code for WiFi must explicitly call `esp_netif_create_default_wifi_sta();` or `esp_netif_create_default_wifi_ap();`.

Please refer to the example initialization code for these three interfaces:

- WiFi Station: [wifi/getting_started/station/main/station_example_main.c](#)
- WiFi Access Point: [wifi/getting_started/softAP/main/softap_example_main.c](#)
- Ethernet: [ethernet/basic/main/ethernet_example_main.c](#)

Other tcpip_adapter API Replacement All the `tcpip_adapter` functions have their `esp-netif` counter-part. Please refer to the `esp_netif.h` grouped into these sections:

- [Setters/Getters](#)
- [DHCP](#)
- [DNS](#)
- [IP address](#)

The TCP/IP Adapter API `tcpip_adapter_get_sta_list()` that was used to acquire a list of associated Wi-Fi stations to the Software Access Point (softAP) has been moved to the Wi-Fi component and renamed to `esp_wifi_ap_get_sta_list_with_ip()`, which is a special case of the ESP-NETIF API [esp_netif_dhcps_get_clients_by_mac\(\)](#) that could be used more generally to provide a list of clients connected to a DHCP server no matter which network interface the server is running on.

Default Event Handlers Event handlers are moved from `tcpip_adapter` to appropriate driver code. There is no change from application code perspective, as all events should be handled in the same way. Please note that for IP-related event handlers, application code usually receives IP addresses in the form of an `esp-netif` specific struct instead of the LwIP structs. However, both structs are binary compatible.

This is the preferred way to print the address:

```
ESP_LOGI(TAG, "got ip:" IPSTR, IP2STR(&event->ip_info.ip));
```

Instead of

```
ESP_LOGI(TAG, "got ip:%s", ip4addr_ntoa(&event->ip_info.ip));
```

Since `ip4addr_ntoa()` is a LwIP API, the `esp-netif` provides `esp_ip4addr_ntoa()` as a replacement. However, the above method using `IP2STR()` is generally preferred.

IP Addresses You are advised to use `esp-netif` defined IP structures. Please note that with default compatibility enabled, the LwIP structs still work.

- [esp-netif IP address definitions](#)

Peripherals

Peripheral Clock Gating As usual, peripheral clock gating is still handled by driver itself, users do not need to take care of the peripheral module clock gating.

However, for advanced users who implement their own drivers based on `hal` and `soc` components, the previous clock gating include path has been changed from `driver/periph_ctrl.h` to `esp_private/periph_ctrl.h`.

RTC Subsystem Control RTC control APIs have been moved from `driver/rtc_cntl.h` to `esp_private/rtc_ctrl.h`.

ADC

ADC Oneshot & Continuous Mode Drivers The ADC oneshot mode driver has been redesigned.

- The new driver is in `esp_adc` component and the include path is `esp_adc/adc_oneshot.h`.
- The legacy driver is still available in the previous include path `driver/adc.h`.

The ADC continuous mode driver has been moved from `driver` component to `esp_adc` component.

- The include path has been changed from `driver/adc.h` to `esp_adc/adc_continuous.h`.

Attempting to use the legacy include path `driver/adc.h` of either driver triggers the build warning below by default. However, the warning can be suppressed by enabling the [`CONFIG\_ADC\_SUPPRESS\_DEPRECATED\_WARN`](#) Kconfig option.

```
legacy adc driver is deprecated, please migrate to use esp_adc/adc_oneshot.h and
↳ esp_adc/adc_continuous.h for oneshot mode and continuous mode drivers
↳ respectively
```

ADC Calibration Driver The ADC calibration driver has been redesigned.

- The new driver is in `esp_adc` component and the include path is `esp_adc/adc_cali.h` and `esp_adc/adc_cali_scheme.h`.

Legacy driver is still available by including `esp_adc_cal.h`. However, if users still would like to use the include path of the legacy driver, users should add `esp_adc` component to the list of component requirements in `CMakeLists.txt`.

Attempting to use the legacy include path `esp_adc_cal.h` triggers the build warning below by default. However, the warning can be suppressed by enabling the [`CONFIG\_ADC\_CALI\_SUPPRESS\_DEPRECATED\_WARN`](#) Kconfig option.

```
legacy adc calibration driver is deprecated, please migrate to use esp_adc/adc_
↳ cali.h and esp_adc/adc_cali_scheme.h
```

API Changes

- The ADC power management APIs `adc_power_acquire` and `adc_power_release` have made private and moved to `esp_private/adc_share_hw_ctrl.h`.
 - The two APIs were previously made public due to a HW errata workaround.
 - Now, ADC power management is completely handled internally by drivers.
 - Users who still require this API can include `esp_private/adc_share_hw_ctrl.h` to continue using these functions.
- `driver/adc2_wifi_private.h` has been moved to `esp_private/adc_share_hw_ctrl.h`.
- Enums `ADC_UNIT_BOTH`, `ADC_UNIT_ALTER`, and `ADC_UNIT_MAX` in `adc_unit_t` have been removed.
- The following enumerations have been removed as some of their enumeration values are not supported on all chips. This would lead to the driver triggering a runtime error if an unsupported value is used.
 - Enum `ADC_CHANNEL_MAX`
 - Enum `ADC_ATTEN_MAX`

- Enum `ADC_CONV_UNIT_MAX`
- API `hall_sensor_read` on ESP32 has been removed. Hall sensor is no longer supported on ESP32.
- API `adc_set_i2s_data_source` and `adc_i2s_mode_init` have been deprecated. Related enum `adc_i2s_source_t` has been deprecated. Please migrate to use `esp_adc/adc_continuous.h`.
- API `adc_digi_filter_reset`, `adc_digi_filter_set_config`, `adc_digi_filter_get_config` and `adc_digi_filter_enable` have been removed. These APIs behaviours are not guaranteed. Enum `adc_digi_filter_idx_t`, `adc_digi_filter_mode_t` and structure `adc_digi_iir_filter_t` have been removed as well.
- API `esp_adc_cal_characterize` has been deprecated, please migrate to `adc_cali_create_scheme_curve_fitting` or `adc_cali_create_scheme_line_fitting` instead.
- API `esp_adc_cal_raw_to_voltage` has been deprecated, please migrate to `adc_cali_raw_to_voltage` instead.
- API `esp_adc_cal_get_voltage` has been deprecated, please migrate to `adc_one_shot_get_calibrated_result` instead.

GPIO

- The previous Kconfig option `RTCIO_SUPPORT_RTC_GPIO_DESC` has been removed, thus the `rtc_gpio_desc` array is unavailable. Please use `rtc_io_desc` array instead.
- The user callback of a GPIO interrupt should no longer read the GPIO interrupt status register to get the GPIO's pin number of triggering the interrupt. You should use the callback argument to determine the GPIO's pin number instead.
 - Previously, when a GPIO interrupt occurs, the GPIO's interrupt status register is cleared after calling the user callbacks. Thus, it was possible for users to read the GPIO's interrupt status register inside the callback to determine which GPIO was used to trigger the interrupt.
 - However, clearing the interrupt status register after calling the user callbacks can potentially cause edge-triggered interrupts to be lost. For example, if an edge-triggered interrupt is triggered/retriggered while the user callbacks are being called, that interrupt will be cleared without its registered user callback being handled.
 - Now, the GPIO's interrupt status register is cleared **before** invoking the user callbacks. Thus, users can no longer read the GPIO interrupt status register to determine which pin has triggered the interrupt. Instead, users should use the callback argument to pass the pin number.

Timer Group Driver Timer Group driver has been redesigned into *GPTimer*, which aims to unify and simplify the usage of general purpose timer.

Although it is recommended to use the new driver APIs, the legacy driver is still available in the previous include path `driver/timer.h`. However, by default, including `driver/timer.h` triggers the build warning below. The warning can be suppressed by the Kconfig option `CONFIG_GPTIMER_SUPPRESS_DEPRECATED_WARN`.

legacy timer group driver is deprecated, please migrate to `driver/gptimer.h`

The major breaking changes in concept and usage are listed as follows:

Breaking Changes in Concepts

- `timer_group_t` and `timer_idx_t` which used to identify the hardware timer are removed from user's code. In the new driver, a timer is represented by `gptimer_handle_t`.
- Definition of timer clock source is moved to `gptimer_clock_source_t`, the previous `timer_src_clk_t` is not used.
- Definition of timer count direction is moved to `gptimer_count_direction_t`, the previous `timer_count_dir_t` is not used.
- Only level interrupt is supported, `timer_intr_t` and `timer_intr_mode_t` are not used.
- Auto-reload is enabled by set the `gptimer_alarm_config_t::auto_reload_on_alarm` flag. `timer_autoreload_t` is not used.

Breaking Changes in Usage

- Timer initialization is done by creating a timer instance from `gptimer_new_timer()`. Basic configurations like clock source, resolution and direction should be set in `gptimer_config_t`. Note that, specific configurations of alarm events are not needed during the installation stage of the driver.
- Alarm event is configured by `gptimer_set_alarm_action()`, with parameters set in the `gptimer_alarm_config_t`.
- Setting and getting count value are done by `gptimer_get_raw_count()` and `gptimer_set_raw_count()`. The driver does not help convert the raw value into UTC time-stamp. Instead, the conversion should be done from user's side as the timer resolution is also known to the user.
- The driver will install the interrupt service as well if `gptimer_event_callbacks_t::on_alarm` is set to a valid callback function. In the callback, users do not have to deal with the low level registers (like "clear interrupt status", "re-enable alarm event" and so on). So functions like `timer_group_get_intr_status_in_isr` and `timer_group_get_auto_reload_in_isr` are not used anymore.
- To update the alarm configurations when alarm event happens, one can call `gptimer_set_alarm_action()` in the interrupt callback, then the alarm will be re-enabled again.
- Alarm will always be re-enabled by the driver if `gptimer_alarm_config_t::auto_reload_on_alarm` is set to true.

UART

Removed/Deprecated items	Replacement	Remarks
<code>uart_isr_register()</code>	None	UART interrupt handling is implemented by driver itself.
<code>uart_isr_free()</code>	None	UART interrupt handling is implemented by driver itself.
<code>use_ref_tick</code> in <code>uart_config_t</code>	<code>uart_config_t::source_clk</code>	Select the clock source.
<code>uart_enable_pattern_detection()</code>	<code>uart_enable_pattern_detection()</code>	Enable pattern detection interrupt.

I2C

Removed/Deprecated items	Replacement	Remarks
<code>i2c_isr_register()</code>	None	I2C interrupt handling is implemented by driver itself.
<code>i2c_isr_unregister()</code>	None	I2C interrupt handling is implemented by driver itself.
<code>i2c_opmode_t</code>	None	It is not used anywhere in ESP-IDF.

SPI

Removed/Deprecated items	Replacement	Remarks
<code>spi_cal_clock()</code>	<code>spi_get_actual_clock()</code>	Get SPI real working frequency.

- The internal header file `spi_common_internal.h` has been moved to `esp_private/spi_common_internal.h`.

SDMMC

Removed/Deprecated items	Replacement	Remarks
<code>sdmmc_host_pullup_enable()</code>	set <code>SDMMC_SLOT_FLAG_INTERNAL_PULLUP</code> flag in <code>sdmmc_slot_config_t::flags</code>	Enable internal pull up.

LEDC

Removed/Deprecated items	Replacement	Remarks
bit_num in <code>ledc_timer_config_t</code>	<code>ledc_timer_config_t::duty_resolution</code>	Set resolution of the duty cycle.

Pulse Counter Driver Pulse counter driver has been redesigned (see *PCNT*), which aims to unify and simplify the usage of PCNT peripheral.

Although it is recommended to use the new driver APIs, the legacy driver is still available in the previous include path `driver/pcnt.h`. However, including `driver/pcnt.h` triggers the build warning below by default. The warning can be suppressed by the Kconfig option `CONFIG_PCNT_SUPPRESS_DEPRECATED_WARN`.

legacy pcnt driver is deprecated, please migrate to use `driver/pulse_cnt.h`

The major breaking changes in concept and usage are listed as follows:

Breaking Changes in Concepts

- `pcnt_port_t`, `pcnt_unit_t` and `pcnt_channel_t` which used to identify the hardware unit and channel are removed from user's code. In the new driver, PCNT unit is represented by `pcnt_unit_handle_t`, likewise, PCNT channel is represented by `pcnt_channel_handle_t`. Both of them are opaque pointers.
- `pcnt_evt_type_t` is not used any more, they have been replaced by a universal **Watch Point Event**. In the event callback `pcnt_watch_cb_t`, it is still possible to distinguish different watch points from `pcnt_watch_event_data_t`.
- `pcnt_count_mode_t` is replaced by `pcnt_channel_edge_action_t`, and `pcnt_ctrl_mode_t` is replaced by `pcnt_channel_level_action_t`.

Breaking Changes in Usage

- Previously, the PCNT unit configuration and channel configuration were combined into a single function: `pcnt_unit_config`. They are now split into the two factory APIs: `pcnt_new_unit()` and `pcnt_new_channel()` respectively.
 - Only the count range is necessary for initializing a PCNT unit. GPIO number assignment has been moved to `pcnt_new_channel()`.
 - High/Low control mode and positive/negative edge count mode are set by stand-alone functions: `pcnt_channel_set_edge_action()` and `pcnt_channel_set_level_action()`.
- `pcnt_get_counter_value` is replaced by `pcnt_unit_get_count()`.
- `pcnt_counter_pause` is replaced by `pcnt_unit_stop()`.
- `pcnt_counter_resume` is replaced by `pcnt_unit_start()`.
- `pcnt_counter_clear` is replaced by `pcnt_unit_clear_count()`.
- `pcnt_intr_enable` and `pcnt_intr_disable` are removed. In the new driver, the interrupt is enabled by registering event callbacks `pcnt_unit_register_event_callbacks()`.
- `pcnt_event_enable` and `pcnt_event_disable` are removed. In the new driver, the PCNT events are enabled/disabled by adding/removing watch points `pcnt_unit_add_watch_point()`, `pcnt_unit_remove_watch_point()`.
- `pcnt_set_event_value` is removed. In the new driver, event value is also set when adding watch point by `pcnt_unit_add_watch_point()`.
- `pcnt_get_event_value` and `pcnt_get_event_status` are removed. In the new driver, these information are provided by event callback `pcnt_watch_cb_t` in the `pcnt_watch_event_data_t`.
- `pcnt_isr_register` and `pcnt_isr_unregister` are removed. Register of the ISR handler from user's code is no longer permitted. Users should register event callbacks instead by calling `pcnt_unit_register_event_callbacks()`.
- `pcnt_set_pin` is removed and the new driver no longer allows the switching of the GPIO at runtime. If users want to change to other GPIOs, please delete the existing PCNT channel by `pcnt_del_channel()` and reinstall with the new GPIO number by `pcnt_new_channel()`.

- `pcnt_filter_enable`, `pcnt_filter_disable` and `pcnt_set_filter_value` are replaced by `pcnt_unit_set_glitch_filter()`. Meanwhile, `pcnt_get_filter_value` has been removed.
- `pcnt_set_mode` is replaced by `pcnt_channel_set_edge_action()` and `pcnt_channel_set_level_action()`.
- `pcnt_isr_service_install`, `pcnt_isr_service_uninstall`, `pcnt_isr_handler_add` and `pcnt_isr_handler_remove` are replaced by `pcnt_unit_register_event_callbacks()`. The default ISR handler is lazy installed in the new driver.

RMT Driver RMT driver has been redesigned (see [RMT transceiver](#)), which aims to unify and extend the usage of RMT peripheral.

Although it is recommended to use the new driver APIs, the legacy driver is still available in the previous include path `driver/rmt.h`. However, including `driver/rmt.h` triggers the build warning below by default. The warning can be suppressed by the Kconfig option `CONFIG_RMT_SUPPRESS_DEPRECATED_WARN`.

The legacy RMT driver is deprecated, please use `driver/rmt_tx.h` and/or `driver/rmt_rx.h`

The major breaking changes in concept and usage are listed as follows:

Breaking Changes in Concepts

- `rmt_channel_t` which used to identify the hardware channel are removed from user space. In the new driver, RMT channel is represented by `rmt_channel_handle_t`. The channel is dynamically allocated by the driver, instead of designated by user.
- `rmt_item32_t` is replaced by `rmt_symbol_word_t`, which avoids a nested union inside a struct.
- `rmt_mem_t` is removed, as we do not allow users to access RMT memory block (a.k.a. RMTMEM) directly. Direct access to RMTMEM does not make sense but make mistakes, especially when the RMT channel also connected with a DMA channel.
- `rmt_mem_owner_t` is removed, as the ownership is controlled by driver, not by user anymore.
- `rmt_source_clk_t` is replaced by `rmt_clock_source_t`, and note they are not binary compatible.
- `rmt_data_mode_t` is removed, the RMT memory access mode is configured to always use Non-FIFO and DMA mode.
- `rmt_mode_t` is removed, as the driver has stand alone install functions for TX and RX channels.
- `rmt_idle_level_t` is removed, setting IDLE level for TX channel is available in `rmt_transmit_config_t::eot_level`.
- `rmt_carrier_level_t` is removed, setting carrier polarity is available in `rmt_carrier_config_t::polarity_active_low`.
- `rmt_channel_status_t` and `rmt_channel_status_result_t` are removed, they are not used anywhere.
- Transmitting by RMT channel does not expect user to prepare the RMT symbols, instead, user needs to provide an RMT Encoder to tell the driver how to convert user data into RMT symbols.

Breaking Changes in Usage

- Channel installation has been separated for TX and RX channels into `rmt_new_tx_channel()` and `rmt_new_rx_channel()`.
- `rmt_set_clk_div` and `rmt_get_clk_div` are removed. Channel clock configuration can only be done during channel installation.
- `rmt_set_rx_idle_thresh` and `rmt_get_rx_idle_thresh` are removed. In the new driver, the RX channel IDLE threshold is redesigned into a new concept `rmt_receive_config_t::signal_range_max_ns`.
- `rmt_set_mem_block_num` and `rmt_get_mem_block_num` are removed. In the new driver, the memory block number is determined by `rmt_tx_channel_config_t::mem_block_symbols` and `rmt_rx_channel_config_t::mem_block_symbols`.

- `rmt_set_tx_carrier` is removed, the new driver uses `rmt_apply_carrier()` to set carrier behavior.
- `rmt_set_mem_pd` and `rmt_get_mem_pd` are removed. The memory power is managed by the driver automatically.
- `rmt_memory_rw_rst`, `rmt_tx_memory_reset` and `rmt_rx_memory_reset` are removed. Memory reset is managed by the driver automatically.
- `rmt_tx_start` and `rmt_rx_start` are merged into a single function `rmt_enable()`, for both TX and RX channels.
- `rmt_tx_stop` and `rmt_rx_stop` are merged into a single function `rmt_disable()`, for both TX and RX channels.
- `rmt_set_memory_owner` and `rmt_get_memory_owner` are removed. RMT memory owner guard is added automatically by the driver.
- `rmt_set_tx_loop_mode` and `rmt_get_tx_loop_mode` are removed. In the new driver, the loop mode is configured in `rmt_transmit_config_t::loop_count`.
- `rmt_set_source_clk` and `rmt_get_source_clk` are removed. Configuring clock source is only possible during channel installation by `rmt_tx_channel_config_t::clk_src` and `rmt_rx_channel_config_t::clk_src`.
- `rmt_set_rx_filter` is removed. In the new driver, the filter threshold is redesigned into a new concept `rmt_receive_config_t::signal_range_min_ns`.
- `rmt_set_idle_level` and `rmt_get_idle_level` are removed. Setting IDLE level for TX channel is available in `rmt_transmit_config_t::eot_level`.
- `rmt_set_rx_intr_en`, `rmt_set_err_intr_en`, `rmt_set_tx_intr_en`, `rmt_set_tx_thr_intr_en` and `rmt_set_rx_thr_intr_en` are removed. The new driver does not allow user to turn on/off interrupt from user space. Instead, it provides callback functions.
- `rmt_set_gpio` and `rmt_set_pin` are removed. The new driver does not support to switch GPIO dynamically at runtime.
- `rmt_config` is removed. In the new driver, basic configuration is done during the channel installation stage.
- `rmt_isr_register` and `rmt_isr_deregister` are removed, the interrupt is allocated by the driver itself.
- `rmt_driver_install` is replaced by `rmt_new_tx_channel()` and `rmt_new_rx_channel()`.
- `rmt_driver_uninstall` is replaced by `rmt_del_channel()`.
- `rmt_fill_tx_items`, `rmt_write_items` and `rmt_write_sample` are removed. In the new driver, user needs to provide an encoder to "translate" the user data into RMT symbols.
- `rmt_get_counter_clock` is removed, as the channel clock resolution is configured by user from `rmt_tx_channel_config_t::resolution_hz`.
- `rmt_wait_tx_done` is replaced by `rmt_tx_wait_all_done()`.
- `rmt_translator_init`, `rmt_translator_set_context` and `rmt_translator_get_context` are removed. In the new driver, the translator has been replaced by the RMT encoder.
- `rmt_get_ringbuf_handle` is removed. The new driver does not use Ringbuffer to save RMT symbols. Instead, the incoming data are saved to the user provided buffer directly. The user buffer can even be mounted to DMA link internally.
- `rmt_register_tx_end_callback` is replaced by `rmt_tx_register_event_callbacks()`, where user can register `rmt_tx_event_callbacks_t::on_trans_done` event callback.
- `rmt_set_intr_enable_mask` and `rmt_clr_intr_enable_mask` are removed, as the interrupt is handled by the driver, user does not need to take care of it.
- `rmt_add_channel_to_group` and `rmt_remove_channel_from_group` are replaced by RMT sync manager. Please refer to `rmt_new_sync_manager()`.
- `rmt_set_tx_loop_count` is removed. The loop count in the new driver is configured in `rmt_transmit_config_t::loop_count`.
- `rmt_enable_tx_loop_autostop` is removed. In the new driver, TX loop auto stop is always enabled if available, it is not configurable anymore.

LCD

- The LCD panel initialization flow is slightly changed. Now the `esp_lcd_panel_init()` will not turn on the display automatically. User needs to call `esp_lcd_panel_disp_on_off()` to manually turn on the display. Note, this is different from turning on backlight. With this breaking change, user can flash a

predefined pattern to the screen before turning on the screen. This can help avoid random noise on the screen after a power on reset.

- `esp_lcd_panel_disp_off()` is deprecated, please use `esp_lcd_panel_disp_on_off()` instead.
- `dc_as_cmd_phase` is removed. The SPI LCD driver currently does not support a 9-bit SPI LCD. Please always use a dedicated GPIO to control the LCD D/C line.
- The way to register RGB panel event callbacks has been moved from the `esp_lcd_rgb_panel_config_t` into a separate API `esp_lcd_rgb_panel_register_event_callbacks()`. However, the event callback signature is not changed.
- Previous `relax_on_idle` flag in `esp_lcd_rgb_panel_config_t` has been renamed into `esp_lcd_rgb_panel_config_t::refresh_on_demand`, which expresses the same meaning but with a clear name.
- If the RGB LCD is created with the `refresh_on_demand` flag enabled, the driver will not start a refresh in the `esp_lcd_panel_draw_bitmap()`. Now users have to call `esp_lcd_rgb_panel_refresh()` to refresh the screen by themselves.
- `esp_lcd_color_space_t` is deprecated, please use `lcd_color_space_t` to describe the color space, and use `lcd_rgb_element_order_t` to describe the data order of RGB color.

MCPWM MCPWM driver was redesigned (see [MCPWM](#)), meanwhile, the legacy driver is deprecated.

The new driver's aim is to make each MCPWM submodule independent to each other, and give the freedom of resource connection back to users.

Although it is recommended to use the new driver APIs, the legacy driver is still available in the previous include path `driver/mcpwm.h`. However, using legacy driver triggers the build warning below by default. This warning can be suppressed by the Kconfig option `CONFIG_MCPWM_SUPPRESS_DEPRECATED_WARN`.

```
legacy MCPWM driver is deprecated, please migrate to the new driver (include_
↪driver/mcpwm_prelude.h)
```

The major breaking changes in concept and usage are listed as follows:

Breaking Changes in Concepts The new MCPWM driver is object-oriented, where most of the MCPWM submodule has a driver object associated with it. The driver object is created by factory function like `mcpwm_new_timer()`. IO control function always needs an object handle, in the first place.

The legacy driver has an inappropriate assumption, that is the MCPWM operator should be connected to different MCPWM timer. In fact, the hardware does not have such limitation. In the new driver, a MCPWM timer can be connected to multiple operators, so that the operators can achieve the best synchronization performance.

The legacy driver presets the way to generate a PWM waveform into a so called `mcpwm_duty_type_t`. However, the duty cycle modes listed there are far from sufficient. Likewise, legacy driver has several preset `mcpwm_deadtime_type_t`, which also does not cover all the use cases. What is more, user usually gets confused by the name of the duty cycle mode and dead-time mode. In the new driver, there are no such limitation, but user has to construct the generator behavior from scratch.

In the legacy driver, the ways to synchronize the MCPWM timer by GPIO, software and other timer module are not unified. It increased learning costs for users. In the new driver, the synchronization APIs are unified.

The legacy driver has mixed the concepts of "Fault detector" and "Fault handler". Which make the APIs very confusing to users. In the new driver, the fault object just represents a failure source, and we introduced a new concept -- **brake** to express the concept of "Fault handler". What is more, the new driver supports software fault.

The legacy drive only provides callback functions for the capture submodule. The new driver provides more useful callbacks for various MCPWM submodules, like timer stop, compare match, fault enter, brake, etc.

- `mcpwm_io_signals_t` and `mcpwm_pin_config_t` are not used. GPIO configuration has been moved into submodule's configuration structure.
- `mcpwm_timer_t`, `mcpwm_generator_t` are not used. Timer and generator are represented by `mcpwm_timer_handle_t` and `mcpwm_gen_handle_t`.

- `mcpwm_fault_signal_t` and `mcpwm_sync_signal_t` are not used. Fault and sync source are represented by `mcpwm_fault_handle_t` and `mcpwm_sync_handle_t`.
- `mcpwm_capture_signal_t` is not used. A capture channel is represented by `mcpwm_cap_channel_handle_t`.

Breaking Changes in Usage


- `mcpwm_gpio_init` and `mcpwm_set_pin`: GPIO configurations are moved to submodule's own configuration. e.g., set the PWM GPIO in `mcpwm_generator_config_t::gen_gpio_num`.
- `mcpwm_init`: To get an expected PWM waveform, users need to allocate at least one MCPWM timer and MCPWM operator, then connect them by calling `mcpwm_operator_connect_timer()`. After that, users should set the generator's actions on various events by calling e.g., `mcpwm_generator_set_actions_on_timer_event()`, `mcpwm_generator_set_actions_on_compare_event()`.
- `mcpwm_group_set_resolution`: in the new driver, the group resolution is fixed to the maximum, usually it is 80 MHz.
- `mcpwm_timer_set_resolution`: MCPWM Timer resolution is set in `mcpwm_timer_config_t::resolution_hz`.
- `mcpwm_set_frequency`: PWM frequency is determined by `mcpwm_timer_config_t::resolution_hz`, `mcpwm_timer_config_t::count_mode` and `mcpwm_timer_config_t::period_ticks`.
- `mcpwm_set_duty`: To set the PWM duty cycle, users should call `mcpwm_comparator_set_compare_value()` to change comparator's threshold.
- `mcpwm_set_duty_type`: There is no preset duty cycle types. The duty cycle type is configured by setting different generator actions. e.g., `mcpwm_generator_set_actions_on_timer_event()`.
- `mcpwm_set_signal_high` and `mcpwm_set_signal_low` are replaced by `mcpwm_generator_set_force_level()`. In the new driver, it is implemented by setting force action for the generator, instead of changing the duty cycle to 0% or 100% at the background.
- `mcpwm_start` and `mcpwm_stop` are replaced by `mcpwm_timer_start_stop()`. You have more modes to start and stop the MCPWM timer, see `mcpwm_timer_start_stop_cmd_t`.
- `mcpwm_carrier_init` is replaced by `mcpwm_operator_apply_carrier()`.
- `mcpwm_carrier_enable` and `mcpwm_carrier_disable`: Enabling and disabling carrier submodule is done automatically by checking whether the carrier configuration structure `mcpwm_carrier_config_t` is NULL.
- `mcpwm_carrier_set_period` is replaced by `mcpwm_carrier_config_t::frequency_hz`.
- `mcpwm_carrier_set_duty_cycle` is replaced by `mcpwm_carrier_config_t::duty_cycle`.
- `mcpwm_carrier_oneshot_mode_enable` is replaced by `mcpwm_carrier_config_t::first_pulse_duration`.
- `mcpwm_carrier_oneshot_mode_disable` is removed. Disabling the first pulse (a.k.a the one-shot pulse) in the carrier is never supported by the hardware.
- `mcpwm_carrier_output_invert` is replaced by `mcpwm_carrier_config_t::invert_before_modulate` and `mcpwm_carrier_config_t::invert_after_modulate`.
- `mcpwm_deadtime_enable` and `mcpwm_deadtime_disable` are replaced by `mcpwm_generator_set_dead_time()`.
- `mcpwm_fault_init` is replaced by `mcpwm_new_gpio_fault()`.
- `mcpwm_fault_set_oneshot_mode`, `mcpwm_fault_set_cyc_mode` are replaced by `mcpwm_operator_set_brake_on_fault()` and `mcpwm_generator_set_actions_on_brake_event()`.
- `mcpwm_capture_enable` is removed. It is duplicated to `mcpwm_capture_enable_channel()`.
- `mcpwm_capture_disable` is removed. It is duplicated to `mcpwm_capture_disable_channel()`.
- `mcpwm_capture_enable_channel` and `mcpwm_capture_disable_channel` are replaced by `mcpwm_capture_channel_enable()` and `mcpwm_capture_channel_disable()`.
- `mcpwm_capture_signal_get_value` and `mcpwm_capture_signal_get_edge`: Capture timer count value and capture edge are provided in the capture event callback, via `mcpwm_capture_event_data_t`. Capture data are only valuable when capture event happens. Providing single API to fetch capture data is meaningless.
- `mcpwm_sync_enable` is removed. It is duplicated to `mcpwm_sync_configure()`.
- `mcpwm_sync_configure` is replaced by `mcpwm_timer_set_phase_on_sync()`.
- `mcpwm_sync_disable` is equivalent to setting `mcpwm_timer_sync_phase_config_t::sync_src` to NULL.

- `mcpwm_set_timer_sync_output` is replaced by `mcpwm_new_timer_sync_src()`.
- `mcpwm_timer_trigger_soft_sync` is replaced by `mcpwm_soft_sync_activate()`.
- `mcpwm_sync_invert_gpio_synchro` is equivalent to setting `mcpwm_gpio_sync_src_config_t::active_ne`.
- `mcpwm_isr_register` is removed. You can register various event callbacks instead. For example, to register capture event callback, users can use `mcpwm_capture_channel_register_event_callbacks()`.

I2S Driver The I2S driver has been redesigned (see [I2S Driver](#)), which aims to rectify the shortcomings of the driver that were exposed when supporting all the new features of ESP32-C3 & ESP32-S3. The new driver's APIs are available by including corresponding I2S mode's header files `driver/i2s/include/driver/i2s_std.h`, `driver/i2s/include/driver/i2s_pdm.h`, or `driver/i2s/include/driver/i2s_tdm.h`.

Meanwhile, the old driver's APIs in `driver/deprecated/driver/i2s.h` are still supported for backward compatibility. But there will be warnings if users keep using the old APIs in their projects, these warnings can be suppressed by the Kconfig option `CONFIG_I2S_SUPPRESS_DEPRECATED_WARN`.

Here is the general overview of the current I2S files:



Breaking changes in Concepts

Independent TX/RX channels The minimum control unit in new I2S driver are now individual TX/RX channels instead of an entire I2S controller (that consists of multiple channels).

- The TX and RX channels of the same I2S controller can be controlled separately, meaning that they are configured such that they can be started or stopped separately.
- The `i2s_chan_handle_t` handle type is used to uniquely identify I2S channels. All the APIs require the channel handle and users need to maintain the channel handles by themselves.
- On the ESP32-C3 and ESP32-S3, TX and RX channels in the same controller can be configured to different clocks or modes.
- However, on the ESP32 and ESP32-S2, the TX and RX channels of the same controller still share some hardware resources. Thus, configurations may cause one channel to affect another channel in the same controller.
- The channels can be registered to an available I2S controller automatically by setting `i2s_port_t::I2S_NUM_AUTO` as I2S port ID which causes the driver to search for the available TX/RX channels. However, the driver also supports registering channels to a specific port.
- In order to distinguish between TX/RX channels and sound channels, the term "channel" in the context of the I2S driver only refers to TX/RX channels. Meanwhile, sound channels are referred to as "slots".

I2S Mode Categorization I2S communication modes are categorized into the following three modes. Note that:

- **Standard mode:** Standard mode always has two slots, it can support Philips, MSB, and PCM (short frame sync) formats. Please refer to [driver/i2s/include/driver/i2s_std.h](#) for more details.
- **PDM mode:** PDM mode only supports two slots with 16-bit data width, but the configurations of PDM TX and PDM RX are slightly different. For PDM TX, the sample rate can be set by `i2s_pdm_tx_clk_config_t::sample_rate`, and its clock frequency depends on the up-sampling configuration. For PDM RX, the sample rate can be set by `i2s_pdm_rx_clk_config_t::sample_rate`, and its clock frequency depends on the down-sampling configuration. Please refer to [driver/i2s/include/driver/i2s_pdm.h](#) for details.
- **TDM mode:** TDM mode can support up to 16 slots. It can work in Philips, MSB, PCM (short frame sync), and PCM (long frame sync) formats. Please refer to [driver/i2s/include/driver/i2s_tdm.h](#) for details.

When allocating a new channel in a specific mode, users should initialize that channel by its corresponding function. It is strongly recommended to use the helper macros to generate the default configurations in case the default values are changed in the future.

Independent Slot and Clock Configuration The slot configurations and clock configurations can be configured separately.

- Call `i2s_channel_init_std_mode()`, `i2s_channel_init_pdm_rx_mode()`, `i2s_channel_init_pdm_tx_mode()`, or `i2s_channel_init_tdm_mode()` to initialize the slot/clock/gpio_pin configurations.
- Calling `i2s_channel_reconfig_std_slot()`, `i2s_channel_reconfig_pdm_rx_slot()`, `i2s_channel_reconfig_pdm_tx_slot()`, or `i2s_channel_reconfig_tdm_slot()` can change the slot configurations after initialization.
- Calling `i2s_channel_reconfig_std_clock()`, `i2s_channel_reconfig_pdm_rx_clock()`, `i2s_channel_reconfig_pdm_tx_clock()`, or `i2s_channel_reconfig_tdm_clock()` can change the clock configurations after initialization.
- Calling `i2s_channel_reconfig_std_gpio()`, `i2s_channel_reconfig_pdm_rx_gpio()`, `i2s_channel_reconfig_pdm_tx_gpio()`, or `i2s_channel_reconfig_tdm_gpio()` can change the GPIO configurations after initialization.

Misc

- States and state-machine are adopted in the new I2S driver to avoid APIs called in wrong state.
- ADC and DAC modes are removed. They are only supported in their own drivers and the legacy I2S driver.

Breaking Changes in Usage To use the new I2S driver, please follow these steps:

1. Call `i2s_new_channel()` to acquire channel handles. We should specify the work role and I2S port in this step. Besides, the TX or RX channel handle will be generated by the driver. Inputting both two TX and RX channel handles is not necessary but at least one handle is needed. In the case of inputting both two handles, the driver will work at the duplex mode. Both TX and RX channels will be available on a same port, and they will share the MCLK, BCLK and WS signal. But if only one of the TX or RX channel handle is inputted, this channel will only work in the simplex mode.
2. Call `i2s_channel_init_std_mode()`, `i2s_channel_init_pdm_rx_mode()`, `i2s_channel_init_pdm_tx_mode()` or `i2s_channel_init_tdm_mode()` to initialize the channel to the specified mode. Corresponding slot, clock and GPIO configurations are needed in this step.
3. (Optional) Call `i2s_channel_register_event_callback()` to register the ISR event callback functions. I2S events now can be received by the callback function synchronously, instead of from the event queue asynchronously.
4. Call `i2s_channel_enable()` to start the hardware of I2S channel. In the new driver, I2S does not start automatically after installed, and users are supposed to know clearly whether the channel has started or not.
5. Read or write data by `i2s_channel_read()` or `i2s_channel_write()`. Certainly, only the RX channel handle is supposed to be inputted in `i2s_channel_read()` and the TX channel handle in `i2s_channel_write()`.
6. (Optional) The slot, clock and GPIO configurations can be changed by corresponding 'reconfig' functions, but `i2s_channel_disable()` must be called before updating the configurations.

7. Call `i2s_channel_disable()` to stop the hardware of I2S channel.
8. Call `i2s_del_channel()` to delete and release the resources of the channel if it is not needed any more, but the channel must be disabled before deleting it.

Register Access Macros Previously, all register access macros could be used as expressions, so the following was allowed:

```
uint32_t val = REG_SET_BITS(reg, bits, mask);
```

In ESP-IDF v5.0, register access macros which write or read-modify-write the register can no longer be used as expressions, and can only be used as statements. This applies to the following macros: `REG_WRITE`, `REG_SET_BIT`, `REG_CLR_BIT`, `REG_SET_BITS`, `REG_SET_FIELD`, `WRITE_PERI_REG`, `CLEAR_PERI_REG_MASK`, `SET_PERI_REG_MASK`, `SET_PERI_REG_BITS`.

To store the value which would have been written into the register, split the operation as follows:

```
uint32_t new_val = REG_READ(reg) | mask;
REG_WRITE(reg, new_val);
```

To get the value of the register after modification (which may be different from the value written), add an explicit read:

```
REG_SET_BITS(reg, bits, mask);
uint32_t new_val = REG_READ(reg);
```

Protocols

Mbed TLS For ESP-IDF v5.0, [Mbed TLS](#) has been updated from v2.x to v3.1.0.

For more details about Mbed TLS's migration from version 2.x to version 3.0 or greater, please refer to the [official guide](#).

Breaking Changes (Summary)

Most Structure Fields Are Now Private

- Direct access to fields of structures (`struct` types) declared in public headers is no longer supported.
- Appropriate accessor functions (getter/setter) must be used for the same. A temporary workaround would be to use `MBEDTLS_PRIVATE` macro (**not recommended**).
- For more details, refer to the [official guide](#).

SSL

- Removed support for TLS 1.0, 1.1, and DTLS 1.0
- Removed support for SSL 3.0

Deprecated Functions Were Removed from Cryptography Modules

- The functions `MBEDTLS_*_ret()` (related to MD, SHA, RIPEMD, RNG, HMAC modules) was renamed to replace the corresponding functions without `_ret` appended and updated return value.
- For more details, refer to the [official guide](#).

Deprecated Config Options Following are some of the important config options deprecated by this update. The configs related to and/or dependent on these have also been deprecated.

- `MBEDTLS_SSL_PROTO_SSL3` : Support for SSL 3.0
- `MBEDTLS_SSL_PROTO_TLS1` : Support for TLS 1.0
- `MBEDTLS_SSL_PROTO_TLS1_1` : Support for TLS 1.1
- `MBEDTLS_SSL_PROTO_DTLS` : Support for DTLS 1.1 (Only DTLS 1.2 is supported now)
- `MBEDTLS_DES_C` : Support for 3DES ciphersuites
- `MBEDTLS_RC4_MODE` : Support for RC4-based ciphersuites

Note: This list includes only major options configurable through `idf.py menuconfig`. For more details on deprecated options, refer to the [official guide](#).

Miscellaneous

Disabled Diffie-Hellman Key Exchange Modes The Diffie-Hellman Key Exchange modes have now been disabled by default due to security risks (see warning text [here](#)). Related configs are given below:

- `MBEDTLS_DHM_C` : Support for the Diffie-Hellman-Merkle module
- `MBEDTLS_KEY_EXCHANGE_DHE_PSK` : Support for Diffie-Hellman PSK (pre-shared-key) TLS authentication modes
- `MBEDTLS_KEY_EXCHANGE_DHE_RSA` : Support for cipher suites with the prefix `TLS-DHE-RSA-WITH-`

Note: During the initial step of the handshake (i.e., `client_hello`), the server selects a cipher from the list that the client publishes. As the `DHE_PSK/DHE_RSA` ciphers have now been disabled by the above change, the server would fall back to an alternative cipher; if in a rare case, it does not support any other cipher, the handshake would fail. To retrieve the list of ciphers supported by the server, one must attempt to connect with the server with a specific cipher from the client-side. Few utilities can help do this, e.g., `sslsan`.

Remove `certs` Module from X509 Library

- The `mbdttls/certs.h` header is no longer available in mbedtlls 3.1. Most applications can safely remove it from the list of includes.

Breaking Change for `esp_crt_bundle_set` API

- The `esp_crt_bundle_set()` API now requires one additional argument named `bundle_size`. The return type of the API has also been changed to `esp_err_t` from `void`.

Breaking Change for `esp_ds_rsa_sign` API

- The `esp_ds_rsa_sign()` API now requires one less argument. The argument `mode` is no longer required.

HTTPS Server

Breaking Changes (Summary) Names of variables holding different certs in `httpd_ssl_config_t` structure have been updated.

- `httpd_ssl_config::servercert` variable inherits role of `cacert_pem` variable.
- `httpd_ssl_config::servercert_len` variable inherits role of `cacert_len` variable
- `httpd_ssl_config::cacert_pem` variable inherits role of `client_verify_cert_pem` variable

- `httpd_ssl_config::cacert_len` variable inherits role of `client_verify_cert_len` variable

The return type of the `httpd_ssl_stop()` API has been changed to `esp_err_t` from `void`.

ESP HTTPS OTA

Breaking Changes (Summary)

- The function `esp_https_ota()` now requires pointer to `esp_https_ota_config_t` as argument instead of pointer to `esp_http_client_config_t`.

ESP-TLS

Breaking Changes (Summary)

esp_tls_t Structure Is Now Private The `esp_tls_t` has now been made completely private. You cannot access its internal structures directly. Any necessary data that needs to be obtained from the ESP-TLS handle can be done through respective getter/setter functions. If there is a requirement of a specific getter/setter function, please raise an [issue](#) on ESP-IDF.

The list of newly added getter/setter function is as follows:

- `esp_tls_get_ssl_context()` - Obtain the ssl context of the underlying ssl stack from the ESP-TLS handle.

Function Deprecations And Recommended Alternatives Following table summarizes the deprecated functions removed and their alternatives to be used from ESP-IDF v5.0 onwards.

Deprecated Function	Alternative
<code>esp_tls_conn_new()</code>	<code>esp_tls_conn_new_sync()</code>
<code>esp_tls_conn_delete()</code>	<code>esp_tls_conn_destroy()</code>

- The function `esp_tls_conn_http_new()` has now been termed as deprecated. Please use the alternative function `esp_tls_conn_http_new_sync()` (or its asynchronous `esp_tls_conn_http_new_async()`). Note that the alternatives need an additional parameter `esp_tls_t`, which has to be initialized using the `esp_tls_init()` function.

HTTP Server

Breaking Changes (Summary)

- `http_server.h` header is no longer available in `esp_http_server`. Please use `esp_http_server.h` instead.

ESP HTTP Client

Breaking Changes (Summary)

- The functions `esp_http_client_read()` and `esp_http_client_fetch_headers()` now return an additional return value `-ESP_ERR_HTTP_EAGAIN` for timeout errors - call timed-out before any data was ready.

TCP Transport

Breaking Changes (Summary)

- The function `esp_transport_read()` now returns 0 for a connection timeout and < 0 for other errors. Please refer `esp_tcp_transport_err_t` for all possible return values.

MQTT Client

Breaking Changes (Summary)

- `esp_mqtt_client_config_t` have all fields grouped in sub structs.

Most common configurations are listed below:

- Broker address now is set in `esp_mqtt_client_config_t::broker::address::uri`
- Security related to broker verification in `esp_mqtt_client_config_t::broker::verification`
- Client username is set in `esp_mqtt_client_config_t::credentials::username`
- `esp_mqtt_client_config_t` no longer supports the `user_context` field. Please use `esp_mqtt_client_register_event()` instead for registering an event handler; the last argument `event_handler_arg` can be used to pass user context to the handler.

ESP-Modbus

Breaking Changes (Summary) The ESP-IDF component `freemodbus` has been removed from ESP-IDF and is supported as a separate component. Additional information for the ESP-Modbus component can be found in the separate repository:

- [ESP-Modbus component on GitHub](#)

The main component folder of the new application shall include the component manager manifest file `idf_component.yml` as in the example below:

```
dependencies:
  espressif/esp-modbus:
    version: "^1.0"
```

The `esp-modbus` component can be found in [component manager registry](#). Refer to [component manager documentation](#) for more information on how to set up the component manager.

For applications targeting v4.x releases of ESP-IDF that need to use new `esp-modbus` component, adding the component manager manifest file `idf_component.yml` will be sufficient to pull in the new component. However, users should also exclude the legacy `freemodbus` component from the build. This can be achieved using the statement below in the project's `CMakeLists.txt`:

```
set(EXCLUDE_COMPONENTS freemodbus)
```

Provisioning

Protocomm The `pop` field in the `protocomm_set_security()` API is now deprecated. Please use the `sec_params` field instead of `pop`. This parameter should contain the structure (including the security parameters) as required by the protocol version used.

For example, when using security version 2, the `sec_params` parameter should contain the pointer to the structure of type `protocomm_security2_params_t`.

Wi-Fi Provisioning

- The `pop` field in the `wifi_prov_mgr_start_provisioning()` API is now deprecated. For backward compatibility, `pop` can be still passed as a string for security version 1. However, for security version 2, the `wifi_prov_sec_params` argument needs to be passed instead of `pop`. This parameter should contain the structure (containing the security parameters) as required by the protocol version used. For example, when using security version 2, the `wifi_prov_sec_params` parameter should contain the pointer to the structure of type `wifi_prov_security2_params_t`. For security 1, the behaviour and the usage of the API remain the same.
- The API `wifi_prov_mgr_is_provisioned()` does not return `ESP_ERR_INVALID_STATE` error any more. This API now works without any dependency on provisioning manager initialization state.

ESP Local Control The `pop` field in the `esp_local_ctrl_proto_sec_cfg_t` API is now deprecated. Please use the `sec_params` field instead of `pop`. This field should contain the structure (containing the security parameters) as required by the protocol version used.

For example, when using security version 2, the `sec_params` field should contain pointer to the structure of type `esp_local_ctrl_security2_params_t`.

Removed or Deprecated Components

Components Moved to ESP-IDF Component Registry Following components are removed from ESP-IDF and moved to [ESP-IDF Component Registry](#):

- [libsodium](#)
- [cbor](#)
- [jsmn](#)
- [esp_modem](#)
- [nghttp](#)
- [mdns](#)
- [esp_websocket_client](#)
- [asio](#)
- [freemodbus](#)
- [sh2lib](#)
- [expat](#)
- [coap](#)
- [esp-cryptoauthlib](#)
- [qrcode](#)
- [tjpgd](#)
- [esp_serial_slave_link](#)
- [tinyusb](#)

Note: Please note that http parser functionality which was previously part of `nghttp` component is now part of [http_parser](#) component.

These components can be installed using `idf.py add-dependency` command.

For example, to install `libsodium` component with the exact version X.Y, run `idf.py add-dependency libsodium==X.Y`.

To install `libsodium` component with the latest version compatible to X.Y according to [semver](#) rules, run `idf.py add-dependency libsodium~X.Y`.

To find out which versions of each component are available, open <https://components.espressif.com>, search for the component by its name and check the versions listed on the component page.

Deprecated Components The following components are removed since they were deprecated in ESP-IDF v4.x:

- `tcpip_adapter`. Please use the [ESP-NETIF](#) component instead; you can follow the [TCP/IP Adapter](#).

Note: OpenSSL-API component is no longer supported. It is not available in the IDF Component Registry, either. Please use [ESP-TLS](#) or [mbedtls](#) API directly.

Note: `esp_adc_cal` component is no longer supported. New adc calibration driver is in `esp_adc` component. Legacy adc calibration driver has been moved into `esp_adc` component. To use legacy `esp_adc_cal` driver APIs, you should add `esp_adc` component to the list of component requirements in `CMakeLists.txt`. Also check [Peripherals Migration Guide](#) for more details.

The targets components are no longer necessary after refactoring and have been removed:

- `esp32`
- `esp32s2`
- `esp32s3`
- `esp32c2`
- `esp32c3`
- `esp32h2`

Storage

New Component for the Partition APIs Breaking change: all the Partition API code has been moved to a new component [esp_partition](#). For the complete list of affected functions and data-types, see header file [esp_partition.h](#).

These API functions and data-types were previously a part of the [spi_flash](#) component, and thus possible dependencies on the `spi_flash` in existing applications may cause the build failure, in case of direct `esp_partition_*` APIs/data-types use (for instance, fatal error: `esp_partition.h`: No such file or directory at lines with `#include "esp_partition.h"`). If you encounter such an issue, please update your project's `CMakeLists.txt` file as follows:

Original dependency setup:

```
idf_component_register(...  
                        REQUIRES spi_flash)
```

Updated dependency setup:

```
idf_component_register(...  
                        REQUIRES spi_flash esp_partition)
```

Note: Please update relevant `REQUIRES` or `PRIV_REQUIRES` section according to your project. The above-presented code snippet is just an example.

If the issue persists, please let us know and we will assist you with your code migration.

SDMMC/SDSPI SD card frequency on SDMMC/SDSPI interface can be now configured through `sdmmc_host_t.max_freq_khz` to a specific value, not only `SDMMC_FREQ_PROBING` (400 kHz), `SDMMC_FREQ_DEFAULT` (20 MHz), or `SDMMC_FREQ_HIGHSPEED` (40 MHz). Previously, in case you have specified a custom frequency other than any of the above-mentioned values, the closest lower-or-equal one was selected anyway.

Now, the underlying drivers calculate the nearest fitting value, given by available frequency dividers instead of an enumeration item selection. This could cause troubles in communication with your SD card without a change of the existing application code. If you encounter such an issue, please, keep trying different frequencies around your desired

value unless you find the one working well. To check the frequency value calculated and actually applied, use `void sdmmc_card_print_info(FILE* stream, const sdmmc_card_t* card)` function.

FatFs FatFs is now updated to v0.14. As a result, the function signature of `f_mkfs()` has changed. The new signature is `FRESULT f_mkfs(const TCHAR* path, const MKFS_PARM* opt, void* work, UINT len);` which uses `MKFS_PARM` struct as a second argument.

Partition Table The partition table generator no longer supports misaligned partitions. When generating a partition table, ESP-IDF only accepts partitions with offsets that align to 4 KB. This change only affects generating new partition tables. Reading and writing to already existing partitions remains unchanged.

VFS The `esp_vfs_semihost_register()` function signature is changed as follows:

- The new signature is `esp_err_t esp_vfs_semihost_register(const char* base_path);`
- The `host_path` parameter of the old signature no longer exists. Instead, the OpenOCD command `ESP_SEMIHOST_BASEDIR` should be used to set the full path on the host.

Function Signature Changes The following functions now return `esp_err_t` instead of `void` or `nvs_iterator_t`. Previously, when parameters were invalid or when something goes wrong internally, these functions would assert() or return a `nullptr`. With an `esp_err_t` returned, you can get better error reporting.

- `nvs_entry_find()`
- `nvs_entry_next()`
- `nvs_entry_info()`

Because the `esp_err_t` return type changes, the usage patterns of `nvs_entry_find()` and `nvs_entry_next()` become different. Both functions now modify iterators via parameters instead of returning an iterator.

The old programming pattern to iterate over an NVS partition was as follows:

```
nvs_iterator_t it = nvs_entry_find(<nvs_partition_name>, <namespace>, NVS_TYPE_
↪ANY);
while (it != NULL) {
    nvs_entry_info_t info;
    nvs_entry_info(it, &info);
    it = nvs_entry_next(it);
    printf("key '%s', type '%d'", info.key, info.type);
};
```

The new programming pattern to iterate over an NVS partition is now:

```
nvs_iterator_t it = nullptr;
esp_err_t res = nvs_entry_find(<nvs_partition_name>, <namespace>, NVS_TYPE_ANY, &
↪it);
while(res == ESP_OK) {
    nvs_entry_info_t info;
    nvs_entry_info(it, &info); // Can omit error check if parameters are_
↪guaranteed to be non-NULL
    printf("key '%s', type '%d'", info.key, info.type);
    res = nvs_entry_next(&it);
}
nvs_release_iterator(it);
```

Iterator Validity Note that because the function signature changes, if there is a parameter error, you may get an invalid iterator from `nvs_entry_find()`. Hence, it is important to initialize the iterator

to `NULL` before using `nvs_entry_find()`, so that you can avoid complex error checking before calling `nvs_release_iterator()`. A good example is the programming pattern above.

Removed SDSPI Deprecated API Structure `sdspi_slot_config_t` and function `sdspi_host_init_slot()` are removed, and replaced by structure `sdspi_device_config_t` and function `sdspi_host_init_device()` respectively.

ROM SPI Flash In versions before v5.0, ROM SPI flash functions were included via `esp32**/rom/spi_flash.h`. Thus, code written to support different ESP chips might be filled with ROM headers of different targets. Furthermore, not all of the APIs could be used on all ESP chips.

Now, the common APIs are extracted to `esp_rom_spiflash.h`. Although it is not a breaking change, you are strongly recommended to only use the functions from this header (i.e., prefixed with `esp_rom_spiflash` and included by `esp_rom_spiflash.h`) for better cross-compatibility between ESP chips.

To make ROM SPI flash APIs clearer, the following functions are also renamed:

- `esp_rom_spiflash_lock()` to `esp_rom_spiflash_set_bp()`
- `esp_rom_spiflash_unlock()` to `esp_rom_spiflash_clear_bp()`

SPI Flash Driver The `esp_flash_speed_t` enum type is now deprecated. Instead, you may now directly pass the real clock frequency value to the flash configuration structure. The following example demonstrates how to configure a flash frequency of 80MHz:

```
esp_flash_spi_device_config_t dev_cfg = {
    // Other members
    .freq_mhz = 80,
    // Other members
};
```

Legacy SPI Flash Driver To make SPI flash drivers more stable, the legacy SPI flash driver is removed from v5.0. The legacy SPI flash driver refers to default `spi_flash` driver since v3.0, and the SPI flash driver with configuration option `CONFIG_SPI_FLASH_USE_LEGACY_IMPL` enabled since v4.0. The major breaking change here is that the legacy `spi_flash` driver is no longer supported from v5.0. Therefore, the legacy driver APIs and the `CONFIG_SPI_FLASH_USE_LEGACY_IMPL` configuration option are both removed. Please use the new `spi_flash` driver's APIs instead.

Removed items	Replacement
<code>spi_flash_erase_sector()</code>	<code>esp_flash_erase_region()</code>
<code>spi_flash_erase_range()</code>	<code>esp_flash_erase_region()</code>
<code>spi_flash_write()</code>	<code>esp_flash_write()</code>
<code>spi_flash_read()</code>	<code>esp_flash_read()</code>
<code>spi_flash_write_encrypted()</code>	<code>esp_flash_write_encrypted()</code>
<code>spi_flash_read_encrypted()</code>	<code>esp_flash_read_encrypted()</code>

Note: New functions with prefix `esp_flash` accept an additional `esp_flash_t*` parameter. You can simply set it to `NULL`. This will make the function to run the main flash (`esp_flash_default_chip`).

The `esp_spi_flash.h` header is deprecated as system functions are no longer public. To use flash memory mapping APIs, you may include `spi_flash_mmap.h` instead.

System

Inter-Processor Call IPC (Inter-Processor Call) feature is no longer a stand-alone component and has been integrated into the `esp_system` component.

Thus, any project presenting a `CMakeLists.txt` file with the parameters `PRIV_REQUIRES esp_ipc` or `REQUIRES esp_ipc` should be modified to simply remove these options as the `esp_system` component is included by default.

ESP Clock The ESP Clock API (functions/types/macros prefixed with `esp_clk`) has been made into a private API. Thus, the previous include paths `#include "ESP32-P4/clk.h"` and `#include "esp_clk.h"` have been removed. If users still require usage of the ESP Clock API (though this is not recommended), it can be included via `#include "esp_private/esp_clk.h"`.

Note: Private APIs are not stable and are no longer subject to the ESP-IDF versioning scheme's breaking change rules. Thus, it is not recommended for users to continue calling private APIs in their applications.

Cache Error Interrupt The Cache Error Interrupt API (functions/types/macros prefixed with `esp_cache_err`) has been made into a private API. Thus, the previous include path `#include "ESP32-P4/cache_err_int.h"` has been removed. If users still require usage of the Cache Error Interrupt API (though this is not recommended), it can be included via `#include "esp_private/cache_err_int.h"`.

bootloader_support

- The function `bootloader_common_get_reset_reason()` has been removed. Please use the function `esp_rom_get_reset_reason()` in the ROM component.
- The functions `esp_secure_boot_verify_sbv2_signature_block()` and `esp_secure_boot_verify_rsa_signature_block()` have been removed without replacement. We do not expect users to use these directly. If they are indeed still necessary, please open a feature request on [GitHub](#) explaining why these functions are necessary to you.

Brownout The Brownout API (functions/types/macros prefixed with `esp_brownout`) has been made into a private API. Thus, the previous include path `#include "brownout.h"` has been removed. If users still require usage of the Brownout API (though this is not recommended), it can be included via `#include "esp_private/brownout.h"`.

Trax The Trax API (functions/types/macros prefixed with `trax_`) has been made into a private API. Thus, the previous include path `#include "trax.h"` has been removed. If users still require usage of the Trax API (though this is not recommended), it can be included via `#include "esp_private/trax.h"`.

ROM The previously deprecated ROM-related header files located in `components/esp32/rom/` (old include path: `rom/*.h`) have been moved. Please use the new target-specific path from `components/esp_rom/include/ESP32-P4/` (new include path: `ESP32-P4/rom/*.h`).

esp_hw_support

- The header files `soc/cpu.h` have been deleted and deprecated CPU util functions have been removed. ESP-IDF developers should include `esp_cpu.h` instead for equivalent functions.
- The header files `hal/cpu_ll.h`, `hal/cpu_hal.h`, `hal/soc_ll.h`, `hal/soc_hal.h` and `interrupt_controller_hal.h` CPU API functions have been deprecated. ESP-IDF developers should include `esp_cpu.h` instead for equivalent functions.
- The header file `compare_set.h` have been deleted. ESP-IDF developers should use `esp_cpu_compare_and_set()` function provided in `esp_cpu.h` instead.

- `esp_cpu_get_ccount()`, `esp_cpu_set_ccount()` and `esp_cpu_in_ocd_debug_mode()` were removed from `esp_cpu.h`. ESP-IDF developers should use respectively `esp_cpu_get_cycle_count()`, `esp_cpu_set_cycle_count()` and `esp_cpu_dbg_is_attached()` instead.
- The header file `esp_intr.h` has been deleted. Please include `esp_intr_alloc.h` to allocate and manipulate interrupts.
- The Panic API (functions/types/macros prefixed with `esp_panic`) has been made into a private API. Thus, the previous include path `#include "esp_panic.h"` has been removed. If users still require usage of the Trax API (though this is not recommended), it can be included via `#include "esp_private/panic_reason.h"`. Besides, developers should include `esp_debug_helpers.h` instead to use any debug-related helper functions, e.g., `print_backtrace`.
- The header file `soc_log.h` is now renamed to `esp_hw_log.h` and has been made private. Users are encouraged to use logging APIs provided under `esp_log.h` instead.
- The header files `spinlock.h`, `clk_ctrl_os.h`, and `rtc_wdt.h` must now be included without the `soc` prefix. For example, `#include "spinlock.h"`.
- `esp_chip_info()` returns the chip version in the format `= 100 * major eFuse version + minor eFuse version`. Thus, the revision in the `esp_chip_info_t` structure is expanded to `uint16_t` to fit the new format.

PSRAM

- The target-specific header file `spiram.h` and the header file `esp_spiram.h` have been removed. A new component `esp_psram` is created instead. For PSRAM/SPIRAM-related functions, users now include `esp_psram.h` and set the `esp_psram` component as a component requirement in their `CMakeLists.txt` project files.
- `esp_spiram_get_chip_size` and `esp_spiram_get_size` have been deleted. You should use `esp_psram_get_size` instead.

eFuse

- The parameter type of function `esp_secure_boot_read_key_digests()` changed from `ets_secure_boot_key_digests_t*` to `esp_secure_boot_key_digests_t*`. The new type is the same as the old one, except that the `allow_key_revoke` flag has been removed. The latter was always set to `true` in current code, not providing additional information.
- Added eFuse wafer revisions: major and minor. The `esp_efuse_get_chip_ver()` API is not compatible with these changes, so it was removed. Instead, please use the following APIs: `efuse_hal_get_major_chip_version()`, `efuse_hal_get_minor_chip_version()` or `efuse_hal_chip_revision()`.

esp_common `EXT_RAM_ATTR` is deprecated. Use the new macro `EXT_RAM_BSS_ATTR` to put `.bss` on PSRAM.

esp_system

- The header files `esp_random.h`, `esp_mac.h`, and `esp_chip_info.h`, which were all previously indirectly included via the header file `esp_system.h`, must now be included directly. These indirect inclusions from `esp_system.h` have been removed.
- The Backtrace Parser API (functions/types/macros prefixed with `esp_eh_frame_`) has been made into a private API. Thus, the previous include path `#include "eh_frame_parser.h"` has been removed. If users still require usage of the Backtrace Parser API (though this is not recommended), it can be included via `#include "esp_private/eh_frame_parser.h"`.
- The Interrupt Watchdog API (functions/types/macros prefixed with `esp_int_wdt_`) has been made into a private API. Thus, the previous include path `#include "esp_int_wdt.h"` has been removed. If users still require usage of the Interrupt Watchdog API (though this is not recommended), it can be included via `#include "esp_private/esp_int_wdt.h"`.

SoC Dependency

- Public API headers listed in the Doxyfiles will not expose unstable and unnecessary SoC header files, such as `soc/soc.h` and `soc/rtc.h`. That means the user has to explicitly include them in their code if these "missing" header files are still wanted.
- Kconfig option `LEGACY_INCLUDE_COMMON_HEADERS` is also removed.
- The header file `soc/soc_memory_types.h` has been deprecated. Users should use the `esp_memory_utils.h` instead. Including `soc/soc_memory_types.h` will bring a build warning like `soc_memory_types.h` is deprecated, please migrate to `esp_memory_utils.h`.

APP Trace One of the timestamp sources has changed from the legacy timer group driver to the new *GPTimer*. Kconfig choices like `APPTRACE_SV_TS_SOURCE_TIMER00` has been changed to `APPTRACE_SV_TS_SOURCE_GPTIMER`. User no longer need to choose the group and timer ID.

esp_timer The FRC2-based legacy implementation of `esp_timer` available on ESP32 has been removed. The simpler and more efficient implementation based on the LAC timer is now the only option.

ESP Image The image SPI speed enum definitions have been renamed.

- Enum `ESP_IMAGE_SPI_SPEED_80M` has been renamed to `ESP_IMAGE_SPI_SPEED_DIV_1`.
- Enum `ESP_IMAGE_SPI_SPEED_40M` has been renamed to `ESP_IMAGE_SPI_SPEED_DIV_2`.
- Enum `ESP_IMAGE_SPI_SPEED_26M` has been renamed to `ESP_IMAGE_SPI_SPEED_DIV_3`.
- Enum `ESP_IMAGE_SPI_SPEED_20M` has been renamed to `ESP_IMAGE_SPI_SPEED_DIV_4`.

Task Watchdog Timers

- The API for `esp_task_wdt_init()` has changed as follows:
 - Configuration is now passed as a configuration structure.
 - The function will now handle subscribing of the idle tasks if configured to do so.
- The former `CONFIG_ESP_TASK_WDT` configuration option has been renamed to *CONFIG_ESP_TASK_WDT_INIT* and a new *CONFIG_ESP_TASK_WDT_EN* option has been introduced.

FreeRTOS

Legacy API and Data Types Previously, the `configENABLE_BACKWARD_COMPATIBILITY` option was set by default, thus allowing pre FreeRTOS v8.0.0 function names and data types to be used. The `configENABLE_BACKWARD_COMPATIBILITY` is now disabled by default, thus legacy FreeRTOS names/types are no longer supported by default. Users should do one of the followings:

- Update their code to remove usage of legacy FreeRTOS names/types.
- Enable the *CONFIG_FREERTOS_ENABLE_BACKWARD_COMPATIBILITY* to explicitly allow the usage of legacy names/types.

Tasks Snapshot The header `task_snapshot.h` has been removed from `freertos/task.h`. ESP-IDF developers should include `freertos/task_snapshot.h` if they need tasks snapshot API.

The function `vTaskGetSnapshot()` now returns `BaseType_t`. Return value shall be `pdTRUE` on success and `pdFALSE` otherwise.

FreeRTOS Asserts Previously, FreeRTOS asserts were configured separately from the rest of the system using the `FREERTOS_ASSERT` kconfig option. This option has now been removed and the configuration is now done through `COMPILER_OPTIMIZATION_ASSERTION_LEVEL`.

Port Macro API The file `portmacro_deprecated.h` which was added to maintain backward compatibility for deprecated APIs is removed. Users are advised to use the alternate functions for the deprecated APIs as listed below:

- `portENTER_CRITICAL_NESTED()` is removed. Users should use the `portSET_INTERRUPT_MASK_FROM_ISR()` macro instead.
- `portEXIT_CRITICAL_NESTED()` is removed. Users should use the `portCLEAR_INTERRUPT_MASK_FROM_ISR()` macro instead.
- `vPortCPUInitializeMutex()` is removed. Users should use the `spinlock_initialize()` function instead.
- `vPortCPUAcquireMutex()` is removed. Users should use the `spinlock_acquire()` function instead.
- `vPortCPUAcquireMutexTimeout()` is removed. Users should use the `spinlock_acquire()` function instead.
- `vPortCPUReleaseMutex()` is removed. Users should use the `spinlock_release()` function instead.

App Update

- The functions `esp_ota_get_app_description()` and `esp_ota_get_app_elf_sha256()` have been termed as deprecated. Please use the alternative functions `esp_app_get_description()` and `esp_app_get_elf_sha256()` respectively. These functions have now been moved to a new component `esp_app_format`. Please refer to the header file `esp_app_desc.h`.

Bootloader Support

- The `esp_app_desc_t` structure, which used to be declared in `esp_app_format.h`, is now declared in `esp_app_desc.h`.
- The function `bootloader_common_get_partition_description()` has now been made private. Please use the alternative function `esp_ota_get_partition_description()`. Note that this function takes `esp_partition_t` as its first argument instead of `esp_partition_pos_t`.

Chip Revision The bootloader checks the chip revision at the beginning of the application loading. The application can only be loaded if the version is $\geq$ `CONFIG_ESP32P4_REV_MIN` and $<$ `CONFIG_ESP32P4_REV_MAX_FULL`.

During the OTA upgrade, the version requirements and chip revision in the application header are checked for compatibility. The application can only be updated if the version is $\geq$ `CONFIG_ESP32P4_REV_MIN` and $<$ `CONFIG_ESP32P4_REV_MAX_FULL`.

Tools

ESP-IDF Monitor ESP-IDF Monitor makes the following changes regarding baud-rate:

- ESP-IDF monitor now uses the custom console baud-rate (`CONFIG_ESP_CONSOLE_UART_BAUDRATE`) by default instead of 115200.
- Setting a custom baud from menuconfig is no longer supported.
- A custom baud-rate can be specified from command line with the `idf.py monitor -b <baud>` command or through setting environment variables.
- Please note that the baud-rate argument has been renamed from `-B` to `-b` in order to be consistent with the global baud-rate `idf.py -b <baud>`. Run `idf.py monitor --help` for more information.

Deprecated Commands `idf.py` sub-commands and `cmake` target names have been unified to use hyphens (–) instead of underscores (_). Using a deprecated sub-command or target name will produce a warning. Users are advised to migrate to using the new sub-commands and target names. The following changes have been made:

Table 1: Deprecated Sub-command and Target Names

Old Name	New Name
<code>efuse_common_table</code>	<code>efuse-common-table</code>
<code>efuse_custom_table</code>	<code>efuse-custom-table</code>
<code>erase_flash</code>	<code>erase-flash</code>
<code>partition_table</code>	<code>partition-table</code>
<code>partition_table-flash</code>	<code>partition-table-flash</code>
<code>post_debug</code>	<code>post-debug</code>
<code>show_efuse_table</code>	<code>show-efuse-table</code>
<code>erase_otadata</code>	<code>erase-otadata</code>
<code>read_otadata</code>	<code>read-otadata</code>

Esptool The `CONFIG_ESPTOOLPY_FLASHSIZE_DETECT` option has been renamed to `CONFIG_ESPTOOLPY_HEADER_FLASHSIZE_UPDATE` and has been disabled by default. New and existing projects migrated to ESP-IDF v5.0 have to set `CONFIG_ESPTOOLPY_FLASHSIZE`. If this is not possible due to an unknown flash size at build time, then `CONFIG_ESPTOOLPY_HEADER_FLASHSIZE_UPDATE` can be enabled. However, once enabled, to keep the digest valid, an SHA256 digest is no longer appended to the image when updating the binary header with the flash size during flashing.

Windows Environment The Msys/Mingw-based Windows environment support got deprecated in ESP-IDF v4.0 and was entirely removed in v5.0. Please use [ESP-IDF Tools Installer](#) to set up a compatible environment. The options include Windows Command Line, Power Shell and the graphical user interface based on Eclipse IDE. In addition, a VS Code-based environment can be set up with the supported plugin: <https://github.com/espressif/vscode-esp-idf-extension>.

6.1.2 Migration from 5.0 to 5.1

GCC

GCC Version The previous GCC version was GCC 11.2.0. This has now been upgraded to GCC 12.2.0 on all targets. Users that need to port their code from GCC 11.2.0 to 12.2.0 should refer to the series of official GCC porting guides listed below:

- [Porting to GCC 12](#)

Warnings The upgrade to GCC 12.2.0 has resulted in the addition of new warnings, or enhancements to existing warnings. The full details of all GCC warnings can be found in [GCC Warning Options](#). Users are advised to double-check their code, then fix the warnings if possible. Unfortunately, depending on the warning and the complexity of the user's code, some warnings will be false positives that require non-trivial fixes. In such cases, users can choose to suppress the warning in multiple ways. This section outlines some common warnings that users are likely to encounter and ways to fix them.

-Wuse-after-free Typically, this warning should not produce false-positives for release-level code. But this may appear in test cases. There is an example of how it was fixed in ESP-IDF's `test_realloc.c`.

```
void *x = malloc(64);
void *y = realloc(x, 48);
TEST_ASSERT_EQUAL_PTR(x, y);
```

Pointers may be converted to `int` to avoid warning `-Wuse-after-free`.

```
int x = (int) malloc(64);
int y = (int) realloc((void *) x, 48);
TEST_ASSERT_EQUAL_UINT32((uint32_t) x, (uint32_t) y);
```

-Waddress GCC 12.2.0 introduces an enhanced version of the `-Waddress` warning option, which is now more eager in detecting the checking of pointers to an array in `if`-statements.

The following code triggers the warning:

```
char array[8];
...
if (array)
    memset(array, 0xff, sizeof(array));
```

Eliminating unnecessary checks resolves the warning.

```
char array[8];
...
memset(array, 0xff, sizeof(array));
```

RISC-V Builds Outside of ESP-IDF The RISC-V extensions `zicsr` and `zifencei` have been separated from the `I` extension. GCC 12 reflects this change, and as a result, when building for RISC-V ESP32 chips outside of the ESP-IDF framework, you must include the `_zicsr_zifencei` postfix when specifying the `-march` option in your build system.

Example:

```
riscv32-esp-elf-gcc main.c -march=rv32imac
```

Now is replaced with:

```
riscv32-esp-elf-gcc main.c -march=rv32imac_zicsr_zifencei
```

Peripherals

GPSPI Following items are deprecated. Since ESP-IDF v5.1, GPSPI clock source is configurable.

- `spi_get_actual_clock` is deprecated, you should use `spi_device_get_actual_freq()` instead.

LEDC

- `soc_periph_ledc_clk_src_legacy_t::LEDC_USE_RTC8M_CLK` is deprecated. Please use `LEDC_USE_RC_FAST_CLK` instead.

Storage

FatFs `esp_vfs_fat_sdmmc_unmount()` is now deprecated, and you can use `esp_vfs_fat_sdcard_unmount()` instead. This API is deprecated in previous ESP-IDF versions, but without a deprecation warning or migration guide. Since ESP-IDF v5.1, calling this `esp_vfs_fat_sdmmc_unmount()` API will generate a deprecation warning.

SPI_FLASH

- `spi_flash_get_counters()` is deprecated, please use `esp_flash_get_counters()` instead.
- `spi_flash_dump_counters()` is deprecated, please use `esp_flash_dump_counters()` instead.
- `spi_flash_reset_counters()` is deprecated, please use `esp_flash_reset_counters()` instead.

Networking

SNTP SNTP module now provides thread safe APIs to access lwIP functionality. It is recommended to use `ESP_NETIF` API. Please refer to the chapter [SNTP API](#) for more details.

System

FreeRTOS

Dynamic Memory Allocation

In the past, FreeRTOS commonly utilized the function `malloc()` to allocate dynamic memory. As a result, if an application allowed `malloc()` to allocate memory from external RAM (by configuring the [CONFIG_SPIRAM_USE](#) option as `CONFIG_SPIRAM_USE_MALLOC`), FreeRTOS had the potential to allocate dynamic memory from external RAM, and the specific location was determined by the heap allocator.

Note: Dynamic memory allocation for tasks (which are likely to consume the most memory) were an exception to the scenario above. FreeRTOS would use a separate memory allocation function to guarantee that dynamic memory allocated for a task was always placed in internal RAM.

Allowing FreeRTOS objects (such as queues and semaphores) to be placed in external RAM becomes an issue if those objects are accessed while the cache is disabled (such as during SPI flash write operations) and would lead to a cache access errors (see [Fatal Errors](#) for more details).

Therefore, FreeRTOS has been updated to always use internal memory (i.e., DRAM) for dynamic memory allocation. Calling FreeRTOS creation functions (e.g., `xTaskCreate()`, `xQueueCreate()`) guarantees that the memory allocated for those tasks/objects is from internal memory (see [FreeRTOS Heap](#) for more details).

Warning: If you previously relied on [CONFIG_SPIRAM_USE](#) to place FreeRTOS objects into external memory, this change will lead to increased usage of internal memory due the FreeRTOS objects now being allocated there.

To place a FreeRTOS task/object into external memory, it is now necessary to do so explicitly. The following methods can be employed:

- Allocate the task/object using one of the `...CreateWithCaps()` API such as `xTaskCreateWithCaps()` and `xQueueCreateWithCaps()` (see *IDF Additional API* for more details).
- Manually allocate external memory for those objects using `heap_caps_malloc()`, then create the objects from the allocated memory using one of the `...CreateStatic()` FreeRTOS functions.

Power Management

- `esp_pm_config_esp32xx_t` is deprecated, use `esp_pm_config_t` instead.
- `esp32xx/pm.h` is deprecated, use `esp_pm.h` instead.

6.1.3 Migration from 5.1 to 5.2

GCC

GCC Version The previous GCC version was GCC 12.2.0. This has now been upgraded to GCC 13.2.0 on all targets. Users that need to port their code from GCC 12.2.0 to 13.2.0 should refer to the series of official GCC porting guides listed below:

- [Porting to GCC 13](#)

Common Porting Problems and Fixes

`stdio.h` No Longer Includes `sys/types.h`

Issue Compilation errors may occur in code that previously worked with the old toolchain. For example:

```
#include <stdio.h>
clock_t var; // error: expected specifier-qualifier-list before 'clock_t'
```

Solution To resolve this issue, the correct header must be included. Refactor the code like this:

```
#include <time.h>
clock_t var;
```

Peripherals

UART

- `UART_FIFO_LEN` is deprecated. Please use `UART_HW_FIFO_LEN` instead.

I2C I2C driver has been redesigned (see [I2C API Reference](#)), which aims to unify the interface and extend the usage of I2C peripheral. Although it is recommended to use the new driver APIs, the legacy driver is still available in the previous include path `driver/i2c.h`.

The major breaking changes in concept and usage are listed as follows:

Major Changes in Concepts

- `i2c_config_t` which was used to configure the I2C bus, but it doesn't really tell whether to configure master or slave. So in the new design, master and slave initialization are separate, user can call `i2c_master_bus_config_t` or `i2c_slave_config_t`.
- `i2c_mode_t` which was used to tell whether I2C controller works in slave mode or master mode. This enumerator has been deprecated. In the new driver, users don't need to manually set the mode anymore since master and slave APIs are different.
- `i2c_rw_t` which was used to tell whether I2C master controller is performing a *write* or a *read* operation. This is now deprecated.
- `i2c_addr_mode_t` was renamed to `i2c_addr_bit_len_t`.
- In the legacy driver, operations needed to be chained with a command list (dynamically or statically created). The new driver now handles this internally, making the operations more size and space efficient.
- Capability flags like `I2C_SCLK_SRC_FLAG_FOR_NOMAL` are used to select clock source in the legacy driver. In the new driver, users can select clock source directly.

Major Changes in Usage

- I2C bus initialization is done in two parts: first, initialization of the bus with `i2c_new_master_bus()`, then, initialization of the I2C device with `i2c_master_bus_add_device()`.
- `i2c_reset_tx_fifo` and `i2c_reset_rx_fifo` have been removed, since it is never required to reset the fifo by users. Whole bus can still be reset by calling `i2c_master_bus_reset`.
- `i2c_cmd_link_xxx` functions have been removed, user doesn't need to use link to link commands on its own.
- `i2c_master_write_to_device` has been renamed to `i2c_master_transmit`.
- `i2c_master_read_from_device` has been renamed to `i2c_master_receive`.
- `i2c_master_write_read_device` has been renamed to `i2c_master_transmit_receive`.
- `i2c_slave_write_buffer` has been renamed to `i2c_slave_transmit`.
- `i2c_slave_read_buffer` has been renamed to `i2c_slave_receive`.

Protocols

CoAP CoAP examples have been moved to [idf-extra-components](#) repository.

HTTP2 `http2_request` example has been moved to [idf-extra-components](#) repository.

Storage

NVS Encryption

- For SoCs with the HMAC peripheral (`SOC_HMAC_SUPPORTED`), turning on [Flash Encryption](#) will no longer automatically turn on [NVS Encryption](#).
- You will need to explicitly turn on NVS encryption and select the required scheme (flash encryption-based or HMAC peripheral-based). You can select the HMAC peripheral-based scheme ([CONFIG_NVS_SEC_KEY_PROTECTION_SCHEME](#)), even if flash encryption is not enabled.

- SoCs without the HMAC peripheral will still automatically turn on NVS encryption when flash encryption is enabled.

SPI Flash Driver XMC-C series flash suspend support has been removed. According to feedback from the flash manufacturer, in some situations the XMC-C flash would require a 1ms interval between resume and next command. This is too long for a software request. Based on the above reason, in order to use suspend safely, we decide to remove flash suspend support from XMC-C series. But you can still force enable it via `CONFIG_SPI_FLASH_FORCE_ENABLE_XMC_C_SUSPEND`. If you have any questions, please contact espressif business support.

System

FreeRTOS

IDF FreeRTOS Upgrade The IDF FreeRTOS kernel (which is a dual-core SMP implementation of FreeRTOS) has been upgraded to be based on Vanilla FreeRTOS v10.5.1. With this upgrade, the design and implementation of IDF FreeRTOS has also been changed significantly. As a result, users should take note of the following changes to kernel behavior and API:

- When enabling single-core mode via the `CONFIG_FREERTOS_UNICORE` option, the kernel's behavior will now be identical to Vanilla FreeRTOS (see [Single-Core Mode](#) for more details).
- For SMP related APIs that were added by IDF FreeRTOS, checks on `xCoreID` arguments are now stricter. Providing out of range values for `xCoreID` arguments will now trigger an assert.
- The following SMP related APIs are now deprecated and replaced due to naming consistency reasons:
 - `xTaskGetAffinity()` is deprecated, call `xTaskGetCoreID()` instead.
 - `xTaskGetIdleTaskHandleForCPU()` is deprecated, call `xTaskGetIdleTaskHandleForCore()` instead.
 - `xTaskGetCurrentTaskHandleForCPU()` is deprecated, call `xTaskGetCurrentTaskHandleForCore()` instead.

Task Snapshot The Task Snapshot API has been made private due to a lack of a practical way for the API to be used from user code (the scheduler must be halted before the API can be called).

Panic Handler Behavior The choice `CONFIG_ESP_SYSTEM_PANIC_GDBSTUB` in the configuration option `CONFIG_ESP_SYSTEM_PANIC` has been made dependent on whether the `esp_gdbstub` component is included in the build. When trimming the list of components in the build using `set (COMPONENTS main), esp_gdbstub` component has to be added to this list of components to make the `CONFIG_ESP_SYSTEM_PANIC_GDBSTUB` option available.

Chapter 7

Libraries and Frameworks

7.1 Cloud Frameworks

ESP32-P4 supports multiple cloud frameworks using agents built on top of ESP-IDF. Here are the pointers to various supported cloud frameworks' agents and examples:

7.1.1 ESP RainMaker

ESP RainMaker is a complete solution for accelerated AIoT development. [ESP RainMaker on GitHub](#).

7.1.2 AWS IoT

<https://github.com/espressif/esp-aws-iot> is an open source repository for ESP32-P4 based on Amazon Web Services' [aws-iot-device-sdk-embedded-C](#).

7.1.3 Azure IoT

<https://github.com/espressif/esp-azure> is an open source repository for ESP32-P4 based on Microsoft Azure's [azure-iot-sdk-c](#) SDK.

7.1.4 Google IoT Core

<https://github.com/espressif/esp-google-iot> is an open source repository for ESP32-P4 based on Google's [iot-device-sdk-embedded-c](#) SDK.

7.1.5 Aliyun IoT

<https://github.com/espressif/esp-aliyun> is an open source repository for ESP32-P4 based on Aliyun's [iotkit-embedded](#) SDK.

7.1.6 Joylink IoT

<https://github.com/espressif/esp-joylink> is an open source repository for ESP32-P4 based on Joylink's [joylink_dev_sdk](#) SDK.

7.1.7 Tencent IoT

<https://github.com/espressif/esp-welink> is an open source repository for ESP32-P4 based on Tencent's [welink SDK](#).

7.1.8 Tencentyun IoT

<https://github.com/espressif/esp-qcloud> is an open source repository for ESP32-P4 based on Tencentyun's [qcloud-iot-sdk-embedded-c SDK](#).

7.1.9 Baidu IoT

<https://github.com/espressif/esp-baidu-iot> is an open source repository for ESP32-P4 based on Baidu's [iot-sdk-c SDK](#).

7.2 Espressif's Frameworks

Here you will find a collection of the official Espressif libraries and frameworks.

7.2.1 Espressif Audio Development Framework

The ESP-ADF is a comprehensive framework for audio applications including:

- CODEC's HAL
- Music players and recorders
- Audio processing
- Bluetooth speakers
- Internet radios
- Hands-free devices
- Speech decognition

This framework is available on GitHub: [ESP-ADF](#).

7.2.2 ESP-CSI

ESP-CSI is an experimental implementation that uses the Wi-Fi Channel State Information to detect the presence of a human body.

See the [ESP-CSI](#) project for more information.

7.2.3 Espressif DSP Library

The library provides algorithms optimized specifically for digital signal processing applications. This library supports:

- Matrix multiplication
- Dot product
- FFT (Fast Fourier Transform)
- IIR (Infinite Impulse Response)
- FIR (Finite Impulse Response)
- Vector math operations

This library is available on Github: [ESP-DSP library](#).

7.2.4 ESP-WIFI-MESH Development Framework

This framework is based on the ESP-WIFI-MESH protocol with the following features:

- Fast network configuration
- Stable upgrade
- Efficient debugging
- LAN control
- Various application demos

This framework is available on Github: [ESP-MDF](#).

7.2.5 ESP-WHO

The ESP-WHO is a face detection and recognition framework using the ESP32 and camera.

This framework is available on Github: [ESP-WHO](#).

7.2.6 ESP RainMaker

[ESP RainMaker](#) is a complete solution for accelerated AIoT development. Using ESP RainMaker, you can create AIoT devices from the firmware to the integration with voice-assistant, phone apps and cloud backend.

This project is available on Github: [ESP RainMaker on GitHub](#).

7.2.7 ESP-IoT-Solution

[ESP-IoT-Solution](#) contains commonly used device drivers and code frameworks when developing IoT systems. The device drivers and code frameworks within the ESP-IoT-Solution are organized as separate components, allowing them to be easily integrated into an ESP-IDF project.

ESP-IoT-Solution includes:

- Device drivers for sensors, display, audio, GUI, input, actuators, etc.
- Framework and documentation for low power, security, storage, etc.
- Guide for Espressif open source solutions from practical application point.

This solution is available on Github: [ESP-IoT-Solution on GitHub](#).

7.2.8 ESP-Protocols

The [ESP-Protocols](#) repository contains a collection of protocol components for ESP-IDF. The code within ESP-Protocols is organized into separate components, allowing them to be easily integrated into an ESP-IDF project. Additionally, each component is available in [IDF Component Registry](#).

ESP-Protocols components:

- [esp_modem](#) enables connectivity with GSM/LTE modems using AT commands or PPP protocol. See the [esp_modem documentation](#).
- [mdns](#) (mDNS) is a multicast UDP service that is used to provide local network service and host discovery. See the [mdns documentation](#).
- [esp_websocket_client](#) is a managed component for ESP-IDF that contains implementation of WebSocket protocol client for ESP32. See the [esp_websocket_client documentation](#). For details of WebSocket protocol client, see [WebSocket_protocol_client](#).
- [asio](#) is a cross-platform C++ library, see <https://think-async.com/Asio/>. It provides a consistent asynchronous model using a modern C++ approach. See the [asio documentation](#).

7.2.9 ESP-BSP

The [ESP-BSP](#) repository contains Board Support Packages (BSPs) for various Espressif's and third-party development boards. BSPs help to quickly get started with a supported board. Usually they contain pinout definition and helper functions that will initialize peripherals for the specific board. Additionally, the BSPs contain drivers for external chips populated on the development board, such as sensors, displays, audio codecs, etc.

7.2.10 ESP-IDF-CXX

[ESP-IDF-CXX](#) contains C++ wrappers for part of ESP-IDF. The focuses are on ease of use, safety, automatic resource management. They also move error checking from runtime to compile time to prevent running failure. There are C++ classes for ESP-Timer, I2C, SPI, GPIO and other peripherals or features of ESP-IDF. ESP-IDF-CXX is [available as a component](#) from the component registry. Please check the project's [README.md](#) for more information.

Chapter 8

Contributions Guide

We welcome contributions to the ESP-IDF project!

8.1 How to Contribute

Contributions to ESP-IDF - fixing bugs, adding features, adding documentation - are welcome. We accept contributions via [Github Pull Requests](#).

8.2 Before Contributing

Before sending us a Pull Request, please consider this list of points:

- Is the contribution entirely your own work, or already licensed under an Apache License 2.0 compatible Open Source License? If not then we unfortunately cannot accept it. Please check the [Copyright Header Guide](#) for additional information.
- Does any new code conform to the ESP-IDF [Style Guide](#)?
- Have you installed the [pre-commit hook](#) for ESP-IDF project?
- Does the code documentation follow requirements in [Documenting Code](#)?
- Is the code adequately commented for people to understand how it is structured?
- Is there documentation or examples that go with code contributions? There are additional suggestions for writing good examples in [examples](#) readme.
- Are comments and documentation written in clear English, with no spelling or grammar errors?
- Example contributions are also welcome. Please check the [Creating Examples](#) guide for these.
- If the contribution contains multiple commits, are they grouped together into logical changes (one major change per pull request)? Are any commits with names like "fixed typo" [squashed into previous commits](#)?
- If you are unsure about any of these points, please open the Pull Request anyhow and then ask us for feedback.

8.3 Pull Request Process

After you open the Pull Request, there will probably be some discussion in the comments field of the request itself.

Once the Pull Request is ready to merge, it will first be merged into our internal git system for in-house automated testing.

If this process passes, it will be merged into the public GitHub repository.

8.4 Legal Part

Before a contribution can be accepted, you will need to sign our [Contributor Agreement](#). You will be prompted for this automatically as part of the Pull Request process.

8.5 Related Documents

8.5.1 Espressif IoT Development Framework Style Guide

About This Guide

Purpose of this style guide is to encourage use of common coding practices within the ESP-IDF.

Style guide is a set of rules which are aimed to help create readable, maintainable, and robust code. By writing code which looks the same way across the code base we help others read and comprehend the code. By using same conventions for spaces and newlines we reduce chances that future changes will produce huge unreadable diffs. By following common patterns for module structure and by using language features consistently we help others understand code behavior.

We try to keep rules simple enough, which means that they can not cover all potential cases. In some cases one has to bend these simple rules to achieve readability, maintainability, or robustness.

When doing modifications to third-party code used in ESP-IDF, follow the way that particular project is written. That will help propose useful changes for merging into upstream project.

C Code Formatting

Naming

- Any variable or function which is only used in a single source file should be declared `static`.
- Public names (non-static variables and functions) should be namespaced with a per-component or per-unit prefix, to avoid naming collisions. ie `esp_vfs_register()` or `esp_console_run()`. Starting the prefix with `esp_` for Espressif-specific names is optional, but should be consistent with any other names in the same component.
- Static variables should be prefixed with `s_` for easy identification. For example, `static bool s_invert`.
- Avoid unnecessary abbreviations (ie shortening `data` to `dat`), unless the resulting name would otherwise be very long.

Indentation Use 4 spaces for each indentation level. Do not use tabs for indentation. Configure the editor to emit 4 spaces each time you press tab key.

Vertical Space Place one empty line between functions. Do not begin or end a function with an empty line.

```
void function1()
{
    do_one_thing();
    do_another_thing();
}
// INCORRECT, do not place empty line here
// place empty line here
void function2()
{
    // INCORRECT, do not use an empty line here
    int var = 0;
    while (var < SOME_CONSTANT) {
        do_stuff(&var);
    }
}
```

(continues on next page)

(continued from previous page)

```

    }
}

```

The maximum line length is 120 characters as long as it does not seriously affect the readability.

Horizontal Space Always add single space after conditional and loop keywords:

```

if (condition) {      // correct
    // ...
}

switch (n) {          // correct
    case 0:
        // ...
}

for(int i = 0; i < CONST; ++i) {    // INCORRECT
    // ...
}

```

Add single space around binary operators. No space is necessary for unary operators. It is okay to drop space around multiply and divide operators:

```

const int y = y0 + (x - x0) * (y1 - y0) / (x1 - x0);    // correct

const int y = y0 + (x - x0)*(y1 - y0)/(x1 - x0);        // also okay

int y_cur = -y;                                          // correct
++y_cur;

const int y = y0+(x-x0)*(y1-y0)/(x1-x0);                // INCORRECT

```

No space is necessary around `.` and `->` operators.

Sometimes adding horizontal space within a line can help make code more readable. For example, you can add space to align function arguments:

```

esp_rom_gpio_connect_in_signal(PIN_CAM_D6,    I2S0I_DATA_IN14_IDX, false);
esp_rom_gpio_connect_in_signal(PIN_CAM_D7,    I2S0I_DATA_IN15_IDX, false);
esp_rom_gpio_connect_in_signal(PIN_CAM_HREF,   I2S0I_H_ENABLE_IDX,  false);
esp_rom_gpio_connect_in_signal(PIN_CAM_PCLK,   I2S0I_DATA_IN15_IDX, false);

```

Note however that if someone goes to add new line with a longer identifier as first argument (e.g., `PIN_CAM_VSYNC`), it will not fit. So other lines would have to be realigned, adding meaningless changes to the commit.

Therefore, use horizontal alignment sparingly, especially if you expect new lines to be added to the list later.

Never use TAB characters for horizontal alignment.

Never add trailing whitespace at the end of the line.

Braces

- Function definition should have a brace on a separate line:

```

// This is correct:
void function(int arg)
{
    // ...
}

```

(continues on next page)

(continued from previous page)

```
// NOT like this:
void function(int arg) {
}
```

- Within a function, place opening brace on the same line with conditional and loop statements:

```
if (condition) {
    do_one();
} else if (other_condition) {
    do_two();
}
```

Comments Use `//` for single line comments. For multi-line comments it is okay to use either `//` on each line or a `/* */` block.

Although not directly related to formatting, here are a few notes about using comments effectively.

- Do not use single comments to disable some functionality:

```
void init_something()
{
    setup_dma();
    // load_resources();           // WHY is this thing commented, asks_
    ↪the reader?
    start_timer();
}
```

- If some code is no longer required, remove it completely. If you need it you can always look it up in git history of this file. If you disable some call because of temporary reasons, with an intention to restore it in the future, add explanation on the adjacent line:

```
void init_something()
{
    setup_dma();
    // TODO: we should load resources here, but loader is not fully integrated_
    ↪yet.
    // load_resources();
    start_timer();
}
```

- Same goes for `#if 0 ... #endif` blocks. Remove code block completely if it is not used. Otherwise, add comment explaining why the block is disabled. Do not use `#if 0 ... #endif` or comments to store code snippets which you may need in the future.
- Do not add trivial comments about authorship and change date. You can always look up who modified any given line using git. E.g., this comment adds clutter to the code without adding any useful information:

```
void init_something()
{
    setup_dma();
    // XXX add 2016-09-01
    init_dma_list();
    fill_dma_item(0);
    // end XXX add
    start_timer();
}
```

Line Endings Commits should only contain files with LF (Unix style) endings.

Windows users can configure git to check out CRLF (Windows style) endings locally and commit LF endings by setting the `core.autocrlf` setting. *Github has a document about setting this option* <github-line-endings>.

If you accidentally have some commits in your branch that add LF endings, you can convert them to Unix by running this command in an MSYS2 or Unix terminal (change directory to the IDF working directory and check the correct branch is currently checked out, beforehand):

```
git rebase --exec 'git diff-tree --no-commit-id --name-only -r HEAD | xargs_
↪dos2unix && git commit -a --amend --no-edit --allow-empty' master
```

(Note that this line rebases on master, change the branch name at the end to rebase on another branch.)

For updating a single commit, it is possible to run `dos2unix FILENAME` and then run `git commit --amend`

Formatting Your Code ESP-IDF uses Astyle to format source code. The configuration is stored in [tools/ci/astyle-rules.yml](#) file.

Initially, all components are excluded from formatting checks. You can enable formatting checks for the component by removing it from `components_not_formatted_temporary` list. Then run:

```
pre-commit run --files <path_to_files> astyle_py
```

Alternatively, you can run `astyle_py` manually. You can install it with `pip install astyle_py==VERSION`. Make sure you have the same version installed as the one specified in [.pre-commit-config.yml](#) file. With `astyle_py` installed, run:

```
astyle_py --rules=$IDF_PATH/tools/ci/astyle-rules.yml <path-to-file>
```

Type Definitions Should be `snake_case`, ending with `_t` suffix:

```
typedef int signed_32_bit_t;
```

Enum Enums should be defined through the *typedef* and be namespaced:

```
typedef enum
{
    MODULE_FOO_ONE,
    MODULE_FOO_TWO,
    MODULE_FOO_THREE
} module_foo_t;
```

Assertions The standard C `assert()` function, defined in `assert.h` should be used to check conditions that should be true in source code. In the default configuration, an assert condition that returns `false` or `0` will call `abort()` and trigger a *Fatal Error*.

`assert()` should only be used to detect unrecoverable errors due to a serious internal logic bug or corruption, where it is not possible for the program to continue. For recoverable errors, including errors that are possible due to invalid external input, an *error value should be returned*.

Note: When asserting a value of type `esp_err_t` is equal to `ESP_OK`, use the [ESP_ERROR_CHECK Macro](#) instead of an `assert()`.

It is possible to configure ESP-IDF projects with assertions disabled (see [CONFIG_COMPILER_OPTIMIZATION_ASSERTION_LEVEL](#)). Therefore, functions called in an `assert()` statement should not have side-effects.

It is also necessary to use particular techniques to avoid "variable set but not used" warnings when assertions are disabled, due to code patterns such as:

```
int res = do_something();
assert(res == 0);
```

Once the `assert` is optimized out, the `res` value is unused and the compiler will warn about this. However the function `do_something()` must still be called, even if assertions are disabled.

When the variable is declared and initialized in a single statement, a good strategy is to cast it to `void` on a new line. The compiler will not produce a warning, and the variable can still be optimized out of the final binary:

```
int res = do_something();
assert(res == 0);
(void)res;
```

If the variable is declared separately, for example if it is used for multiple assertions, then it can be declared with the GCC attribute `__attribute__((unused))`. The compiler will not produce any unused variable warnings, but the variable can still be optimized out:

```
int res __attribute__((unused));

res = do_something();
assert(res == 0);

res = do_something_else();
assert(res != 0);
```

Header File Guards

All public facing header files should have preprocessor guards. A `pragma` is preferred:

```
#pragma once
```

over the following pattern:

```
#ifndef FILE_NAME_H
#define FILE_NAME_H
...
#endif // FILE_NAME_H
```

In addition to guard macros, all C header files should have `extern "C"` guards to allow the header to be used from C++ code. Note that the following order should be used: `pragma once`, then any `#include` statements, then `extern "C"` guards:

```
#pragma once

#include <stdint.h>

#ifdef __cplusplus
extern "C" {
#endif

/* declarations go here */

#ifdef __cplusplus
}
#endif
```

Include Statements

When writing `#include` statements, try to maintain the following order:

- C standard library headers.
- Other POSIX standard headers and common extensions to them (such as `sys/queue.h`.)
- Common IDF headers (`esp_log.h`, `esp_system.h`, `esp_timer.h`, `esp_sleep.h`, etc.)
- Headers of other components, such as FreeRTOS.
- Public headers of the current component.
- Private headers.

Use angle brackets for C standard library headers and other POSIX headers (`#include <stdio.h>`).

Use double quotes for all other headers (`#include "esp_log.h"`).

C++ Code Formatting

The same rules as for C apply. Where they are not enough, apply the following rules.

File Naming C++ Header files have the extension `.hpp`. C++ source files have the extension `.cpp`. The latter is important for the compiler to distinguish them from normal C source files.

Naming

- **Class and struct** names shall be written in `CamelCase` with a capital letter as beginning. Member variables and methods shall be in `snake_case`. An exception from `CamelCase` is if the readability is severely decreased, e.g., in `GPIOOutput`, then an underscore `_` is allowed to make it more readable: `GPIO_Output`.
- **Namespaces** shall be in lower `snake_case`.
- **Templates** are specified in the line above the function declaration.
- Interfaces in terms of Object-Oriented Programming shall be named without the suffix `...Interface`. Later, this makes it easier to extract interfaces from normal classes and vice versa without making a breaking change.

Member Order in Classes In order of precedence:

- First put the public members, then the protected, then private ones. Omit public, protected or private sections without any members.
- First put constructors/destructors, then member functions, then member variables.

For example:

```
class ForExample {
public:
    // first constructors, then default constructor, then destructor
    ForExample(double example_factor_arg);
    ForExample();
    ~ForExample();

    // then remaining public methods
    set_example_factor(double example_factor_arg);

    // then public member variables
    uint32_t public_data_member;

private:
    // first private methods
    void internal_method();

    // then private member variables
    double example_factor;
};
```

Spacing

- Do not indent inside namespaces.
- Put public, protected and private labels at the same indentation level as the corresponding class label.

Simple Example

```
// file spaceship.h
#ifndef SPACESHIP_H_
#define SPACESHIP_H_
#include <cstdlib>

namespace spaceships {

class SpaceShip {
public:
    SpaceShip(size_t crew);
    size_t get_crew_size() const;

private:
    const size_t crew;
};

class SpaceShuttle : public SpaceShip {
public:
    SpaceShuttle();
};

class Sojuz : public SpaceShip {
public:
    Sojuz();
};

template <typename T>
class CargoShip {
public:
    CargoShip(const T &cargo);

private:
    T cargo;
};

} // namespace spaceships

#endif // SPACESHIP_H_

// file spaceship.cpp
#include "spaceship.h"

namespace spaceships {

// Putting the curly braces in the same line for constructors is OK if it only
// initializes
// values in the initializer list
SpaceShip::SpaceShip(size_t crew) : crew(crew) { }

size_t SpaceShip::get_crew_size() const
{
    return crew;
}
```

(continues on next page)

(continued from previous page)

```

SpaceShuttle::SpaceShuttle() : SpaceShip(7)
{
    // doing further initialization
}

Sojuz::Sojuz() : SpaceShip(3)
{
    // doing further initialization
}

template <typename T>
CargoShip<T>::CargoShip(const T &cargo) : cargo(cargo) { }

} // namespace spaceships

```

CMake Code Style

- Indent with four spaces.
- Maximum line length 120 characters. When splitting lines, try to focus on readability where possible (for example, by pairing up keyword/argument pairs on individual lines).
- Do not put anything in the optional parentheses after `endforeach()`, `endif()`, etc.
- Use lowercase (`with_underscores`) for command, function, and macro names.
- For locally scoped variables, use lowercase (`with_underscores`).
- For globally scoped variables, use uppercase (`WITH_UNDERSCORES`).
- Otherwise follow the defaults of the [cmake-lint](#) project.

Configuring the Code Style for a Project Using EditorConfig

EditorConfig helps developers define and maintain consistent coding styles between different editors and IDEs. The EditorConfig project consists of a file format for defining coding styles and a collection of text editor plugins that enable editors to read the file format and adhere to defined styles. EditorConfig files are easily readable and they work nicely with version control systems.

For more information, see [EditorConfig Website](#).

Third Party Component Code Styles

ESP-IDF integrates a number of third party components where these components may have differing code styles.

FreeRTOS The code style adopted by FreeRTOS is described in the [FreeRTOS style guide](#). Formatting of FreeRTOS source code is automated using [Uncrustify](#), thus a copy of the FreeRTOS code style's Uncrustify configuration (`uncrustify.cfg`) is stored within ESP-IDF FreeRTOS component.

If a FreeRTOS source file is modified, the updated file can be formatted again by following the steps below:

1. Ensure that Uncrustify (v0.69.0) is installed on your system
2. Run the following command on the update FreeRTOS source file (where `source.c` is the path to the source file that requires formatting).

```

uncrustify -c $IDF_PATH/components/freertos/FreeRTOS-Kernel/uncrustify.cfg --
↪replace source.c --no-backup

```

Documenting Code

Please see the guide here: [Documenting Code](#).

Structure

To be written.

Language Features

To be written.

8.5.2 Install Pre-commit Hook for ESP-IDF Project

Required Dependency

Python 3.8.* or above. This is our recommended Python version for ESP-IDF developers.

If you still have Python versions not compatible, update your Python versions before installing the pre-commit hook.

Install pre-commit

Run `pip install pre-commit`

Install pre-commit Hook

1. Go to the ESP-IDF project directory.
2. Run `pre-commit install --allow-missing-config -t pre-commit -t commit-msg`. Install hook by this approach will let you commit successfully even in branches without the `.pre-commit-config.yaml`
3. pre-commit hook will run automatically when you are running `git commit` command

Uninstall pre-commit Hook

Run `pre-commit uninstall`

What Is More?

For detailed usage, please refer to the documentation of [pre-commit](#).

Common Problems For Windows Users

`/usr/bin/env: python: Permission denied.`

If you are in Git Bash, please check the python executable location by run `which python`.

If the executable is under `~/AppData/Local/Microsoft/WindowsApps/`, then it is a link to Windows AppStore, not a real one.

Please install Python manually and update this in your `PATH` environment variable.

Your `USERPROFILE` contains non-ASCII characters

`pre-commit` may fail when initializing an environment for a particular hook when the path of `pre-commit`'s cache contains non-ASCII characters. The solution is to set `PRE_COMMIT_HOME` to a path containing only standard characters before running `pre-commit`.

- **CMD:** `set PRE_COMMIT_HOME=C:\somepath\pre-commit`
- **PowerShell:** `$Env:PRE_COMMIT_HOME = "C:\somepath\pre-commit"`
- **git bash:** `export PRE_COMMIT_HOME="/c/somepath/pre-commit"`

8.5.3 Documenting Code

The purpose of this description is to provide a quick summary of the documentation style used in [espressif/esp-idf](#) repository and how to add new documentation.


Introduction

When documenting code for this repository, please follow [Doxygen style](#). You are doing it by inserting special commands, for instance `@param`, into standard comments blocks, for example:

```
/**
 * @param ratio this is oxygen to air ratio
 */
```

Doxygen can phrase the code, extract the commands together with subsequent text, and build documentation out of it.

Typical comment block, that contains documentation of a function, looks like below:



Doxygen supports a couple of formatting styles. It also gives you great flexibility on the level of details to include in documentation. To get familiar with available features, please check data-rich and very well-organized [Doxygen Manual](#).

Why We Need Doxygen?

The ultimate goal is to ensure that all the code is consistently documented, so we can use tools like [Sphinx](#) and [Breathe](#) to aid preparation and automatic updates of API documentation when the code changes.

With these tools, the above piece of code renders like below:

```

348
349 /**
350  * @brief disable GPIO wake-up function.
351  * @param gpio_num: GPIO number
352  * @return ESP_OK: success
353  *         ESP_ERR_INVALID_ARG: parameter error
354  */
355 esp_err_t gpio_wakeup_disable(gpio_num_t gpio_num);
356

```

`esp_err_t gpio_wakeup_disable(gpio_num_t gpio_num)`

disable GPIO wake-up function.

Return

ESP_OK: success ESP_ERR_INVALID_ARG: parameter error

Parameters

- `gpio_num` - GPIO number

Go for It!

When writing code for this repository, please follow guidelines below:

1. Document all building blocks of code, including functions, structs, typedefs, enums, macros, etc. Provide enough information about purpose, functionality, and limitations of documented items, as you would like to see them documented when reading the code by others.
2. Documentation of function should describe what this function does. If it accepts input parameters and returns some value, all of them should be explained.
3. Do not add a data type before parameter or any other characters besides spaces. All spaces and line breaks are compressed into a single space. If you like to break a line, then break it twice.

do not add data type

white spaces are compressed

a line break that will render

this line break will not render

```

41 /**
42  * @brief Set log level for given tag
43  *
44  * If logging for given component has already been enabled, changes previous setting.
45  *
46  * @param tag Tag of the log entries to enable. Must be a non-NULL zero terminated string.
47  *           Value "" resets log level for all tags to the given value.
48  *
49  * @param level Selects log level to enable.
50  *             Only logs at this and lower levels will be shown.
51  */
52 void esp_log_level_set(const char* tag, esp_log_level_t level);

```

`void esp_log_level_set(const char* tag, esp_log_level_t level)`

Set log level for given tag.

If logging for given component has already been enabled, changes previous setting.

Parameters

- `tag` - Tag of the log entries to enable. Must be a non-NULL zero terminated string. Value "" resets log level for all tags to the given value.
- `level` - Selects log level to enable. Only logs at this and lower levels will be shown.

4. If function has void input or does not return any value, then skip @param or @return.

```

26@ /**
27  * @brief Initialize BT controller
28  *
29  * This function should be called only once,
30  * before any other BT functions are called.
31  */
32 void bt_controller_init(void);

```

```
void bt_controller_init(void)
```

Initialize BT controller.

This function should be called only once, before any other BT functions are called.

- When documenting a define as well as members of a struct or enum, place specific comment like below after each member.

```

45@ /**
46  * Mode of opening the non-volatile storage
47  *
48  */
49@ typedef enum {
50     NVS_READONLY, /*!< Read only */
51     NVS_READWRITE /*!< Read and write */
52 } nvs_open_mode;

```

/*!< how to documented members */

```
enum nvs_open_mode
```

Mode of opening the non-volatile storage.

Values:

```
NVS_READONLY
```

Read only

```
NVS_READWRITE
```

Read and write

- To provide well-formatted lists, break the line after command (like @return in the example below).

```

*
* @return
*
* - ESP_OK if erase operation was successful
*
* - ESP_ERR_NVS_INVALID_HANDLE if handle has been closed or is NULL
*
* - ESP_ERR_NVS_READ_ONLY if handle was opened as read only
*
* - ESP_ERR_NVS_NOT_FOUND if the requested key does not exist
*
* - other error codes from the underlying storage driver
*

```

- Overview of functionality of documented header file, or group of files that make a library, should be placed in a separate README.rst file of the same directory. If this directory contains header files for different APIs, then the file name should be apiname-readme.rst.

Go One Extra Mile

Here are a couple of tips on how you can make your documentation even better and more useful to the reader and writer.

When writing code, please follow the guidelines below:

- Add code snippets to illustrate implementation. To do so, enclose snippet using @code{c} and @endcode commands.

```

*
* @code{c}
* // Example of using nvs_get_i32:
* int32_t max_buffer_size = 4096; // default value
* esp_err_t err = nvs_get_i32(my_handle, "max_buffer_size", &max_buffer_size);
* assert(err == ESP_OK || err == ESP_ERR_NVS_NOT_FOUND);
* // if ESP_ERR_NVS_NOT_FOUND was returned, max_buffer_size will still
* // have its default value.

```

(continues on next page)

(continued from previous page)

```
* @endcode
*
```

The code snippet should be enclosed in a comment block of the function that it illustrates.

2. To highlight some important information use command `@attention` or `@note`.

```
*
* @attention
* 1. This API only impact WIFI_MODE_STA or WIFI_MODE_APSTA mode
* 2. If the ESP32 is connected to an AP, call esp_wifi_disconnect to
  ↳ disconnect.
*
```

Above example also shows how to use a numbered list.

3. To provide common description to a group of similar functions, enclose them using `/**@{ */` and `/**@} */` markup commands.

```
/**@{ */
/**
 * @brief common description of similar functions
 *
 */
void first_similar_function (void);
void second_similar_function (void);
/**@} */
```

For practical example see [nvs_flash/include/nvs.h](#).

4. You may want to go even further and skip some code like repetitive defines or enumerations. In such case, enclose the code within `/** @cond */` and `/** @endcond */` commands. Example of such implementation is provided in [driver/gpio/include/driver/gpio.h](#).
5. Use markdown to make your documentation even more readable. You will add headers, links, tables and more.

```
*
* [ESP32-P4 Technical Reference Manual] (https://www.espressif.com/sites/default/files/documentation/esp32-p4\_technical\_reference\_manual\_en.pdf)
*
```


Note: Code snippets, notes, links, etc., will not make it to the documentation, if not enclosed in a comment block associated with one of the documented objects.

6. Prepare one or more complete code examples together with description. Place description to a separate file `README.md` in specific folder of [examples](#) directory.

Standardize Document Format

When it comes to text, please follow guidelines below to provide well-formatted Markdown (.md) or reST (.rst) documents.

1. Please ensure that one paragraph is written in one line. Do not break lines like below. Breaking lines to enhance readability is only suitable for writing code. To make the text easier to read, it is recommended to place an empty line to separate the paragraph.
2. Please make the line number of CN and EN documents consistent like below. The benefit of this approach is that it can save time for both writers and translators. When non-bilingual writers need to update text, they only need to update the same line in the corresponding CN or EN document. For translators, if documents are updated in English, then translators can quickly locate where to update in the corresponding CN document later. Besides, by comparing the total number of lines in EN and CN documents, you can quickly find out whether the CN version lags behind the EN version.




```

11 SPI Bus Lock
12 ~~~~~
13
14 To realize the multiplexing of different devices from different drivers (SPI Master, SPI Flash, etc.), an SPI bus lock is applied on
15 each SPI bus. Drivers can attach their devices onto the bus with the arbitration of the lock.
16 Each bus lock is initialized with a BG (background) service registered. All devices request to do transactions on the bus should wait
17 until the BG to be successfully disabled.
18 - For SPI1 bus, the BG is the cache, the bus lock will help to disable the cache before device operations starts, and enable it again
  after device releasing the lock. No devices on SPI1 is allowed using ISR (it's meaningless for the task to yield to other tasks when
  the cache is disabled).

```

Fig. 1: One line for one paragraph (click to enlarge)



```

11 SPI Bus Lock
12 ~~~~~
13
14 To realize the multiplexing of different devices from different drivers (SPI Master, SPI Flash, etc.), an SPI bus lock is applied on
15 each SPI bus. Drivers can attach their devices onto the bus with the arbitration of the lock.
16 Each bus lock is initialized with a BG (background) service registered. All devices request to do transactions on the bus should wait
17 until the BG to be successfully disabled.
18 - For SPI1 bus, the BG is the cache, the bus lock will help to disable the cache before device operations starts, and enable it again
19 after device releasing the lock. No devices on SPI1 is allowed using ISR (it's meaningless for the task to yield to other tasks when
20 the cache is disabled).
21
22
23
24

```

Fig. 2: No line breaks within the same paragraph (click to enlarge)



```

1 *****
2 Getting Started with VS Code IDE
3 *****
4 :link_to_translation:'zh_CN:[中文]'
5
6 We have official support for VS Code and we aim to provide complete
7 end to end support for all actions related to ESP-IDF namely build,
8 flash, monitor, debug, tracing, core-dump, System Trace Viewer, etc.
9
10 Quick Install Guide
11 =====
12 Recommended way to install ESP-IDF Visual Studio Code Extension is by
13 downloading it from 'VS Code Marketplace <https://marketplace.
14 visualstudio.com/items?itemName=espressif.esp-idf-extension>' or
15 following 'Quick Installation Guide <https://github.com/espressif/
16 vscode-esp-idf-extension/blob/master/docs/tutorial/install.md>'.
17
18 Review the 'tutorials <https://github.com/espressif/
19 vscode-esp-idf-extension/blob/master/docs/tutorial/toc.md>' for
20 ESP-IDF Visual Studio Code Extension to learn how to use all features.
21
22 Supported Features
23 =====
24

```

Fig. 3: Keep the line number for EN and CN documents consistent (click to enlarge)

Building Documentation

The documentation is built with the *esp-docs* Python package, which is a wrapper around [Sphinx](#).

To install it simply do:

```
pip install esp-docs
```

After a successful install then the documentation can be built from the docs folder with:

```
build-docs build
```

or for specific target and language with:

```
build-docs -t esp32 -l en build
```

For more in-depth documentation about *esp-docs* features please see the documentation at [esp-docs](#).

Wrap Up

We love good code that is doing cool things. We love it even better, if it is well-documented, so we can quickly make it run and also do the cool things.

Go ahead, contribute your code and documentation!

Related Documents

- [API Documentation Template](#)

8.5.4 Creating Examples

Each ESP-IDF example is a complete project that someone else can copy and adapt the code to solve their own problem. Examples should demonstrate ESP-IDF functionality, while keeping this purpose in mind.

Structure

- The `main` directory should contain a source file named `(something)_example_main.c` with the main functionality.
- If the example has additional functionality, split it logically into separate C or C++ source files under `main` and place a corresponding header file in the same directory.
- If the example has a lot of additional functionality, consider adding a `components` directory to the example project and make some example-specific components with library functionality. Only do this if the components are specific to the example, if they are generic or common functionality then they should be added to ESP-IDF itself.
- The example should have a `README.md` file. Use the [template example README](#) and adapt it for your particular example.
- Examples should have a `pytest_<example name>.py` file for running an automated example test. If submitting a GitHub Pull Request which includes an example, it is OK not to include this file initially. The details can be discussed as part of the [Pull Request](#). Please refer to [IDF Tests with Pytest Guide](#) for details.

General Guidelines

Example code should follow the [Espressif IoT Development Framework Style Guide](#).

Checklist

Checklist before submitting a new example:

- Example project name (in `README.md`) uses the word "example". Use "example" instead of "demo", "test" or similar words.
- Example does one distinct thing. If the example does more than one thing at a time, split it into two or more examples.
- Example has a `README.md` file which is similar to the [template example README](#).
- Functions and variables in the example are named according to [naming section of the style guide](#). For non-static names which are only specific to the example's source files, you can use `example` or something similar as a prefix.
- All code in the example is well structured and commented.
- Any unnecessary code (old debugging logs, commented-out code, etc.) is removed from the example.
- Options in the example (like network names, addresses, etc) are not hard-coded. Use configuration items if possible, or otherwise declare macros or constants.
- Configuration items are provided in a `KConfig.projbuild` file with a menu named "Example Configuration". See existing example projects to see how this is done.
- All original example code has a license header saying it is "in the public domain / CC0", and a warranty disclaimer clause. Alternatively, the example is licensed under Apache License 2.0. See existing examples for headers to adapt from.
- Any adapted or third party example code has the original license header on it. This code must be licensed compatible with Apache License 2.0.

8.5.5 API Documentation Template

Note: *INSTRUCTIONS*

1. Use this file ([docs/en/api-reference/template.rst](#)) as a template to document API.
 2. Change the file name to the name of the header file that represents the documented API.
 3. Include respective files with descriptions from the API folder using `..include::`:
 - `README.rst`
 - `example.rst`
 - ...
 4. Optionally provide description right in this file.
 5. Once done, remove all instructions like this one and any superfluous headers.
-

Overview

Note: *INSTRUCTIONS*

1. Provide overview where and how this API may be used.
 2. Include code snippets to illustrate functionality of particular functions when applicable.
 3. To distinguish between sections, use the following [heading levels](#):
 - `#` with overline, for parts
 - `*` with overline, for chapters
 - `=` for sections
 - `-` for subsections
 - `^` for subsubsections
 - `"` for paragraphs
-

Application Example

Note: INSTRUCTIONS

1. Prepare one or more practical examples to demonstrate functionality of this API.
2. Each example should follow pattern of projects located in `esp-idf/examples/` folder.
3. Place example in this folder, and add `README.md` file.
4. Provide overview of demonstrated functionality in `README.md`.
5. With good overview readers should be able to understand what example does without opening the source code.
6. Depending on complexity of example, break down description of code into parts and provide overview of functionality of each part.
7. Include flow diagram and screenshots of application output if applicable.
8. Finally add in this section synopsis of each example together with link to respective folder in `esp-idf/examples/`.

API Reference

Note: INSTRUCTIONS

1. ESP-IDF repository provides automatic update of API reference documentation using *code markup retrieved by Doxygen from header files*.
2. Update is done on each documentation build by invoking Sphinx extension `esp_extensions/run_doxygen.py` for all header files listed in the `INPUT` statement of [docs/doxygen/Doxyfile](#).
3. Each line of the `INPUT` statement (other than a comment that begins with `##`) contains a path to header file `*.h` that is used to generate corresponding `*.inc` files:

```
##
## Wi-Fi - API Reference
##
../components/esp32/include/esp_wifi.h \
../components/esp32/include/esp_smartconfig.h \
```

4. When the headers are expanded, any macros defined by default in `sdkconfig.h` as well as any macros defined in SOC-specific `include/soc/*_caps.h` headers will be expanded. This allows the headers to include or exclude material based on the `IDF_TARGET` value.
5. The `*.inc` files contain formatted reference of API members generated automatically on each documentation build. All `*.inc` files are placed in `Sphinx_build` directory. To see directives generated, e.g., `esp_wifi.h`, run `python gen-dxd.py esp32/include/esp_wifi.h`.
6. To show contents of `*.inc` file in documentation, include it as follows:

```
.. include-build-file:: inc/esp_wifi.inc
```

For example see [docs/en/api-reference/network/esp_wifi.rst](#)

7. Optionally, rather than using `*.inc` files, you may want to describe API in your own way. See [docs/en/api-reference/storage/fatfs.rst](#) for example.

Below is the list of common `.. doxygen...::` directives:

- Functions - `.. doxygenfunction:: name_of_function`
- Unions - `.. doxygenunion:: name_of_union`
- Structures - `.. doxygenstruct:: name_of_structure` together with `:members:`
- Macros - `.. doxygendefine:: name_of_define`
- Type Definitions - `.. doxygentypedef:: name_of_type`
- Enumerations - `.. doxygenenum:: name_of_enumeration`

See [Breathe documentation](#) for additional information.

To provide a link to header file, use the `link custom role` directive as follows:

```
* :component_file:`path_to/header_file.h`
```

8. In any case, to generate API reference, the file [docs/doxygen/Doxyfile](#) should be updated with paths to *.h headers that are being documented.
 9. When changes are committed and documentation is built, check how this section has been rendered. *Correct annotations* in respective header files, if required.
-

8.5.6 Contributor Agreement

Individual Contributor Non-Exclusive License Agreement Including the Traditional Patent License OPTION

Thank you for your interest in contributing to this Espressif project hosted on GitHub ("We" or "Us").

The purpose of this contributor agreement ("Agreement") is to clarify and document the rights granted by contributors to Us. To make this document effective, please follow the instructions in the [Contributions Guide](#).

1. DEFINITIONS **You** means the Individual Copyright owner who submits a Contribution to Us. If You are an employee and submit the Contribution as part of your employment, You must have had Your employer approve this Agreement or sign the Entity version of this Agreement.

Contribution means any original work of authorship (software and/or documentation) including any modifications or additions to an existing work, Submitted by You to Us, in which You own the Copyright. If You do not own the Copyright in the entire work of authorship, please contact Us by submitting a comment on GitHub.

Copyright means all rights protecting works of authorship owned or controlled by You, including copyright, moral and neighboring rights, as appropriate, for the full term of their existence including any extensions by You.

Material means the software or documentation made available by Us to third parties. When this Agreement covers more than one software project, the Material means the software or documentation to which the Contribution was Submitted. After You Submit the Contribution, it may be included in the Material.

Submit means any form of physical, electronic, or written communication sent to Us, including but not limited to electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, Us, but excluding communication that is conspicuously marked or otherwise designated in writing by You as "Not a Contribution."

Submission Date means the date You Submit a Contribution to Us.

Documentation means any non-software portion of a Contribution.

2. COPYRIGHT LICENSE 2.1 Grant of Copyright License to Us

Subject to the terms and conditions of this Agreement, You hereby grant to Us a worldwide, royalty-free, NON-exclusive, perpetual and irrevocable license, with the right to transfer an unlimited number of non-exclusive licenses or to grant sublicenses to third parties, under the Copyright covering the Contribution to use the Contribution by all means, including, but not limited to:

- to publish the Contribution
- to modify the Contribution, to prepare derivative works based upon or containing the Contribution and to combine the Contribution with other software code
- to reproduce the Contribution in original or modified form
- to distribute, to make the Contribution available to the public, display and publicly perform the Contribution in original or modified form

2.2 Moral Rights remain unaffected to the extent they are recognized and not waivable by applicable law. Notwithstanding, You may add your name in the header of the source code files of Your Contribution and We will respect this attribution when using Your Contribution.

3. PATENT LICENSE 3.1 Grant of Patent License to US

Subject to the terms and conditions of this Agreement, You hereby grant to Us a worldwide, royalty-free, non-exclusive, perpetual and irrevocable (except as stated in Section 3.2) patent license, with the right to transfer an unlimited number of non-exclusive licenses or to grant sublicenses to third parties, to make, have made, use, sell, offer for sale, import and otherwise transfer the Contribution and the Contribution in combination with the Material (and portions of such combination). This license applies to all patents owned or controlled by You, whether already acquired or hereafter acquired, that would be infringed by making, having made, using, selling, offering for sale, importing or otherwise transferring of Your Contribution(s) alone or by combination of Your Contribution(s) with the Material.

3.2 Revocation of Patent License

You reserve the right to revoke the patent license stated in section 3.1 if We make any infringement claim that is targeted at your Contribution and not asserted for a Defensive Purpose. An assertion of claims of the Patents shall be considered for a "Defensive Purpose" if the claims are asserted against an entity that has filed, maintained, threatened, or voluntarily participated in a patent infringement lawsuit against Us or any of Our licensees.

4. DISCLAIMER THE CONTRIBUTION IS PROVIDED "AS IS". MORE PARTICULARLY, ALL EXPRESSED OR IMPLIED WARRANTIES INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT ARE EXPRESSLY DISCLAIMED BY YOU TO US AND BY US TO YOU. TO THE EXTENT THAT ANY SUCH WARRANTIES CANNOT BE DISCLAIMED, SUCH WARRANTY IS LIMITED IN DURATION TO THE MINIMUM PERIOD PERMITTED BY LAW.

5. Consequential Damage Waiver TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IN NO EVENT WILL YOU OR US BE LIABLE FOR ANY LOSS OF PROFITS, LOSS OF ANTICIPATED SAVINGS, LOSS OF DATA, INDIRECT, SPECIAL, INCIDENTAL, CONSEQUENTIAL AND EXEMPLARY DAMAGES ARISING OUT OF THIS AGREEMENT REGARDLESS OF THE LEGAL OR EQUITABLE THEORY (CONTRACT, TORT OR OTHERWISE) UPON WHICH THE CLAIM IS BASED.

6. Approximation of Disclaimer and Damage Waiver IF THE DISCLAIMER AND DAMAGE WAIVER MENTIONED IN SECTION 4 AND SECTION 5 CANNOT BE GIVEN LEGAL EFFECT UNDER APPLICABLE LOCAL LAW, REVIEWING COURTS SHALL APPLY LOCAL LAW THAT MOST CLOSELY APPROXIMATES AN ABSOLUTE WAIVER OF ALL CIVIL LIABILITY IN CONNECTION WITH THE CONTRIBUTION.

7. Term 7.1 This Agreement shall come into effect upon Your acceptance of the terms and conditions.

7.2 In the event of a termination of this Agreement, sections 4, 5, 6, 7 and 8 shall survive such termination and shall remain in full force thereafter. For the avoidance of doubt, Contributions that are already licensed under a free and open source license at the date of the termination shall remain in full force after the termination of this Agreement.

8. Miscellaneous 8.1 This Agreement and all disputes, claims, actions, suits or other proceedings arising out of this agreement or relating in any way to it shall be governed by the laws of People's Republic of China excluding its private international law provisions.

8.2 This Agreement sets out the entire agreement between You and Us for Your Contributions to Us and overrides all other agreements or understandings.

8.3 If any provision of this Agreement is found void and unenforceable, such provision will be replaced to the extent possible with a provision that comes closest to the meaning of the original provision and that is enforceable. The terms and conditions set forth in this Agreement shall apply notwithstanding any failure of essential purpose of this Agreement or any limited remedy to the maximum extent possible under law.

8.4 You agree to notify Us of any facts or circumstances of which you become aware that would make this Agreement inaccurate in any respect.

You

Date	
Name	
Title	
Address	

Us

Date	
Name	
Title	
Address	

8.5.7 Copyright Header Guide

ESP-IDF is released under [the Apache License 2.0](#) with some additional third-party copyrighted code released under various licenses. For further information please refer to [the list of copyrights and licenses](#).

This page explains how the source code should be properly marked with a copyright header. ESP-IDF uses the [Software Package Data Exchange \(SPDX\)](#) format which is short and can be easily read by humans or processed by automated tools for copyright checks.

How to Check the Copyright Headers

Please make sure you have installed the [pre-commit hooks](#) which contain a copyright header checker as well. The checker can suggest a header if it is not able to detect a properly formatted SPDX header.

What If the Checker's Suggestion Is Incorrect?

No automated checker (no matter how good is) can replace humans. So the developer's responsibility is to modify the offered header to be in line with the law and the license restrictions of the original code on which the work is based on. Certain licenses are not compatible between each other. Such corner cases will be covered by the following examples.

The checker can be configured with the `tools/ci/check_copyright_config.yaml` configuration file. Please check the options it offers and consider updating it in order to match the headers correctly.

Common Examples of Copyright Headers

The simplest case is when the code is not based on any licensed previous work, e.g., it was written completely from scratch. Such code can be decorated with the following copyright header and put under the license of ESP-IDF:

```
/*
 * SPDX-FileCopyrightText: 2015-2023 Espressif Systems (Shanghai) CO LTD
 *
 * SPDX-License-Identifier: Apache-2.0
 */
```

Less Restrictive Parts of ESP-IDF Some parts of ESP-IDF are deliberately under less restrictive licenses in order to ease their re-use in commercial closed source projects. This is the case for [ESP-IDF examples](#) which are in Public domain or under the Creative Commons Zero Universal (CC0) license. The following header can be used in such source files:

```
/*
 * SPDX-FileCopyrightText: 2015-2023 Espressif Systems (Shanghai) CO LTD
 *
 * SPDX-License-Identifier: Unlicense OR CC0-1.0
 */
```

The option allowing multiple licenses joined with the OR keyword from the above example can be achieved with the definition of multiple allowed licenses in the `tools/ci/check_copyright_config.yaml` configuration file. Please use this option with care and only selectively for a limited part of ESP-IDF.

Third Party Licenses Code licensed under different licenses, modified by Espressif Systems and included in ESP-IDF cannot be licensed under Apache License 2.0 not even if the checker suggests it. It is advised to keep the original copyright header and add an SPDX before it.

The following example is a suitable header for a code licensed under the "GNU General Public License v2.0 or later" held by John Doe with some additional modifications done by Espressif Systems:

```
/*
 * SPDX-FileCopyrightText: 1991 John Doe
 *
 * SPDX-License-Identifier: GPL-2.0-or-later
 *
 * SPDX-FileContributor: 2019-2023 Espressif Systems (Shanghai) CO LTD
 */
```

The licenses can be identified and the short SPDX identifiers can be found in the official [SPDX license list](#). Other very common licenses are the GPL-2.0-only, the BSD-3-Clause, and the BSD-2-Clause.

In exceptional case, when a license is not present on the [SPDX license list](#), it can be expressed by using the [LicenseRef-\[idString\]](#) custom license identifier, for example `LicenseRef-Special-License`. The full license text must be added into the `LICENSES` directory under `Special-License` filename.

```
/*
 * SPDX-FileCopyrightText: 2015-2023 Espressif Systems (Shanghai) CO LTD
 *
 * SPDX-License-Identifier: LicenseRef-Special-License
 */
```

Dedicated `LicenseRef-Included` custom license identifier can be used to express a situation when the custom license is included directly in the source file.

```
/*
 * SPDX-FileCopyrightText: 2015-2023 Espressif Systems (Shanghai) CO LTD
 *
 * SPDX-License-Identifier: LicenseRef-Included
 *
 * <Full custom license text>
 */
```

The configuration stored in `tools/ci/check_copyright_config.yaml` offers features useful for third party licenses:

- A different license can be defined for the files part of a third party library.
- The check for a selected set of files can be permanently disabled. Please use this option with care and only in cases when none of the other options are suitable.

8.5.8 pytest in ESP-IDF

ESP-IDF has numerous types of tests that are meant to be executed on an ESP chip (known as **on target testing**). Target tests are usually compiled as part of an IDF project used for testing (known as a **test app**), where test apps follows the same build, flash, and monitor process of any other standard IDF project.

Typically, on target testing will require a connected host (e.g., a PC) that is responsible for triggering a particular test case, providing test data, and inspecting test results.

ESP-IDF uses the pytest framework (and some pytest plugins) on the host side to automate on target testing. This guide introduces pytest in ESP-IDF and covers the following concepts:

1. The different types of test apps in ESP-IDF.
2. Using the pytest framework in Python scripts to automate target testing.
3. ESP-IDF Continuous Integration (CI) target testing process.
4. How to run target tests locally with pytest.
5. pytest tips and tricks.

Note: In ESP-IDF, we use the following pytest plugins by default:

- [pytest-embedded](#) with default services `esp`, `idf`
- [pytest-rerunfailures](#)

All the concepts and usages introduced in this guide are based on the default behavior of these plugins, thus may not be available in vanilla pytest.

Installation

All dependencies could be installed by running the install script with the `--enable-pytest` argument:

```
$ install.sh --enable-pytest
```

Common Issues During Installation

No Package 'dbus-1' Found

```
configure: error: Package requirements (dbus-1 >= 1.8) were not met:

No package 'dbus-1' found

Consider adjusting the PKG_CONFIG_PATH environment variable if you
installed software in a non-standard prefix.
```

If you encounter the error message above, you may need to install some missing packages.

If you are using Ubuntu, you may need to run:

```
sudo apt-get install libdbus-glib-1-dev
```

or

```
sudo apt-get install libdbus-1-dev
```

For other Linux distributions, please Google the error message above and find which missing packages need to be installed for your particular distribution.

Invalid Command 'bdist_wheel'

```
error: invalid command 'bdist_wheel'
```

If you encounter the error message above, you may need to install some missing Python packages such as:

```
python -m pip install -U pip
```

or

```
python -m pip install wheel
```

Note: Before running the pip commands, please make sure you are using the IDF Python virtual environment.

Test Apps

ESP-IDF contains different types of test apps that can be automated using pytest.

Component Tests ESP-IDF components typically contain component specific test apps that execute component specific unit tests. Component test apps are the recommended way to test components. All the test apps should be located under `${IDF_PATH}/components/<COMPONENT_NAME>/test_apps`, for example:

```
components/
├── my_component/
│   ├── include/
│   │   └── ...
│   ├── test_apps/
│   │   ├── test_app_1
│   │   │   ├── main/
│   │   │   │   └── ...
│   │   │   ├── CMakeLists.txt
│   │   │   └── pytest_my_component_app_1.py
│   │   ├── test_app_2
│   │   │   ├── ...
│   │   │   └── pytest_my_component_app_2.py
│   │   └── parent_folder
│   │       ├── test_app_3
│   │       │   ├── ...
│   │       │   └── pytest_my_component_app_3.py
│   │       └── ...
│   ├── my_component.c
│   └── CMakeLists.txt
```

Example Tests The purpose of ESP-IDF examples is to demonstrate parts of ESP-IDF functionality to users (refer to [Examples Readme](#) for more information).

However, to ensure that these examples operate correctly, examples can be treated as test apps and executed automatically by using pytest. All examples should be located under `${IDF_PATH}/examples`, with tested example including a Python test script, for example:

```
examples/
├── parent_folder/
│   └── example_1/
│       ├── main/
│       │   └── ...
│       ├── CMakeLists.txt
│       └── pytest_example_1.py
```


Custom Tests Custom Tests are tests that aim to test some arbitrary functionality of ESP-IDF, thus are not intended to demonstrate IDF functionality to users in any way.

All custom test apps are located under `${IDF_PATH}/tools/test_apps`. For more information please refer to the [Custom Test Readme](#).

pytest in ESP-IDF

pytest Execution Process

1. Bootstrapping Phase

Create session-scoped caches:

- port-target cache
- port-app cache

2. Collection Phase

- A. Gather all Python files with the prefix `pytest_`.
- B. Gather all test functions with the prefix `test_`.
- C. Apply the [params](#), and duplicate the test functions.
- D. Filter the test cases with CLI options. For the detailed usages, see [Filter the Test Cases](#).

3. Execution Phase

- A. Construct the [fixtures](#). In ESP-IDF, the common fixtures are initialized in this order:

- a. `pexpect_proc`: [pexpect](#) instance
- b. `app`: [IdfApp](#) instance

The test app's information (e.g., `sdkconfig`, `flash_files`, `partition_table`, etc) would be parsed at this phase.

- c. `serial`: [IdfSerial](#) instance

The port of the host to which the target is connected is auto-detected. In the case of multiple targets connected to the host, the test target's type is parsed from the app. The test app binary files are flashed to the test target automatically.

- d. `dut`: [IdfDut](#) instance

- B. Run the real test function.

- C. Deconstruct the fixtures in this order:

- a. `dut`
 - i. close the `serial` port.
 - ii. (Only for apps with [Unity test framework](#)) generate JUnit report of the Unity test cases.
- b. `serial`
- c. `app`
- d. `pexpect_proc`: Close the file descriptor

- D. (Only for apps with [Unity test framework](#))

If `dut.expect_from_unity_output()` is called, an `AssertionError` is raised upon detection of a Unity test failure.

4. Reporting Phase

- A. Generate JUnit report of the test functions.
- B. Modify the JUnit report test case name into ESP-IDF test case ID format: `<target>.<config>.<test function name>`.

5. Finalization Phase (Only for apps with [Unity test framework](#))

Combine the JUnit reports if the JUnit reports of the Unity test cases are generated.

Basic Example This following Python test script example is taken from `pytest_console_basic.py`.

```
@pytest.mark.esp32
@pytest.mark.esp32c3
@pytest.mark.generic
@pytest.mark.parametrize('config', [
    'history',
    'nohistory',
], indirect=True)
def test_console_advanced(config: str, dut: IdfDut) -> None:
    if config == 'history':
        dut.expect('Command history enabled')
    elif config == 'nohistory':
        dut.expect('Command history disabled')
```

To demonstrate how pytest is typically used in an ESP-IDF test script, let us go through this simple test script line by line in the following subsections.

Target Markers Pytest markers can be used to indicate which targets (i.e., which ESP chip) a particular test case should run on. For example:

```
@pytest.mark.esp32      # <-- support esp32
@pytest.mark.esp32c3    # <-- support esp32c3
@pytest.mark.generic    # <-- test env "generic"
```

The example above indicates that a particular test case is supported on the ESP32 and ESP32-C3. Furthermore, the target's board type should be generic. For more details regarding the generic type, you may run `pytest --markers` to get detailed information regarding all markers.

Note: If the test case can be run on all targets officially supported by ESP-IDF (call `idf.py --list-targets` for more details), you can use a special marker `supported_targets` to apply all of them in one line.

Parameterized Markers You can use `pytest.mark.parametrize` with `config` to apply the same test to different apps with different `sdkconfig.ci.xxx` files. For more information about `sdkconfig.ci.xxx` files, please refer to the Configuration Files section under [this readme](#).

```
@pytest.mark.parametrize('config', [
    'history',      # <-- run with app built by sdkconfig.ci.history
    'nohistory',    # <-- run with app built by sdkconfig.ci.nohistory
], indirect=True) # <-- `indirect=True` is required
```

Overall, this test function would be replicated to 4 test cases:

- `esp32.history.test_console_advanced`
- `esp32.nohistory.test_console_advanced`
- `esp32c3.history.test_console_advanced`
- `esp32c3.nohistory.test_console_advanced`

Testing Serial Output To ensure that test has executed successfully on target, the test script can test that serial output of the target using the `dut.expect()` function, for example:

```
def test_console_advanced(config: str, dut: IdfDut) -> None: # The value of_
↪argument ``config`` is assigned by the parameterization.
    if config == 'history':
        dut.expect('Command history enabled')
    elif config == 'nohistory':
        dut.expect('Command history disabled')
```

The `dut.expect(...)` will first compile the expected string into regex, which in turn is then used to seek through the serial output until the compiled regex is matched, or until a timeout occurs.

Please pay extra attention to the expected string when it contains regex keyword characters (e.g., parentheses, square brackets). Alternatively, you may use `dut.expect_exact(...)` that will attempt to match the string without converting it into regex.

For more information regarding the different types of `expect` functions, please refer to the [pytest-embedded Expecting documentation](#).

Advanced Examples

Multi-Target Tests with the Same App In some cases a test may involve multiple targets running the same test app. In this case, multiple DUTs can be instantiated using `parameterize`, for example:

```
@pytest.mark.esp32s2
@pytest.mark.esp32s3
@pytest.mark.usb_host
@pytest.mark.parametrize('count', [
    2,
], indirect=True)
def test_usb_host(dut: Tuple[IdfDut, IdfDut]) -> None:
    device = dut[0] # <-- assume the first dut is the device
    host = dut[1]   # <-- and the second dut is the host
    ...
```

After setting the param `count` to 2, all these fixtures are changed into tuples.

Multi-Target Tests with Different Apps In some cases (in particular protocol tests), a test may involve multiple targets running different test apps (e.g., separate targets to act as master and slave). In this case, multiple DUTs with different test apps can be instantiated using `parameterize`.

This code example is taken from `pytest_wifi_getting_started.py`.

```
@pytest.mark.esp32
@pytest.mark.multi_dut_generic
@pytest.mark.parametrize(
    'count, app_path', [
        (2,
         f'{os.path.join(os.path.dirname(__file__), "softAP")}|{os.path.join(os.
→path.dirname(__file__), "station")}',
        ], indirect=True
)
def test_wifi_getting_started(dut: Tuple[IdfDut, IdfDut]) -> None:
    softap = dut[0]
    station = dut[1]
    ...
```

Here the first DUT was flashed with the app `softAP`, and the second DUT was flashed with the app `station`.

Note: Here the `app_path` should be set with absolute path. The `__file__` macro in Python would return the absolute path of the test script itself.

Multi-Target Tests with Different Apps and Targets This code example is taken from `pytest_wifi_getting_started.py`. As the comment says, for now it is not running in the ESP-IDF CI.

```

@pytest.mark.parametrize(
    'count, app_path, target', [
        (2,
         f'{os.path.join(os.path.dirname(__file__), "softAP")}|{os.path.join(os.
→path.dirname(__file__), "station")}',
         'esp32|esp32s2'),
        (2,
         f'{os.path.join(os.path.dirname(__file__), "softAP")}|{os.path.join(os.
→path.dirname(__file__), "station")}',
         'esp32s2|esp32'),
    ],
    indirect=True,
)
def test_wifi_getting_started(dut: Tuple[IdfDut, IdfDut]) -> None:
    softap = dut[0]
    station = dut[1]
    ...

```

Overall, this test function would be replicated to 2 test cases:

- softAP with ESP32 target, and station with ESP32-S2 target
- softAP with ESP32-S2 target, and station with ESP32 target

Support Different Targets with Different sdkconfig Files This code example is taken from `pytest_panic.py` as an advanced example.

```

CONFIGS = [
    pytest.param('coredump_flash_bin_crc', marks=[pytest.mark.esp32, pytest.mark.
→esp32s2]),
    pytest.param('coredump_flash_elf_sha', marks=[pytest.mark.esp32]), # sha256
→only supported on esp32
    pytest.param('coredump_uart_bin_crc', marks=[pytest.mark.esp32, pytest.mark.
→esp32s2]),
    pytest.param('coredump_uart_elf_crc', marks=[pytest.mark.esp32, pytest.mark.
→esp32s2]),
    pytest.param('gdbstub', marks=[pytest.mark.esp32, pytest.mark.esp32s2]),
    pytest.param('panic', marks=[pytest.mark.esp32, pytest.mark.esp32s2]),
]

@pytest.mark.parametrize('config', CONFIGS, indirect=True)
...

```

Custom Classes Usually, you may want to write a custom class under these conditions:

1. Add more reusable functions for a certain number of DUTs.
2. Add custom setup and teardown functions in different phases described in Section *pytest Execution Process*.

This code example is taken from `panic/confest.py`.

```

class PanicTestDut (IdfDut) :
    ...

@pytest.fixture(scope='module')
def monkeypatch_module(request: FixtureRequest) -> MonkeyPatch:
    mp = MonkeyPatch()
    request.addfinalizer(mp.undo)
    return mp

@pytest.fixture(scope='module', autouse=True)

```

(continues on next page)

(continued from previous page)

```
def replace_dut_class(monkeypatch_module: MonkeyPatch) -> None:
    monkeypatch_module.setattr('pytest_embedded_idf.dut.IdfDut', PanicTestDut)
```

`monkeypatch_module` provides a [module-scoped monkeypatch](#) fixture.

`replace_dut_class` is a [module-scoped autouse](#) fixture. This function replaces the `IdfDut` class with your custom class.

Mark Flaky Tests Certain test cases are based on Ethernet or Wi-Fi. However, the test may be flaky due to networking issues. Thus, it is possible to mark a particular test case as flaky.

This code example is taken from `pytest_esp_eth.py`.

```
@pytest.mark.flaky(reruns=3, reruns_delay=5)
def test_esp_eth_ip101(dut: IdfDut) -> None:
    ...
```

This flaky marker means that if the test function failed, the test case would rerun for a maximum of 3 times with 5 seconds delay.

Mark Known Failures Sometimes, a test can consistently fail for the following reasons:

- The feature under test (or the test itself) has a bug.
- The test environment is unstable (e.g., due to network issues) leading to a high failure ratio.

Now you may mark this test case with marker `xfail` with a user-friendly readable reason.

This code example is taken from `pytest_panic.py`

```
@pytest.mark.xfail('config.getvalue("target") == "esp32s2"', reason='raised_
↳ IllegalInstruction instead')
def test_cache_error(dut: PanicTestDut, config: str, test_func_name: str) -> None:
```

This marker means that test is a known failure on the ESP32-S2.

Mark Nightly Run Test Cases Some test cases are only triggered in nightly run pipelines due to a lack of runners.

```
@pytest.mark.nightly_run
```

This marker means that the test case would only be run with env var `NIGHTLY_RUN` or `INCLUDE_NIGHTLY_RUN`.

Mark Temporarily Disabled in CI Some test cases which can pass locally may need to be temporarily disabled in CI due to a lack of runners.

```
@pytest.mark.temp_skip_ci(targets=['esp32', 'esp32s2'], reason='lack of runners')
```

This marker means that the test case could still be run locally with `pytest --target esp32`, but will not run in CI.

Run Unity Test Cases For component-based unit test apps, all single-board test cases (including normal test cases and multi-stage test cases) can be run using the following command:

```
def test_component_ut(dut: IdfDut):
    dut.run_all_single_board_cases()
```

Using this command will skip all the test cases containing the `[ignore]` tag.

If you need to run a group of test cases, you may run:

```
def test_component_ut(dut: IdfDut):  
    dut.run_all_single_board_cases(group='psram')
```

It would trigger all test cases with the `[psram]` tag.

You may also see that there are some test scripts with the following statements, which are deprecated. Please use the suggested one as above.

```
def test_component_ut(dut: IdfDut):  
    dut.expect_exact('Press ENTER to see the list of tests')  
    dut.write('*')  
    dut.expect_unity_test_output()
```

For further reading about our unit testing in ESP-IDF, please refer to [our unit testing guide](#).

Running Tests in CI

The workflow in CI is simple, build jobs > target test jobs.

Build Jobs

Build Job Names

- Component-based Unit Tests: `build_pytest_components_<target>`
- Example Tests: `build_pytest_examples_<target>`
- Custom Tests: `build_pytest_test_apps_<target>`

Build Job Commands The command used by CI to build all the relevant tests is: `python $IDF_PATH/tools/ci/ci_build_apps.py <parent_dir> --target <target> -vv --pytest-apps`

All apps which supported the specified target would be built with all supported `sdkconfig` files under `build_<target>_<config>`.

For example, If you run `python $IDF_PATH/tools/ci/ci_build_apps.py $IDF_PATH/examples/system/console/basic --target esp32 --pytest-apps`, the folder structure would be like this:

```
basic  
├── build_esp32_history/  
│   └── ...  
├── build_esp32_nohistory/  
│   └── ...  
├── main/  
├── CMakeLists.txt  
├── pytest_console_basic.py  
└── ...
```

All the build folders would be uploaded as artifacts under the same directories.

Target Test Jobs

Target Test Job Names

- Component-based Unit Tests: `component_ut_pytest_<target>_<test_env>`
- Example Tests: `example_test_pytest_<target>_<test_env>`
- Custom Tests: `test_app_test_pytest_<target>_<test_env>`

Target Test Job Commands The command used by CI to run all the relevant tests is: `pytest <parent_dir> --target <target> -m <test_env_marker>`

All test cases with the specified target marker and the test env marker under the parent folder would be executed.

The binaries in the target test jobs are downloaded from build jobs. the artifacts would be placed under the same directories.

Running Tests Locally

First you need to install ESP-IDF with additional Python requirements:

```
$ cd $IDF_PATH
$ bash install.sh --enable-pytest
$ . ./export.sh
```

By default, the pytest script will look for the build directory in this order:

- `build_<target>_<sdkconfig>`
- `build_<target>`
- `build_<sdkconfig>`
- `build`

Which means, the simplest way to run pytest is calling `idf.py build`.

For example, if you want to run all the esp32 tests under the `$IDF_PATH/examples/get-started/hello_world` folder, you should run:

```
$ cd examples/get-started/hello_world
$ idf.py build
$ pytest --target esp32
```

If you have multiple `sdkconfig` files in your test app, like those `sdkconfig.ci.*` files, the simple `idf.py build` won't apply the extra `sdkconfig` files. Let us take `$IDF_PATH/examples/system/console/basic` as an example.

If you want to test this app with config history, and build with `idf.py build`, you should run

```
$ cd examples/system/console/basic
$ idf.py -DSDKCONFIG_DEFAULTS="sdkconfig.defaults;sdkconfig.ci.history" build
$ pytest --target esp32 --sdkconfig history
```

If you want to build and test with all `sdkconfig` files at the same time, you should use our CI script as an helper script:

```
$ cd examples/system/console/basic
$ python $IDF_PATH/tools/ci/ci_build_apps.py . --target esp32 -vv --pytest-apps
$ pytest --target esp32
```

The app with `sdkconfig.ci.history` will be built in `build_esp32_history`, and the app with `sd-
kconfig.ci.nohistory` will be built in `build_esp32_nohistory`. `pytest --target esp32` will run tests on both apps.

Tips and Tricks

Filter the Test Cases

- Filter by target with `pytest --target <target>`
`pytest` would run all the test cases that support specified target.
- Filter by `sdkconfig` file with `pytest --sdkconfig <sdkconfig>`
If `<sdkconfig>` is default, `pytest` would run all the test cases with the `sdkconfig` file `sdkconfig.defaults`.
In other cases, `pytest` would run all the test cases with `sdkconfig` file `sdkconfig.ci.<sdkconfig>`.

Add New Markers We are using two types of custom markers, target markers which indicate that the test cases should support this target, and env markers which indicate that the test cases should be assigned to runners with these tags in CI.

You can add new markers by adding one line under the `${IDF_PATH}/conftest.py`. If it is a target marker, it should be added into `TARGET_MARKERS`. If it is a marker that specifies a type of test environment, it should be added into `ENV_MARKERS`. The syntax should be: `<marker_name>: <marker_description>`.

Generate JUnit Report You can call pytest with `--junitxml <filepath>` to generate the JUnit report. In ESP-IDF, the test case name would be unified as `<target>.<config>.<function_name>`.

Skip Auto Flash Binary Skipping auto-flash binary every time would be useful when you are debugging your test script.

You can call pytest with `--skip-autoflash y` to achieve it.

Record Statistics Sometimes you may need to record some statistics while running the tests, like the performance test statistics.

You can use `record_xml_attribute` fixture in your test script, and the statistics would be recorded as attributes in the JUnit report.

Logging System Sometimes you may need to add some extra logging lines while running the test cases.

You can use [Python logging module](#) to achieve this.

Useful Logging Functions (as Fixture)

log_performance

```
def test_hello_world(
    dut: IdfDut,
    log_performance: Callable[[str, object], None],
) -> None:
    log_performance('test', 1)
```

The above example would log the performance item with pre-defined format: `[performance][test]: 1` and record it under the `properties` tag in the JUnit report if `--junitxml <filepath>` is specified. The JUnit test case node would look like:

```
<testcase classname="examples.get-started.hello_world.pytest_hello_world" file=
↳ "examples/get-started/hello_world/pytest_hello_world.py" line="13" name="esp32.
↳ default.test_hello_world" time="8.389">
  <properties>
    <property name="test" value="1"/>
  </properties>
</testcase>
```

check_performance We provide C macros `TEST_PERFORMANCE_LESS_THAN` and `TEST_PERFORMANCE_GREATER_THAN` to log the performance item and check if the value is in the valid range. Sometimes the performance item value could not be measured in C code, so we also provide a Python function for the same purpose. Please note that using C macros is the preferred approach, since the Python function could not recognize the threshold values of the same performance item under different `ifdef` blocks well.


```
def test_hello_world(  
    dut: IdfDut,  
    check_performance: Callable[[str, float, str], None],  
) -> None:  
    check_performance('RSA_2048KEY_PUBLIC_OP', 123, 'esp32')  
    check_performance('RSA_2048KEY_PUBLIC_OP', 19001, 'esp32')
```

The above example would first get the threshold values of the performance item `RSA_2048KEY_PUBLIC_OP` from `components/idf_test/include/idf_performance.h` and the target-specific one `components/idf_test/include/esp32/idf_performance_target.h`, then check if the value reached the minimum limit or exceeded the maximum limit.

Let us assume the value of `IDF_PERFORMANCE_MAX_RSA_2048KEY_PUBLIC_OP` is 19000. so the first `check_performance` line would pass and the second one would fail with warning: [Performance] `RSA_2048KEY_PUBLIC_OP` value is 19001, doesn't meet pass standard 19000.0.

Further Readings

- pytest documentation: <https://docs.pytest.org/en/latest/contents.html>
- pytest-embedded documentation: <https://docs.espressif.com/projects/pytest-embedded/en/latest/>

Chapter 9

ESP-IDF Versions

The ESP-IDF GitHub repository is updated regularly, especially the master branch where new development takes place.

For production use, there are also stable releases available.

9.1 Releases

The documentation for the current stable release version can always be found at this URL:

<https://docs.espressif.com/projects/esp-idf/en/stable/>

Documentation for the latest version (master branch) can always be found at this URL:

<https://docs.espressif.com/projects/esp-idf/en/latest/>

The full history of releases can be found on the GitHub repository [Releases page](#). There you can find release notes, links to each version of the documentation, and instructions for obtaining each version.

9.2 Which Version Should I Start With?

- For production purposes, use the [current stable version](#). Stable versions have been manually tested, and are updated with "bugfix releases" which fix bugs without changing other functionality (see [Versioning Scheme](#) for more details). Every stable release version can be found on the [Releases page](#). Also refer to [Compatibility Between ESP-IDF Releases and Revisions of Espressif SoCs](#) to make sure the ESP-IDF version you selected is compatible with the chip revision you are going to produce with.
- For prototyping, experimentation or for developing new ESP-IDF features, use the [latest version \(master branch in Git\)](#). The latest version in the master branch has all the latest features and has passed automated testing, but has not been completely manually tested ("bleeding edge").
- If a required feature is not yet available in a stable release, but you do not want to use the master branch, it is possible to check out a pre-release version or a release branch. It is recommended to start from a stable version and then follow the instructions for [Updating to a Pre-Release Version](#) or [Updating to a Release Branch](#).
- If you plan to use another project which is based on ESP-IDF, please check the documentation of that project to determine the version(s) of ESP-IDF it is compatible with.

See [Updating ESP-IDF](#) if you already have a local copy of ESP-IDF and wish to update it.

9.3 Versioning Scheme

ESP-IDF uses [Semantic Versioning](#). This means that:

- Major Releases, like `v3.0`, add new functionality and may change functionality. This includes removing deprecated functionality.
If updating to a new major release (for example, from `v2.1` to `v3.0`), some of your project's code may need updating and functionality may need to be re-tested. The release notes on the [Releases page](#) include lists of Breaking Changes to refer to.
- Minor Releases like `v3.1` add new functionality and fix bugs but will not change or remove documented functionality, or make incompatible changes to public APIs.
If updating to a new minor release (for example, from `v3.0` to `v3.1`), your project's code does not require updating, but you should re-test your project. Pay particular attention to the items mentioned in the release notes on the [Releases page](#).
- Bugfix Releases like `v3.0.1` only fix bugs and do not add new functionality.
If updating to a new bugfix release (for example, from `v3.0` to `v3.0.1`), you do not need to change any code in your project, and you only need to re-test the functionality directly related to bugs listed in the release notes on the [Releases page](#).

9.4 Support Periods

Each ESP-IDF major and minor release version has an associated support period. After this period, the release is End of Life and no longer supported.

The [ESP-IDF Support Period Policy](#) explains this in detail, and describes how the support periods for each release are determined.

Each release on the [Releases page](#) includes information about the support period for that particular release.

As a general guideline:

- If starting a new project, use the latest stable release.
- If you have a GitHub account, click the "Watch" button in the top-right of the [Releases page](#) and choose "Releases only". GitHub will notify you whenever a new release is available. Whenever a bug fix release is available for the version you are using, plan to update to it.
- If possible, periodically update the project to a new major or minor ESP-IDF version (for example, once a year.) The update process should be straightforward for Minor updates, but may require some planning and checking of the release notes for Major updates.
- Always plan to update to a newer release before the release you are using becomes End of Life.

Each ESP-IDF major and minor release (V4.1, V4.2, etc) is supported for 30 months after the initial stable release date.

Supported means that the ESP-IDF team will continue to apply bug fixes, security fixes, etc to the release branch on GitHub, and periodically make new bugfix releases as needed.

Support period is divided into "Service" and "Maintenance" period:

Period	Duration	Recommended for new projects?
Service	12 months	Yes
Maintenance	18 months	No

During the Service period, bugfixes releases are more frequent. In some cases, support for new features may be added during the Service period (this is reserved for features which are needed to meet particular regulatory requirements or standards for new products, and which carry a very low risk of introducing regressions.)


During the Maintenance period, the version is still supported but only bugfixes for high severity issues or security issues will be applied.

Using an "In Service" version is recommended when starting a new project.

Users are encouraged to upgrade all projects to a newer ESP-IDF release before the support period finishes and the release becomes End of Life (EOL). It is our policy to not continue fixing bugs in End of Life releases.

Pre-release versions (betas, previews, `-rc` and `-dev` versions, etc) are not covered by any support period. Sometimes a particular feature is marked as "Preview" in a release, which means it is also not covered by the support period.

The ESP-IDF Programming Guide has information about the [different versions of ESP-IDF](#) (major, minor, bugfix, etc).



9.5 Checking the Current Version

The local ESP-IDF version can be checked by using `idf.py`:

```
idf.py --version
```

The ESP-IDF version is also compiled into the firmware and can be accessed (as a string) via the macro `IDF_VER`. The default ESP-IDF bootloader will print the version on boot (the version information is not always updated if the code in the GitHub repo is updated, it only changes if there is a clean build or if that particular source file is recompiled).

If writing code that needs to support multiple ESP-IDF versions, the version can be checked at compile time using [compile-time macros](#).

Examples of ESP-IDF versions:

Version String	Meaning
v3.2-dev-306-gbeb3611ca	Master branch pre-release. - v3.2-dev - in development for version 3.2. - 306 - number of commits after v3.2 development started. - beb3611ca - commit identifier.
v3.0.2	Stable release, tagged v3.0.2.
v3.1-beta1-75-g346d6b0ea	Beta version in development (on a release branch). - v3.1-beta1 - pre-release tag. - 75 - number of commits after the pre-release beta tag was assigned. - 346d6b0ea - commit identifier.
v3.0.1-dirty	Stable release, tagged v3.0.1. - dirty means that there are modifications in the local ESP-IDF directory.

9.6 Git Workflow

The development (Git) workflow of the Espressif ESP-IDF team is as follows:

- New work is always added on the master branch (latest version) first. The ESP-IDF version on `master` is always tagged with `-dev` (for "in development"), for example `v3.1-dev`.
- Changes are first added to an internal Git repository for code review and testing but are pushed to GitHub after automated testing passes.
- When a new version (developed on `master`) becomes feature complete and "beta" quality, a new branch is made for the release, for example `release/v3.1`. A pre-release tag is also created, for example `v3.1-beta1`. You can see a full [list of branches](#) and a [list of tags](#) on GitHub. Beta pre-releases have release notes which may include a significant number of Known Issues.
- As testing of the beta version progresses, bug fixes will be added to both the `master` branch and the release branch. New features for the next release may start being added to `master` at the same time.
- Once testing is nearly complete a new release candidate is tagged on the release branch, for example `v3.1-rc1`. This is still a pre-release version.
- If no more significant bugs are found or reported, then the final Major or Minor Version is tagged, for example `v3.1`. This version appears on the [Releases page](#).
- As bugs are reported in released versions, the fixes will continue to be committed to the same release branch.
- Regular bugfix releases are made from the same release branch. After manual testing is complete, a bugfix release is tagged (i.e., `v3.1.1`) and appears on the [Releases page](#).

9.7 Updating ESP-IDF

Updating ESP-IDF depends on which version(s) you wish to follow:

- [Updating to Stable Release](#) is recommended for production use.

- [Updating to Master Branch](#) is recommended for the latest features, development use, and testing.
- [Updating to a Release Branch](#) is a compromise between the first two.

Note: These guides assume that you already have a local copy of ESP-IDF cloned. To get one, check Step 2 in the [Getting Started](#) guide for any ESP-IDF version.

9.7.1 Updating to Stable Release

To update to a new ESP-IDF release (recommended for production use), this is the process to follow:

- Check the [Releases page](#) regularly for new releases.
- When a bugfix release for the version you are using is released (for example, if using v3.0.1 and v3.0.2 is released), check out the new bugfix version into the existing ESP-IDF directory.
- In Linux or macOS system, please run the following commands to update the local branch to vX.Y.Z:

```
cd $IDF_PATH
git fetch
git checkout vX.Y.Z
git submodule update --init --recursive
```

- In the Windows system, please replace `cd $IDF_PATH` with `cd %IDF_PATH%`.
- When major or minor updates are released, check the Release Notes on the releases page and decide if you want to update or to stay with your current release. Updating is via the same Git commands shown above.

Note: If you installed the stable release via zip file instead of using git, it might not be possible to update versions using the commands. In this case, update by downloading a new zip file and replacing the entire `IDF_PATH` directory with its contents.

9.7.2 Updating to a Pre-Release Version

It is also possible to `git checkout` a tag corresponding to a pre-release version or release candidate, the process is the same as [Updating to Stable Release](#).

Pre-release tags are not always found on the [Releases page](#). Consult the [list of tags](#) on GitHub for a full list. Caveats for using a pre-release are similar to [Updating to a Release Branch](#).

9.7.3 Updating to Master Branch

Note: Using Master branch means living "on the bleeding edge" with the latest ESP-IDF code.

To use the latest version on the ESP-IDF master branch, this is the process to follow:

- In Linux or macOS system, please run the following commands to check out to the master branch locally:

```
cd $IDF_PATH
git checkout master
git pull
git submodule update --init --recursive
```

- In the Windows system, please replace `cd $IDF_PATH` with `cd %IDF_PATH%`.
- Periodically, re-run `git pull` to pull the latest version of master. Note that you may need to change your project or report bugs after updating your master branch.

- To switch from master to a release branch or stable version, run `git checkout` as shown in the other sections.

Important: It is strongly recommended to regularly run `git pull` and then `git submodule update --init --recursive` so a local copy of master does not get too old. Arbitrary old master branch revisions are effectively unsupportable "snapshots" that may have undocumented bugs. For a semi-stable version, try [Updating to a Release Branch](#) instead.

9.7.4 Updating to a Release Branch

In terms of stability, using a release branch is part-way between using the master branch and only using stable releases. A release branch is always beta quality or better, and receives bug fixes before they appear in each stable release.

You can find a [list of branches](#) on GitHub.

For example, in Linux or macOS system, you can execute the following commands to follow the branch for ESP-IDF v3.1, including any bugfixes for future releases like v3.1.1, etc:

```
cd $IDF_PATH
git fetch
git checkout release/v3.1
git pull
git submodule update --init --recursive
```

In the Windows system, please replace `cd $IDF_PATH` with `cd %IDF_PATH%`.

Each time you `git pull` this branch, ESP-IDF will be updated with fixes for this release.

Note: There is no dedicated documentation for release branches. It is recommended to use the documentation for the closest version to the branch which is currently checked out.

Chapter 10

Resources

10.1 PlatformIO



- *[What Is PlatformIO?](#)*
- *[Installation](#)*
- *[Configuration](#)*
- *[Tutorials](#)*
- *[Project Examples](#)*
- *[Next Steps](#)*

10.1.1 What Is PlatformIO?

PlatformIO is a cross-platform embedded development environment with out-of-the-box support for ESP-IDF.

Since ESP-IDF support within PlatformIO is not maintained by the Espressif team, please report any issues with PlatformIO directly to its developers in [the official PlatformIO repositories](#).

A detailed overview of the PlatformIO ecosystem and its philosophy can be found in [the official PlatformIO documentation](#).

10.1.2 Installation

- [PlatformIO IDE](#) is a toolset for embedded C/C++ development available on Windows, macOS and Linux platforms.
- [PlatformIO Core \(CLI\)](#) is a command-line tool that consists of multi-platform build system, platform and library managers and other integration components. It can be used with a variety of code development environments and allows integration with cloud platforms and web services

10.1.3 Configuration

Please go through [the official PlatformIO configuration guide for ESP-IDF](#).

10.1.4 Tutorials

- [ESP-IDF and ESP32-DevKitC: debugging, unit testing, project analysis](#)

10.1.5 Project Examples

Please check ESP-IDF page in [the official PlatformIO documentation](#)

10.1.6 Next Steps

Here are some useful links for exploring the PlatformIO ecosystem:

- Learn more about [integrations with other IDEs or Text Editors](#)
- Get help from [PlatformIO community](#)

10.2 CLion

10.2.1 What Is CLion?

[CLion](#) is a cross-platform integrated Development Environment (IDE) for C and C++ programming. CLion also provides dedicated support for ESP-IDF, allowing developers to seamlessly work with the ESP-IDF framework.

10.2.2 Installation

To install CLion, please follow the instructions provided in [Install CLion](#) for your operating system (Windows, macOS, or Linux).

10.2.3 Configuration

To configure an ESP-IDF project in CLion, please refer to the guide on [Configure an ESP CMake project in CLion](#). This guide will walk you through the necessary steps to set up your project properly.

10.2.4 Resources

For more information about CLion and ESP-IDF integration, please refer to the following resource:

- [CLion Documentation](#): The official documentation for CLion provides detailed information on various aspects of the IDE, including ESP-IDF integration.

10.3 VisualGDB

10.3.1 What Is VisualGDB?

[VisualGDB](#) is a powerful extension for Microsoft Visual Studio that provides advanced development tools and features for embedded systems, including support for the ESP-IDF framework. VisualGDB allows you to leverage the familiar

and feature-rich Visual Studio environment for your ESP-IDF projects, enabling efficient coding, debugging, and deployment.

10.3.2 Installation

Please download and install VisualGDB by following the steps stated in [VisualGDB download and installation](#).

10.3.3 Configuration

[Creating Advanced ESP32 Projects with ESP-IDF](#) provide basic steps about how to configure an ESP-IDF project in VisualGDB.

You can also refer to [Advanced ESP-IDF Project Structure](#) to get a more comprehensive impression for developing ESP-IDF projects using VisualGDB.

10.3.4 Resources

For more information about VisualGDB and ESP-IDF integration, refer to the following resources:

- [VisualGDB Documentation](#): The official documentation for VisualGDB provides comprehensive guides and tutorials on using VisualGDB with ESP-IDF.

For inquiries related to these third-party tools, we recommend seeking assistance from the respective tool's support channels or user communities.

10.4 Useful Links

- The [esp32.com forum](#) is a place to ask questions and find community resources.
- Check the [Issues](#) section on GitHub if you find a bug or have a feature request. Please check existing [Issues](#) before opening a new one.
- A comprehensive collection of [solutions](#), [practical applications](#), [components and drivers](#) based on ESP-IDF is available in [ESP IoT Solution](#) repository. In most of cases descriptions are provided both in English and in 中文.
- To develop applications using Arduino platform, refer to [Arduino core for the ESP32, ESP32-S2 and ESP32-C3](#).
- Several [books](#) have been written about ESP32 and they are listed on [Espressif](#) web site.
- If you're interested in contributing to ESP-IDF, please check the [Contributions Guide](#).
- For additional ESP32-P4 product related information, please refer to [documentation](#) section of [Espressif](#) site.
- [Download](#) latest and previous versions of this documentation in PDF and HTML format.

Chapter 11

Copyrights and Licenses

11.1 Software Copyrights

All original source code in this repository is Copyright (C) 2015-2023 Espressif Systems. This source code is licensed under the Apache License 2.0 as described in the file LICENSE.

Additional third party copyrighted code is included under the following licenses.

Where source code headers specify Copyright & License information, this information takes precedence over the summaries made here.

Some examples use external components which are not Apache licensed, please check the copyright description in each example source code.

11.1.1 Firmware Components

These third party libraries can be included into the application (firmware) produced by ESP-IDF.

- [Newlib](#) is licensed under the BSD License and is Copyright of various parties, as described in [COPYING.NEWLIB](#).
- [Xtensa header files](#) are Copyright (C) 2013 Tensilica Inc and are licensed under the MIT License as reproduced in the individual header files.
- Original parts of [FreeRTOS](#) (components/freertos) are Copyright (C) 2017 Amazon.com, Inc. or its affiliates are licensed under the MIT License, as described in [license.txt](#).
- Original parts of [LWIP](#) (components/lwip) are Copyright (C) 2001, 2002 Swedish Institute of Computer Science and are licensed under the BSD License as described in [COPYING file](#).
- [wpa_supplicant](#) Copyright (c) 2003-2022 Jouni Malinen <j@w1.fi> and contributors and licensed under the BSD license.
- [Fast PBKDF2](#) Copyright (c) 2015 Joseph Birr-Pixton and licensed under CC0 Public Domain Dedication license.
- [FreeBSD net80211](#) Copyright (c) 2004-2008 Sam Leffler, Errno Consulting and licensed under the BSD license.
- [argtable3](#) argument parsing library Copyright (C) 1998-2001, 2003-2011, 2013 Stewart Heitmann and licensed under 3-clause BSD license. [argtable3](#) also includes the following software components. For details, please see [argtable3 LICENSE file](#).
 - C Hash Table library, Copyright (c) 2002, Christopher Clark and licensed under 3-clause BSD license.
 - The Better String library, Copyright (c) 2014, Paul Hsieh and licensed under 3-clause BSD license.
 - TCL library, Copyright the Regents of the University of California, Sun Microsystems, Inc., Scriptics Corporation, ActiveState Corporation and other parties, and licensed under TCL/TK License.
- [linenoise](#) line editing library Copyright (c) 2010-2014 Salvatore Sanfilippo, Copyright (c) 2010-2013 Pieter Noordhuis, licensed under 2-clause BSD license.
- [FatFS](#) library, Copyright (C) 2017 ChaN, is licensed under [a BSD-style license](#).

- [cJSON](#) library, Copyright (c) 2009-2017 Dave Gamble and cJSON contributors, is licensed under MIT license as described in [LICENSE file](#).
- [micro-ecc](#) library, Copyright (c) 2014 Kenneth MacKay, is licensed under 2-clause BSD license.
- [Mbed TLS](#) library, Copyright (C) 2006-2018 ARM Limited, is licensed under Apache License 2.0 as described in [LICENSE file](#).
- [SPIFFS](#) library, Copyright (c) 2013-2017 Peter Andersson, is licensed under MIT license as described in [LICENSE file](#).
- [SD/MMC driver](#) is derived from [OpenBSD SD/MMC driver](#), Copyright (c) 2006 Uwe Stuehler, and is licensed under BSD license.
- [ESP-MQTT](#) MQTT Package (contiki-mqtt) - Copyright (c) 2014, Stephen Robinson, MQTT-ESP - Tuan PM <tuanpm at live dot com> is licensed under Apache License 2.0 as described in [LICENSE file](#).
- [BLE Mesh](#) is adapted from Zephyr Project, Copyright (c) 2017-2018 Intel Corporation and licensed under Apache License 2.0.
- [mynewt-nimble](#) Apache Mynewt NimBLE, Copyright 2015-2018, The Apache Software Foundation, is licensed under Apache License 2.0 as described in [LICENSE file](#).
- [TLSF allocator](#) Two Level Segregated Fit memory allocator, Copyright (c) 2006-2016, Matthew Conte, and licensed under the BSD 3-clause license.
- [openthread](#), Copyright (c) The OpenThread Authors, is licensed under BSD License as described in [LICENSE file](#).
- [UBSAN runtime](#) — Copyright (c) 2016, Linaro Limited and Jiří Závěručky, licensed under the BSD 2-clause license.
- [HTTP Parser](#) Based on src/http/nginx_http_parse.c from NGINX copyright Igor Sysoev. Additional changes are licensed under the same terms as NGINX and Joyent, Inc. and other Node contributors. For details please check [LICENSE file](#).
- [SEGGER SystemView](#) target-side library, Copyright (c) 1995-2021 SEGGER Microcontroller GmbH, is licensed under BSD 1-clause license.

11.1.2 Documentation

- HTML version of the [ESP-IDF Programming Guide](#) uses the Sphinx theme [sphinx_idf_theme](#), which is Copyright (c) 2013-2020 Dave Snider, Read the Docs, Inc. & contributors, and Espressif Systems (Shanghai) CO., LTD. It is based on [sphinx_rtd_theme](#). Both are licensed under MIT license.

11.2 ROM Source Code Copyrights

Espressif SoCs mask ROM hardware includes binaries compiled from portions of the following third party software:

- [Newlib](#), licensed under the BSD License and is Copyright of various parties, as described in [COPYING.NEWLIB](#).
- Xtensa libhal, Copyright (c) Tensilica Inc and licensed under the MIT license (see below).
- [TinyBasic](#) Plus, Copyright Mike Field & Scott Lawrence and licensed under the MIT license (see below).
- [miniz](#), by Rich Geldreich - placed into the public domain.
- [TJpgDec](#) Copyright (C) 2011, ChaN, all right reserved. See below for license.
- **Parts of Zephyr RTOS USB stack:**
 - [DesignWare USB device driver](#) Copyright (c) 2016 Intel Corporation and licensed under Apache 2.0 license.
 - [Generic USB device driver](#) Copyright (c) 2006 Bertrik Sikken (bertrik@sikken.nl), 2016 Intel Corporation and licensed under BSD 3-clause license.
 - [USB descriptors functionality](#) Copyright (c) 2017 PHYTEC Messtechnik GmbH, 2017-2018 Intel Corporation and licensed under Apache 2.0 license.
 - [USB DFU class driver](#) Copyright (c) 2015-2016 Intel Corporation, 2017 PHYTEC Messtechnik GmbH and licensed under BSD 3-clause license.
 - [USB CDC ACM class driver](#) Copyright (c) 2015-2016 Intel Corporation and licensed under Apache 2.0 license.

11.3 Xtensa libhal MIT License

Copyright (c) 2003, 2006, 2010 Tensilica Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

11.4 TinyBasic Plus MIT License

Copyright (c) 2012-2013

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

11.5 TJpgDec License

TJpgDec - Tiny JPEG Decompressor R0.01 (C) ChaN, 2011 The TJpgDec is a generic JPEG decompressor module for tiny embedded systems. This is a free software that opened for education, research and commercial developments under license policy of following terms.

Copyright (C) 2011, ChaN, all right reserved.

- The TJpgDec module is a free software and there is NO WARRANTY.
- No restriction on use. You can use, modify and redistribute it for personal, non-profit or commercial products UNDER YOUR RESPONSIBILITY.
- Redistributions of source code must retain the above copyright notice.

Chapter 12

About

This is documentation of [ESP-IDF](#), the framework to develop applications for ESP32-P4.

The ESP32-P4 is a high-performance MCU that supports large internal memory and has powerful image and voice processing capabilities. The MCU consists of a High Performance (HP) system and a Low Power (LP) system. The HP system contains a RISC-V dual-core CPU running up to 400 MHz and rich peripherals, while the LP system contains a low-power RISC-V single-core CPU running up to 40 MHz and various peripherals optimized for low-power applications.




Fig. 1: Espressif IoT Integrated Development Framework

The ESP-IDF, Espressif IoT Development Framework, provides toolchain, API, components and workflows to develop applications for ESP32-P4 using Windows, Linux and macOS operating systems.

Chapter 13

Switch Between Languages

The ESP-IDF Programming Guide is now available in two languages. Please refer to the English version if there is any discrepancy.

- English
- Chinese

You can easily change from one language to another by clicking the language link you can find at the top of every document that has a translation.

 » [API Guides](#) » Fatal Errors [Edit on GitHub](#)

Fatal Errors

[\[中文\]](#)

Overview

In certain situations, execution of the program can not be continued in a well defined way. In ESP-IDF, these situations include:

Index

Symbols

_ESP_LOG_EARLY_ENABLED (*C macro*), 1496

_ip_addr (*C++ struct*), 248

_ip_addr::ip4 (*C++ member*), 248

_ip_addr::ip6 (*C++ member*), 248

_ip_addr::type (*C++ member*), 248

_ip_addr::u_addr (*C++ member*), 248

[anonymous] (*C++ enum*), 601, 1554

[anonymous]::ESP_ERR_FLASH_NO_RESPONSE
(*C++ enumerator*), 602

[anonymous]::ESP_ERR_FLASH_SIZE_NOT_MATCH
(*C++ enumerator*), 601

[anonymous]::ESP_ERR_SLEEP_REJECT
(*C++ enumerator*), 1554

[anonymous]::ESP_ERR_SLEEP_TOO_SHORT_SLEEP_DURATION
(*C++ enumerator*), 1554

A

ana_cmpr_channel_type_t (*C++ enum*), 264

ana_cmpr_channel_type_t::ANA_CMPR_EXT_REF_CHAN
(*C++ enumerator*), 264

ana_cmpr_channel_type_t::ANA_CMPR_SOURCE_CHAN
(*C++ enumerator*), 264

ana_cmpr_clk_src_t (*C++ type*), 263

ana_cmpr_config_t (*C++ struct*), 261

ana_cmpr_config_t::clk_src (*C++ member*),
261

ana_cmpr_config_t::cross_type (*C++
member*), 261

ana_cmpr_config_t::flags (*C++ member*),
262

ana_cmpr_config_t::intr_priority (*C++
member*), 261

ana_cmpr_config_t::io_loop_back (*C++
member*), 261

ana_cmpr_config_t::ref_src (*C++ member*),
261

ana_cmpr_config_t::unit (*C++ member*), 261

ana_cmpr_cross_cb_t (*C++ type*), 263

ana_cmpr_cross_event_data_t (*C++ struct*),
263

ana_cmpr_cross_event_data_t::cross_type
(*C++ member*), 263

ana_cmpr_cross_type_t (*C++ enum*), 264

ana_cmpr_cross_type_t::ANA_CMPR_CROSS_ANY
(*C++ enumerator*), 264

ana_cmpr_cross_type_t::ANA_CMPR_CROSS_DISABLE
(*C++ enumerator*), 264

ana_cmpr_cross_type_t::ANA_CMPR_CROSS_NEG
(*C++ enumerator*), 264

ana_cmpr_cross_type_t::ANA_CMPR_CROSS_POS
(*C++ enumerator*), 264

ana_cmpr_debounce_config_t (*C++ struct*),
262

ana_cmpr_debounce_config_t::wait_us
(*C++ member*), 262

ana_cmpr_del_unit (*C++ function*), 259

ana_cmpr_disable (*C++ function*), 261

ana_cmpr_enable (*C++ function*), 260

ana_cmpr_event_callbacks_t (*C++ struct*),
262

ana_cmpr_event_callbacks_t::on_cross
(*C++ member*), 262

ana_cmpr_get_gpio (*C++ function*), 261

ana_cmpr_handle_t (*C++ type*), 263

ana_cmpr_internal_ref_config_t (*C++
struct*), 262

ana_cmpr_internal_ref_config_t::ref_volt
(*C++ member*), 262

ana_cmpr_new_unit (*C++ function*), 259

ana_cmpr_ref_source_t (*C++ enum*), 263

ana_cmpr_ref_source_t::ANA_CMPR_REF_SRC_EXTERNAL
(*C++ enumerator*), 264

ana_cmpr_ref_source_t::ANA_CMPR_REF_SRC_INTERNAL
(*C++ enumerator*), 263

ana_cmpr_ref_voltage_t (*C++ enum*), 264

ana_cmpr_ref_voltage_t::ANA_CMPR_REF_VOLT_0_PCT_V
(*C++ enumerator*), 264

ana_cmpr_ref_voltage_t::ANA_CMPR_REF_VOLT_10_PCT_V
(*C++ enumerator*), 264

ana_cmpr_ref_voltage_t::ANA_CMPR_REF_VOLT_20_PCT_V
(*C++ enumerator*), 264

ana_cmpr_ref_voltage_t::ANA_CMPR_REF_VOLT_30_PCT_V
(*C++ enumerator*), 264

ana_cmpr_ref_voltage_t::ANA_CMPR_REF_VOLT_40_PCT_V
(*C++ enumerator*), 264

ana_cmpr_ref_voltage_t::ANA_CMPR_REF_VOLT_50_PCT_V
(*C++ enumerator*), 265

ana_cmpr_ref_voltage_t::ANA_CMPR_REF_VOLT_60_PCT_V
(*C++ enumerator*), 265

ana_cmpr_ref_voltage_t::ANA_CMPR_REF_VOLT_70_PCT_V
(*C++ enumerator*), 265

ana_cmpr_register_event_callbacks

- (C++ function), 260
- ana_cmpr_set_cross_type (C++ function), 260
- ana_cmpr_set_debounce (C++ function), 259
- ana_cmpr_set_internal_reference (C++ function), 259
- ANA_CMPR_UNIT_0 (C macro), 263
- ana_cmpr_unit_t (C++ type), 263
- async_memcpy_config_t (C++ struct), 1580
- async_memcpy_config_t::backlog (C++ member), 1580
- async_memcpy_config_t::flags (C++ member), 1580
- async_memcpy_config_t::psram_trans_align (C++ member), 1580
- async_memcpy_config_t::sram_trans_align (C++ member), 1580
- ASYNC_MEMCPY_DEFAULT_CONFIG (C macro), 1580
- async_memcpy_event_t (C++ struct), 1579
- async_memcpy_event_t::data (C++ member), 1580
- async_memcpy_handle_t (C++ type), 1580
- async_memcpy_isr_cb_t (C++ type), 1580
- ## B
- BLE_ADDR_LEN (C macro), 1096
- BLE_UUID128_VAL_LENGTH (C macro), 1095
- bootloader_fill_random (C++ function), 1541
- bootloader_random_disable (C++ function), 1541
- bootloader_random_enable (C++ function), 1541
- bridgeif_config (C++ struct), 241
- bridgeif_config::max_fdb_dyn_entries (C++ member), 241
- bridgeif_config::max_fdb_sta_entries (C++ member), 241
- bridgeif_config::max_ports (C++ member), 241
- bridgeif_config_t (C++ type), 244
- ## C
- CHIP_FEATURE_BLE (C macro), 1507
- CHIP_FEATURE_BT (C macro), 1507
- CHIP_FEATURE_EMB_FLASH (C macro), 1507
- CHIP_FEATURE_EMB_PSRAM (C macro), 1507
- CHIP_FEATURE_IEEE802154 (C macro), 1507
- CHIP_FEATURE_WIFI_BGN (C macro), 1507
- CONFIG_HEAP_TRACING_STACK_DEPTH (C macro), 1468
- ## D
- DEFAULT_HTTP_BUF_SIZE (C macro), 90
- ## E
- EFD_SUPPORT_ISR (C macro), 1200
- efuse_hal_blk_version (C++ function), 1222
- efuse_hal_chip_revision (C++ function), 1222
- efuse_hal_flash_encryption_enabled (C++ function), 1223
- efuse_hal_get_disable_blk_version_major (C++ function), 1223
- efuse_hal_get_disable_wafer_version_major (C++ function), 1223
- efuse_hal_get_mac (C++ function), 1222
- efuse_hal_get_major_chip_version (C++ function), 1223
- efuse_hal_get_minor_chip_version (C++ function), 1223
- efuse_hal_set_ecdsa_key (C++ function), 1223
- emac_rmii_clock_gpio_t (C++ enum), 192
- emac_rmii_clock_gpio_t::EMAC_APPL_CLK_OUT_GPIO (C++ enumerator), 192
- emac_rmii_clock_gpio_t::EMAC_CLK_IN_GPIO (C++ enumerator), 192
- emac_rmii_clock_gpio_t::EMAC_CLK_OUT_180_GPIO (C++ enumerator), 192
- emac_rmii_clock_gpio_t::EMAC_CLK_OUT_GPIO (C++ enumerator), 192
- emac_rmii_clock_mode_t (C++ enum), 191
- emac_rmii_clock_mode_t::EMAC_CLK_DEFAULT (C++ enumerator), 191
- emac_rmii_clock_mode_t::EMAC_CLK_EXT_IN (C++ enumerator), 191
- emac_rmii_clock_mode_t::EMAC_CLK_OUT (C++ enumerator), 192
- eNotifyAction (C++ enum), 1324
- eNotifyAction::eIncrement (C++ enumerator), 1324
- eNotifyAction::eNoAction (C++ enumerator), 1324
- eNotifyAction::eSetBits (C++ enumerator), 1324
- eNotifyAction::eSetValueWithoutOverwrite (C++ enumerator), 1324
- eNotifyAction::eSetValueWithOverwrite (C++ enumerator), 1324
- eSleepModeStatus (C++ enum), 1324
- eSleepModeStatus::eAbortSleep (C++ enumerator), 1325
- eSleepModeStatus::eStandardSleep (C++ enumerator), 1325
- esp_alloc_failed_hook_t (C++ type), 1439
- ESP_APP_DESC_MAGIC_WORD (C macro), 1515
- esp_app_desc_t (C++ struct), 1514
- esp_app_desc_t::app_elf_sha256 (C++ member), 1515
- esp_app_desc_t::date (C++ member), 1515
- esp_app_desc_t::idf_ver (C++ member), 1515
- esp_app_desc_t::magic_word (C++ member), 1515
- esp_app_desc_t::max_efuse_blk_rev_full

- (C++ member), 1515
- esp_app_desc_t::min_efuse_blk_rev_full (C++ member), 1515
- esp_app_desc_t::project_name (C++ member), 1515
- esp_app_desc_t::reserv1 (C++ member), 1515
- esp_app_desc_t::reserv2 (C++ member), 1515
- esp_app_desc_t::secure_version (C++ member), 1515
- esp_app_desc_t::time (C++ member), 1515
- esp_app_desc_t::version (C++ member), 1515
- esp_app_get_description (C++ function), 1514
- esp_app_get_elf_sha256 (C++ function), 1514
- esp_app_get_elf_sha256_str (C++ function), 1514
- esp_apptrace_buffer_get (C++ function), 1213
- esp_apptrace_buffer_put (C++ function), 1213
- esp_apptrace_dest_t (C++ enum), 1216
- esp_apptrace_dest_t::ESP_APPTRACE_DEST_JTAG (C++ enumerator), 1216
- esp_apptrace_dest_t::ESP_APPTRACE_DEST_MAX (C++ enumerator), 1216
- esp_apptrace_dest_t::ESP_APPTRACE_DEST_NUM (C++ enumerator), 1216
- esp_apptrace_dest_t::ESP_APPTRACE_DEST_UART (C++ enumerator), 1216
- esp_apptrace_dest_t::ESP_APPTRACE_DEST_USB (C++ enumerator), 1216
- esp_apptrace_down_buffer_config (C++ function), 1213
- esp_apptrace_down_buffer_get (C++ function), 1214
- esp_apptrace_down_buffer_put (C++ function), 1215
- esp_apptrace_fclose (C++ function), 1215
- esp_apptrace_feof (C++ function), 1216
- esp_apptrace_flush (C++ function), 1214
- esp_apptrace_flush_nolock (C++ function), 1214
- esp_apptrace_fopen (C++ function), 1215
- esp_apptrace_fread (C++ function), 1215
- esp_apptrace_fseek (C++ function), 1216
- esp_apptrace_fstop (C++ function), 1216
- esp_apptrace_ftell (C++ function), 1216
- esp_apptrace_fwrite (C++ function), 1215
- esp_apptrace_host_is_connected (C++ function), 1215
- esp_apptrace_init (C++ function), 1213
- esp_apptrace_read (C++ function), 1214
- esp_apptrace_vprintf (C++ function), 1214
- esp_apptrace_vprintf_to (C++ function), 1214
- esp_apptrace_write (C++ function), 1213
- esp_async_memcpy (C++ function), 1579
- esp_async_memcpy_install (C++ function), 1578
- esp_async_memcpy_install_gdma_ahb (C++ function), 1578
- esp_async_memcpy_install_gdma_axi (C++ function), 1578
- esp_async_memcpy_uninstall (C++ function), 1579
- esp_base_mac_addr_get (C++ function), 1504
- esp_base_mac_addr_set (C++ function), 1504
- ESP_BOOTLOADER_DESC_MAGIC_BYTE (C macro), 1212
- esp_bootloader_desc_t (C++ struct), 1211
- esp_bootloader_desc_t::date_time (C++ member), 1212
- esp_bootloader_desc_t::idf_ver (C++ member), 1212
- esp_bootloader_desc_t::magic_byte (C++ member), 1211
- esp_bootloader_desc_t::reserved (C++ member), 1211
- esp_bootloader_desc_t::reserved2 (C++ member), 1212
- esp_bootloader_desc_t::version (C++ member), 1212
- esp_bootloader_get_description (C++ function), 1211
- esp_cache_msync (C++ function), 1453
- ESP_CACHE_MSINC_FLAG_DIR_C2M (C macro), 1454
- ESP_CACHE_MSINC_FLAG_DIR_M2C (C macro), 1455
- ESP_CACHE_MSINC_FLAG_INVALIDATE (C macro), 1454
- ESP_CACHE_MSINC_FLAG_TYPE_DATA (C macro), 1455
- ESP_CACHE_MSINC_FLAG_TYPE_INST (C macro), 1455
- ESP_CACHE_MSINC_FLAG_UNALIGNED (C macro), 1454
- esp_chip_id_t (C++ enum), 1208
- esp_chip_id_t::ESP_CHIP_ID_ESP32 (C++ enumerator), 1208
- esp_chip_id_t::ESP_CHIP_ID_ESP32C2 (C++ enumerator), 1208
- esp_chip_id_t::ESP_CHIP_ID_ESP32C3 (C++ enumerator), 1208
- esp_chip_id_t::ESP_CHIP_ID_ESP32C6 (C++ enumerator), 1208
- esp_chip_id_t::ESP_CHIP_ID_ESP32H2 (C++ enumerator), 1208
- esp_chip_id_t::ESP_CHIP_ID_ESP32P4 (C++ enumerator), 1208
- esp_chip_id_t::ESP_CHIP_ID_ESP32S2 (C++ enumerator), 1208
- esp_chip_id_t::ESP_CHIP_ID_ESP32S3 (C++ enumerator), 1208

- (C++ *enumerator*), 1208
- esp_chip_id_t::ESP_CHIP_ID_INVALID (C++ *enumerator*), 1208
- esp_chip_info (C++ *function*), 1507
- esp_chip_info_t (C++ *struct*), 1507
- esp_chip_info_t::cores (C++ *member*), 1507
- esp_chip_info_t::features (C++ *member*), 1507
- esp_chip_info_t::model (C++ *member*), 1507
- esp_chip_info_t::revision (C++ *member*), 1507
- esp_chip_model_t (C++ *enum*), 1508
- esp_chip_model_t::CHIP_ESP32 (C++ *enumerator*), 1508
- esp_chip_model_t::CHIP_ESP32C2 (C++ *enumerator*), 1508
- esp_chip_model_t::CHIP_ESP32C3 (C++ *enumerator*), 1508
- esp_chip_model_t::CHIP_ESP32C6 (C++ *enumerator*), 1508
- esp_chip_model_t::CHIP_ESP32H2 (C++ *enumerator*), 1508
- esp_chip_model_t::CHIP_ESP32P4 (C++ *enumerator*), 1508
- esp_chip_model_t::CHIP_ESP32S2 (C++ *enumerator*), 1508
- esp_chip_model_t::CHIP_ESP32S3 (C++ *enumerator*), 1508
- esp_chip_model_t::CHIP_POSIX_LINUX (C++ *enumerator*), 1508
- esp_clk_tree_src_freq_precision_t (C++ *enum*), 278
- esp_clk_tree_src_freq_precision_t::ESP_CLK_TREE_SRC_FREQ_PRECISION_APPROX (C++ *enumerator*), 278
- esp_clk_tree_src_freq_precision_t::ESP_CLK_TREE_SRC_FREQ_PRECISION_CACHED (C++ *enumerator*), 278
- esp_clk_tree_src_freq_precision_t::ESP_CLK_TREE_SRC_FREQ_PRECISION_EXACT (C++ *enumerator*), 278
- esp_clk_tree_src_freq_precision_t::ESP_CLK_TREE_SRC_FREQ_PRECISION_INVALID (C++ *enumerator*), 278
- esp_clk_tree_src_get_freq_hz (C++ *function*), 277
- esp_console_cmd_func_t (C++ *type*), 1231
- esp_console_cmd_register (C++ *function*), 1227
- esp_console_cmd_t (C++ *struct*), 1230
- esp_console_cmd_t::argtable (C++ *member*), 1231
- esp_console_cmd_t::command (C++ *member*), 1231
- esp_console_cmd_t::func (C++ *member*), 1231
- esp_console_cmd_t::help (C++ *member*), 1231
- esp_console_cmd_t::hint (C++ *member*), 1231
- ESP_CONSOLE_CONFIG_DEFAULT (C *macro*), 1231
- esp_console_config_t (C++ *struct*), 1229
- esp_console_config_t::heap_alloc_caps (C++ *member*), 1229
- esp_console_config_t::hint_bold (C++ *member*), 1230
- esp_console_config_t::hint_color (C++ *member*), 1229
- esp_console_config_t::max_cmdline_args (C++ *member*), 1229
- esp_console_config_t::max_cmdline_length (C++ *member*), 1229
- esp_console_deinit (C++ *function*), 1227
- ESP_CONSOLE_DEV_UART_CONFIG_DEFAULT (C *macro*), 1231
- esp_console_dev_uart_config_t (C++ *struct*), 1230
- esp_console_dev_uart_config_t::baud_rate (C++ *member*), 1230
- esp_console_dev_uart_config_t::channel (C++ *member*), 1230
- esp_console_dev_uart_config_t::rx_gpio_num (C++ *member*), 1230
- esp_console_dev_uart_config_t::tx_gpio_num (C++ *member*), 1230
- esp_console_get_completion (C++ *function*), 1228
- esp_console_get_hint (C++ *function*), 1228
- esp_console_init (C++ *function*), 1227
- esp_console_new_repl_uart (C++ *function*), 1228
- esp_console_register_help_command (C++ *function*), 1228
- ESP_CONSOLE_REPL_CONFIG_DEFAULT (C *macro*), 1231
- esp_console_repl_config_t (C++ *struct*), 1230
- esp_console_repl_config_t::history_save_path (C++ *member*), 1230
- esp_console_repl_config_t::max_cmdline_length (C++ *member*), 1230
- esp_console_repl_config_t::max_history_len (C++ *member*), 1230
- esp_console_repl_config_t::prompt (C++ *member*), 1230
- esp_console_repl_config_t::task_priority (C++ *member*), 1230
- esp_console_repl_config_t::task_stack_size (C++ *member*), 1230
- esp_console_repl_s (C++ *struct*), 1231
- esp_console_repl_s::del (C++ *member*), 1231
- esp_console_repl_t (C++ *type*), 1232
- esp_console_run (C++ *function*), 1227
- esp_console_split_argv (C++ *function*), 1227
- esp_console_start_repl (C++ *function*), 1229
- esp_cpu_branch_prediction_enable (C++ *function*), 1512
- esp_cpu_clear_breakpoint (C++ *function*), 1511

- `esp_cpu_clear_watchpoint` (C++ function), 1512
- `esp_cpu_compare_and_set` (C++ function), 1512
- `esp_cpu_configure_region_protection` (C++ function), 1511
- `esp_cpu_cycle_count_t` (C++ type), 1513
- `esp_cpu_dbggr_break` (C++ function), 1512
- `esp_cpu_dbggr_is_attached` (C++ function), 1512
- `esp_cpu_get_call_addr` (C++ function), 1512
- `esp_cpu_get_core_id` (C++ function), 1509
- `esp_cpu_get_cycle_count` (C++ function), 1509
- `esp_cpu_get_sp` (C++ function), 1509
- `ESP_CPU_INTR_DESC_FLAG_RESVD` (C macro), 1513
- `ESP_CPU_INTR_DESC_FLAG_SPECIAL` (C macro), 1513
- `esp_cpu_intr_desc_t` (C++ struct), 1513
- `esp_cpu_intr_desc_t::flags` (C++ member), 1513
- `esp_cpu_intr_desc_t::priority` (C++ member), 1513
- `esp_cpu_intr_desc_t::type` (C++ member), 1513
- `esp_cpu_intr_disable` (C++ function), 1511
- `esp_cpu_intr_edge_ack` (C++ function), 1511
- `esp_cpu_intr_enable` (C++ function), 1511
- `esp_cpu_intr_get_desc` (C++ function), 1509
- `esp_cpu_intr_get_enabled_mask` (C++ function), 1511
- `esp_cpu_intr_get_handler_arg` (C++ function), 1510
- `esp_cpu_intr_get_priority` (C++ function), 1510
- `esp_cpu_intr_get_type` (C++ function), 1510
- `esp_cpu_intr_handler_t` (C++ type), 1513
- `esp_cpu_intr_has_handler` (C++ function), 1510
- `esp_cpu_intr_set_handler` (C++ function), 1510
- `esp_cpu_intr_set_ivt_addr` (C++ function), 1509
- `esp_cpu_intr_set_mtv_t_addr` (C++ function), 1509
- `esp_cpu_intr_set_priority` (C++ function), 1510
- `esp_cpu_intr_set_type` (C++ function), 1510
- `esp_cpu_intr_type_t` (C++ enum), 1513
- `esp_cpu_intr_type_t::ESP_CPU_INTR_TYPE_EDGE` (C++ enumerator), 1513
- `esp_cpu_intr_type_t::ESP_CPU_INTR_TYPE_LEVEL` (C++ enumerator), 1513
- `esp_cpu_intr_type_t::ESP_CPU_INTR_TYPE_BOTH` (C++ enumerator), 1513
- `esp_cpu_pc_to_addr` (C++ function), 1509
- `esp_cpu_reset` (C++ function), 1508
- `esp_cpu_set_breakpoint` (C++ function), 1511
- `esp_cpu_set_cycle_count` (C++ function), 1509
- `esp_cpu_set_watchpoint` (C++ function), 1511
- `esp_cpu_stall` (C++ function), 1508
- `esp_cpu_unstall` (C++ function), 1508
- `esp_cpu_wait_for_intr` (C++ function), 1509
- `esp_cpu_watchpoint_trigger_t` (C++ enum), 1513
- `esp_cpu_watchpoint_trigger_t::ESP_CPU_WATCHPOINT_TRIGGER_DISABLE` (C++ enumerator), 1514
- `esp_cpu_watchpoint_trigger_t::ESP_CPU_WATCHPOINT_TRIGGER_ENABLE` (C++ enumerator), 1514
- `esp_cpu_watchpoint_trigger_t::ESP_CPU_WATCHPOINT_TRIGGER_TOGGLE` (C++ enumerator), 1514
- `esp_crt_bundle_attach` (C++ function), 121
- `esp_crt_bundle_detach` (C++ function), 121
- `esp_crt_bundle_in_use` (C++ function), 121
- `esp_crt_bundle_set` (C++ function), 121
- `esp_deep_sleep` (C++ function), 1550
- `esp_deep_sleep_cb_t` (C++ type), 1552
- `esp_deep_sleep_deregister_hook` (C++ function), 1550
- `esp_deep_sleep_disable_rom_logging` (C++ function), 1551
- `esp_deep_sleep_enable_gpio_wakeup` (C++ function), 1546
- `esp_deep_sleep_register_hook` (C++ function), 1550
- `esp_deep_sleep_start` (C++ function), 1549
- `esp_deep_sleep_try` (C++ function), 1549
- `esp_deep_sleep_try_to_start` (C++ function), 1549
- `esp_deep_sleep_wake_stub_fn_t` (C++ type), 1552
- `esp_deepsleep_gpio_wake_up_mode_t` (C++ enum), 1552
- `esp_deepsleep_gpio_wake_up_mode_t::ESP_GPIO_WAKEUP_MODE_GPIO` (C++ enumerator), 1552
- `esp_deepsleep_gpio_wake_up_mode_t::ESP_GPIO_WAKEUP_MODE_WAKEUP_PIN` (C++ enumerator), 1552
- `esp_default_wake_deep_sleep` (C++ function), 1551
- `esp_deregister_freertos_idle_hook` (C++ function), 1424
- `esp_deregister_freertos_idle_hook_for_cpu` (C++ function), 1424
- `esp_deregister_freertos_tick_hook` (C++ function), 1424
- `esp_deregister_freertos_tick_hook_for_cpu` (C++ function), 1424
- `esp_derive_local_mac` (C++ function), 1505
- `esp_digital_signature_data` (C++ struct), 335
- `esp_digital_signature_data::c` (C++ member), 335
- `esp_digital_signature_data::iv` (C++ member), 335

esp_digital_signature_data::rsa_length (C++ member), 335
 esp_digital_signature_length_t (C++ enum), 336
 esp_digital_signature_length_t::ESP_DS_SIGNATURE_LENGTH_T (C++ enumerator), 336
 esp_digital_signature_length_t::ESP_DS_SIGNATURE_LENGTH_T (C++ enumerator), 336
 esp_digital_signature_length_t::ESP_DS_SIGNATURE_LENGTH_T (C++ enumerator), 336
 esp_digital_signature_length_t::ESP_DS_SIGNATURE_LENGTH_T (C++ enumerator), 336
 esp_digital_signature_length_t::ESP_DS_SIGNATURE_LENGTH_T (C++ enumerator), 336
 esp_dma_buf_location_t (C++ enum), 1456
 esp_dma_buf_location_t::ESP_DMA_BUF_LOCATION_EXTERNAL (C++ enumerator), 1456
 esp_dma_buf_location_t::ESP_DMA_BUF_LOCATION_INTERNAL (C++ enumerator), 1456
 esp_dma_calloc (C++ function), 1455
 esp_dma_is_buffer_aligned (C++ function), 1456
 esp_dma_malloc (C++ function), 1455
 ESP_DMA_MALLOC_FLAG_PSRAM (C macro), 1456
 ESP_DRAM_LOGD (C macro), 1497
 ESP_DRAM_LOGE (C macro), 1496
 ESP_DRAM_LOGI (C macro), 1497
 ESP_DRAM_LOGV (C macro), 1497
 ESP_DRAM_LOGW (C macro), 1497
 ESP_DS_C_LEN (C macro), 336
 esp_ds_context_t (C++ type), 336
 esp_ds_data_t (C++ type), 336
 esp_ds_encrypt_params (C++ function), 334
 esp_ds_finish_sign (C++ function), 334
 esp_ds_is_busy (C++ function), 334
 ESP_DS_IV_BIT_LEN (C macro), 336
 ESP_DS_IV_LEN (C macro), 336
 esp_ds_p_data_t (C++ struct), 335
 esp_ds_p_data_t::length (C++ member), 335
 esp_ds_p_data_t::M (C++ member), 335
 esp_ds_p_data_t::M_prime (C++ member), 335
 esp_ds_p_data_t::Rb (C++ member), 335
 esp_ds_p_data_t::Y (C++ member), 335
 esp_ds_sign (C++ function), 332
 ESP_DS_SIGNATURE_L_BIT_LEN (C macro), 336
 ESP_DS_SIGNATURE_M_PRIME_BIT_LEN (C macro), 336
 ESP_DS_SIGNATURE_MAX_BIT_LEN (C macro), 336
 ESP_DS_SIGNATURE_MD_BIT_LEN (C macro), 336
 ESP_DS_SIGNATURE_PADDING_BIT_LEN (C macro), 336
 esp_ds_start_sign (C++ function), 333
 ESP_EARLY_LOGD (C macro), 1495
 ESP_EARLY_LOGE (C macro), 1495
 ESP_EARLY_LOGI (C macro), 1495
 ESP_EARLY_LOGV (C macro), 1495
 ESP_EARLY_LOGW (C macro), 1495
 esp_ecdsa_load_pubkey (C++ function), 279
 esp_ecdsa_pk_conf_t (C++ struct), 280
 esp_ecdsa_pk_conf_t::efuse_block (C++ member), 281
 esp_ecdsa_pk_conf_t::grp_id (C++ member), 281
 esp_ecdsa_pk_conf_t::load_pubkey (C++ member), 281
 esp_ecdsa_privkey_load_mpi (C++ function), 280
 esp_ecdsa_privkey_load_pk_context (C++ function), 280
 esp_ecdsa_set_pk_context (C++ function), 280
 esp_efuse_batch_write_begin (C++ function), 1255
 esp_efuse_batch_write_cancel (C++ function), 1256
 esp_efuse_batch_write_commit (C++ function), 1256
 esp_efuse_block_is_empty (C++ function), 1256
 esp_efuse_block_t (C++ enum), 1247
 esp_efuse_block_t::EFUSE_BLK0 (C++ enumerator), 1247
 esp_efuse_block_t::EFUSE_BLK1 (C++ enumerator), 1247
 esp_efuse_block_t::EFUSE_BLK10 (C++ enumerator), 1248
 esp_efuse_block_t::EFUSE_BLK2 (C++ enumerator), 1247
 esp_efuse_block_t::EFUSE_BLK3 (C++ enumerator), 1247
 esp_efuse_block_t::EFUSE_BLK4 (C++ enumerator), 1247
 esp_efuse_block_t::EFUSE_BLK5 (C++ enumerator), 1248
 esp_efuse_block_t::EFUSE_BLK6 (C++ enumerator), 1248
 esp_efuse_block_t::EFUSE_BLK7 (C++ enumerator), 1248
 esp_efuse_block_t::EFUSE_BLK8 (C++ enumerator), 1248
 esp_efuse_block_t::EFUSE_BLK9 (C++ enumerator), 1248
 esp_efuse_block_t::EFUSE_BLK_KEY0 (C++ enumerator), 1247
 esp_efuse_block_t::EFUSE_BLK_KEY1 (C++ enumerator), 1248
 esp_efuse_block_t::EFUSE_BLK_KEY2 (C++ enumerator), 1248
 esp_efuse_block_t::EFUSE_BLK_KEY3 (C++ enumerator), 1248
 esp_efuse_block_t::EFUSE_BLK_KEY4 (C++ enumerator), 1248
 esp_efuse_block_t::EFUSE_BLK_KEY5 (C++ enumerator), 1248
 esp_efuse_block_t::EFUSE_BLK_KEY_MAX

- (C++ *enumerator*), 1248
- esp_efuse_block_t::EFUSE_BLK_MAX (C++ *enumerator*), 1248
- esp_efuse_block_t::EFUSE_BLK_SYS_DATA_PART1 (C++ *enumerator*), 1247
- esp_efuse_block_t::EFUSE_BLK_SYS_DATA_PART2 (C++ *enumerator*), 1248
- esp_efuse_block_t::EFUSE_BLK_USER_DATA (C++ *enumerator*), 1247
- esp_efuse_check_errors (C++ *function*), 1260
- esp_efuse_check_secure_version (C++ *function*), 1254
- esp_efuse_coding_scheme_t (C++ *enum*), 1248
- esp_efuse_coding_scheme_t::EFUSE_CODING_SCHEME_NONE (C++ *enumerator*), 1248
- esp_efuse_coding_scheme_t::EFUSE_CODING_SCHEME_CRC8 (C++ *enumerator*), 1248
- esp_efuse_count_unused_key_blocks (C++ *function*), 1258
- esp_efuse_desc_t (C++ *struct*), 1260
- esp_efuse_desc_t::bit_count (C++ *member*), 1260
- esp_efuse_desc_t::bit_start (C++ *member*), 1260
- esp_efuse_desc_t::efuse_block (C++ *member*), 1260
- esp_efuse_disable_rom_download_mode (C++ *function*), 1253
- esp_efuse_enable_rom_secure_download_mode (C++ *function*), 1254
- esp_efuse_find_purpose (C++ *function*), 1257
- esp_efuse_find_unused_key_block (C++ *function*), 1258
- esp_efuse_get_coding_scheme (C++ *function*), 1252
- esp_efuse_get_digest_revoke (C++ *function*), 1258
- esp_efuse_get_field_size (C++ *function*), 1252
- esp_efuse_get_key (C++ *function*), 1257
- esp_efuse_get_key_dis_read (C++ *function*), 1256
- esp_efuse_get_key_dis_write (C++ *function*), 1256
- esp_efuse_get_key_purpose (C++ *function*), 1257
- esp_efuse_get_keypurpose_dis_write (C++ *function*), 1257
- esp_efuse_get_pkg_ver (C++ *function*), 1253
- esp_efuse_get_purpose_field (C++ *function*), 1257
- esp_efuse_get_write_protect_of_digest_revoke (C++ *function*), 1258
- esp_efuse_is_ecdsa_p192_curve_supported (C++ *function*), 1260
- esp_efuse_is_ecdsa_p256_curve_supported (C++ *function*), 1260
- esp_efuse_key_block_unused (C++ *function*), 1257
- esp_efuse_mac_get_custom (C++ *function*),
- esp_efuse_mac_get_default (C++ *function*),
- esp_efuse_purpose_t (C++ *enum*), 1249
- esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_ECDSA (C++ *enumerator*), 1249
- esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_HMAC_D (C++ *enumerator*), 1249
- esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_HMAC_D (C++ *enumerator*), 1249
- esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_HMAC_D (C++ *enumerator*), 1249
- esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_HMAC_U (C++ *enumerator*), 1249
- esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_KM_INI (C++ *enumerator*), 1249
- esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_MAX (C++ *enumerator*), 1249
- esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_SECURE (C++ *enumerator*), 1249
- esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_SECURE (C++ *enumerator*), 1249
- esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_SECURE (C++ *enumerator*), 1249
- esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_USER (C++ *enumerator*), 1249
- esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_XTS_AE (C++ *enumerator*), 1249
- esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_XTS_AE (C++ *enumerator*), 1249
- esp_efuse_purpose_t::ESP_EFUSE_KEY_PURPOSE_XTS_AE (C++ *enumerator*), 1249
- esp_efuse_read_block (C++ *function*), 1253
- esp_efuse_read_field_bit (C++ *function*), 1250
- esp_efuse_read_field_blob (C++ *function*), 1250
- esp_efuse_read_field_cnt (C++ *function*), 1250
- esp_efuse_read_reg (C++ *function*), 1252
- esp_efuse_read_secure_version (C++ *function*), 1254
- esp_efuse_reset (C++ *function*), 1253
- esp_efuse_rom_log_scheme_t (C++ *enum*), 1261
- esp_efuse_rom_log_scheme_t::ESP_EFUSE_ROM_LOG_ALW (C++ *enumerator*), 1261
- esp_efuse_rom_log_scheme_t::ESP_EFUSE_ROM_LOG_ALW (C++ *enumerator*), 1261
- esp_efuse_rom_log_scheme_t::ESP_EFUSE_ROM_LOG_ON (C++ *enumerator*), 1261
- esp_efuse_rom_log_scheme_t::ESP_EFUSE_ROM_LOG_ON (C++ *enumerator*), 1261
- esp_efuse_set_digest_revoke (C++ *function*), 1258

- [esp_efuse_set_key_dis_read \(C++ function\), 1256](#)
[esp_efuse_set_key_dis_write \(C++ function\), 1256](#)
[esp_efuse_set_key_purpose \(C++ function\), 1257](#)
[esp_efuse_set_keypurpose_dis_write \(C++ function\), 1258](#)
[esp_efuse_set_read_protect \(C++ function\), 1252](#)
[esp_efuse_set_rom_log_scheme \(C++ function\), 1254](#)
[esp_efuse_set_write_protect \(C++ function\), 1251](#)
[esp_efuse_set_write_protect_of_digest \(C++ function\), 1258](#)
[esp_efuse_update_secure_version \(C++ function\), 1255](#)
[esp_efuse_write_block \(C++ function\), 1253](#)
[esp_efuse_write_field_bit \(C++ function\), 1251](#)
[esp_efuse_write_field_blob \(C++ function\), 1251](#)
[esp_efuse_write_field_cnt \(C++ function\), 1251](#)
[esp_efuse_write_key \(C++ function\), 1259](#)
[esp_efuse_write_keys \(C++ function\), 1259](#)
[esp_efuse_write_reg \(C++ function\), 1252](#)
[ESP_ERR_CODING \(C macro\), 1261](#)
[ESP_ERR_DAMAGED_READING \(C macro\), 1261](#)
[ESP_ERR_EFUSE \(C macro\), 1261](#)
[ESP_ERR_EFUSE_CNT_IS_FULL \(C macro\), 1261](#)
[ESP_ERR_EFUSE_REPEATED_PROG \(C macro\), 1261](#)
[ESP_ERR_ESP_NETIF_BASE \(C macro\), 243](#)
[ESP_ERR_ESP_NETIF_DHCP_ALREADY_STARTED \(C macro\), 243](#)
[ESP_ERR_ESP_NETIF_DHCP_ALREADY_STOPPED \(C macro\), 243](#)
[ESP_ERR_ESP_NETIF_DHCP_NOT_STOPPED \(C macro\), 243](#)
[ESP_ERR_ESP_NETIF_DHCPC_START_FAILED \(C macro\), 243](#)
[ESP_ERR_ESP_NETIF_DHCPS_START_FAILED \(C macro\), 244](#)
[ESP_ERR_ESP_NETIF_DNS_NOT_CONFIGURED \(C macro\), 244](#)
[ESP_ERR_ESP_NETIF_DRIVER_ATTACH_FAILED \(C macro\), 243](#)
[ESP_ERR_ESP_NETIF_IF_NOT_READY \(C macro\), 243](#)
[ESP_ERR_ESP_NETIF_INIT_FAILED \(C macro\), 243](#)
[ESP_ERR_ESP_NETIF_INVALID_PARAMS \(C macro\), 243](#)
[ESP_ERR_ESP_NETIF_IP6_ADDR_FAILED \(C macro\), 244](#)
[ESP_ERR_ESP_NETIF_MLD6_FAILED \(C macro\), 244](#)
[ESP_ERR_ESP_NETIF_NO_MEM \(C macro\), 243](#)
[ESP_ERR_ESP_TLS_BASE \(C macro\), 74](#)
[ESP_ERR_ESP_TLS_CANNOT_CREATE_SOCKET \(C macro\), 74](#)
[ESP_ERR_ESP_TLS_CANNOT_RESOLVE_HOSTNAME \(C macro\), 74](#)
[ESP_ERR_ESP_TLS_CONNECTION_TIMEOUT \(C macro\), 74](#)
[ESP_ERR_ESP_TLS_FAILED_CONNECT_TO_HOST \(C macro\), 74](#)
[ESP_ERR_ESP_TLS_SE_FAILED \(C macro\), 74](#)
[ESP_ERR_ESP_TLS_SOCKET_SETOPT_FAILED \(C macro\), 74](#)
[ESP_ERR_ESP_TLS_TCP_CLOSED_FIN \(C macro\), 74](#)
[ESP_ERR_ESP_TLS_UNSUPPORTED_PROTOCOL_FAMILY \(C macro\), 74](#)
[ESP_ERR_FLASH_BASE \(C macro\), 1264](#)
[ESP_ERR_FLASH_NOT_INITIALISED \(C macro\), 601](#)
[ESP_ERR_FLASH_OP_FAIL \(C macro\), 594](#)
[ESP_ERR_FLASH_OP_TIMEOUT \(C macro\), 594](#)
[ESP_ERR_FLASH_PROTECTED \(C macro\), 601](#)
[ESP_ERR_FLASH_UNSUPPORTED_CHIP \(C macro\), 601](#)
[ESP_ERR_FLASH_UNSUPPORTED_HOST \(C macro\), 601](#)
[ESP_ERR_HTTP_BASE \(C macro\), 90](#)
[ESP_ERR_HTTP_CONNECT \(C macro\), 90](#)
[ESP_ERR_HTTP_CONNECTING \(C macro\), 90](#)
[ESP_ERR_HTTP_CONNECTION_CLOSED \(C macro\), 91](#)
[ESP_ERR_HTTP_EAGAIN \(C macro\), 90](#)
[ESP_ERR_HTTP_FETCH_HEADER \(C macro\), 90](#)
[ESP_ERR_HTTP_INVALID_TRANSPORT \(C macro\), 90](#)
[ESP_ERR_HTTP_MAX_REDIRECT \(C macro\), 90](#)
[ESP_ERR_HTTP_WRITE_DATA \(C macro\), 90](#)
[ESP_ERR_HTTPD_ALLOC_MEM \(C macro\), 146](#)
[ESP_ERR_HTTPD_BASE \(C macro\), 145](#)
[ESP_ERR_HTTPD_HANDLER_EXISTS \(C macro\), 145](#)
[ESP_ERR_HTTPD_HANDLERS_FULL \(C macro\), 145](#)
[ESP_ERR_HTTPD_INVALID_REQ \(C macro\), 145](#)
[ESP_ERR_HTTPD_RESP_HDR \(C macro\), 145](#)
[ESP_ERR_HTTPD_RESP_SEND \(C macro\), 146](#)
[ESP_ERR_HTTPD_RESULT_TRUNC \(C macro\), 145](#)
[ESP_ERR_HTTPD_TASK \(C macro\), 146](#)
[ESP_ERR_HTTPS_OTA_BASE \(C macro\), 1271](#)
[ESP_ERR_HTTPS_OTA_IN_PROGRESS \(C macro\), 1271](#)
[ESP_ERR_HW_CRYPT0_BASE \(C macro\), 1264](#)
[ESP_ERR_INVALID_ARG \(C macro\), 1263](#)
[ESP_ERR_INVALID_CRC \(C macro\), 1264](#)
[ESP_ERR_INVALID_MAC \(C macro\), 1264](#)
[ESP_ERR_INVALID_RESPONSE \(C macro\), 1264](#)

- ESP_ERR_INVALID_SIZE (*C macro*), 1263
- ESP_ERR_INVALID_STATE (*C macro*), 1263
- ESP_ERR_INVALID_VERSION (*C macro*), 1264
- ESP_ERR_MBEDTLS_CERT_PARTLY_OK (*C macro*), 74
- ESP_ERR_MBEDTLS_CTR_DRBG_SEED_FAILED (*C macro*), 75
- ESP_ERR_MBEDTLS_PK_PARSE_KEY_FAILED (*C macro*), 75
- ESP_ERR_MBEDTLS_SSL_CONF_ALPN_PROTOCOLS_FAILED (*C macro*), 75
- ESP_ERR_MBEDTLS_SSL_CONF_OWN_CERT_FAILED (*C macro*), 75
- ESP_ERR_MBEDTLS_SSL_CONF_PSK_FAILED (*C macro*), 75
- ESP_ERR_MBEDTLS_SSL_CONFIG_DEFAULTS_FAILED (*C macro*), 75
- ESP_ERR_MBEDTLS_SSL_HANDSHAKE_FAILED (*C macro*), 75
- ESP_ERR_MBEDTLS_SSL_SET_HOSTNAME_FAILED (*C macro*), 75
- ESP_ERR_MBEDTLS_SSL_SETUP_FAILED (*C macro*), 75
- ESP_ERR_MBEDTLS_SSL_TICKET_SETUP_FAILED (*C macro*), 75
- ESP_ERR_MBEDTLS_SSL_WRITE_FAILED (*C macro*), 75
- ESP_ERR_MBEDTLS_X509_CRT_PARSE_FAILED (*C macro*), 75
- ESP_ERR_MEMPROT_BASE (*C macro*), 1264
- ESP_ERR_MESH_BASE (*C macro*), 1264
- ESP_ERR_NO_MEM (*C macro*), 1263
- ESP_ERR_NOT_ALLOWED (*C macro*), 1264
- ESP_ERR_NOT_ENOUGH_UNUSED_KEY_BLOCKS (*C macro*), 1261
- ESP_ERR_NOT_FINISHED (*C macro*), 1264
- ESP_ERR_NOT_FOUND (*C macro*), 1263
- ESP_ERR_NOT_SUPPORTED (*C macro*), 1263
- ESP_ERR_NVS_BASE (*C macro*), 1139
- ESP_ERR_NVS_CONTENT_DIFFERS (*C macro*), 1140
- ESP_ERR_NVS_CORRUPT_KEY_PART (*C macro*), 1140
- ESP_ERR_NVS_ENCR_NOT_SUPPORTED (*C macro*), 1140
- ESP_ERR_NVS_INVALID_HANDLE (*C macro*), 1139
- ESP_ERR_NVS_INVALID_LENGTH (*C macro*), 1139
- ESP_ERR_NVS_INVALID_NAME (*C macro*), 1139
- ESP_ERR_NVS_INVALID_STATE (*C macro*), 1139
- ESP_ERR_NVS_KEY_TOO_LONG (*C macro*), 1139
- ESP_ERR_NVS_KEYS_NOT_INITIALIZED (*C macro*), 1140
- ESP_ERR_NVS_NEW_VERSION_FOUND (*C macro*), 1140
- ESP_ERR_NVS_NO_FREE_PAGES (*C macro*), 1139
- ESP_ERR_NVS_NOT_ENOUGH_SPACE (*C macro*), 1139
- ESP_ERR_NVS_NOT_FOUND (*C macro*), 1139
- ESP_ERR_NVS_NOT_INITIALIZED (*C macro*), 1139
- ESP_ERR_NVS_PAGE_FULL (*C macro*), 1139
- ESP_ERR_NVS_PART_NOT_FOUND (*C macro*), 1140
- ESP_ERR_NVS_READ_ONLY (*C macro*), 1139
- ESP_ERR_NVS_REMOVE_FAILED (*C macro*), 1139
- ESP_ERR_NVS_SEC_BASE (*C macro*), 1148
- ESP_ERR_NVS_SEC_HMAC_KEY_BLK_ALREADY_USED (*C macro*), 1148
- ESP_ERR_NVS_SEC_HMAC_KEY_GENERATION_FAILED (*C macro*), 1148
- ESP_ERR_NVS_SEC_HMAC_KEY_NOT_FOUND (*C macro*), 1148
- ESP_ERR_NVS_SEC_HMAC_XTS_KEYS_DERIV_FAILED (*C macro*), 1148
- ESP_ERR_NVS_TYPE_MISMATCH (*C macro*), 1139
- ESP_ERR_NVS_VALUE_TOO_LONG (*C macro*), 1140
- ESP_ERR_NVS_WRONG_ENCRYPTION (*C macro*), 1140
- ESP_ERR_NVS_XTS_CFG_FAILED (*C macro*), 1140
- ESP_ERR_NVS_XTS_CFG_NOT_FOUND (*C macro*), 1140
- ESP_ERR_NVS_XTS_DECR_FAILED (*C macro*), 1140
- ESP_ERR_NVS_XTS_ENCR_FAILED (*C macro*), 1140
- ESP_ERR_OTA_BASE (*C macro*), 1526
- ESP_ERR_OTA_PARTITION_CONFLICT (*C macro*), 1527
- ESP_ERR_OTA_ROLLBACK_FAILED (*C macro*), 1527
- ESP_ERR_OTA_ROLLBACK_INVALID_STATE (*C macro*), 1527
- ESP_ERR_OTA_SELECT_INFO_INVALID (*C macro*), 1527
- ESP_ERR_OTA_SMALL_SEC_VER (*C macro*), 1527
- ESP_ERR_OTA_VALIDATE_FAILED (*C macro*), 1527
- esp_err_t (*C++ type*), 1265
- ESP_ERR_TIMEOUT (*C macro*), 1264
- esp_err_to_name (*C++ function*), 1263
- esp_err_to_name_r (*C++ function*), 1263
- ESP_ERR_WIFI_BASE (*C macro*), 1264
- ESP_ERR_WOLFSSL_CERT_VERIFY_SETUP_FAILED (*C macro*), 75
- ESP_ERR_WOLFSSL_CTX_SETUP_FAILED (*C macro*), 76
- ESP_ERR_WOLFSSL_KEY_VERIFY_SETUP_FAILED (*C macro*), 75
- ESP_ERR_WOLFSSL_SSL_CONF_ALPN_PROTOCOLS_FAILED (*C macro*), 75
- ESP_ERR_WOLFSSL_SSL_HANDSHAKE_FAILED (*C macro*), 76
- ESP_ERR_WOLFSSL_SSL_SET_HOSTNAME_FAILED (*C macro*), 75
- ESP_ERR_WOLFSSL_SSL_SETUP_FAILED (*C macro*), 76
- ESP_ERR_WOLFSSL_SSL_WRITE_FAILED (*C macro*), 76

- macro*), 76
- ESP_ERROR_CHECK (*C macro*), 1264
- ESP_ERROR_CHECK_WITHOUT_ABORT (*C macro*), 1264
- esp_eth_config_t (*C++ struct*), 179
- esp_eth_config_t::check_link_period_ms (*C++ member*), 179
- esp_eth_config_t::mac (*C++ member*), 179
- esp_eth_config_t::on_lowlevel_deinit_done (*C++ member*), 180
- esp_eth_config_t::on_lowlevel_init_done (*C++ member*), 180
- esp_eth_config_t::phy (*C++ member*), 179
- esp_eth_config_t::read_phy_reg (*C++ member*), 180
- esp_eth_config_t::stack_input (*C++ member*), 180
- esp_eth_config_t::write_phy_reg (*C++ member*), 180
- esp_eth_decrease_reference (*C++ function*), 179
- esp_eth_del_netif_glue (*C++ function*), 204
- esp_eth_driver_install (*C++ function*), 176
- esp_eth_driver_uninstall (*C++ function*), 176
- esp_eth_handle_t (*C++ type*), 181
- esp_eth_increase_reference (*C++ function*), 179
- esp_eth_io_cmd_t (*C++ enum*), 181
- esp_eth_io_cmd_t::ETH_CMD_CUSTOM_MAC_CMDS (*C++ enumerator*), 182
- esp_eth_io_cmd_t::ETH_CMD_CUSTOM_PHY_CMDS (*C++ enumerator*), 182
- esp_eth_io_cmd_t::ETH_CMD_G_AUTONEGO (*C++ enumerator*), 181
- esp_eth_io_cmd_t::ETH_CMD_G_DUPLEX_MODE (*C++ enumerator*), 182
- esp_eth_io_cmd_t::ETH_CMD_G_MAC_ADDR (*C++ enumerator*), 181
- esp_eth_io_cmd_t::ETH_CMD_G_PHY_ADDR (*C++ enumerator*), 181
- esp_eth_io_cmd_t::ETH_CMD_G_SPEED (*C++ enumerator*), 182
- esp_eth_io_cmd_t::ETH_CMD_READ_PHY_REG (*C++ enumerator*), 182
- esp_eth_io_cmd_t::ETH_CMD_S_AUTONEGO (*C++ enumerator*), 182
- esp_eth_io_cmd_t::ETH_CMD_S_DUPLEX_MODE (*C++ enumerator*), 182
- esp_eth_io_cmd_t::ETH_CMD_S_FLOW_CTRL (*C++ enumerator*), 182
- esp_eth_io_cmd_t::ETH_CMD_S_MAC_ADDR (*C++ enumerator*), 181
- esp_eth_io_cmd_t::ETH_CMD_S_PHY_ADDR (*C++ enumerator*), 181
- esp_eth_io_cmd_t::ETH_CMD_S_PHY_LOOPBACK (*C++ enumerator*), 182
- esp_eth_io_cmd_t::ETH_CMD_S_PROMISCUOUS (*C++ enumerator*), 182
- (*C++ enumerator*), 182
- esp_eth_io_cmd_t::ETH_CMD_S_SPEED (*C++ enumerator*), 182
- esp_eth_io_cmd_t::ETH_CMD_WRITE_PHY_REG (*C++ enumerator*), 182
- esp_eth_ioctl (*C++ function*), 178
- esp_eth_mac_s (*C++ struct*), 185
- esp_eth_mac_s::custom_ioctl (*C++ member*), 189
- esp_eth_mac_s::deinit (*C++ member*), 186
- esp_eth_mac_s::del (*C++ member*), 189
- esp_eth_mac_s::enable_flow_ctrl (*C++ member*), 189
- esp_eth_mac_s::get_addr (*C++ member*), 188
- esp_eth_mac_s::init (*C++ member*), 185
- esp_eth_mac_s::read_phy_reg (*C++ member*), 187
- esp_eth_mac_s::receive (*C++ member*), 187
- esp_eth_mac_s::set_addr (*C++ member*), 187
- esp_eth_mac_s::set_duplex (*C++ member*), 188
- esp_eth_mac_s::set_link (*C++ member*), 188
- esp_eth_mac_s::set_mediator (*C++ member*), 185
- esp_eth_mac_s::set_peer_pause_ability (*C++ member*), 189
- esp_eth_mac_s::set_promiscuous (*C++ member*), 188
- esp_eth_mac_s::set_speed (*C++ member*), 188
- esp_eth_mac_s::start (*C++ member*), 186
- esp_eth_mac_s::stop (*C++ member*), 186
- esp_eth_mac_s::transmit (*C++ member*), 186
- esp_eth_mac_s::transmit_vargs (*C++ member*), 186
- esp_eth_mac_s::write_phy_reg (*C++ member*), 187
- esp_eth_mac_t (*C++ type*), 191
- esp_eth_mediator_s (*C++ struct*), 183
- esp_eth_mediator_s::on_state_changed (*C++ member*), 183
- esp_eth_mediator_s::phy_reg_read (*C++ member*), 183
- esp_eth_mediator_s::phy_reg_write (*C++ member*), 183
- esp_eth_mediator_s::stack_input (*C++ member*), 183
- esp_eth_mediator_t (*C++ type*), 184
- esp_eth_netif_glue_handle_t (*C++ type*), 204
- esp_eth_new_netif_glue (*C++ function*), 204
- esp_eth_phy_802_3_advertise_pause_ability (*C++ function*), 199
- esp_eth_phy_802_3_autonego_ctrl (*C++ function*), 198
- esp_eth_phy_802_3_basic_phy_deinit (*C++ function*), 200
- esp_eth_phy_802_3_basic_phy_init (*C++*

- function*), 200
- esp_eth_phy_802_3_deinit (C++ *function*), 200
- esp_eth_phy_802_3_del (C++ *function*), 200
- esp_eth_phy_802_3_detect_phy_addr (C++ *function*), 200
- esp_eth_phy_802_3_get_addr (C++ *function*), 199
- esp_eth_phy_802_3_get_mmd_addr (C++ *function*), 201
- esp_eth_phy_802_3_init (C++ *function*), 200
- esp_eth_phy_802_3_loopback (C++ *function*), 199
- esp_eth_phy_802_3_mmd_func_t (C++ *enum*), 204
- esp_eth_phy_802_3_mmd_func_t::MMD_FUNC_ADDRESS (C++ *enumerator*), 204
- esp_eth_phy_802_3_mmd_func_t::MMD_FUNC_DATA_LINK (C++ *enumerator*), 204
- esp_eth_phy_802_3_mmd_func_t::MMD_FUNC_DATA_PHY (C++ *enumerator*), 204
- esp_eth_phy_802_3_mmd_func_t::MMD_FUNC_DATA_WINDOW (C++ *enumerator*), 204
- esp_eth_phy_802_3_obj_config_init (C++ *function*), 203
- esp_eth_phy_802_3_pwrctl (C++ *function*), 198
- esp_eth_phy_802_3_read_manufac_info (C++ *function*), 201
- esp_eth_phy_802_3_read_mmd_data (C++ *function*), 202
- esp_eth_phy_802_3_read_mmd_register (C++ *function*), 202
- esp_eth_phy_802_3_read_oui (C++ *function*), 201
- esp_eth_phy_802_3_reset (C++ *function*), 198
- esp_eth_phy_802_3_reset_hw (C++ *function*), 200
- esp_eth_phy_802_3_set_addr (C++ *function*), 198
- esp_eth_phy_802_3_set_duplex (C++ *function*), 199
- esp_eth_phy_802_3_set_link (C++ *function*), 199
- esp_eth_phy_802_3_set_mediator (C++ *function*), 198
- esp_eth_phy_802_3_set_mmd_addr (C++ *function*), 201
- esp_eth_phy_802_3_set_speed (C++ *function*), 199
- esp_eth_phy_802_3_write_mmd_data (C++ *function*), 202
- esp_eth_phy_802_3_write_mmd_register (C++ *function*), 202
- ESP_ETH_PHY_ADDR_AUTO (C *macro*), 197
- esp_eth_phy_into_phy_802_3 (C++ *function*), 203
- esp_eth_phy_new_dp83848 (C++ *function*), 193
- esp_eth_phy_new_ip101 (C++ *function*), 193
- esp_eth_phy_new_ksz80xx (C++ *function*), 193
- esp_eth_phy_new_lan87xx (C++ *function*), 193
- esp_eth_phy_new_rtl8201 (C++ *function*), 193
- esp_eth_phy_reg_rw_data_t (C++ *struct*), 181
- esp_eth_phy_reg_rw_data_t::reg_addr (C++ *member*), 181
- esp_eth_phy_reg_rw_data_t::reg_value_p (C++ *member*), 181
- esp_eth_phy_s (C++ *struct*), 193
- esp_eth_phy_s::advertise_pause_ability (C++ *member*), 195
- esp_eth_phy_s::autonego_ctrl (C++ *member*), 194
- esp_eth_phy_s::custom_ioctl (C++ *member*), 196
- esp_eth_phy_s::deinit (C++ *member*), 194
- esp_eth_phy_s::del (C++ *member*), 196
- esp_eth_phy_s::get_addr (C++ *member*), 195
- esp_eth_phy_s::get_link (C++ *member*), 194
- esp_eth_phy_s::init (C++ *member*), 194
- esp_eth_phy_s::loopback (C++ *member*), 195
- esp_eth_phy_s::pwrctl (C++ *member*), 195
- esp_eth_phy_s::reset (C++ *member*), 194
- esp_eth_phy_s::reset_hw (C++ *member*), 194
- esp_eth_phy_s::set_addr (C++ *member*), 195
- esp_eth_phy_s::set_duplex (C++ *member*), 196
- esp_eth_phy_s::set_link (C++ *member*), 195
- esp_eth_phy_s::set_mediator (C++ *member*), 193
- esp_eth_phy_s::set_speed (C++ *member*), 196
- esp_eth_phy_t (C++ *type*), 197
- esp_eth_start (C++ *function*), 177
- esp_eth_state_t (C++ *enum*), 184
- esp_eth_state_t::ETH_STATE_DEINIT (C++ *enumerator*), 184
- esp_eth_state_t::ETH_STATE_DUPLEX (C++ *enumerator*), 184
- esp_eth_state_t::ETH_STATE_LINK (C++ *enumerator*), 184
- esp_eth_state_t::ETH_STATE_LLINIT (C++ *enumerator*), 184
- esp_eth_state_t::ETH_STATE_PAUSE (C++ *enumerator*), 184
- esp_eth_state_t::ETH_STATE_SPEED (C++ *enumerator*), 184
- esp_eth_stop (C++ *function*), 177
- esp_eth_transmit (C++ *function*), 177
- esp_eth_transmit_vargs (C++ *function*), 178
- esp_eth_update_input_path (C++ *function*), 177
- esp_etm_channel_config_t (C++ *struct*), 286
- esp_etm_channel_connect (C++ *function*), 285
- esp_etm_channel_disable (C++ *function*), 285
- esp_etm_channel_enable (C++ *function*), 284
- esp_etm_channel_handle_t (C++ *type*), 286

- `esp_etm_del_channel` (C++ function), 284
- `esp_etm_del_event` (C++ function), 285
- `esp_etm_del_task` (C++ function), 285
- `esp_etm_dump` (C++ function), 286
- `esp_etm_event_handle_t` (C++ type), 286
- `esp_etm_new_channel` (C++ function), 284
- `esp_etm_task_handle_t` (C++ type), 286
- `ESP_EVENT_ANY_BASE` (C macro), 1284
- `ESP_EVENT_ANY_ID` (C macro), 1285
- `ESP_EVENT_DECLARE_BASE` (C macro), 1284
- `ESP_EVENT_DEFINE_BASE` (C macro), 1284
- `esp_event_dump` (C++ function), 1283
- `esp_event_handler_instance_register` (C++ function), 1279
- `esp_event_handler_instance_register_with` (C++ function), 1278
- `esp_event_handler_instance_t` (C++ type), 1285
- `esp_event_handler_instance_unregister` (C++ function), 1281
- `esp_event_handler_instance_unregister_with` (C++ function), 1280
- `esp_event_handler_register` (C++ function), 1277
- `esp_event_handler_register_with` (C++ function), 1278
- `esp_event_handler_t` (C++ type), 1285
- `esp_event_handler_unregister` (C++ function), 1280
- `esp_event_handler_unregister_with` (C++ function), 1280
- `esp_event_isr_post` (C++ function), 1282
- `esp_event_isr_post_to` (C++ function), 1282
- `esp_event_loop_args_t` (C++ struct), 1284
- `esp_event_loop_args_t::queue_size` (C++ member), 1284
- `esp_event_loop_args_t::task_core_id` (C++ member), 1284
- `esp_event_loop_args_t::task_name` (C++ member), 1284
- `esp_event_loop_args_t::task_priority` (C++ member), 1284
- `esp_event_loop_args_t::task_stack_size` (C++ member), 1284
- `esp_event_loop_create` (C++ function), 1276
- `esp_event_loop_create_default` (C++ function), 1276
- `esp_event_loop_delete` (C++ function), 1276
- `esp_event_loop_delete_default` (C++ function), 1277
- `esp_event_loop_handle_t` (C++ type), 1285
- `esp_event_loop_run` (C++ function), 1277
- `esp_event_post` (C++ function), 1281
- `esp_event_post_to` (C++ function), 1281
- `ESP_EXECUTE_EXPRESSION_WITH_STACK` (C macro), 1219
- `esp_execute_shared_stack_function` (C++ function), 1218
- `ESP_FAIL` (C macro), 1263
- `esp_fill_random` (C++ function), 1540
- `esp_flash_chip_driver_initialized` (C++ function), 585
- `esp_flash_counter_t` (C++ struct), 602
- `esp_flash_counter_t::bytes` (C++ member), 602
- `esp_flash_counter_t::count` (C++ member), 602
- `esp_flash_counter_t::time` (C++ member), 602
- `esp_flash_counters_t` (C++ struct), 602
- `esp_flash_counters_t::erase` (C++ member), 603
- `esp_flash_counters_t::read` (C++ member), 603
- `esp_flash_counters_t::write` (C++ member), 603
- `esp_flash_dump_counters` (C++ function), 602
- `esp_flash_enc_mode_t` (C++ enum), 605
- `esp_flash_enc_mode_t::ESP_FLASH_ENC_MODE_DEVELOPMENT` (C++ enumerator), 605
- `esp_flash_enc_mode_t::ESP_FLASH_ENC_MODE_DISABLED` (C++ enumerator), 605
- `esp_flash_enc_mode_t::ESP_FLASH_ENC_MODE_RELEASE` (C++ enumerator), 605
- `esp_flash_encrypt_check_and_update` (C++ function), 603
- `esp_flash_encrypt_contents` (C++ function), 604
- `esp_flash_encrypt_enable` (C++ function), 604
- `esp_flash_encrypt_init` (C++ function), 604
- `esp_flash_encrypt_initialized_once` (C++ function), 603
- `esp_flash_encrypt_is_write_protected` (C++ function), 604
- `esp_flash_encrypt_region` (C++ function), 604
- `esp_flash_encrypt_state` (C++ function), 603
- `esp_flash_encryption_cfg_verify_release_mode` (C++ function), 605
- `esp_flash_encryption_enable_secure_features` (C++ function), 605
- `esp_flash_encryption_enabled` (C++ function), 603
- `esp_flash_encryption_init_checks` (C++ function), 604
- `esp_flash_encryption_set_release_mode` (C++ function), 605
- `esp_flash_erase_chip` (C++ function), 586
- `esp_flash_erase_region` (C++ function), 586
- `esp_flash_get_chip_write_protect` (C++ function), 587
- `esp_flash_get_counters` (C++ function), 602
- `esp_flash_get_physical_size` (C++ function), 586
- `esp_flash_get_protectable_regions`

- (C++ function), 587
- esp_flash_get_protected_region (C++ function), 588
- esp_flash_get_size (C++ function), 585
- esp_flash_init (C++ function), 585
- esp_flash_io_mode_t (C++ enum), 600
- esp_flash_io_mode_t::SPI_FLASH_DIO (C++ enumerator), 601
- esp_flash_io_mode_t::SPI_FLASH_DOUT (C++ enumerator), 601
- esp_flash_io_mode_t::SPI_FLASH_FASTRD (C++ enumerator), 600
- esp_flash_io_mode_t::SPI_FLASH_OPI_DTR (C++ enumerator), 601
- esp_flash_io_mode_t::SPI_FLASH_OPI_STR (C++ enumerator), 601
- esp_flash_io_mode_t::SPI_FLASH_QIO (C++ enumerator), 601
- esp_flash_io_mode_t::SPI_FLASH_QOUT (C++ enumerator), 601
- esp_flash_io_mode_t::SPI_FLASH_READ_MODE_MAX (C++ enumerator), 601
- esp_flash_io_mode_t::SPI_FLASH_SLOWRD (C++ enumerator), 600
- esp_flash_is_quad_mode (C++ function), 590
- esp_flash_os_functions_t (C++ struct), 590
- esp_flash_os_functions_t::check_yield (C++ member), 591
- esp_flash_os_functions_t::delay_us (C++ member), 591
- esp_flash_os_functions_t::end (C++ member), 591
- esp_flash_os_functions_t::get_system_time (C++ member), 591
- esp_flash_os_functions_t::get_temp_buffer (C++ member), 591
- esp_flash_os_functions_t::region_protected (C++ member), 591
- esp_flash_os_functions_t::release_temp_buffer (C++ member), 591
- esp_flash_os_functions_t::set_flash_op_status (C++ member), 591
- esp_flash_os_functions_t::start (C++ member), 591
- esp_flash_os_functions_t::yield (C++ member), 591
- esp_flash_read (C++ function), 588
- esp_flash_read_encrypted (C++ function), 590
- esp_flash_read_id (C++ function), 585
- esp_flash_read_unique_chip_id (C++ function), 586
- esp_flash_region_t (C++ struct), 590
- esp_flash_region_t::offset (C++ member), 590
- esp_flash_region_t::size (C++ member), 590
- esp_flash_reset_counters (C++ function), 602
- esp_flash_set_chip_write_protect (C++ function), 587
- esp_flash_set_protected_region (C++ function), 588
- esp_flash_speed_s (C++ enum), 600
- esp_flash_speed_s::ESP_FLASH_10MHZ (C++ enumerator), 600
- esp_flash_speed_s::ESP_FLASH_120MHZ (C++ enumerator), 600
- esp_flash_speed_s::ESP_FLASH_20MHZ (C++ enumerator), 600
- esp_flash_speed_s::ESP_FLASH_26MHZ (C++ enumerator), 600
- esp_flash_speed_s::ESP_FLASH_40MHZ (C++ enumerator), 600
- esp_flash_speed_s::ESP_FLASH_5MHZ (C++ enumerator), 600
- esp_flash_speed_s::ESP_FLASH_80MHZ (C++ enumerator), 600
- esp_flash_speed_s::ESP_FLASH_SPEED_MAX (C++ enumerator), 600
- esp_flash_speed_t (C++ type), 600
- esp_flash_spi_device_config_t (C++ struct), 584
- esp_flash_spi_device_config_t::cs_id (C++ member), 584
- esp_flash_spi_device_config_t::cs_io_num (C++ member), 584
- esp_flash_spi_device_config_t::freq_mhz (C++ member), 584
- esp_flash_spi_device_config_t::host_id (C++ member), 584
- esp_flash_spi_device_config_t::input_delay_ns (C++ member), 584
- esp_flash_spi_device_config_t::io_mode (C++ member), 584
- esp_flash_spi_device_config_t::speed (C++ member), 584
- esp_flash_t (C++ struct), 591
- esp_flash_t::busy (C++ member), 592
- esp_flash_t::chip_drv (C++ member), 591
- esp_flash_t::chip_id (C++ member), 592
- esp_flash_t::host (C++ member), 591
- esp_flash_t::hpm_dummy_ena (C++ member), 592
- esp_flash_t::os_func (C++ member), 592
- esp_flash_t::os_func_data (C++ member), 592
- esp_flash_t::read_mode (C++ member), 592
- esp_flash_t::reserved_flags (C++ member), 592
- esp_flash_t::size (C++ member), 592
- esp_flash_write (C++ function), 589
- esp_flash_write_encrypted (C++ function), 589
- esp_flash_write_protect_crypt_cnt (C++ function), 604

- `esp_freertos_idle_cb_t` (C++ type), 1424
- `esp_freertos_tick_cb_t` (C++ type), 1424
- `esp_gcov_dump` (C++ function), 1216
- `esp_get_deep_sleep_wake_stub` (C++ function), 1551
- `esp_get_flash_encryption_mode` (C++ function), 604
- `esp_get_free_heap_size` (C++ function), 1501
- `esp_get_free_internal_heap_size` (C++ function), 1501
- `esp_get_idf_version` (C++ function), 1503
- `esp_get_minimum_free_heap_size` (C++ function), 1502
- `ESP_GOTO_ON_ERROR` (C macro), 1262
- `ESP_GOTO_ON_ERROR_ISR` (C macro), 1262
- `ESP_GOTO_ON_FALSE` (C macro), 1262
- `ESP_GOTO_ON_FALSE_ISR` (C macro), 1262
- `esp_hmac_calculate` (C++ function), 329
- `esp_hmac_jtag_disable` (C++ function), 330
- `esp_hmac_jtag_enable` (C++ function), 330
- `esp_http_client_add_auth` (C++ function), 85
- `esp_http_client_addr_type_t` (C++ enum), 94
- `esp_http_client_addr_type_t::HTTP_ADDR_TYPE_INET` (C++ enumerator), 94
- `esp_http_client_addr_type_t::HTTP_ADDR_TYPE_INET6` (C++ enumerator), 94
- `esp_http_client_addr_type_t::HTTP_ADDR_TYPE_UNSPECIFIED` (C++ enumerator), 94
- `esp_http_client_auth_type_t` (C++ enum), 93
- `esp_http_client_auth_type_t::HTTP_AUTH_TYPE_BASIC` (C++ enumerator), 93
- `esp_http_client_auth_type_t::HTTP_AUTH_TYPE_DIGEST` (C++ enumerator), 93
- `esp_http_client_auth_type_t::HTTP_AUTH_TYPE_NONE` (C++ enumerator), 93
- `esp_http_client_cancel_request` (C++ function), 80
- `esp_http_client_cleanup` (C++ function), 84
- `esp_http_client_close` (C++ function), 84
- `esp_http_client_config_t` (C++ struct), 87
- `esp_http_client_config_t::addr_type` (C++ member), 90
- `esp_http_client_config_t::auth_type` (C++ member), 88
- `esp_http_client_config_t::buffer_size` (C++ member), 89
- `esp_http_client_config_t::buffer_size_bytes` (C++ member), 89
- `esp_http_client_config_t::cert_len` (C++ member), 88
- `esp_http_client_config_t::cert_pem` (C++ member), 88
- `esp_http_client_config_t::client_cert_len` (C++ member), 88
- `esp_http_client_config_t::client_cert_pem` (C++ member), 88
- `esp_http_client_config_t::client_key_len` (C++ member), 88
- `esp_http_client_config_t::client_key_password` (C++ member), 88
- `esp_http_client_config_t::client_key_password_len` (C++ member), 88
- `esp_http_client_config_t::client_key_pem` (C++ member), 88
- `esp_http_client_config_t::common_name` (C++ member), 89
- `esp_http_client_config_t::crt_bundle_attach` (C++ member), 89
- `esp_http_client_config_t::disable_auto_redirect` (C++ member), 89
- `esp_http_client_config_t::ds_data` (C++ member), 90
- `esp_http_client_config_t::event_handler` (C++ member), 89
- `esp_http_client_config_t::host` (C++ member), 88
- `esp_http_client_config_t::if_name` (C++ member), 90
- `esp_http_client_config_t::is_async` (C++ member), 89
- `esp_http_client_config_t::keep_alive_count` (C++ member), 90
- `esp_http_client_config_t::keep_alive_enable` (C++ member), 90
- `esp_http_client_config_t::keep_alive_idle` (C++ member), 90
- `esp_http_client_config_t::keep_alive_interval` (C++ member), 90
- `esp_http_client_config_t::max_authorization_retries` (C++ member), 89
- `esp_http_client_config_t::max_redirection_count` (C++ member), 89
- `esp_http_client_config_t::method` (C++ member), 89
- `esp_http_client_config_t::password` (C++ member), 88
- `esp_http_client_config_t::path` (C++ member), 88
- `esp_http_client_config_t::port` (C++ member), 88
- `esp_http_client_config_t::query` (C++ member), 88
- `esp_http_client_config_t::skip_cert_common_name_check` (C++ member), 89
- `esp_http_client_config_t::timeout_ms` (C++ member), 89
- `esp_http_client_config_t::tls_version` (C++ member), 88
- `esp_http_client_config_t::transport_type` (C++ member), 89
- `esp_http_client_config_t::url` (C++ member), 88
- `esp_http_client_config_t::use_global_ca_store` (C++ member), 89

esp_http_client_config_t::user_agent (C++ member), 89
 esp_http_client_config_t::user_data (C++ member), 89
 esp_http_client_config_t::username (C++ member), 88
 esp_http_client_delete_header (C++ function), 83
 esp_http_client_event (C++ struct), 86
 esp_http_client_event::client (C++ member), 87
 esp_http_client_event::data (C++ member), 87
 esp_http_client_event::data_len (C++ member), 87
 esp_http_client_event::event_id (C++ member), 87
 esp_http_client_event::header_key (C++ member), 87
 esp_http_client_event::header_value (C++ member), 87
 esp_http_client_event::user_data (C++ member), 87
 esp_http_client_event_handle_t (C++ type), 91
 esp_http_client_event_id_t (C++ enum), 91
 esp_http_client_event_id_t::HTTP_EVENT_DISCONNECTED (C++ enumerator), 91
 esp_http_client_event_id_t::HTTP_EVENT_ERROR (C++ enumerator), 91
 esp_http_client_event_id_t::HTTP_EVENT_HEADER_SENT (C++ enumerator), 91
 esp_http_client_event_id_t::HTTP_EVENT_HEADERS_SENT (C++ enumerator), 91
 esp_http_client_event_id_t::HTTP_EVENT_ON_CONNECTED (C++ enumerator), 91
 esp_http_client_event_id_t::HTTP_EVENT_ON_DATA (C++ enumerator), 91
 esp_http_client_event_id_t::HTTP_EVENT_ON_FINISH (C++ enumerator), 91
 esp_http_client_event_id_t::HTTP_EVENT_ON_HEADERS (C++ enumerator), 91
 esp_http_client_event_id_t::HTTP_EVENT_REDIRECT (C++ enumerator), 92
 esp_http_client_event_t (C++ type), 91
 esp_http_client_fetch_headers (C++ function), 84
 esp_http_client_flush_response (C++ function), 86
 esp_http_client_get_chunk_length (C++ function), 86
 esp_http_client_get_content_length (C++ function), 84
 esp_http_client_get_errno (C++ function), 83
 esp_http_client_get_header (C++ function), 81
 esp_http_client_get_password (C++ function), 82
 esp_http_client_get_post_field (C++ function), 81
 esp_http_client_get_status_code (C++ function), 84
 esp_http_client_get_transport_type (C++ function), 85
 esp_http_client_get_url (C++ function), 86
 esp_http_client_get_user_data (C++ function), 82
 esp_http_client_get_username (C++ function), 81
 esp_http_client_handle_t (C++ type), 91
 esp_http_client_init (C++ function), 80
 esp_http_client_is_chunked_response (C++ function), 84
 esp_http_client_is_complete_data_received (C++ function), 86
 esp_http_client_method_t (C++ enum), 92
 esp_http_client_method_t::HTTP_METHOD_COPY (C++ enumerator), 93
 esp_http_client_method_t::HTTP_METHOD_DELETE (C++ enumerator), 92
 esp_http_client_method_t::HTTP_METHOD_GET (C++ enumerator), 92
 esp_http_client_method_t::HTTP_METHOD_HEAD (C++ enumerator), 92
 esp_http_client_method_t::HTTP_METHOD_LOCK (C++ enumerator), 93
 esp_http_client_method_t::HTTP_METHOD_MAX (C++ enumerator), 93
 esp_http_client_method_t::HTTP_METHOD_MKCOL (C++ enumerator), 93
 esp_http_client_method_t::HTTP_METHOD_MOVE (C++ enumerator), 93
 esp_http_client_method_t::HTTP_METHOD_NOTIFY (C++ enumerator), 93
 esp_http_client_method_t::HTTP_METHOD_OPTIONS (C++ enumerator), 93
 esp_http_client_method_t::HTTP_METHOD_PATCH (C++ enumerator), 92
 esp_http_client_method_t::HTTP_METHOD_POST (C++ enumerator), 92
 esp_http_client_method_t::HTTP_METHOD_PROPFIND (C++ enumerator), 93
 esp_http_client_method_t::HTTP_METHOD_PROPPATCH (C++ enumerator), 93
 esp_http_client_method_t::HTTP_METHOD_PUT (C++ enumerator), 92
 esp_http_client_method_t::HTTP_METHOD_SUBSCRIBE (C++ enumerator), 93
 esp_http_client_method_t::HTTP_METHOD_UNLOCK (C++ enumerator), 93
 esp_http_client_method_t::HTTP_METHOD_UNSUBSCRIBE (C++ enumerator), 93
 esp_http_client_on_data (C++ struct), 87
 esp_http_client_on_data::client (C++ member), 87

- esp_http_client_on_data::data_process (C++ member), 87
 esp_http_client_on_data_t (C++ type), 91
 esp_http_client_open (C++ function), 83
 esp_http_client_perform (C++ function), 80
 esp_http_client_proto_ver_t (C++ enum), 92
 esp_http_client_proto_ver_t::ESP_HTTP_CLIENT_148_VER_ANY (C++ enumerator), 92
 esp_http_client_proto_ver_t::ESP_HTTP_CLIENT_148_VER_ANY (C++ enumerator), 150
 esp_http_client_proto_ver_t::ESP_HTTP_CLIENT_148_VER_ANY (C++ enumerator), 149
 esp_http_client_proto_ver_t::ESP_HTTP_CLIENT_148_VER_ANY (C++ enumerator), 139
 esp_http_client_read (C++ function), 84
 esp_http_client_read_response (C++ function), 86
 esp_http_client_redirect_event_data (C++ struct), 87
 esp_http_client_redirect_event_data::client (C++ member), 87
 esp_http_client_redirect_event_data::status (C++ member), 87
 esp_http_client_redirect_event_data_t (C++ type), 91
 esp_http_client_reset_redirect_counter (C++ function), 85
 esp_http_client_set_auth_data (C++ function), 85
 esp_http_client_set_authtype (C++ function), 82
 esp_http_client_set_header (C++ function), 81
 esp_http_client_set_method (C++ function), 83
 esp_http_client_set_password (C++ function), 82
 esp_http_client_set_post_field (C++ function), 81
 esp_http_client_set_redirection (C++ function), 85
 esp_http_client_set_timeout_ms (C++ function), 83
 esp_http_client_set_url (C++ function), 80
 esp_http_client_set_user_data (C++ function), 82
 esp_http_client_set_username (C++ function), 82
 esp_http_client_transport_t (C++ enum), 92
 esp_http_client_transport_t::HTTP_TRANSPORT_OVER_TCP (C++ enumerator), 92
 esp_http_client_transport_t::HTTP_TRANSPORT_OVER_TCP (C++ enumerator), 92
 esp_http_client_transport_t::HTTP_TRANSPORT_UNKNOWN (C++ enumerator), 92
 esp_http_client_write (C++ function), 83
 esp_http_server_event_data (C++ struct), 140
 esp_http_server_event_data::data_len (C++ member), 141
 esp_http_server_event_data::fd (C++ member), 140
 esp_http_server_event_id_t (C++ enum),
 esp_http_server_event_id_t::HTTP_SERVER_EVENT_DISCONNECT (C++ enumerator), 150
 esp_http_server_event_id_t::HTTP_SERVER_EVENT_ERROR (C++ enumerator), 149
 esp_http_server_event_id_t::HTTP_SERVER_EVENT_HEADER (C++ enumerator), 149
 esp_http_server_event_id_t::HTTP_SERVER_EVENT_ON_DATA (C++ enumerator), 149
 esp_http_server_event_id_t::HTTP_SERVER_EVENT_ON_DATA (C++ enumerator), 150
 esp_http_server_event_id_t::HTTP_SERVER_EVENT_ON_DATA (C++ enumerator), 149
 esp_http_server_event_id_t::HTTP_SERVER_EVENT_SEND (C++ enumerator), 150
 esp_http_server_event_id_t::HTTP_SERVER_EVENT_STATUS (C++ enumerator), 149
 esp_http_server_event_id_t::HTTP_SERVER_EVENT_STOP (C++ enumerator), 150
 ESP_HTTPD_DEF_CTRL_PORT (C macro), 145
 esp_https_ota (C++ function), 1267
 esp_https_ota_abort (C++ function), 1269
 esp_https_ota_begin (C++ function), 1268
 esp_https_ota_config_t (C++ struct), 1270
 esp_https_ota_config_t::bulk_flash_erase (C++ member), 1271
 esp_https_ota_config_t::http_client_init_cb (C++ member), 1271
 esp_https_ota_config_t::http_config (C++ member), 1271
 esp_https_ota_config_t::max_http_request_size (C++ member), 1271
 esp_https_ota_config_t::partial_http_download (C++ member), 1271
 esp_https_ota_event_t (C++ enum), 1271
 esp_https_ota_event_t::ESP_HTTPS_OTA_ABORT (C++ enumerator), 1272
 esp_https_ota_event_t::ESP_HTTPS_OTA_CONNECTED (C++ enumerator), 1271
 esp_https_ota_event_t::ESP_HTTPS_OTA_DECRYPT_CB (C++ enumerator), 1271
 esp_https_ota_event_t::ESP_HTTPS_OTA_FINISH (C++ enumerator), 1272
 esp_https_ota_event_t::ESP_HTTPS_OTA_GET_IMG_DESC (C++ enumerator), 1271
 esp_https_ota_event_t::ESP_HTTPS_OTA_START (C++ enumerator), 1271
 esp_https_ota_event_t::ESP_HTTPS_OTA_UPDATE_BOOT (C++ enumerator), 1272
 esp_https_ota_event_t::ESP_HTTPS_OTA_VERIFY_CHIP (C++ enumerator), 1271

esp_https_ota_event_t::ESP_HTTPS_OTA_WRITE_FLAG (C++ enumerator), 1272
 esp_https_ota_finish (C++ function), 1269
 esp_https_ota_get_image_len_read (C++ function), 1270
 esp_https_ota_get_image_size (C++ function), 1270
 esp_https_ota_get_img_desc (C++ function), 1269
 esp_https_ota_get_status_code (C++ function), 1270
 esp_https_ota_handle_t (C++ type), 1271
 esp_https_ota_is_complete_data_received (C++ function), 1269
 esp_https_ota_perform (C++ function), 1268
 esp_https_server_user_cb (C++ type), 153
 esp_https_server_user_cb_arg (C++ struct), 151
 esp_https_server_user_cb_arg::tls (C++ member), 152
 esp_https_server_user_cb_arg::user_cb_state (C++ member), 152
 esp_https_server_user_cb_arg_t (C++ type), 153
 ESP_IDF_VERSION (C macro), 1504
 ESP_IDF_VERSION_MAJOR (C macro), 1503
 ESP_IDF_VERSION_MINOR (C macro), 1503
 ESP_IDF_VERSION_PATCH (C macro), 1503
 ESP_IDF_VERSION_VAL (C macro), 1504
 esp_iface_mac_addr_set (C++ function), 1505
 esp_image_flash_size_t (C++ enum), 1209
 esp_image_flash_size_t::ESP_IMAGE_FLASH_SIZE_128KB (C++ enumerator), 1209
 esp_image_flash_size_t::ESP_IMAGE_FLASH_SIZE_16MB (C++ enumerator), 1209
 esp_image_flash_size_t::ESP_IMAGE_FLASH_SIZE_1MB (C++ enumerator), 1209
 esp_image_flash_size_t::ESP_IMAGE_FLASH_SIZE_256KB (C++ enumerator), 1209
 esp_image_flash_size_t::ESP_IMAGE_FLASH_SIZE_4KB (C++ enumerator), 1209
 esp_image_flash_size_t::ESP_IMAGE_FLASH_SIZE_4MB (C++ enumerator), 1209
 esp_image_flash_size_t::ESP_IMAGE_FLASH_SIZE_64KB (C++ enumerator), 1209
 esp_image_flash_size_t::ESP_IMAGE_FLASH_SIZE_8KB (C++ enumerator), 1209
 esp_image_flash_size_t::ESP_IMAGE_FLASH_SIZE_8MB (C++ enumerator), 1209
 ESP_IMAGE_HEADER_MAGIC (C macro), 1207
 esp_image_header_t (C++ struct), 1206
 esp_image_header_t::chip_id (C++ member), 1207
 esp_image_header_t::entry_addr (C++ member), 1206
 esp_image_header_t::hash_appended (C++ member), 1207
 esp_image_header_t::magic (C++ member), 1206
 esp_image_header_t::max_chip_rev_full (C++ member), 1207
 esp_image_header_t::min_chip_rev (C++ member), 1207
 esp_image_header_t::min_chip_rev_full (C++ member), 1207
 esp_image_header_t::reserved (C++ member), 1207
 esp_image_header_t::segment_count (C++ member), 1206
 esp_image_header_t::spi_mode (C++ member), 1206
 esp_image_header_t::spi_pin_drv (C++ member), 1207
 esp_image_header_t::spi_size (C++ member), 1206
 esp_image_header_t::spi_speed (C++ member), 1206
 esp_image_header_t::wp_pin (C++ member), 1206
 ESP_IMAGE_MAX_SEGMENTS (C macro), 1207
 esp_image_segment_header_t (C++ struct), 1207
 esp_image_segment_header_t::data_len (C++ member), 1207
 esp_image_segment_header_t::load_addr (C++ member), 1207
 esp_image_spi_freq_t (C++ enum), 1209
 esp_image_spi_freq_t::ESP_IMAGE_SPI_SPEED_DIV_1 (C++ enumerator), 1209
 esp_image_spi_freq_t::ESP_IMAGE_SPI_SPEED_DIV_2 (C++ enumerator), 1209
 esp_image_spi_freq_t::ESP_IMAGE_SPI_SPEED_DIV_3 (C++ enumerator), 1209
 esp_image_spi_freq_t::ESP_IMAGE_SPI_SPEED_DIV_4 (C++ enumerator), 1209
 esp_image_spi_mode_t (C++ enum), 1208
 esp_image_spi_mode_t::ESP_IMAGE_SPI_MODE_DIO (C++ enumerator), 1208
 esp_image_spi_mode_t::ESP_IMAGE_SPI_MODE_DOUT (C++ enumerator), 1208
 esp_image_spi_mode_t::ESP_IMAGE_SPI_MODE_FAST_READ (C++ enumerator), 1208
 esp_image_spi_mode_t::ESP_IMAGE_SPI_MODE_QIO (C++ enumerator), 1208
 esp_image_spi_mode_t::ESP_IMAGE_SPI_MODE_QOUT (C++ enumerator), 1208
 esp_image_spi_mode_t::ESP_IMAGE_SPI_MODE_SLOW_READ (C++ enumerator), 1209
 esp_intr_alloc (C++ function), 1485
 esp_intr_alloc_intrstatus (C++ function), 1486
 esp_intr_cpu_affinity_t (C++ enum), 1484
 esp_intr_cpu_affinity_t::ESP_INTR_CPU_AFFINITY_0 (C++ enumerator), 1485
 esp_intr_cpu_affinity_t::ESP_INTR_CPU_AFFINITY_1 (C++ enumerator), 1485

esp_intr_cpu_affinity_t::ESP_INTR_CPU_AFFINITY_ADDR::addr (C++ member), 248
 (C++ enumerator), 1484
 ESP_INTR_CPU_AFFINITY_TO_CORE_ID (C macro), 1484
 ESP_INTR_DISABLE (C macro), 1489
 esp_intr_disable (C++ function), 1487
 esp_intr_disable_source (C++ function), 1488
 esp_intr_dump (C++ function), 1488
 ESP_INTR_ENABLE (C macro), 1489
 esp_intr_enable (C++ function), 1487
 esp_intr_enable_source (C++ function), 1487
 ESP_INTR_FLAG_EDGE (C macro), 1488
 ESP_INTR_FLAG_HIGH (C macro), 1489
 ESP_INTR_FLAG_INTRDISABLED (C macro), 1489
 ESP_INTR_FLAG_IRAM (C macro), 1488
 ESP_INTR_FLAG_LEVEL1 (C macro), 1488
 ESP_INTR_FLAG_LEVEL2 (C macro), 1488
 ESP_INTR_FLAG_LEVEL3 (C macro), 1488
 ESP_INTR_FLAG_LEVEL4 (C macro), 1488
 ESP_INTR_FLAG_LEVEL5 (C macro), 1488
 ESP_INTR_FLAG_LEVEL6 (C macro), 1488
 ESP_INTR_FLAG_LEVELMASK (C macro), 1489
 ESP_INTR_FLAG_LOWMED (C macro), 1489
 ESP_INTR_FLAG_NMI (C macro), 1488
 ESP_INTR_FLAG_SHARED (C macro), 1488
 esp_intr_flags_to_level (C++ function), 1488
 esp_intr_free (C++ function), 1486
 esp_intr_get_cpu (C++ function), 1486
 esp_intr_get_intno (C++ function), 1487
 esp_intr_level_to_flags (C++ function), 1488
 esp_intr_mark_shared (C++ function), 1485
 esp_intr_noniram_disable (C++ function), 1487
 esp_intr_noniram_enable (C++ function), 1487
 esp_intr_reserve (C++ function), 1485
 esp_intr_set_in_iram (C++ function), 1487
 esp_ip4_addr (C++ struct), 248
 esp_ip4_addr1 (C macro), 249
 esp_ip4_addr1_16 (C macro), 249
 esp_ip4_addr2 (C macro), 249
 esp_ip4_addr2_16 (C macro), 249
 esp_ip4_addr3 (C macro), 249
 esp_ip4_addr3_16 (C macro), 249
 esp_ip4_addr4 (C macro), 249
 esp_ip4_addr4_16 (C macro), 249
 esp_ip4_addr::addr (C++ member), 248
 esp_ip4_addr_get_byte (C macro), 249
 esp_ip4_addr_t (C++ type), 250
 esp_ip4addr_aton (C++ function), 235
 ESP_IP4ADDR_INIT (C macro), 249
 esp_ip4addr_ntoa (C++ function), 234
 ESP_IP4TOADDR (C macro), 249
 ESP_IP4TOUINT32 (C macro), 249
 esp_ip6_addr (C++ struct), 248
 esp_ip6_addr::addr (C++ member), 248
 esp_ip6_addr::zone (C++ member), 248
 ESP_IP6_ADDR_BLOCK1 (C macro), 249
 ESP_IP6_ADDR_BLOCK2 (C macro), 249
 ESP_IP6_ADDR_BLOCK3 (C macro), 249
 ESP_IP6_ADDR_BLOCK4 (C macro), 249
 ESP_IP6_ADDR_BLOCK5 (C macro), 249
 ESP_IP6_ADDR_BLOCK6 (C macro), 249
 ESP_IP6_ADDR_BLOCK7 (C macro), 249
 ESP_IP6_ADDR_BLOCK8 (C macro), 249
 esp_ip6_addr_t (C++ type), 250
 esp_ip6_addr_type_t (C++ enum), 250
 esp_ip6_addr_type_t::ESP_IP6_ADDR_IS_GLOBAL (C++ enumerator), 250
 esp_ip6_addr_type_t::ESP_IP6_ADDR_IS_IPV4_MAPPED (C++ enumerator), 250
 esp_ip6_addr_type_t::ESP_IP6_ADDR_IS_LINK_LOCAL (C++ enumerator), 250
 esp_ip6_addr_type_t::ESP_IP6_ADDR_IS_SITE_LOCAL (C++ enumerator), 250
 esp_ip6_addr_type_t::ESP_IP6_ADDR_IS_UNIQUE_LOCAL (C++ enumerator), 250
 esp_ip6_addr_type_t::ESP_IP6_ADDR_IS_UNKNOWN (C++ enumerator), 250
 ESP_IP6ADDR_INIT (C macro), 249
 esp_ip_addr_t (C++ type), 250
 ESP_IP_IS_ANY (C macro), 249
 ESP_IPADDR_TYPE_ANY (C macro), 249
 ESP_IPADDR_TYPE_V4 (C macro), 249
 ESP_IPADDR_TYPE_V6 (C macro), 249
 esp_ipc_call (C++ function), 1479
 esp_ipc_call_blocking (C++ function), 1479
 esp_ipc_func_t (C++ type), 1480
 esp_ipc_isr_asm_call (C macro), 1482
 esp_ipc_isr_asm_call_blocking (C macro), 1482
 esp_ipc_isr_call (C++ function), 1480
 esp_ipc_isr_call_blocking (C++ function), 1480
 esp_ipc_isr_func_t (C++ type), 1482
 esp_ipc_isr_release_other_cpu (C++ function), 1481
 esp_ipc_isr_stall_abort (C++ function), 1481
 esp_ipc_isr_stall_other_cpu (C++ function), 1481
 esp_ipc_isr_stall_pause (C++ function), 1481
 esp_ipc_isr_stall_resume (C++ function), 1481
 esp_lcd_i2c_bus_handle_t (C++ type), 425
 esp_lcd_i80_bus_handle_t (C++ type), 425
 esp_lcd_new_panel_io_i2c (C macro), 425
 esp_lcd_new_panel_io_i2c_v1 (C++ function), 422
 esp_lcd_new_panel_io_i2c_v2 (C++ function), 422
 esp_lcd_new_panel_io_spi (C++ function),

- 421
- `esp_lcd_new_panel_nt35510` (C++ function), 428
- `esp_lcd_new_panel_ssd1306` (C++ function), 429
- `esp_lcd_new_panel_st7789` (C++ function), 428
- `esp_lcd_panel_del` (C++ function), 426
- `esp_lcd_panel_dev_config_t` (C++ struct), 429
- `esp_lcd_panel_dev_config_t::bits_per_pixel` (C++ member), 429
- `esp_lcd_panel_dev_config_t::color_space` (C++ member), 429
- `esp_lcd_panel_dev_config_t::data_endian` (C++ member), 429
- `esp_lcd_panel_dev_config_t::flags` (C++ member), 430
- `esp_lcd_panel_dev_config_t::reset_active_high` (C++ member), 429
- `esp_lcd_panel_dev_config_t::reset_gpio_num` (C++ member), 429
- `esp_lcd_panel_dev_config_t::rgb_ele_order` (C++ member), 429
- `esp_lcd_panel_dev_config_t::rgb_endian` (C++ member), 429
- `esp_lcd_panel_dev_config_t::vendor_config` (C++ member), 430
- `esp_lcd_panel_disp_off` (C++ function), 427
- `esp_lcd_panel_disp_on_off` (C++ function), 427
- `esp_lcd_panel_disp_sleep` (C++ function), 428
- `esp_lcd_panel_draw_bitmap` (C++ function), 426
- `esp_lcd_panel_handle_t` (C++ type), 420
- `esp_lcd_panel_init` (C++ function), 426
- `esp_lcd_panel_invert_color` (C++ function), 427
- `esp_lcd_panel_io_callbacks_t` (C++ struct), 422
- `esp_lcd_panel_io_callbacks_t::on_color_trans_done` (C++ member), 422
- `esp_lcd_panel_io_color_trans_done_cb_t` (C++ type), 425
- `esp_lcd_panel_io_del` (C++ function), 421
- `esp_lcd_panel_io_event_data_t` (C++ struct), 422
- `esp_lcd_panel_io_handle_t` (C++ type), 420
- `esp_lcd_panel_io_i2c_config_t` (C++ struct), 424
- `esp_lcd_panel_io_i2c_config_t::control_byte` (C++ member), 424
- `esp_lcd_panel_io_i2c_config_t::dc_bit_offset` (C++ member), 424
- `esp_lcd_panel_io_i2c_config_t::dc_low_on_data` (C++ member), 424
- `esp_lcd_panel_io_i2c_config_t::dev_addr` (C++ member), 424
- `esp_lcd_panel_io_i2c_config_t::disable_control_phase` (C++ member), 424
- `esp_lcd_panel_io_i2c_config_t::flags` (C++ member), 424
- `esp_lcd_panel_io_i2c_config_t::lcd_cmd_bits` (C++ member), 424
- `esp_lcd_panel_io_i2c_config_t::lcd_param_bits` (C++ member), 424
- `esp_lcd_panel_io_i2c_config_t::on_color_trans_done` (C++ member), 424
- `esp_lcd_panel_io_i2c_config_t::scl_speed_hz` (C++ member), 424
- `esp_lcd_panel_io_i2c_config_t::user_ctx` (C++ member), 424
- `esp_lcd_panel_io_register_event_callbacks` (C++ function), 421
- `esp_lcd_panel_io_rx_param` (C++ function), 422
- `esp_lcd_panel_io_spi_config_t` (C++ struct), 422
- `esp_lcd_panel_io_spi_config_t::cs_gpio_num` (C++ member), 423
- `esp_lcd_panel_io_spi_config_t::cs_high_active` (C++ member), 424
- `esp_lcd_panel_io_spi_config_t::dc_gpio_num` (C++ member), 423
- `esp_lcd_panel_io_spi_config_t::dc_high_on_cmd` (C++ member), 423
- `esp_lcd_panel_io_spi_config_t::dc_low_on_data` (C++ member), 423
- `esp_lcd_panel_io_spi_config_t::dc_low_on_param` (C++ member), 423
- `esp_lcd_panel_io_spi_config_t::flags` (C++ member), 424
- `esp_lcd_panel_io_spi_config_t::lcd_cmd_bits` (C++ member), 423
- `esp_lcd_panel_io_spi_config_t::lcd_param_bits` (C++ member), 423
- `esp_lcd_panel_io_spi_config_t::lsb_first` (C++ member), 424
- `esp_lcd_panel_io_spi_config_t::octal_mode` (C++ member), 423
- `esp_lcd_panel_io_spi_config_t::on_color_trans_done` (C++ member), 423
- `esp_lcd_panel_io_spi_config_t::pclk_hz` (C++ member), 423
- `esp_lcd_panel_io_spi_config_t::quad_mode` (C++ member), 423
- `esp_lcd_panel_io_spi_config_t::sio_mode` (C++ member), 423
- `esp_lcd_panel_io_spi_config_t::spi_mode` (C++ member), 423
- `esp_lcd_panel_io_spi_config_t::trans_queue_depth` (C++ member), 423
- `esp_lcd_panel_io_spi_config_t::user_ctx` (C++ member), 423
- `esp_lcd_panel_io_tx_color` (C++ function), 424

- 421
- `esp_lcd_panel_io_tx_param` (C++ function), 420
- `esp_lcd_panel_mirror` (C++ function), 426
- `esp_lcd_panel_reset` (C++ function), 425
- `esp_lcd_panel_set_gap` (C++ function), 427
- `esp_lcd_panel_swap_xy` (C++ function), 426
- `esp_lcd_spi_bus_handle_t` (C++ type), 425
- `esp_light_sleep_start` (C++ function), 1549
- `esp_local_ctrl_add_property` (C++ function), 99
- `esp_local_ctrl_config` (C++ struct), 103
- `esp_local_ctrl_config::handlers` (C++ member), 103
- `esp_local_ctrl_config::max_properties` (C++ member), 103
- `esp_local_ctrl_config::proto_sec` (C++ member), 103
- `esp_local_ctrl_config::transport` (C++ member), 103
- `esp_local_ctrl_config::transport_config` (C++ member), 103
- `esp_local_ctrl_config_t` (C++ type), 104
- `esp_local_ctrl_get_property` (C++ function), 99
- `esp_local_ctrl_get_transport_ble` (C++ function), 99
- `esp_local_ctrl_get_transport_httpd` (C++ function), 99
- `esp_local_ctrl_handlers` (C++ struct), 101
- `esp_local_ctrl_handlers::get_prop_values` (C++ member), 101
- `esp_local_ctrl_handlers::set_prop_values` (C++ member), 102
- `esp_local_ctrl_handlers::usr_ctx` (C++ member), 102
- `esp_local_ctrl_handlers::usr_ctx_free_fn` (C++ member), 102
- `esp_local_ctrl_handlers_t` (C++ type), 103
- `esp_local_ctrl_prop` (C++ struct), 100
- `esp_local_ctrl_prop::ctx` (C++ member), 101
- `esp_local_ctrl_prop::ctx_free_fn` (C++ member), 101
- `esp_local_ctrl_prop::flags` (C++ member), 100
- `esp_local_ctrl_prop::name` (C++ member), 100
- `esp_local_ctrl_prop::size` (C++ member), 100
- `esp_local_ctrl_prop::type` (C++ member), 100
- `esp_local_ctrl_prop_t` (C++ type), 103
- `esp_local_ctrl_prop_val` (C++ struct), 101
- `esp_local_ctrl_prop_val::data` (C++ member), 101
- `esp_local_ctrl_prop_val::free_fn` (C++ member), 101
- `esp_local_ctrl_prop_val::size` (C++ member), 101
- `esp_local_ctrl_prop_val_t` (C++ type), 103
- `esp_local_ctrl_proto_sec` (C++ enum), 104
- `esp_local_ctrl_proto_sec::PROTOCOLCOM_SEC0` (C++ enumerator), 104
- `esp_local_ctrl_proto_sec::PROTOCOLCOM_SEC1` (C++ enumerator), 104
- `esp_local_ctrl_proto_sec::PROTOCOLCOM_SEC2` (C++ enumerator), 104
- `esp_local_ctrl_proto_sec::PROTOCOLCOM_SEC_CUSTOM` (C++ enumerator), 104
- `esp_local_ctrl_proto_sec_cfg` (C++ struct), 102
- `esp_local_ctrl_proto_sec_cfg::custom_handle` (C++ member), 102
- `esp_local_ctrl_proto_sec_cfg::pop` (C++ member), 102
- `esp_local_ctrl_proto_sec_cfg::sec_params` (C++ member), 102
- `esp_local_ctrl_proto_sec_cfg::version` (C++ member), 102
- `esp_local_ctrl_proto_sec_cfg_t` (C++ type), 104
- `esp_local_ctrl_proto_sec_t` (C++ type), 104
- `esp_local_ctrl_remove_property` (C++ function), 99
- `esp_local_ctrl_security1_params_t` (C++ type), 104
- `esp_local_ctrl_security2_params_t` (C++ type), 104
- `esp_local_ctrl_set_handler` (C++ function), 99
- `esp_local_ctrl_start` (C++ function), 99
- `esp_local_ctrl_stop` (C++ function), 99
- `ESP_LOCAL_CTRL_TRANSPORT_BLE` (C macro), 103
- `esp_local_ctrl_transport_config_ble_t` (C++ type), 103
- `esp_local_ctrl_transport_config_httpd_t` (C++ type), 104
- `esp_local_ctrl_transport_config_t` (C++ union), 100
- `esp_local_ctrl_transport_config_t::ble` (C++ member), 100
- `esp_local_ctrl_transport_config_t::httpd` (C++ member), 100
- `ESP_LOCAL_CTRL_TRANSPORT_HTTPD` (C macro), 103
- `esp_local_ctrl_transport_t` (C++ type), 103
- `ESP_LOG_BUFFER_CHAR` (C macro), 1495
- `ESP_LOG_BUFFER_CHAR_LEVEL` (C macro), 1494
- `ESP_LOG_BUFFER_HEX` (C macro), 1494
- `ESP_LOG_BUFFER_HEX_LEVEL` (C macro), 1494
- `ESP_LOG_BUFFER_HEXDUMP` (C macro), 1494
- `ESP_LOG_EARLY_IMPL` (C macro), 1496
- `esp_log_early_timestamp` (C++ function), 1493

- `esp_log_get_level_master` (C++ function), 1492
- `ESP_LOG_LEVEL` (C macro), 1496
- `esp_log_level_get` (C++ function), 1493
- `ESP_LOG_LEVEL_LOCAL` (C macro), 1496
- `esp_log_level_set` (C++ function), 1492
- `esp_log_level_t` (C++ enum), 1497
- `esp_log_level_t::ESP_LOG_DEBUG` (C++ enumerator), 1498
- `esp_log_level_t::ESP_LOG_ERROR` (C++ enumerator), 1497
- `esp_log_level_t::ESP_LOG_INFO` (C++ enumerator), 1498
- `esp_log_level_t::ESP_LOG_NONE` (C++ enumerator), 1497
- `esp_log_level_t::ESP_LOG_VERBOSE` (C++ enumerator), 1498
- `esp_log_level_t::ESP_LOG_WARN` (C++ enumerator), 1497
- `esp_log_set_level_master` (C++ function), 1492
- `esp_log_set_vprintf` (C++ function), 1493
- `esp_log_system_timestamp` (C++ function), 1493
- `esp_log_timestamp` (C++ function), 1493
- `esp_log_write` (C++ function), 1493
- `esp_log_writev` (C++ function), 1494
- `ESP_LOGD` (C macro), 1496
- `ESP_LOGE` (C macro), 1496
- `ESP_LOGI` (C macro), 1496
- `ESP_LOGV` (C macro), 1496
- `ESP_LOGW` (C macro), 1496
- `esp_mac_addr_len_get` (C++ function), 1506
- `esp_mac_type_t` (C++ enum), 1506
- `esp_mac_type_t::ESP_MAC_BASE` (C++ enumerator), 1506
- `esp_mac_type_t::ESP_MAC_BT` (C++ enumerator), 1506
- `esp_mac_type_t::ESP_MAC_EFUSE_CUSTOM` (C++ enumerator), 1506
- `esp_mac_type_t::ESP_MAC_EFUSE_EXT` (C++ enumerator), 1506
- `esp_mac_type_t::ESP_MAC_EFUSE_FACTORY` (C++ enumerator), 1506
- `esp_mac_type_t::ESP_MAC_ETH` (C++ enumerator), 1506
- `esp_mac_type_t::ESP_MAC_IEEE802154` (C++ enumerator), 1506
- `esp_mac_type_t::ESP_MAC_WIFI_SOFTAP` (C++ enumerator), 1506
- `esp_mac_type_t::ESP_MAC_WIFI_STA` (C++ enumerator), 1506
- `esp_mmu_map` (C++ function), 1448
- `esp_mmu_map_dump_mapped_blocks` (C++ function), 1449
- `esp_mmu_map_get_max_consecutive_free_block_size` (C++ function), 1449
- `ESP_MMU_MMAP_FLAG_PADDR_SHARED` (C macro), 1450
- `esp_mmu_paddr_find_caps` (C++ function), 1450
- `esp_mmu_paddr_to_vaddr` (C++ function), 1449
- `esp_mmu_unmap` (C++ function), 1449
- `esp_mmu_vaddr_to_paddr` (C++ function), 1449
- `esp_mqtt_client_config_t` (C++ struct), 50
- `esp_mqtt_client_config_t` (C++ type), 57
- `esp_mqtt_client_config_t::broker` (C++ member), 50
- `esp_mqtt_client_config_t::broker_t` (C++ struct), 50
- `esp_mqtt_client_config_t::broker_t::address` (C++ member), 50
- `esp_mqtt_client_config_t::broker_t::address_t` (C++ struct), 51
- `esp_mqtt_client_config_t::broker_t::address_t::host` (C++ member), 51
- `esp_mqtt_client_config_t::broker_t::address_t::password` (C++ member), 51
- `esp_mqtt_client_config_t::broker_t::address_t::port` (C++ member), 51
- `esp_mqtt_client_config_t::broker_t::address_t::topic` (C++ member), 51
- `esp_mqtt_client_config_t::broker_t::address_t::url` (C++ member), 51
- `esp_mqtt_client_config_t::broker_t::verification` (C++ member), 50
- `esp_mqtt_client_config_t::broker_t::verification_t` (C++ struct), 51
- `esp_mqtt_client_config_t::broker_t::verification_t::method` (C++ member), 52
- `esp_mqtt_client_config_t::broker_t::verification_t::payload` (C++ member), 51
- `esp_mqtt_client_config_t::broker_t::verification_t::payload_t` (C++ struct), 51
- `esp_mqtt_client_config_t::broker_t::verification_t::payload_t::buffer` (C++ member), 50
- `esp_mqtt_client_config_t::buffer_t` (C++ struct), 52
- `esp_mqtt_client_config_t::buffer_t::out_size` (C++ member), 52
- `esp_mqtt_client_config_t::buffer_t::size` (C++ member), 52
- `esp_mqtt_client_config_t::credentials` (C++ member), 50
- `esp_mqtt_client_config_t::credentials_t` (C++ struct), 52

Espressif Systems **1943** [Submit Document Feedback](#) Release v5.2.7

- tion*), 225
- `esp_netif_action_got_ip` (C++ function), 225
- `esp_netif_action_join_ip6_multicast_group` (C++ function), 226
- `esp_netif_action_leave_ip6_multicast_group` (C++ function), 226
- `esp_netif_action_remove_ip6_address` (C++ function), 226
- `esp_netif_action_start` (C++ function), 224
- `esp_netif_action_stop` (C++ function), 225
- `esp_netif_add_ip6_address` (C++ function), 234
- `esp_netif_attach` (C++ function), 224
- `ESP_NETIF_BR_DROP` (C macro), 244
- `ESP_NETIF_BR_FDW_CPU` (C macro), 244
- `ESP_NETIF_BR_FLOOD` (C macro), 244
- `esp_netif_callback_fn` (C++ type), 237
- `esp_netif_config` (C++ struct), 242
- `esp_netif_config::base` (C++ member), 243
- `esp_netif_config::driver` (C++ member), 243
- `esp_netif_config::stack` (C++ member), 243
- `esp_netif_config_t` (C++ type), 244
- `esp_netif_create_ip6_linklocal` (C++ function), 233
- `ESP_NETIF_DEFAULT_OPENTHREAD` (C macro), 215
- `esp_netif_deinit` (C++ function), 223
- `esp_netif_destroy` (C++ function), 224
- `esp_netif_dhcp_option_id_t` (C++ enum), 245
- `esp_netif_dhcp_option_id_t::ESP_NETIF_DOMAIN_NAME_OPTION` (C++ enumerator), 246
- `esp_netif_dhcp_option_id_t::ESP_NETIF_IP_ADDRESS_LEASE_TIME` (C++ enumerator), 246
- `esp_netif_dhcp_option_id_t::ESP_NETIF_IP_REQUEST_RETRY_TIME` (C++ enumerator), 246
- `esp_netif_dhcp_option_id_t::ESP_NETIF_IP_REQUESTED_IP_ADDRESS` (C++ enumerator), 246
- `esp_netif_dhcp_option_id_t::ESP_NETIF_IP_REQUESTED_IP_ADDRESS_RANGE` (C++ enumerator), 246
- `esp_netif_dhcp_option_id_t::ESP_NETIF_IP_REQUESTED_IP_ADDRESS_SUBNET_MASK` (C++ enumerator), 245
- `esp_netif_dhcp_option_id_t::ESP_NETIF_VENDOR_CLASS_IDENTIFIER` (C++ enumerator), 246
- `esp_netif_dhcp_option_id_t::ESP_NETIF_VENDOR_SPECIFIC_INFO` (C++ enumerator), 246
- `esp_netif_dhcp_option_mode_t` (C++ enum), 245
- `esp_netif_dhcp_option_mode_t::ESP_NETIF_OP_GET` (C++ enumerator), 245
- `esp_netif_dhcp_option_mode_t::ESP_NETIF_OP_MAX` (C++ enumerator), 245
- `esp_netif_dhcp_option_mode_t::ESP_NETIF_OP_SET` (C++ enumerator), 245
- `esp_netif_dhcp_option_mode_t::ESP_NETIF_OP_STOP` (C++ enumerator), 245
- `esp_netif_dhcp_status_t` (C++ enum), 245
- `esp_netif_dhcp_status_t::ESP_NETIF_DHCP_INIT` (C++ enumerator), 245
- `esp_netif_dhcp_status_t::ESP_NETIF_DHCP_STARTED` (C++ enumerator), 245
- `esp_netif_dhcp_status_t::ESP_NETIF_DHCP_STATUS_MAX` (C++ enumerator), 245
- `esp_netif_dhcp_status_t::ESP_NETIF_DHCP_STOPPED` (C++ enumerator), 245
- `esp_netif_dhcpc_get_status` (C++ function), 231
- `esp_netif_dhcpc_option` (C++ function), 230
- `esp_netif_dhcpc_start` (C++ function), 231
- `esp_netif_dhcpc_stop` (C++ function), 231
- `esp_netif_dhcps_get_clients_by_mac` (C++ function), 232
- `esp_netif_dhcps_get_status` (C++ function), 231
- `esp_netif_dhcps_option` (C++ function), 230
- `esp_netif_dhcps_start` (C++ function), 232
- `esp_netif_dhcps_stop` (C++ function), 232
- `esp_netif_dns_info_t` (C++ struct), 239
- `esp_netif_dns_info_t::ip` (C++ member), 239
- `esp_netif_dns_type_t` (C++ enum), 245
- `esp_netif_dns_type_t::ESP_NETIF_DNS_BACKUP` (C++ enumerator), 245
- `esp_netif_dns_type_t::ESP_NETIF_DNS_FALLBACK` (C++ enumerator), 245
- `esp_netif_dns_type_t::ESP_NETIF_DNS_MAIN` (C++ enumerator), 245
- `esp_netif_dns_type_t::ESP_NETIF_DNS_MAX` (C++ enumerator), 245
- `esp_netif_driver_base_s` (C++ struct), 242
- `esp_netif_driver_base_s::netif` (C++ member), 242
- `esp_netif_driver_base_s::post_attach` (C++ member), 242
- `esp_netif_driver_ifconfig` (C++ struct), 242
- `esp_netif_driver_ifconfig::driver_free_rx_buffer` (C++ member), 242
- `esp_netif_driver_ifconfig::handle` (C++ member), 242
- `esp_netif_driver_ifconfig::transmit` (C++ member), 242
- `esp_netif_driver_ifconfig::transmit_wrap` (C++ member), 242
- `esp_netif_driver_ifconfig_t` (C++ type), 244
- `esp_netif_find_if` (C++ function), 236
- `esp_netif_find_predicate_t` (C++ type), 237
- `esp_netif_flags` (C++ enum), 247
- `esp_netif_flags::ESP_NETIF_DHCP_CLIENT` (C++ enumerator), 247
- `esp_netif_flags::ESP_NETIF_DHCP_SERVER` (C++ enumerator), 247
- `esp_netif_flags::ESP_NETIF_FLAG_AUTOUP` (C++ enumerator), 247

esp_netif_flags::ESP_NETIF_FLAG_EVENT_ID_MODIFIED (C++ enumerator), 247
 esp_netif_flags::ESP_NETIF_FLAG_GARP (C++ enumerator), 247
 esp_netif_flags::ESP_NETIF_FLAG_IS_BRIDGE (C++ enumerator), 247
 esp_netif_flags::ESP_NETIF_FLAG_IS_PPP (C++ enumerator), 247
 esp_netif_flags::ESP_NETIF_FLAG_MLDV6_REPORT (C++ enumerator), 247
 esp_netif_flags_t (C++ type), 244
 esp_netif_free_rx_buffer (C++ function), 254
 esp_netif_get_all_ip6 (C++ function), 234
 esp_netif_get_all_preferred_ip6 (C++ function), 234
 esp_netif_get_default_netif (C++ function), 227
 esp_netif_get_desc (C++ function), 235
 esp_netif_get_dns_info (C++ function), 233
 esp_netif_get_event_id (C++ function), 236
 esp_netif_get_flags (C++ function), 235
 esp_netif_get_handle_from_ifkey (C++ function), 235
 esp_netif_get_handle_from_netif_impl (C++ function), 253
 esp_netif_get_hostname (C++ function), 228
 esp_netif_get_ifkey (C++ function), 235
 esp_netif_get_io_driver (C++ function), 235
 esp_netif_get_ip6_global (C++ function), 233
 esp_netif_get_ip6_linklocal (C++ function), 233
 esp_netif_get_ip_info (C++ function), 228
 esp_netif_get_mac (C++ function), 227
 esp_netif_get_netif_impl (C++ function), 254
 esp_netif_get_netif_impl_index (C++ function), 229
 esp_netif_get_netif_impl_name (C++ function), 230
 esp_netif_get_nr_of_ifs (C++ function), 236
 esp_netif_get_old_ip_info (C++ function), 228
 esp_netif_get_route_prio (C++ function), 236
 esp_netif_htonl (C macro), 249
 esp_netif_inherent_config (C++ struct), 241
 esp_netif_inherent_config::bridge_info (C++ member), 242
 esp_netif_inherent_config::flags (C++ member), 241
 esp_netif_inherent_config::get_ip_event (C++ member), 241
 esp_netif_inherent_config::if_desc (C++ member), 242
 esp_netif_inherent_config::if_key (C++ member), 242
 esp_netif_inherent_config::ip_info (C++ member), 241
 esp_netif_inherent_config::lost_ip_event (C++ member), 242
 esp_netif_inherent_config::mac (C++ member), 241
 esp_netif_inherent_config::route_prio (C++ member), 242
 esp_netif_inherent_config_t (C++ type), 244
 ESP_NETIF_INHERENT_DEFAULT_OPENTHREAD (C macro), 215
 esp_netif_init (C++ function), 223
 esp_netif_iodriver_handle (C++ type), 244
 esp_netif_ip4_makeu32 (C macro), 249
 esp_netif_ip6_get_addr_type (C++ function), 247
 esp_netif_ip6_info_t (C++ struct), 240
 esp_netif_ip6_info_t::ip (C++ member), 240
 esp_netif_ip_addr_copy (C++ function), 247
 esp_netif_ip_event_type (C++ enum), 247
 esp_netif_ip_event_type::ESP_NETIF_IP_EVENT_GOT_IP (C++ enumerator), 247
 esp_netif_ip_event_type::ESP_NETIF_IP_EVENT_LOST_IP (C++ enumerator), 247
 esp_netif_ip_event_type_t (C++ type), 244
 esp_netif_ip_info_t (C++ struct), 239
 esp_netif_ip_info_t::gw (C++ member), 239
 esp_netif_ip_info_t::ip (C++ member), 239
 esp_netif_ip_info_t::netmask (C++ member), 239
 esp_netif_is_netif_up (C++ function), 228
 esp_netif_join_ip6_multicast_group (C++ function), 227
 esp_netif_leave_ip6_multicast_group (C++ function), 227
 esp_netif_napt_disable (C++ function), 230
 esp_netif_napt_enable (C++ function), 230
 esp_netif_netstack_buf_free (C++ function), 236
 esp_netif_netstack_buf_ref (C++ function), 236
 esp_netif_netstack_config_t (C++ type), 244
 esp_netif_new (C++ function), 224
 esp_netif_next (C++ function), 236
 esp_netif_next_unsafe (C++ function), 236
 esp_netif_pair_mac_ip_t (C++ struct), 243
 esp_netif_pair_mac_ip_t::ip (C++ member), 243
 esp_netif_pair_mac_ip_t::mac (C++ member), 243
 esp_netif_receive (C++ function), 224
 esp_netif_receive_t (C++ type), 244
 esp_netif_remove_ip6_address (C++ function), 234
 esp_netif_set_default_netif (C++ function), 224

- tion*), 227
- `esp_netif_set_dns_info` (C++ function), 232
- `esp_netif_set_driver_config` (C++ function), 224
- `esp_netif_set_hostname` (C++ function), 228
- `esp_netif_set_ip4_addr` (C++ function), 234
- `esp_netif_set_ip_info` (C++ function), 229
- `esp_netif_set_link_speed` (C++ function), 254
- `esp_netif_set_mac` (C++ function), 227
- `esp_netif_set_old_ip_info` (C++ function), 229
- `ESP_NETIF_SNTP_DEFAULT_CONFIG` (C macro), 239
- `ESP_NETIF_SNTP_DEFAULT_CONFIG_MULTIPLE` (C macro), 238
- `esp_netif_sntp_deinit` (C++ function), 237
- `esp_netif_sntp_init` (C++ function), 237
- `esp_netif_sntp_reachability` (C++ function), 237
- `esp_netif_sntp_start` (C++ function), 237
- `esp_netif_sntp_sync_wait` (C++ function), 237
- `esp_netif_str_to_ip4` (C++ function), 235
- `esp_netif_str_to_ip6` (C++ function), 235
- `esp_netif_t` (C++ type), 244
- `esp_netif_tcpip_exec` (C++ function), 237
- `esp_netif_transmit` (C++ function), 254
- `esp_netif_transmit_wrap` (C++ function), 254
- `esp_ng_type_t` (C++ enum), 1091
- `esp_ng_type_t::ESP_NG_3072` (C++ enumerator), 1091
- `ESP_OK` (C macro), 1263
- `ESP_OK_EFUSE_CNT` (C macro), 1261
- `esp_openthread_auto_start` (C++ function), 206
- `esp_openthread_border_router_deinit` (C++ function), 216
- `esp_openthread_border_router_init` (C++ function), 215
- `esp_openthread_compatibility_error_callback` (C++ type), 211
- `esp_openthread_config_t` (C++ struct), 210
- `esp_openthread_config_t::netif_config` (C++ member), 210
- `esp_openthread_config_t::platform_config` (C++ member), 210
- `esp_openthread_coprocessor_reset_failure_callback` (C++ type), 211
- `esp_openthread_dataset_changed_event_t` (C++ struct), 207
- `esp_openthread_dataset_changed_event_t::new_dataset` (C++ member), 207
- `esp_openthread_dataset_changed_event_t::type` (C++ member), 207
- `esp_openthread_dataset_type_t` (C++ enum), 212
- `esp_openthread_dataset_type_t::OPENTHREAD_DATASET_PENDING` (C++ enumerator), 212
- `esp_openthread_dataset_type_t::OPENTHREAD_PENDING` (C++ enumerator), 212
- `esp_openthread_deinit` (C++ function), 206
- `esp_openthread_event_t` (C++ enum), 211
- `esp_openthread_event_t::OPENTHREAD_EVENT_ATTACHED` (C++ enumerator), 211
- `esp_openthread_event_t::OPENTHREAD_EVENT_DATASET` (C++ enumerator), 212
- `esp_openthread_event_t::OPENTHREAD_EVENT_DETACHED` (C++ enumerator), 211
- `esp_openthread_event_t::OPENTHREAD_EVENT_GOT_IP6` (C++ enumerator), 211
- `esp_openthread_event_t::OPENTHREAD_EVENT_IF_DOWN` (C++ enumerator), 211
- `esp_openthread_event_t::OPENTHREAD_EVENT_IF_UP` (C++ enumerator), 211
- `esp_openthread_event_t::OPENTHREAD_EVENT_LOST_IP6` (C++ enumerator), 211
- `esp_openthread_event_t::OPENTHREAD_EVENT_MULTICAST` (C++ enumerator), 211
- `esp_openthread_event_t::OPENTHREAD_EVENT_MULTICAST` (C++ enumerator), 211
- `esp_openthread_event_t::OPENTHREAD_EVENT_PUBLISH` (C++ enumerator), 212
- `esp_openthread_event_t::OPENTHREAD_EVENT_REMOVE_M` (C++ enumerator), 212
- `esp_openthread_event_t::OPENTHREAD_EVENT_ROLE_CHANGE` (C++ enumerator), 211
- `esp_openthread_event_t::OPENTHREAD_EVENT_SET_DNS` (C++ enumerator), 212
- `esp_openthread_event_t::OPENTHREAD_EVENT_START` (C++ enumerator), 211
- `esp_openthread_event_t::OPENTHREAD_EVENT_STOP` (C++ enumerator), 211
- `esp_openthread_event_t::OPENTHREAD_EVENT_TREL_ADD` (C++ enumerator), 211
- `esp_openthread_event_t::OPENTHREAD_EVENT_TREL_MULTICAST` (C++ enumerator), 212
- `esp_openthread_event_t::OPENTHREAD_EVENT_TREL_REMOVE` (C++ enumerator), 212
- `esp_openthread_get_backbone_netif` (C++ function), 216
- `esp_openthread_get_instance` (C++ function), 206
- `esp_openthread_get_meshcop_instance_name` (C++ function), 216
- `esp_openthread_get_netif` (C++ function), 215
- `esp_openthread_host_connection_config_t` (C++ struct), 209
- `esp_openthread_host_connection_config_t::host_connection` (C++ member), 209
- `esp_openthread_host_connection_config_t::host_uart` (C++ member), 209
- `esp_openthread_host_connection_config_t::host_usb` (C++ member), 210
- `esp_openthread_host_connection_config_t::spi_slave` (C++ member), 210

- (C++ member), 210
- esp_openthread_host_connection_mode_t (C++ enum), 212
- esp_openthread_host_connection_mode_t::HOST_CONNECTION_MODE_CLI_UART (C++ enumerator), 213
- esp_openthread_host_connection_mode_t::HOST_CONNECTION_MODE_CLI_USB (C++ enumerator), 213
- esp_openthread_host_connection_mode_t::HOST_CONNECTION_MODE_MAX (C++ enumerator), 213
- esp_openthread_host_connection_mode_t::HOST_CONNECTION_MODE_NONE (C++ enumerator), 213
- esp_openthread_host_connection_mode_t::HOST_CONNECTION_MODE_RCP_SPI (C++ enumerator), 213
- esp_openthread_host_connection_mode_t::HOST_CONNECTION_MODE_RCP_UART (C++ enumerator), 213
- esp_openthread_init (C++ function), 205
- esp_openthread_launch_mainloop (C++ function), 206
- esp_openthread_lock_acquire (C++ function), 213
- esp_openthread_lock_deinit (C++ function), 213
- esp_openthread_lock_init (C++ function), 213
- esp_openthread_lock_release (C++ function), 214
- esp_openthread_mainloop_context_t (C++ struct), 207
- esp_openthread_mainloop_context_t::error_code (C++ member), 208
- esp_openthread_mainloop_context_t::max Esp (C++ member), 208
- esp_openthread_mainloop_context_t::read Esp (C++ member), 208
- esp_openthread_mainloop_context_t::time Esp (C++ member), 208
- esp_openthread_mainloop_context_t::write Esp (C++ member), 208
- esp_openthread_mainloop_exit (C++ function), 206
- esp_openthread_netif_glue_deinit (C++ function), 214
- esp_openthread_netif_glue_init (C++ function), 214
- esp_openthread_platform_config_t (C++ struct), 210
- esp_openthread_platform_config_t::host Esp (C++ member), 210
- esp_openthread_platform_config_t::port Esp (C++ member), 210
- esp_openthread_platform_config_t::radio Esp (C++ member), 210
- esp_openthread_port_config_t (C++ struct), 210
- esp_openthread_port_config_t::netif_queue_size (C++ member), 210
- esp_openthread_port_config_t::storage Esp (C++ member), 210
- esp_openthread_port_config_t::task_queue_size (C++ member), 210
- esp_openthread_radio_config_t (C++ struct), 209
- esp_openthread_radio_config_t::radio_mode (C++ member), 209
- esp_openthread_radio_config_t::radio_spi_config (C++ member), 209
- esp_openthread_radio_config_t::radio_uart_config (C++ member), 209
- esp_openthread_radio_mode_t (C++ enum), 212
- esp_openthread_radio_mode_t::RADIO_MODE_MAX (C++ enumerator), 212
- esp_openthread_radio_mode_t::RADIO_MODE_NATIVE (C++ enumerator), 212
- esp_openthread_radio_mode_t::RADIO_MODE_SPI_RCP (C++ enumerator), 212
- esp_openthread_radio_mode_t::RADIO_MODE_TREL (C++ enumerator), 212
- esp_openthread_radio_mode_t::RADIO_MODE_UART_RCP (C++ enumerator), 212
- esp_openthread_rcp_failure_handler (C++ type), 211
- esp_openthread_register_meshcop_e_handler (C++ function), 215
- esp_openthread_role_changed_event_t (C++ struct), 207
- esp_openthread_role_changed_event_t::current_role (C++ member), 207
- esp_openthread_role_changed_event_t::previous_role (C++ member), 207
- esp_openthread_set_backbone_netif (C++ function), 215
- esp_openthread_set_meshcop_instance_name (C++ function), 216
- esp_openthread_spi_host_config_t (C++ struct), 208
- esp_openthread_spi_host_config_t::dma_channel (C++ member), 208
- esp_openthread_spi_host_config_t::host_device (C++ member), 208
- esp_openthread_spi_host_config_t::intr_pin (C++ member), 209
- esp_openthread_spi_host_config_t::spi_device (C++ member), 208
- esp_openthread_spi_host_config_t::spi_interface (C++ member), 208
- esp_openthread_spi_slave_config_t (C++ struct), 209
- esp_openthread_spi_slave_config_t::bus_config (C++ member), 209
- esp_openthread_spi_slave_config_t::host_device (C++ member), 209
- esp_openthread_spi_slave_config_t::intr_pin (C++ member), 209
- esp_openthread_spi_slave_config_t::slave_config (C++ member), 209

Espressif Systems 1948 Release v5.2.7

- (C++ enumerator), 159
- esp_ping_profile_t::ESP_PING_PROF_SEQN (C++ enumerator), 159
- esp_ping_profile_t::ESP_PING_PROF_SIZE (C++ enumerator), 159
- esp_ping_profile_t::ESP_PING_PROF_TIMEGAP (C++ enumerator), 159
- esp_ping_profile_t::ESP_PING_PROF_TOS (C++ enumerator), 159
- esp_ping_profile_t::ESP_PING_PROF_TTL (C++ enumerator), 159
- esp_ping_start (C++ function), 157
- esp_ping_stop (C++ function), 157
- esp_pm_config_esp32_t (C++ type), 1533
- esp_pm_config_esp32c2_t (C++ type), 1533
- esp_pm_config_esp32c3_t (C++ type), 1533
- esp_pm_config_esp32c6_t (C++ type), 1534
- esp_pm_config_esp32s2_t (C++ type), 1533
- esp_pm_config_esp32s3_t (C++ type), 1533
- esp_pm_config_t (C++ struct), 1533
- esp_pm_config_t::light_sleep_enable (C++ member), 1533
- esp_pm_config_t::max_freq_mhz (C++ member), 1533
- esp_pm_config_t::min_freq_mhz (C++ member), 1533
- esp_pm_configure (C++ function), 1531
- esp_pm_dump_locks (C++ function), 1533
- esp_pm_get_configuration (C++ function), 1531
- esp_pm_lock_acquire (C++ function), 1532
- esp_pm_lock_create (C++ function), 1532
- esp_pm_lock_delete (C++ function), 1532
- esp_pm_lock_handle_t (C++ type), 1534
- esp_pm_lock_release (C++ function), 1532
- esp_pm_lock_type_t (C++ enum), 1534
- esp_pm_lock_type_t::ESP_PM_APB_FREQ_MAX (C++ enumerator), 1534
- esp_pm_lock_type_t::ESP_PM_CPU_FREQ_MAX (C++ enumerator), 1534
- esp_pm_lock_type_t::ESP_PM_NO_LIGHT_SLEEP (C++ enumerator), 1534
- esp_pthread_cfg_t (C++ struct), 1538
- esp_pthread_cfg_t::inherit_cfg (C++ member), 1539
- esp_pthread_cfg_t::pin_to_core (C++ member), 1539
- esp_pthread_cfg_t::prio (C++ member), 1539
- esp_pthread_cfg_t::stack_size (C++ member), 1539
- esp_pthread_cfg_t::thread_name (C++ member), 1539
- esp_pthread_get_cfg (C++ function), 1538
- esp_pthread_get_default_config (C++ function), 1538
- esp_pthread_init (C++ function), 1538
- esp_pthread_set_cfg (C++ function), 1538
- esp_random (C++ function), 1540
- esp_read_mac (C++ function), 1505
- esp_register_freertos_idle_hook (C++ function), 1423
- esp_register_freertos_idle_hook_for_cpu (C++ function), 1423
- esp_register_freertos_tick_hook (C++ function), 1424
- esp_register_freertos_tick_hook_for_cpu (C++ function), 1424
- esp_register_shutdown_handler (C++ function), 1501
- esp_reset_reason (C++ function), 1501
- esp_reset_reason_t (C++ enum), 1502
- esp_reset_reason_t::ESP_RST_BROWNOUT (C++ enumerator), 1503
- esp_reset_reason_t::ESP_RST_CPU_LOCKUP (C++ enumerator), 1503
- esp_reset_reason_t::ESP_RST_DEEPSLEEP (C++ enumerator), 1502
- esp_reset_reason_t::ESP_RST_EFUSE (C++ enumerator), 1503
- esp_reset_reason_t::ESP_RST_EXT (C++ enumerator), 1502
- esp_reset_reason_t::ESP_RST_INT_WDT (C++ enumerator), 1502
- esp_reset_reason_t::ESP_RST_JTAG (C++ enumerator), 1503
- esp_reset_reason_t::ESP_RST_PANIC (C++ enumerator), 1502
- esp_reset_reason_t::ESP_RST_POWERON (C++ enumerator), 1502
- esp_reset_reason_t::ESP_RST_PWR_GLITCH (C++ enumerator), 1503
- esp_reset_reason_t::ESP_RST_SDIO (C++ enumerator), 1503
- esp_reset_reason_t::ESP_RST_SW (C++ enumerator), 1502
- esp_reset_reason_t::ESP_RST_TASK_WDT (C++ enumerator), 1502
- esp_reset_reason_t::ESP_RST_UNKNOWN (C++ enumerator), 1502
- esp_reset_reason_t::ESP_RST_USB (C++ enumerator), 1503
- esp_reset_reason_t::ESP_RST_WDT (C++ enumerator), 1502
- esp_restart (C++ function), 1501
- ESP_RETURN_ON_ERROR (C macro), 1262
- ESP_RETURN_ON_ERROR_ISR (C macro), 1262
- ESP_RETURN_ON_FALSE (C macro), 1262
- ESP_RETURN_ON_FALSE_ISR (C macro), 1262
- esp_rom_delay_us (C++ function), 1476
- esp_rom_get_cpu_ticks_per_us (C++ function), 1477
- esp_rom_get_reset_reason (C++ function), 1476
- esp_rom_install_channel_putc (C++ function), 1476

- `esp_rom_install_uart_printf` (C++ function), 1476
- `esp_rom_printf` (C++ function), 1476
- `esp_rom_route_intr_matrix` (C++ function), 1477
- `esp_rom_set_cpu_ticks_per_us` (C++ function), 1477
- `esp_rom_software_reset_cpu` (C++ function), 1476
- `esp_rom_software_reset_system` (C++ function), 1476
- `esp_secure_boot_key_digests_t` (C++ struct), 1260
- `esp_secure_boot_key_digests_t::key_digests` (C++ member), 1261
- `esp_secure_boot_read_key_digests` (C++ function), 1260
- `esp_set_deep_sleep_wake_stub` (C++ function), 1551
- `esp_set_deep_sleep_wake_stub_default_entry` (C++ function), 1551
- `esp_sleep_config_gpio_isolate` (C++ function), 1551
- `esp_sleep_cpu_retention_deinit` (C++ function), 1551
- `esp_sleep_cpu_retention_init` (C++ function), 1551
- `esp_sleep_disable_bt_wakeup` (C++ function), 1548
- `esp_sleep_disable_wakeup_source` (C++ function), 1546
- `esp_sleep_disable_wifi_beacon_wakeup` (C++ function), 1548
- `esp_sleep_disable_wifi_wakeup` (C++ function), 1548
- `esp_sleep_enable_bt_wakeup` (C++ function), 1548
- `esp_sleep_enable_gpio_switch` (C++ function), 1551
- `esp_sleep_enable_gpio_wakeup` (C++ function), 1547
- `esp_sleep_enable_timer_wakeup` (C++ function), 1546
- `esp_sleep_enable_uart_wakeup` (C++ function), 1547
- `esp_sleep_enable_wifi_beacon_wakeup` (C++ function), 1548
- `esp_sleep_enable_wifi_wakeup` (C++ function), 1548
- `esp_sleep_get_ext1_wakeup_status` (C++ function), 1548
- `esp_sleep_get_gpio_wakeup_status` (C++ function), 1548
- `esp_sleep_get_wakeup_cause` (C++ function), 1550
- `esp_sleep_get_wakeup_causes` (C++ function), 1550
- `esp_sleep_is_valid_wakeup_gpio` (C++ function), 1546
- `esp_sleep_mode_t` (C++ enum), 1554
- `esp_sleep_mode_t::ESP_SLEEP_MODE_DEEP_SLEEP` (C++ enumerator), 1554
- `esp_sleep_mode_t::ESP_SLEEP_MODE_LIGHT_SLEEP` (C++ enumerator), 1554
- `esp_sleep_pd_config` (C++ function), 1548
- `esp_sleep_pd_domain_t` (C++ enum), 1552
- `esp_sleep_pd_domain_t::ESP_PD_DOMAIN_CPU` (C++ enumerator), 1552
- `esp_sleep_pd_domain_t::ESP_PD_DOMAIN_MAX` (C++ enumerator), 1552
- `esp_sleep_pd_domain_t::ESP_PD_DOMAIN_MODEM` (C++ enumerator), 1552
- `esp_sleep_pd_domain_t::ESP_PD_DOMAIN_RC32K` (C++ enumerator), 1552
- `esp_sleep_pd_domain_t::ESP_PD_DOMAIN_RC_FAST` (C++ enumerator), 1552
- `esp_sleep_pd_domain_t::ESP_PD_DOMAIN_TOP` (C++ enumerator), 1552
- `esp_sleep_pd_domain_t::ESP_PD_DOMAIN_VDDSDIO` (C++ enumerator), 1552
- `esp_sleep_pd_domain_t::ESP_PD_DOMAIN_XTAL` (C++ enumerator), 1552
- `esp_sleep_pd_domain_t::ESP_PD_DOMAIN_XTAL32K` (C++ enumerator), 1552
- `esp_sleep_pd_option_t` (C++ enum), 1552
- `esp_sleep_pd_option_t::ESP_PD_OPTION_AUTO` (C++ enumerator), 1553
- `esp_sleep_pd_option_t::ESP_PD_OPTION_OFF` (C++ enumerator), 1553
- `esp_sleep_pd_option_t::ESP_PD_OPTION_ON` (C++ enumerator), 1553
- `esp_sleep_source_t` (C++ enum), 1553
- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_ALL` (C++ enumerator), 1553
- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_BT` (C++ enumerator), 1554
- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_COCPU` (C++ enumerator), 1554
- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_COCPU_TRAP_T` (C++ enumerator), 1554
- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_EXT0` (C++ enumerator), 1553
- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_EXT1` (C++ enumerator), 1553
- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_GPIO` (C++ enumerator), 1553
- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_TIMER` (C++ enumerator), 1553
- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_TOUCHPAD` (C++ enumerator), 1553
- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_UART` (C++ enumerator), 1553
- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_UART1` (C++ enumerator), 1553
- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_UART2` (C++ enumerator), 1553

- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_ULP` (C++ enumerator), 1553
- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_UNDEFINED` (C++ enumerator), 1553
- `esp_sleep_source_t::ESP_SLEEP_WAKEUP_WIFI` (C++ enumerator), 1554
- `esp_sleep_wakeup_cause_t` (C++ type), 1552
- `esp_sntp_config` (C++ struct), 238
- `esp_sntp_config::index_of_first_server` (C++ member), 238
- `esp_sntp_config::ip_event_to_renew` (C++ member), 238
- `esp_sntp_config::num_of_servers` (C++ member), 238
- `esp_sntp_config::renew_servers_after_renew` (C++ member), 238
- `esp_sntp_config::server_from_dhcp` (C++ member), 238
- `esp_sntp_config::servers` (C++ member), 238
- `esp_sntp_config::smooth_sync` (C++ member), 238
- `esp_sntp_config::start` (C++ member), 238
- `esp_sntp_config::sync_cb` (C++ member), 238
- `esp_sntp_config::wait_for_sync` (C++ member), 238
- `esp_sntp_config_t` (C++ type), 239
- `esp_sntp_enabled` (C++ function), 1574
- `esp_sntp_get_sync_interval` (C macro), 1575
- `esp_sntp_get_sync_mode` (C macro), 1575
- `esp_sntp_get_sync_status` (C macro), 1575
- `esp_sntp_getoperatingmode` (C++ function), 1575
- `esp_sntp_getreachability` (C++ function), 1575
- `esp_sntp_getserver` (C++ function), 1574
- `esp_sntp_getservername` (C++ function), 1574
- `esp_sntp_init` (C++ function), 1574
- `esp_sntp_operatingmode_t` (C++ enum), 1576
- `esp_sntp_operatingmode_t::ESP_SNTP_OPMODE_LISTEN_ONLY` (C++ enumerator), 1576
- `esp_sntp_operatingmode_t::ESP_SNTP_OPMODE_POLL` (C++ enumerator), 1576
- `esp_sntp_restart` (C macro), 1575
- `ESP_SNTP_SERVER_LIST` (C macro), 238
- `esp_sntp_set_sync_interval` (C macro), 1575
- `esp_sntp_set_sync_mode` (C macro), 1575
- `esp_sntp_set_sync_status` (C macro), 1575
- `esp_sntp_set_time_sync_notification_cb` (C macro), 1575
- `esp_sntp_setoperatingmode` (C++ function), 1574
- `esp_sntp_setserver` (C++ function), 1574
- `esp_sntp_setservername` (C++ function), 1574
- `esp_sntp_stop` (C++ function), 1574
- `esp_sntp_sync_time` (C macro), 1575
- `esp_sntp_time_cb_t` (C++ type), 239
- `esp_spiffs_check` (C++ function), 1183
- `esp_spiffs_format` (C++ function), 1183
- `esp_spiffs_gc` (C++ function), 1183
- `esp_spiffs_info` (C++ function), 1183
- `esp_spiffs_mounted` (C++ function), 1183
- `esp_srp_exchange_proofs` (C++ function), 1091
- `esp_srp_free` (C++ function), 1089
- `esp_srp_gen_salt_verifier` (C++ function), 1089
- `esp_srp_get_session_key` (C++ function), 1090
- `esp_srp_handle_t` (C++ type), 1091
- `esp_srp_init` (C++ function), 1089
- `esp_srp_set_salt_verifier` (C++ function), 1090
- `esp_srp_srv_pubkey` (C++ function), 1089
- `esp_srp_srv_pubkey_from_salt_verifier` (C++ function), 1090
- `esp_system_abort` (C++ function), 1502
- `esp_systick_new_etm_alarm_event` (C++ function), 290
- `esp_sysview_flush` (C++ function), 1217
- `esp_sysview_heap_trace_alloc` (C++ function), 1217
- `esp_sysview_heap_trace_free` (C++ function), 1217
- `esp_sysview_heap_trace_start` (C++ function), 1217
- `esp_sysview_heap_trace_stop` (C++ function), 1217
- `esp_sysview_vprintf` (C++ function), 1217
- `esp_task_wdt_add` (C++ function), 1584
- `esp_task_wdt_add_user` (C++ function), 1584
- `esp_task_wdt_config_t` (C++ struct), 1586
- `esp_task_wdt_config_t::idle_core_mask` (C++ member), 1586
- `esp_task_wdt_config_t::timeout_ms` (C++ member), 1586
- `esp_task_wdt_config_t::trigger_panic` (C++ member), 1586
- `esp_task_wdt_deinit` (C++ function), 1584
- `esp_task_wdt_delete` (C++ function), 1585
- `esp_task_wdt_delete_user` (C++ function), 1585
- `esp_task_wdt_init` (C++ function), 1583
- `esp_task_wdt_isr_user_handler` (C++ function), 1585
- `esp_task_wdt_print_triggered_tasks` (C++ function), 1585
- `esp_task_wdt_reconfigure` (C++ function), 1583
- `esp_task_wdt_reset` (C++ function), 1584
- `esp_task_wdt_reset_user` (C++ function), 1584
- `esp_task_wdt_status` (C++ function), 1585
- `esp_task_wdt_user_handle_t` (C++ type), 1586

- `esp_timer_cb_t` (C++ type), 1475
- `esp_timer_create` (C++ function), 1472
- `esp_timer_create_args_t` (C++ struct), 1475
- `esp_timer_create_args_t::arg` (C++ member), 1475
- `esp_timer_create_args_t::callback` (C++ member), 1475
- `esp_timer_create_args_t::dispatch_method` (C++ member), 1475
- `esp_timer_create_args_t::name` (C++ member), 1475
- `esp_timer_create_args_t::skip_unhandled_event` (C++ member), 1475
- `esp_timer_deinit` (C++ function), 1471
- `esp_timer_delete` (C++ function), 1473
- `esp_timer_dispatch_t` (C++ enum), 1475
- `esp_timer_dispatch_t::ESP_TIMER_MAX` (C++ enumerator), 1475
- `esp_timer_dispatch_t::ESP_TIMER_TASK` (C++ enumerator), 1475
- `esp_timer_dump` (C++ function), 1474
- `esp_timer_early_init` (C++ function), 1471
- `esp_timer_get_expiry_time` (C++ function), 1473
- `esp_timer_get_next_alarm` (C++ function), 1473
- `esp_timer_get_next_alarm_for_wake_up` (C++ function), 1473
- `esp_timer_get_period` (C++ function), 1473
- `esp_timer_get_time` (C++ function), 1473
- `esp_timer_handle_t` (C++ type), 1475
- `esp_timer_init` (C++ function), 1471
- `esp_timer_is_active` (C++ function), 1474
- `esp_timer_isr_dispatch_need_yield` (C++ function), 1474
- `esp_timer_new_etm_alarm_event` (C++ function), 1474
- `esp_timer_restart` (C++ function), 1472
- `esp_timer_start_once` (C++ function), 1472
- `esp_timer_start_periodic` (C++ function), 1472
- `esp_timer_stop` (C++ function), 1473
- `esp_tls_addr_family` (C++ enum), 73
- `esp_tls_addr_family::ESP_TLS_AF_INET` (C++ enumerator), 73
- `esp_tls_addr_family::ESP_TLS_AF_INET6` (C++ enumerator), 73
- `esp_tls_addr_family::ESP_TLS_AF_UNSPEC` (C++ enumerator), 73
- `esp_tls_addr_family_t` (C++ type), 72
- `esp_tls_cfg` (C++ struct), 69
- `esp_tls_cfg::addr_family` (C++ member), 72
- `esp_tls_cfg::alpn_protos` (C++ member), 70
- `esp_tls_cfg::cacert_buf` (C++ member), 70
- `esp_tls_cfg::cacert_bytes` (C++ member), 70
- `esp_tls_cfg::cacert_pem_buf` (C++ member), 70
- `esp_tls_cfg::cacert_pem_bytes` (C++ member), 70
- `esp_tls_cfg::ciphersuites_list` (C++ member), 72
- `esp_tls_cfg::clientcert_buf` (C++ member), 70
- `esp_tls_cfg::clientcert_bytes` (C++ member), 70
- `esp_tls_cfg::clientcert_pem_buf` (C++ member), 70
- `esp_tls_cfg::clientcert_pem_bytes` (C++ member), 70
- `esp_tls_cfg::clientkey_buf` (C++ member), 70
- `esp_tls_cfg::clientkey_bytes` (C++ member), 70
- `esp_tls_cfg::clientkey_password` (C++ member), 70
- `esp_tls_cfg::clientkey_password_len` (C++ member), 71
- `esp_tls_cfg::clientkey_pem_buf` (C++ member), 70
- `esp_tls_cfg::clientkey_pem_bytes` (C++ member), 70
- `esp_tls_cfg::common_name` (C++ member), 71
- `esp_tls_cfg::crt_bundle_attach` (C++ member), 71
- `esp_tls_cfg::ds_data` (C++ member), 71
- `esp_tls_cfg::ecdsa_key_efuse_blk` (C++ member), 71
- `esp_tls_cfg::if_name` (C++ member), 72
- `esp_tls_cfg::is_plain_tcp` (C++ member), 71
- `esp_tls_cfg::keep_alive_cfg` (C++ member), 71
- `esp_tls_cfg::non_block` (C++ member), 71
- `esp_tls_cfg::psk_hint_key` (C++ member), 71
- `esp_tls_cfg::skip_common_name` (C++ member), 71
- `esp_tls_cfg::timeout_ms` (C++ member), 71
- `esp_tls_cfg::tls_version` (C++ member), 72
- `esp_tls_cfg::use_ecdsa_peripheral` (C++ member), 71
- `esp_tls_cfg::use_global_ca_store` (C++ member), 71
- `esp_tls_cfg::use_secure_element` (C++ member), 71
- `esp_tls_cfg_t` (C++ type), 72
- `esp_tls_conn_destroy` (C++ function), 66
- `esp_tls_conn_http_new` (C++ function), 64
- `esp_tls_conn_http_new_async` (C++ function), 65
- `esp_tls_conn_http_new_sync` (C++ function), 64
- `esp_tls_conn_new_async` (C++ function), 65
- `esp_tls_conn_new_sync` (C++ function), 64
- `esp_tls_conn_read` (C++ function), 65

Espressif Systems

1954

Release v5.2.7

- (C++ member), 1105
- esp_vfs_fat_mount_config_t::max_files (C++ member), 1105
- esp_vfs_fat_register (C++ function), 1100
- esp_vfs_fat_sdcard_format (C++ function), 1103
- esp_vfs_fat_sdcard_unmount (C++ function), 1103
- esp_vfs_fat_sdmmc_mount (C++ function), 1101
- esp_vfs_fat_sdmmc_mount_config_t (C++ type), 1106
- esp_vfs_fat_sdmmc_unmount (C++ function), 1103
- esp_vfs_fat_sdspi_mount (C++ function), 1102
- esp_vfs_fat_spiflash_format_rw_wl (C++ function), 1104
- esp_vfs_fat_spiflash_mount_ro (C++ function), 1104
- esp_vfs_fat_spiflash_mount_rw_wl (C++ function), 1103
- esp_vfs_fat_spiflash_unmount_ro (C++ function), 1105
- esp_vfs_fat_spiflash_unmount_rw_wl (C++ function), 1104
- esp_vfs_fat_unregister_path (C++ function), 1101
- ESP_VFS_FLAG_CONTEXT_PTR (C macro), 1197
- ESP_VFS_FLAG_DEFAULT (C macro), 1197
- ESP_VFS_FLAG_READONLY_FS (C macro), 1197
- esp_vfs_fstat (C++ function), 1189
- esp_vfs_id_t (C++ type), 1197
- esp_vfs_l2tap_eth_filter (C++ function), 251
- esp_vfs_l2tap_intf_register (C++ function), 250
- esp_vfs_l2tap_intf_unregister (C++ function), 250
- esp_vfs_link (C++ function), 1189
- esp_vfs_lseek (C++ function), 1188
- esp_vfs_open (C++ function), 1188
- ESP_VFS_PATH_MAX (C macro), 1197
- esp_vfs_pread (C++ function), 1191
- esp_vfs_pwrite (C++ function), 1191
- esp_vfs_read (C++ function), 1188
- esp_vfs_register (C++ function), 1189
- esp_vfs_register_fd (C++ function), 1190
- esp_vfs_register_fd_range (C++ function), 1189
- esp_vfs_register_fd_with_local_fd (C++ function), 1190
- esp_vfs_register_with_id (C++ function), 1189
- esp_vfs_rename (C++ function), 1189
- esp_vfs_select (C++ function), 1190
- esp_vfs_select_sem_t (C++ struct), 1191
- esp_vfs_select_sem_t::is_sem_local (C++ member), 1191
- esp_vfs_select_sem_t::sem (C++ member), 1191
- esp_vfs_select_triggered (C++ function), 1191
- esp_vfs_select_triggered_isr (C++ function), 1191
- esp_vfs_spiffs_conf_t (C++ struct), 1184
- esp_vfs_spiffs_conf_t::base_path (C++ member), 1184
- esp_vfs_spiffs_conf_t::format_if_mount_failed (C++ member), 1184
- esp_vfs_spiffs_conf_t::max_files (C++ member), 1184
- esp_vfs_spiffs_conf_t::partition_label (C++ member), 1184
- esp_vfs_spiffs_register (C++ function), 1182
- esp_vfs_spiffs_unregister (C++ function), 1182
- esp_vfs_stat (C++ function), 1189
- esp_vfs_t (C++ struct), 1191
- esp_vfs_t::access (C++ member), 1195
- esp_vfs_t::access_p (C++ member), 1195
- esp_vfs_t::close (C++ member), 1193
- esp_vfs_t::close_p (C++ member), 1193
- esp_vfs_t::closedir (C++ member), 1194
- esp_vfs_t::closedir_p (C++ member), 1194
- esp_vfs_t::end_select (C++ member), 1196
- esp_vfs_t::fcntl (C++ member), 1194
- esp_vfs_t::fcntl_p (C++ member), 1194
- esp_vfs_t::flags (C++ member), 1192
- esp_vfs_t::fstat (C++ member), 1193
- esp_vfs_t::fstat_p (C++ member), 1193
- esp_vfs_t::fsync (C++ member), 1195
- esp_vfs_t::fsync_p (C++ member), 1195
- esp_vfs_t::ftruncate (C++ member), 1195
- esp_vfs_t::ftruncate_p (C++ member), 1195
- esp_vfs_t::get_socket_select_semaphore (C++ member), 1196
- esp_vfs_t::ioctl (C++ member), 1195
- esp_vfs_t::ioctl_p (C++ member), 1194
- esp_vfs_t::link (C++ member), 1193
- esp_vfs_t::link_p (C++ member), 1193
- esp_vfs_t::lseek (C++ member), 1192
- esp_vfs_t::lseek_p (C++ member), 1192
- esp_vfs_t::mkdir (C++ member), 1194
- esp_vfs_t::mkdir_p (C++ member), 1194
- esp_vfs_t::open (C++ member), 1193
- esp_vfs_t::open_p (C++ member), 1192
- esp_vfs_t::opendir (C++ member), 1193
- esp_vfs_t::opendir_p (C++ member), 1193
- esp_vfs_t::pread (C++ member), 1192
- esp_vfs_t::pread_p (C++ member), 1192
- esp_vfs_t::pwrite (C++ member), 1192
- esp_vfs_t::pwrite_p (C++ member), 1192
- esp_vfs_t::read (C++ member), 1192
- esp_vfs_t::read_p (C++ member), 1192

- `esp_vfs_t::readdir` (C++ member), 1194
- `esp_vfs_t::readdir_p` (C++ member), 1193
- `esp_vfs_t::readdir_r` (C++ member), 1194
- `esp_vfs_t::readdir_r_p` (C++ member), 1194
- `esp_vfs_t::rename` (C++ member), 1193
- `esp_vfs_t::rename_p` (C++ member), 1193
- `esp_vfs_t::rmdir` (C++ member), 1194
- `esp_vfs_t::rmdir_p` (C++ member), 1194
- `esp_vfs_t::seekdir` (C++ member), 1194
- `esp_vfs_t::seekdir_p` (C++ member), 1194
- `esp_vfs_t::socket_select` (C++ member), 1196
- `esp_vfs_t::start_select` (C++ member), 1196
- `esp_vfs_t::stat` (C++ member), 1193
- `esp_vfs_t::stat_p` (C++ member), 1193
- `esp_vfs_t::stop_socket_select` (C++ member), 1196
- `esp_vfs_t::stop_socket_select_isr` (C++ member), 1196
- `esp_vfs_t::tcdrain` (C++ member), 1196
- `esp_vfs_t::tcdrain_p` (C++ member), 1195
- `esp_vfs_t::tcflow` (C++ member), 1196
- `esp_vfs_t::tcflow_p` (C++ member), 1196
- `esp_vfs_t::tcflush` (C++ member), 1196
- `esp_vfs_t::tcflush_p` (C++ member), 1196
- `esp_vfs_t::tcgetattr` (C++ member), 1195
- `esp_vfs_t::tcgetattr_p` (C++ member), 1195
- `esp_vfs_t::tcgetsid` (C++ member), 1196
- `esp_vfs_t::tcgetsid_p` (C++ member), 1196
- `esp_vfs_t::tcsendbreak` (C++ member), 1196
- `esp_vfs_t::tcsendbreak_p` (C++ member), 1196
- `esp_vfs_t::tcsetattr` (C++ member), 1195
- `esp_vfs_t::tcsetattr_p` (C++ member), 1195
- `esp_vfs_t::telldir` (C++ member), 1194
- `esp_vfs_t::telldir_p` (C++ member), 1194
- `esp_vfs_t::truncate` (C++ member), 1195
- `esp_vfs_t::truncate_p` (C++ member), 1195
- `esp_vfs_t::unlink` (C++ member), 1193
- `esp_vfs_t::unlink_p` (C++ member), 1193
- `esp_vfs_t::utime` (C++ member), 1195
- `esp_vfs_t::utime_p` (C++ member), 1195
- `esp_vfs_t::write` (C++ member), 1192
- `esp_vfs_t::write_p` (C++ member), 1192
- `esp_vfs_unlink` (C++ function), 1189
- `esp_vfs_unregister` (C++ function), 1189
- `esp_vfs_unregister_fd` (C++ function), 1190
- `esp_vfs_unregister_with_id` (C++ function), 1190
- `esp_vfs_usb_serial_jtag_use_driver` (C++ function), 1199
- `esp_vfs_usb_serial_jtag_use_nonblocking` (C++ function), 1199
- `esp_vfs_utime` (C++ function), 1189
- `esp_vfs_write` (C++ function), 1188
- `esp_wake_deep_sleep` (C++ function), 1550
- `essl_clear_intr` (C++ function), 111
- `essl_get_intr` (C++ function), 112
- `essl_get_intr_ena` (C++ function), 112
- `essl_get_packet` (C++ function), 110
- `essl_get_rx_data_size` (C++ function), 110
- `essl_get_tx_buffer_num` (C++ function), 109
- `essl_handle_t` (C++ type), 112
- `essl_init` (C++ function), 109
- `essl_read_reg` (C++ function), 111
- `essl_reset_cnt` (C++ function), 110
- `essl_sdio_config_t` (C++ struct), 113
- `essl_sdio_config_t::card` (C++ member), 113
- `essl_sdio_config_t::recv_buffer_size` (C++ member), 113
- `essl_sdio_deinit_dev` (C++ function), 113
- `essl_sdio_init_dev` (C++ function), 113
- `essl_send_packet` (C++ function), 110
- `essl_send_slave_intr` (C++ function), 112
- `essl_set_intr_ena` (C++ function), 112
- `essl_spi_config_t` (C++ struct), 118
- `essl_spi_config_t::rx_sync_reg` (C++ member), 119
- `essl_spi_config_t::spi` (C++ member), 119
- `essl_spi_config_t::tx_buf_size` (C++ member), 119
- `essl_spi_config_t::tx_sync_reg` (C++ member), 119
- `essl_spi_deinit_dev` (C++ function), 113
- `essl_spi_get_packet` (C++ function), 114
- `essl_spi_init_dev` (C++ function), 113
- `essl_spi_rdbuf` (C++ function), 115
- `essl_spi_rdbuf_polling` (C++ function), 116
- `essl_spi_rddma` (C++ function), 117
- `essl_spi_rddma_done` (C++ function), 117
- `essl_spi_rddma_seg` (C++ function), 117
- `essl_spi_read_reg` (C++ function), 114
- `essl_spi_reset_cnt` (C++ function), 115
- `essl_spi_send_packet` (C++ function), 115
- `essl_spi_wrbuf` (C++ function), 116
- `essl_spi_wrbuf_polling` (C++ function), 116
- `essl_spi_wrdma` (C++ function), 117
- `essl_spi_wrdma_done` (C++ function), 118
- `essl_spi_wrdma_seg` (C++ function), 118
- `essl_spi_write_reg` (C++ function), 114
- `essl_wait_for_ready` (C++ function), 109
- `essl_wait_int` (C++ function), 111
- `essl_write_reg` (C++ function), 111
- `eTaskGetState` (C++ function), 1302
- `eTaskState` (C++ enum), 1324
- `eTaskState::eBlocked` (C++ enumerator), 1324
- `eTaskState::eDeleted` (C++ enumerator), 1324
- `eTaskState::eInvalid` (C++ enumerator), 1324
- `eTaskState::eReady` (C++ enumerator), 1324
- `eTaskState::eRunning` (C++ enumerator), 1324
- `eTaskState::eSuspended` (C++ enumerator), 1324
- `ETH_DEFAULT_CONFIG` (C macro), 181
- `ETH_DEFAULT_SPI` (C macro), 191

eth_event_t (C++ enum), 184
 eth_event_t::ETHERNET_EVENT_CONNECTED (C++ enumerator), 184
 eth_event_t::ETHERNET_EVENT_DISCONNECTED (C++ enumerator), 184
 eth_event_t::ETHERNET_EVENT_START (C++ enumerator), 184
 eth_event_t::ETHERNET_EVENT_STOP (C++ enumerator), 184
 eth_mac_clock_config_t (C++ union), 185
 eth_mac_clock_config_t::clock_gpio (C++ member), 185
 eth_mac_clock_config_t::clock_mode (C++ member), 185
 eth_mac_clock_config_t::mii (C++ member), 185
 eth_mac_clock_config_t::rmii (C++ member), 185
 eth_mac_config_t (C++ struct), 189
 eth_mac_config_t::flags (C++ member), 190
 eth_mac_config_t::rx_task_prio (C++ member), 190
 eth_mac_config_t::rx_task_stack_size (C++ member), 189
 eth_mac_config_t::sw_reset_timeout_ms (C++ member), 189
 ETH_MAC_DEFAULT_CONFIG (C macro), 191
 ETH_MAC_FLAG_PIN_TO_CORE (C macro), 191
 ETH_MAC_FLAG_WORK_WITH_CACHE_DISABLE (C macro), 191
 eth_phy_autoneg_cmd_t (C++ enum), 197
 eth_phy_autoneg_cmd_t::ESP_ETH_PHY_AUTONEGO_DISABLE (C++ enumerator), 197
 eth_phy_autoneg_cmd_t::ESP_ETH_PHY_AUTONEGO_ENABLE (C++ enumerator), 197
 eth_phy_autoneg_cmd_t::ESP_ETH_PHY_AUTONEGO_GSTAT (C++ enumerator), 197
 eth_phy_autoneg_cmd_t::ESP_ETH_PHY_AUTONEGO_RSTART (C++ enumerator), 197
 eth_phy_config_t (C++ struct), 197
 eth_phy_config_t::autonego_timeout_ms (C++ member), 197
 eth_phy_config_t::phy_addr (C++ member), 197
 eth_phy_config_t::reset_gpio_num (C++ member), 197
 eth_phy_config_t::reset_timeout_ms (C++ member), 197
 ETH_PHY_DEFAULT_CONFIG (C macro), 197
 eth_spi_custom_driver_config_t (C++ struct), 190
 eth_spi_custom_driver_config_t::config (C++ member), 190
 eth_spi_custom_driver_config_t::deinit (C++ member), 190
 eth_spi_custom_driver_config_t::init (C++ member), 190
 eth_spi_custom_driver_config_t::read (C++ member), 190
 eth_spi_custom_driver_config_t::write (C++ member), 191
 ETS_INTERNAL_INTR_SOURCE_OFF (C macro), 1489
 ETS_INTERNAL_PROFILING_INTR_SOURCE (C macro), 1489
 ETS_INTERNAL_SW0_INTR_SOURCE (C macro), 1489
 ETS_INTERNAL_SW1_INTR_SOURCE (C macro), 1489
 ETS_INTERNAL_TIMER0_INTR_SOURCE (C macro), 1489
 ETS_INTERNAL_TIMER1_INTR_SOURCE (C macro), 1489
 ETS_INTERNAL_TIMER2_INTR_SOURCE (C macro), 1489
 ETS_INTERNAL_UNUSED_INTR_SOURCE (C macro), 1489
 EventBits_t (C++ type), 1385
 eventfd (C++ function), 1200
 EventGroupHandle_t (C++ type), 1385

F

ff_diskio_impl_t (C++ struct), 1098
 ff_diskio_impl_t::init (C++ member), 1098
 ff_diskio_impl_t::ioctl (C++ member), 1099
 ff_diskio_impl_t::read (C++ member), 1099
 ff_diskio_impl_t::status (C++ member), 1098
 ff_diskio_impl_t::write (C++ member), 1099
 ff_diskio_register (C++ function), 1098
 ff_diskio_register_raw_partition (C++ function), 1099
 ff_diskio_register_sdmmc (C++ function), 1099
 ff_diskio_register_wl_partition (C++ function), 1099

G

gpio_config (C++ function), 293
 gpio_config_t (C++ struct), 300
 gpio_config_t::hys_ctrl_mode (C++ member), 301
 gpio_config_t::intr_type (C++ member), 300
 gpio_config_t::mode (C++ member), 300
 gpio_config_t::pin_bit_mask (C++ member), 300
 gpio_config_t::pull_down_en (C++ member), 300
 gpio_config_t::pull_up_en (C++ member), 300
 gpio_deep_sleep_wakeup_disable (C++ function), 300

- `gpio_deep_sleep_wakeup_enable` (C++ function), 300
- `gpio_drive_cap_t` (C++ enum), 305
- `gpio_drive_cap_t::GPIO_DRIVE_CAP_0` (C++ enumerator), 305
- `gpio_drive_cap_t::GPIO_DRIVE_CAP_1` (C++ enumerator), 306
- `gpio_drive_cap_t::GPIO_DRIVE_CAP_2` (C++ enumerator), 306
- `gpio_drive_cap_t::GPIO_DRIVE_CAP_3` (C++ enumerator), 306
- `gpio_drive_cap_t::GPIO_DRIVE_CAP_DEFAULT` (C++ enumerator), 306
- `gpio_drive_cap_t::GPIO_DRIVE_CAP_MAX` (C++ enumerator), 306
- `gpio_dump_io_configuration` (C++ function), 300
- `gpio_etm_event_bind_gpio` (C++ function), 287
- `gpio_etm_event_config_t` (C++ struct), 289
- `gpio_etm_event_config_t::edge` (C++ member), 289
- `gpio_etm_event_config_t::edges` (C++ member), 289
- `gpio_etm_event_edge_t` (C++ enum), 289
- `gpio_etm_event_edge_t::GPIO_ETM_EVENT_EDGE_ANY` (C++ enumerator), 290
- `gpio_etm_event_edge_t::GPIO_ETM_EVENT_EDGE_FALLING` (C++ enumerator), 289
- `gpio_etm_event_edge_t::GPIO_ETM_EVENT_EDGE_RISING` (C++ enumerator), 289
- `GPIO_ETM_EVENT_EDGE_TYPES` (C macro), 289
- `gpio_etm_task_action_t` (C++ enum), 290
- `gpio_etm_task_action_t::GPIO_ETM_TASK_ACTION_DISABLE` (C++ enumerator), 290
- `gpio_etm_task_action_t::GPIO_ETM_TASK_ACTION_ENABLE` (C++ enumerator), 290
- `gpio_etm_task_action_t::GPIO_ETM_TASK_ACTION_TOGGLE` (C++ enumerator), 290
- `GPIO_ETM_TASK_ACTION_TYPES` (C macro), 289
- `gpio_etm_task_add_gpio` (C++ function), 288
- `gpio_etm_task_config_t` (C++ struct), 289
- `gpio_etm_task_config_t::action` (C++ member), 289
- `gpio_etm_task_config_t::actions` (C++ member), 289
- `gpio_etm_task_rm_gpio` (C++ function), 288
- `gpio_force_hold_all` (C++ function), 298
- `gpio_force_unhold_all` (C++ function), 299
- `gpio_get_drive_capability` (C++ function), 297
- `gpio_get_level` (C++ function), 295
- `gpio_hold_dis` (C++ function), 298
- `gpio_hold_en` (C++ function), 297
- `gpio_hys_ctrl_mode_t` (C++ enum), 306
- `gpio_hys_ctrl_mode_t::GPIO_HYS_SOFT_DISABLE` (C++ enumerator), 306
- `gpio_hys_ctrl_mode_t::GPIO_HYS_SOFT_ENABLE` (C++ enumerator), 306
- `gpio_install_isr_service` (C++ function), 296
- `gpio_int_type_t` (C++ enum), 304
- `gpio_int_type_t::GPIO_INTR_ANYEDGE` (C++ enumerator), 304
- `gpio_int_type_t::GPIO_INTR_DISABLE` (C++ enumerator), 304
- `gpio_int_type_t::GPIO_INTR_HIGH_LEVEL` (C++ enumerator), 304
- `gpio_int_type_t::GPIO_INTR_LOW_LEVEL` (C++ enumerator), 304
- `gpio_int_type_t::GPIO_INTR_MAX` (C++ enumerator), 304
- `gpio_int_type_t::GPIO_INTR_NEGEDGE` (C++ enumerator), 304
- `gpio_int_type_t::GPIO_INTR_POSEDGE` (C++ enumerator), 304
- `gpio_intr_disable` (C++ function), 294
- `gpio_intr_enable` (C++ function), 294
- `gpio_iomux_in` (C++ function), 298
- `gpio_iomux_out` (C++ function), 298
- `GPIO_IS_DEEP_SLEEP_WAKEUP_VALID_GPIO` (C macro), 301
- `GPIO_IS_VALID_DIGITAL_IO_PAD` (C macro), 301
- `GPIO_IS_VALID_GPIO` (C macro), 301
- `GPIO_IS_VALID_OUTPUT_GPIO` (C macro), 301
- `gpio_isr_handle_t` (C++ type), 301
- `gpio_isr_handler_add` (C++ function), 297
- `gpio_isr_handler_remove` (C++ function), 297
- `gpio_isr_register` (C++ function), 296
- `gpio_isr_t` (C++ type), 301
- `gpio_mode_t` (C++ enum), 304
- `gpio_mode_t::GPIO_MODE_DISABLE` (C++ enumerator), 304
- `gpio_mode_t::GPIO_MODE_INPUT` (C++ enumerator), 304
- `gpio_mode_t::GPIO_MODE_INPUT_OUTPUT` (C++ enumerator), 305
- `gpio_mode_t::GPIO_MODE_INPUT_OUTPUT_OD` (C++ enumerator), 305
- `gpio_mode_t::GPIO_MODE_OUTPUT` (C++ enumerator), 304
- `gpio_mode_t::GPIO_MODE_OUTPUT_OD` (C++ enumerator), 305
- `gpio_new_etm_event` (C++ function), 286
- `gpio_new_etm_task` (C++ function), 287
- `GPIO_PIN_COUNT` (C macro), 301
- `GPIO_PIN_REG_0` (C macro), 301
- `GPIO_PIN_REG_1` (C macro), 301
- `GPIO_PIN_REG_10` (C macro), 302
- `GPIO_PIN_REG_11` (C macro), 302
- `GPIO_PIN_REG_12` (C macro), 302
- `GPIO_PIN_REG_13` (C macro), 302
- `GPIO_PIN_REG_14` (C macro), 302
- `GPIO_PIN_REG_15` (C macro), 302
- `GPIO_PIN_REG_16` (C macro), 302

- GPIO_PIN_REG_17 (*C macro*), 302
- GPIO_PIN_REG_18 (*C macro*), 302
- GPIO_PIN_REG_19 (*C macro*), 302
- GPIO_PIN_REG_2 (*C macro*), 301
- GPIO_PIN_REG_20 (*C macro*), 302
- GPIO_PIN_REG_21 (*C macro*), 302
- GPIO_PIN_REG_22 (*C macro*), 302
- GPIO_PIN_REG_23 (*C macro*), 302
- GPIO_PIN_REG_24 (*C macro*), 302
- GPIO_PIN_REG_25 (*C macro*), 302
- GPIO_PIN_REG_26 (*C macro*), 302
- GPIO_PIN_REG_27 (*C macro*), 302
- GPIO_PIN_REG_28 (*C macro*), 302
- GPIO_PIN_REG_29 (*C macro*), 302
- GPIO_PIN_REG_3 (*C macro*), 301
- GPIO_PIN_REG_30 (*C macro*), 302
- GPIO_PIN_REG_31 (*C macro*), 303
- GPIO_PIN_REG_32 (*C macro*), 303
- GPIO_PIN_REG_33 (*C macro*), 303
- GPIO_PIN_REG_34 (*C macro*), 303
- GPIO_PIN_REG_35 (*C macro*), 303
- GPIO_PIN_REG_36 (*C macro*), 303
- GPIO_PIN_REG_37 (*C macro*), 303
- GPIO_PIN_REG_38 (*C macro*), 303
- GPIO_PIN_REG_39 (*C macro*), 303
- GPIO_PIN_REG_4 (*C macro*), 301
- GPIO_PIN_REG_40 (*C macro*), 303
- GPIO_PIN_REG_41 (*C macro*), 303
- GPIO_PIN_REG_42 (*C macro*), 303
- GPIO_PIN_REG_43 (*C macro*), 303
- GPIO_PIN_REG_44 (*C macro*), 303
- GPIO_PIN_REG_45 (*C macro*), 303
- GPIO_PIN_REG_46 (*C macro*), 303
- GPIO_PIN_REG_47 (*C macro*), 303
- GPIO_PIN_REG_48 (*C macro*), 303
- GPIO_PIN_REG_49 (*C macro*), 303
- GPIO_PIN_REG_5 (*C macro*), 301
- GPIO_PIN_REG_50 (*C macro*), 303
- GPIO_PIN_REG_51 (*C macro*), 303
- GPIO_PIN_REG_52 (*C macro*), 303
- GPIO_PIN_REG_53 (*C macro*), 303
- GPIO_PIN_REG_54 (*C macro*), 304
- GPIO_PIN_REG_55 (*C macro*), 304
- GPIO_PIN_REG_56 (*C macro*), 304
- GPIO_PIN_REG_6 (*C macro*), 301
- GPIO_PIN_REG_7 (*C macro*), 301
- GPIO_PIN_REG_8 (*C macro*), 302
- GPIO_PIN_REG_9 (*C macro*), 302
- gpio_port_t (*C++ enum*), 304
- gpio_port_t::GPIO_PORT_0 (*C++ enumerator*), 304
- gpio_port_t::GPIO_PORT_MAX (*C++ enumerator*), 304
- gpio_pull_mode_t (*C++ enum*), 305
- gpio_pull_mode_t::GPIO_FLOATING (*C++ enumerator*), 305
- gpio_pull_mode_t::GPIO_PULLDOWN_ONLY (*C++ enumerator*), 305
- gpio_pull_mode_t::GPIO_PULLUP_ONLY (*C++ enumerator*), 305
- gpio_pull_mode_t::GPIO_PULLUP_PULLDOWN (*C++ enumerator*), 305
- gpio_pulldown_dis (*C++ function*), 296
- gpio_pulldown_en (*C++ function*), 296
- gpio_pulldown_t (*C++ enum*), 305
- gpio_pulldown_t::GPIO_PULLDOWN_DISABLE (*C++ enumerator*), 305
- gpio_pulldown_t::GPIO_PULLDOWN_ENABLE (*C++ enumerator*), 305
- gpio_pullup_dis (*C++ function*), 296
- gpio_pullup_en (*C++ function*), 296
- gpio_pullup_t (*C++ enum*), 305
- gpio_pullup_t::GPIO_PULLUP_DISABLE (*C++ enumerator*), 305
- gpio_pullup_t::GPIO_PULLUP_ENABLE (*C++ enumerator*), 305
- gpio_reset_pin (*C++ function*), 293
- gpio_set_direction (*C++ function*), 295
- gpio_set_drive_capability (*C++ function*), 297
- gpio_set_intr_type (*C++ function*), 294
- gpio_set_level (*C++ function*), 294
- gpio_set_pull_mode (*C++ function*), 295
- gpio_sleep_sel_dis (*C++ function*), 299
- gpio_sleep_sel_en (*C++ function*), 299
- gpio_sleep_set_direction (*C++ function*), 299
- gpio_sleep_set_pull_mode (*C++ function*), 299
- gpio_uninstall_isr_service (*C++ function*), 297
- gpio_wakeup_disable (*C++ function*), 295
- gpio_wakeup_enable (*C++ function*), 295
- gptimer_alarm_cb_t (*C++ type*), 325
- gptimer_alarm_config_t (*C++ struct*), 323
- gptimer_alarm_config_t::alarm_count (*C++ member*), 323
- gptimer_alarm_config_t::auto_reload_on_alarm (*C++ member*), 323
- gptimer_alarm_config_t::flags (*C++ member*), 323
- gptimer_alarm_config_t::reload_count (*C++ member*), 323
- gptimer_alarm_event_data_t (*C++ struct*), 325
- gptimer_alarm_event_data_t::alarm_value (*C++ member*), 325
- gptimer_alarm_event_data_t::count_value (*C++ member*), 325
- gptimer_clock_source_t (*C++ type*), 325
- gptimer_config_t (*C++ struct*), 322
- gptimer_config_t::clk_src (*C++ member*), 322
- gptimer_config_t::direction (*C++ member*), 322
- gptimer_config_t::flags (*C++ member*), 322

gptimer_config_t::intr_priority (C++ member), 322
 gptimer_config_t::intr_shared (C++ member), 322
 gptimer_config_t::resolution_hz (C++ member), 322
 gptimer_count_direction_t (C++ enum), 326
 gptimer_count_direction_t::GPTIMER_COUNT_DOWN (C++ enumerator), 326
 gptimer_count_direction_t::GPTIMER_COUNT_UP (C++ enumerator), 326
 gptimer_del_timer (C++ function), 318
 gptimer_disable (C++ function), 321
 gptimer_enable (C++ function), 320
 gptimer_etm_event_config_t (C++ struct), 324
 gptimer_etm_event_config_t::event_type (C++ member), 324
 gptimer_etm_event_type_t (C++ enum), 326
 gptimer_etm_event_type_t::GPTIMER_ETM_EVENT_MATCH (C++ enumerator), 326
 gptimer_etm_event_type_t::GPTIMER_ETM_EVENT_MAX (C++ enumerator), 326
 gptimer_etm_task_config_t (C++ struct), 324
 gptimer_etm_task_config_t::task_type (C++ member), 324
 gptimer_etm_task_type_t (C++ enum), 326
 gptimer_etm_task_type_t::GPTIMER_ETM_TASK_CAPTURE (C++ enumerator), 326
 gptimer_etm_task_type_t::GPTIMER_ETM_TASK_EN_ALARM (C++ enumerator), 326
 gptimer_etm_task_type_t::GPTIMER_ETM_TASK_MAX (C++ enumerator), 326
 gptimer_etm_task_type_t::GPTIMER_ETM_TASK_RELOAD (C++ enumerator), 326
 gptimer_etm_task_type_t::GPTIMER_ETM_TASK_START_COUNT (C++ enumerator), 326
 gptimer_etm_task_type_t::GPTIMER_ETM_TASK_STOP_COUNT (C++ enumerator), 326
 gptimer_event_callbacks_t (C++ struct), 323
 gptimer_event_callbacks_t::on_alarm (C++ member), 323
 gptimer_get_captured_count (C++ function), 319
 gptimer_get_raw_count (C++ function), 318
 gptimer_get_resolution (C++ function), 319
 gptimer_handle_t (C++ type), 325
 gptimer_new_etm_event (C++ function), 324
 gptimer_new_etm_task (C++ function), 324
 gptimer_new_timer (C++ function), 317
 gptimer_register_event_callbacks (C++ function), 319
 gptimer_set_alarm_action (C++ function), 320
 gptimer_set_raw_count (C++ function), 318
 gptimer_start (C++ function), 321
 gptimer_stop (C++ function), 322

H
 heap_caps_add_region (C++ function), 1440
 heap_caps_add_region_with_caps (C++ function), 1441
 heap_caps_aligned_alloc (C++ function), 1434
 heap_caps_aligned_calloc (C++ function), 1435
 heap_caps_aligned_free (C++ function), 1434
 heap_caps_calloc (C++ function), 1435
 heap_caps_calloc_prefer (C++ function), 1438
 heap_caps_check_integrity (C++ function), 1436
 heap_caps_check_integrity_addr (C++ function), 1437
 heap_caps_check_integrity_all (C++ function), 1436
 heap_caps_dump (C++ function), 1438
 heap_caps_dump_all (C++ function), 1438
 heap_caps_enable_nonos_stack_heaps (C++ function), 1440
 heap_caps_free (C++ function), 1434
 heap_caps_get_allocated_size (C++ function), 1438
 heap_caps_get_free_size (C++ function), 1435
 heap_caps_get_info (C++ function), 1436
 heap_caps_get_largest_free_block (C++ function), 1436
 heap_caps_get_minimum_free_size (C++ function), 1435
 heap_caps_get_total_size (C++ function), 1435
 heap_caps_init (C++ function), 1440
 heap_caps_malloc (C++ function), 1434
 heap_caps_malloc_extmem_enable (C++ function), 1437
 heap_caps_malloc_prefer (C++ function), 1437
 heap_caps_print_heap_info (C++ function), 1436
 heap_caps_realloc (C++ function), 1434
 heap_caps_realloc_prefer (C++ function), 1437
 heap_caps_register_failed_alloc_callback (C++ function), 1434
 HEAP_IRAM_ATTR (C macro), 1438
 heap_trace_dump (C++ function), 1467
 heap_trace_dump_caps (C++ function), 1467
 heap_trace_get (C++ function), 1466
 heap_trace_get_count (C++ function), 1466
 heap_trace_init_standalone (C++ function), 1465
 heap_trace_init_tohost (C++ function), 1466
 heap_trace_mode_t (C++ enum), 1468
 heap_trace_mode_t::HEAP_TRACE_ALL (C++ enumerator), 1468

- heap_trace_mode_t::HEAP_TRACE_LEAKS (C++ *enumerator*), 1468
- heap_trace_record_t (C++ *struct*), 1467
- heap_trace_record_t (C++ *type*), 1468
- heap_trace_record_t::address (C++ *member*), 1467
- heap_trace_record_t::allocated_by (C++ *member*), 1467
- heap_trace_record_t::ccount (C++ *member*), 1467
- heap_trace_record_t::freed_by (C++ *member*), 1468
- heap_trace_record_t::size (C++ *member*), 1467
- heap_trace_resume (C++ *function*), 1466
- heap_trace_start (C++ *function*), 1466
- heap_trace_stop (C++ *function*), 1466
- heap_trace_summary (C++ *function*), 1467
- heap_trace_summary_t (C++ *struct*), 1468
- heap_trace_summary_t::capacity (C++ *member*), 1468
- heap_trace_summary_t::count (C++ *member*), 1468
- heap_trace_summary_t::has_overflowed (C++ *member*), 1468
- heap_trace_summary_t::high_water_mark (C++ *member*), 1468
- heap_trace_summary_t::mode (C++ *member*), 1468
- heap_trace_summary_t::total_allocations (C++ *member*), 1468
- heap_trace_summary_t::total_frees (C++ *member*), 1468
- hmac_key_id_t (C++ *enum*), 330
- hmac_key_id_t::HMAC_KEY0 (C++ *enumerator*), 330
- hmac_key_id_t::HMAC_KEY1 (C++ *enumerator*), 330
- hmac_key_id_t::HMAC_KEY2 (C++ *enumerator*), 330
- hmac_key_id_t::HMAC_KEY3 (C++ *enumerator*), 330
- hmac_key_id_t::HMAC_KEY4 (C++ *enumerator*), 330
- hmac_key_id_t::HMAC_KEY5 (C++ *enumerator*), 330
- hmac_key_id_t::HMAC_KEY_MAX (C++ *enumerator*), 331
- http_client_init_cb_t (C++ *type*), 1271
- http_event_handle_cb (C++ *type*), 91
- HTTPD_200 (C *macro*), 144
- HTTPD_204 (C *macro*), 145
- HTTPD_207 (C *macro*), 145
- HTTPD_400 (C *macro*), 145
- HTTPD_404 (C *macro*), 145
- HTTPD_408 (C *macro*), 145
- HTTPD_500 (C *macro*), 145
- httpd_close_func_t (C++ *type*), 148
- httpd_config (C++ *struct*), 141
- httpd_config::backlog_conn (C++ *member*), 141
- httpd_config::close_fn (C++ *member*), 142
- httpd_config::core_id (C++ *member*), 141
- httpd_config::ctrl_port (C++ *member*), 141
- httpd_config::enable_so_linger (C++ *member*), 142
- httpd_config::global_transport_ctx (C++ *member*), 142
- httpd_config::global_transport_ctx_free_fn (C++ *member*), 142
- httpd_config::global_user_ctx (C++ *member*), 142
- httpd_config::global_user_ctx_free_fn (C++ *member*), 142
- httpd_config::keep_alive_count (C++ *member*), 142
- httpd_config::keep_alive_enable (C++ *member*), 142
- httpd_config::keep_alive_idle (C++ *member*), 142
- httpd_config::keep_alive_interval (C++ *member*), 142
- httpd_config::linger_timeout (C++ *member*), 142
- httpd_config::lru_purge_enable (C++ *member*), 141
- httpd_config::max_open_sockets (C++ *member*), 141
- httpd_config::max_resp_headers (C++ *member*), 141
- httpd_config::max_uri_handlers (C++ *member*), 141
- httpd_config::open_fn (C++ *member*), 142
- httpd_config::recv_wait_timeout (C++ *member*), 141
- httpd_config::send_wait_timeout (C++ *member*), 141
- httpd_config::server_port (C++ *member*), 141
- httpd_config::stack_size (C++ *member*), 141
- httpd_config::task_priority (C++ *member*), 141
- httpd_config::uri_match_fn (C++ *member*), 143
- httpd_config_t (C++ *type*), 148
- HTTPD_DEFAULT_CONFIG (C *macro*), 145
- httpd_err_code_t (C++ *enum*), 149
- httpd_err_code_t::HTTPD_400_BAD_REQUEST (C++ *enumerator*), 149
- httpd_err_code_t::HTTPD_401_UNAUTHORIZED (C++ *enumerator*), 149
- httpd_err_code_t::HTTPD_403_FORBIDDEN (C++ *enumerator*), 149
- httpd_err_code_t::HTTPD_404_NOT_FOUND (C++ *enumerator*), 149

[httpd_err_code_t::HTTPD_405_METHOD_NOT_ALLOWED](#) (C++ enumerator), 149
[httpd_err_code_t::HTTPD_408_REQ_TIMEOUT](#) (C++ enumerator), 149
[httpd_err_code_t::HTTPD_411_LENGTH_REQUIRED](#) (C++ enumerator), 149
[httpd_err_code_t::HTTPD_414_URI_TOO_LONG](#) (C++ enumerator), 149
[httpd_err_code_t::HTTPD_431_REQ_HDR_FIELDS_TOO_LARGE](#) (C++ enumerator), 149
[httpd_err_code_t::HTTPD_500_INTERNAL_SERVER_ERROR](#) (C++ enumerator), 149
[httpd_err_code_t::HTTPD_501_METHOD_NOT_IMPLEMENTED](#) (C++ enumerator), 149
[httpd_err_code_t::HTTPD_505_VERSION_NOT_SUPPORTED](#) (C++ enumerator), 149
[httpd_err_code_t::HTTPD_ERR_CODE_MAX](#) (C++ enumerator), 149
[httpd_err_handler_func_t](#) (C++ type), 147
[httpd_free_ctx_fn_t](#) (C++ type), 148
[httpd_get_client_list](#) (C++ function), 140
[httpd_get_global_transport_ctx](#) (C++ function), 139
[httpd_get_global_user_ctx](#) (C++ function), 139
[httpd_handle_t](#) (C++ type), 147
[HTTPD_MAX_REQ_HDR_LEN](#) (C macro), 144
[HTTPD_MAX_URI_LEN](#) (C macro), 144
[httpd_method_t](#) (C++ type), 147
[httpd_open_func_t](#) (C++ type), 148
[httpd_pending_func_t](#) (C++ type), 147
[httpd_query_key_value](#) (C++ function), 130
[httpd_queue_work](#) (C++ function), 138
[httpd_recv_func_t](#) (C++ type), 146
[httpd_register_err_handler](#) (C++ function), 137
[httpd_register_uri_handler](#) (C++ function), 125
[httpd_req](#) (C++ struct), 143
[httpd_req::aux](#) (C++ member), 143
[httpd_req::content_len](#) (C++ member), 143
[httpd_req::free_ctx](#) (C++ member), 144
[httpd_req::handle](#) (C++ member), 143
[httpd_req::ignore_sess_ctx_changes](#) (C++ member), 144
[httpd_req::method](#) (C++ member), 143
[httpd_req::sess_ctx](#) (C++ member), 143
[httpd_req::uri](#) (C++ member), 143
[httpd_req::user_ctx](#) (C++ member), 143
[httpd_req_async_handler_begin](#) (C++ function), 127
[httpd_req_async_handler_complete](#) (C++ function), 127
[httpd_req_get_cookie_val](#) (C++ function), 130
[httpd_req_get_hdr_value_len](#) (C++ function), 129
[httpd_req_get_hdr_value_str](#) (C++ function), 129
[httpd_req_get_url_query_len](#) (C++ function), 129
[httpd_req_get_url_query_str](#) (C++ function), 130
[httpd_req_recv](#) (C++ function), 128
[httpd_req_t](#) (C++ type), 146
[httpd_req_to_sockfd](#) (C++ function), 128
[httpd_resp_send_404](#) (C++ function), 134
[httpd_resp_send_408](#) (C++ function), 135
[httpd_resp_send_500](#) (C++ function), 135
[httpd_resp_send_chunk](#) (C++ function), 132
[httpd_resp_send_err](#) (C++ function), 134
[httpd_resp_sendstr](#) (C++ function), 132
[httpd_resp_sendstr_chunk](#) (C++ function), 132
[httpd_resp_set_hdr](#) (C++ function), 133
[httpd_resp_set_status](#) (C++ function), 133
[httpd_resp_set_type](#) (C++ function), 133
[HTTPD_RESP_USE_STRLEN](#) (C macro), 146
[httpd_send](#) (C++ function), 135
[httpd_send_func_t](#) (C++ type), 146
[httpd_sess_get_ctx](#) (C++ function), 138
[httpd_sess_get_transport_ctx](#) (C++ function), 139
[httpd_sess_set_ctx](#) (C++ function), 138
[httpd_sess_set_pending_override](#) (C++ function), 127
[httpd_sess_set_recv_override](#) (C++ function), 126
[httpd_sess_set_send_override](#) (C++ function), 126
[httpd_sess_set_transport_ctx](#) (C++ function), 139
[httpd_sess_trigger_close](#) (C++ function), 139
[httpd_sess_update_lru_counter](#) (C++ function), 140
[HTTPD_SOCK_ERR_FAIL](#) (C macro), 144
[HTTPD_SOCK_ERR_INVALID](#) (C macro), 144
[HTTPD_SOCK_ERR_TIMEOUT](#) (C macro), 144
[httpd_socket_recv](#) (C++ function), 136
[httpd_socket_send](#) (C++ function), 136
[httpd_ssl_config](#) (C++ struct), 152
[httpd_ssl_config::alpn_protos](#) (C++ member), 153
[httpd_ssl_config::cacert_len](#) (C++ member), 152
[httpd_ssl_config::cacert_pem](#) (C++ member), 152
[httpd_ssl_config::cert_select_cb](#) (C++ member), 153
[httpd_ssl_config::ecdsa_key_efuse_blk](#) (C++ member), 152
[httpd_ssl_config::httpd](#) (C++ member), 152
[httpd_ssl_config::port_insecure](#) (C++ member), 152

[httpd_ssl_config::port_secure](#) (C++ member), [152](#)
[httpd_ssl_config::prvtkey_len](#) (C++ member), [152](#)
[httpd_ssl_config::prvtkey_pem](#) (C++ member), [152](#)
[httpd_ssl_config::servercert](#) (C++ member), [152](#)
[httpd_ssl_config::servercert_len](#) (C++ member), [152](#)
[httpd_ssl_config::session_tickets](#) (C++ member), [153](#)
[httpd_ssl_config::ssl_userdata](#) (C++ member), [153](#)
[httpd_ssl_config::transport_mode](#) (C++ member), [152](#)
[httpd_ssl_config::use_ecdsa_peripheral](#) (C++ member), [152](#)
[httpd_ssl_config::use_secure_element](#) (C++ member), [153](#)
[httpd_ssl_config::user_cb](#) (C++ member), [153](#)
[HTTPD_SSL_CONFIG_DEFAULT](#) (C macro), [153](#)
[httpd_ssl_config_t](#) (C++ type), [153](#)
[httpd_ssl_start](#) (C++ function), [151](#)
[httpd_ssl_stop](#) (C++ function), [151](#)
[httpd_ssl_transport_mode_t](#) (C++ enum), [153](#)
[httpd_ssl_transport_mode_t::HTTPD_SSL_TRANSPORT_INSECURE](#) (C++ enumerator), [154](#)
[httpd_ssl_transport_mode_t::HTTPD_SSL_TRANSPORT_SECURE](#) (C++ enumerator), [153](#)
[httpd_ssl_user_cb_state_t](#) (C++ enum), [154](#)
[httpd_ssl_user_cb_state_t::HTTPD_SSL_USER_CB_STATE_CLOSE](#) (C++ enumerator), [154](#)
[httpd_ssl_user_cb_state_t::HTTPD_SSL_USER_CB_STATE_CREATE](#) (C++ enumerator), [154](#)
[httpd_start](#) (C++ function), [137](#)
[httpd_stop](#) (C++ function), [138](#)
[HTTPD_TYPE_JSON](#) (C macro), [145](#)
[HTTPD_TYPE_OCTET](#) (C macro), [145](#)
[HTTPD_TYPE_TEXT](#) (C macro), [145](#)
[httpd_unregister_uri](#) (C++ function), [126](#)
[httpd_unregister_uri_handler](#) (C++ function), [126](#)
[httpd_uri](#) (C++ struct), [144](#)
[httpd_uri::handler](#) (C++ member), [144](#)
[httpd_uri::method](#) (C++ member), [144](#)
[httpd_uri::uri](#) (C++ member), [144](#)
[httpd_uri::user_ctx](#) (C++ member), [144](#)
[httpd_uri_match_func_t](#) (C++ type), [148](#)
[httpd_uri_match_wildcard](#) (C++ function), [131](#)
[httpd_uri_t](#) (C++ type), [146](#)
[httpd_work_fn_t](#) (C++ type), [148](#)
[HttpStatus_Code](#) (C++ enum), [94](#)
[HttpStatus_Code::HttpStatus_BadRequest](#) (C++ enumerator), [94](#)
[HttpStatus_Code::HttpStatus_Forbidden](#) (C++ enumerator), [94](#)
[HttpStatus_Code::HttpStatus_Found](#) (C++ enumerator), [94](#)
[HttpStatus_Code::HttpStatus_InternalError](#) (C++ enumerator), [94](#)
[HttpStatus_Code::HttpStatus_MovedPermanently](#) (C++ enumerator), [94](#)
[HttpStatus_Code::HttpStatus_MultipleChoices](#) (C++ enumerator), [94](#)
[HttpStatus_Code::HttpStatus_NotFound](#) (C++ enumerator), [94](#)
[HttpStatus_Code::HttpStatus_Ok](#) (C++ enumerator), [94](#)
[HttpStatus_Code::HttpStatus_PermanentRedirect](#) (C++ enumerator), [94](#)
[HttpStatus_Code::HttpStatus_SeeOther](#) (C++ enumerator), [94](#)
[HttpStatus_Code::HttpStatus_TemporaryRedirect](#) (C++ enumerator), [94](#)
[HttpStatus_Code::HttpStatus_Unauthorized](#) (C++ enumerator), [94](#)
I
[i2c_ack_type_t](#) (C++ enum), [364](#)
[i2c_ack_type_t::I2C_MASTER_ACK](#) (C++ enumerator), [364](#)
[i2c_ack_type_t::I2C_MASTER_ACK_MAX](#) (C++ enumerator), [364](#)
[i2c_ack_type_t::I2C_MASTER_LAST_NACK](#) (C++ enumerator), [364](#)
[i2c_ack_type_t::I2C_MASTER_NACK](#) (C++ enumerator), [364](#)
[i2c_ack_value_t](#) (C++ enum), [361](#)
[i2c_ack_value_t::I2C_ACK_VAL](#) (C++ enumerator), [361](#)
[i2c_ack_value_t::I2C_NACK_VAL](#) (C++ enumerator), [361](#)
[i2c_addr_bit_len_t](#) (C++ enum), [363](#)
[i2c_addr_bit_len_t::I2C_ADDR_BIT_LEN_10](#) (C++ enumerator), [363](#)
[i2c_addr_bit_len_t::I2C_ADDR_BIT_LEN_7](#) (C++ enumerator), [363](#)
[i2c_addr_mode_t](#) (C++ enum), [364](#)
[i2c_addr_mode_t::I2C_ADDR_BIT_10](#) (C++ enumerator), [364](#)
[i2c_addr_mode_t::I2C_ADDR_BIT_7](#) (C++ enumerator), [364](#)
[i2c_addr_mode_t::I2C_ADDR_BIT_MAX](#) (C++ enumerator), [364](#)
[i2c_bus_mode_t](#) (C++ enum), [365](#)
[i2c_bus_mode_t::I2C_BUS_MODE_MASTER](#) (C++ enumerator), [365](#)
[i2c_bus_mode_t::I2C_BUS_MODE_SLAVE](#) (C++ enumerator), [365](#)
[i2c_clock_source_t](#) (C++ type), [362](#)
[i2c_del_master_bus](#) (C++ function), [349](#)
[i2c_del_slave_device](#) (C++ function), [356](#)

- `I2C_DEVICE_ADDRESS_NOT_USED` (C macro), 355
- `i2c_device_config_t` (C++ struct), 353
- `i2c_device_config_t::dev_addr_length` (C++ member), 353
- `i2c_device_config_t::device_address` (C++ member), 354
- `i2c_device_config_t::disable_ack_check` (C++ member), 354
- `i2c_device_config_t::flags` (C++ member), 354
- `i2c_device_config_t::scl_speed_hz` (C++ member), 354
- `i2c_device_config_t::scl_wait_us` (C++ member), 354
- `i2c_hal_clk_config_t` (C++ struct), 362
- `i2c_hal_clk_config_t::clkm_div` (C++ member), 362
- `i2c_hal_clk_config_t::hold` (C++ member), 362
- `i2c_hal_clk_config_t::scl_high` (C++ member), 362
- `i2c_hal_clk_config_t::scl_low` (C++ member), 362
- `i2c_hal_clk_config_t::scl_wait_high` (C++ member), 362
- `i2c_hal_clk_config_t::sda_hold` (C++ member), 362
- `i2c_hal_clk_config_t::sda_sample` (C++ member), 362
- `i2c_hal_clk_config_t::setup` (C++ member), 362
- `i2c_hal_clk_config_t::tout` (C++ member), 362
- `i2c_master_bus_add_device` (C++ function), 349
- `i2c_master_bus_config_t` (C++ struct), 353
- `i2c_master_bus_config_t::clk_source` (C++ member), 353
- `i2c_master_bus_config_t::enable_internal_pullup` (C++ member), 353
- `i2c_master_bus_config_t::flags` (C++ member), 353
- `i2c_master_bus_config_t::glitch_ignore` (C++ member), 353
- `i2c_master_bus_config_t::i2c_port` (C++ member), 353
- `i2c_master_bus_config_t::intr_priority` (C++ member), 353
- `i2c_master_bus_config_t::scl_io_num` (C++ member), 353
- `i2c_master_bus_config_t::sda_io_num` (C++ member), 353
- `i2c_master_bus_config_t::trans_queue_depth` (C++ member), 353
- `i2c_master_bus_handle_t` (C++ type), 359
- `i2c_master_bus_reset` (C++ function), 352
- `i2c_master_bus_rm_device` (C++ function), 349
- `i2c_master_bus_wait_all_done` (C++ function), 352
- `i2c_master_callback_t` (C++ type), 360
- `i2c_master_command_t` (C++ enum), 361
- `i2c_master_command_t::I2C_MASTER_CMD_READ` (C++ enumerator), 361
- `i2c_master_command_t::I2C_MASTER_CMD_START` (C++ enumerator), 361
- `i2c_master_command_t::I2C_MASTER_CMD_STOP` (C++ enumerator), 361
- `i2c_master_command_t::I2C_MASTER_CMD_WRITE` (C++ enumerator), 361
- `i2c_master_dev_handle_t` (C++ type), 359
- `i2c_master_event_callbacks_t` (C++ struct), 355
- `i2c_master_event_callbacks_t::on_trans_done` (C++ member), 355
- `i2c_master_event_data_t` (C++ struct), 359
- `i2c_master_event_data_t::event` (C++ member), 359
- `i2c_master_event_t` (C++ enum), 361
- `i2c_master_event_t::I2C_EVENT_ALIVE` (C++ enumerator), 361
- `i2c_master_event_t::I2C_EVENT_DONE` (C++ enumerator), 361
- `i2c_master_event_t::I2C_EVENT_NACK` (C++ enumerator), 361
- `i2c_master_event_t::I2C_EVENT_TIMEOUT` (C++ enumerator), 361
- `i2c_master_execute_defined_operations` (C++ function), 351
- `i2c_master_get_bus_handle` (C++ function), 352
- `i2c_master_probe` (C++ function), 350
- `i2c_master_receive` (C++ function), 350
- `i2c_master_register_event_callbacks` (C++ function), 351
- `i2c_master_status_t` (C++ enum), 360
- `i2c_master_status_t::I2C_STATUS_ACK_ERROR` (C++ enumerator), 360
- `i2c_master_status_t::I2C_STATUS_DONE` (C++ enumerator), 360
- `i2c_master_status_t::I2C_STATUS_IDLE` (C++ enumerator), 360
- `i2c_master_status_t::I2C_STATUS_READ` (C++ enumerator), 360
- `i2c_master_status_t::I2C_STATUS_START` (C++ enumerator), 360
- `i2c_master_status_t::I2C_STATUS_STOP` (C++ enumerator), 360
- `i2c_master_status_t::I2C_STATUS_TIMEOUT` (C++ enumerator), 360
- `i2c_master_status_t::I2C_STATUS_WRITE` (C++ enumerator), 360
- `i2c_master_transmit` (C++ function), 349
- `i2c_master_transmit_receive` (C++ function), 350

- `i2c_mode_t` (C++ *enum*), 363
- `i2c_mode_t::I2C_MODE_MASTER` (C++ *enumerator*), 363
- `i2c_mode_t::I2C_MODE_MAX` (C++ *enumerator*), 363
- `i2c_mode_t::I2C_MODE_SLAVE` (C++ *enumerator*), 363
- `i2c_new_master_bus` (C++ *function*), 349
- `i2c_new_slave_device` (C++ *function*), 355
- `i2c_operation_job_t` (C++ *struct*), 354
- `i2c_operation_job_t::ack_check` (C++ *member*), 354
- `i2c_operation_job_t::ack_value` (C++ *member*), 354
- `i2c_operation_job_t::command` (C++ *member*), 354
- `i2c_operation_job_t::data` (C++ *member*), 354
- `i2c_operation_job_t::read` (C++ *member*), 355
- `i2c_operation_job_t::total_bytes` (C++ *member*), 354
- `i2c_operation_job_t::write` (C++ *member*), 354
- `i2c_port_num_t` (C++ *type*), 359
- `i2c_port_t` (C++ *enum*), 362
- `i2c_port_t::I2C_NUM_0` (C++ *enumerator*), 363
- `i2c_port_t::I2C_NUM_1` (C++ *enumerator*), 363
- `i2c_port_t::I2C_NUM_MAX` (C++ *enumerator*), 363
- `i2c_rw_t` (C++ *enum*), 363
- `i2c_rw_t::I2C_MASTER_READ` (C++ *enumerator*), 363
- `i2c_rw_t::I2C_MASTER_WRITE` (C++ *enumerator*), 363
- `i2c_slave_config_t` (C++ *struct*), 358
- `i2c_slave_config_t::access_ram_en` (C++ *member*), 358
- `i2c_slave_config_t::addr_bit_len` (C++ *member*), 358
- `i2c_slave_config_t::broadcast_en` (C++ *member*), 358
- `i2c_slave_config_t::clk_source` (C++ *member*), 358
- `i2c_slave_config_t::flags` (C++ *member*), 358
- `i2c_slave_config_t::i2c_port` (C++ *member*), 358
- `i2c_slave_config_t::intr_priority` (C++ *member*), 358
- `i2c_slave_config_t::scl_io_num` (C++ *member*), 358
- `i2c_slave_config_t::sda_io_num` (C++ *member*), 358
- `i2c_slave_config_t::send_buf_depth` (C++ *member*), 358
- `i2c_slave_config_t::slave_addr` (C++ *member*), 358
- `i2c_slave_config_t::slave_unmatch_en` (C++ *member*), 358
- `i2c_slave_dev_handle_t` (C++ *type*), 360
- `i2c_slave_event_callbacks_t` (C++ *struct*), 358
- `i2c_slave_event_callbacks_t::on_recv_done` (C++ *member*), 359
- `i2c_slave_read_ram` (C++ *function*), 357
- `i2c_slave_read_write_status_t` (C++ *enum*), 364
- `i2c_slave_read_write_status_t::I2C_SLAVE_READ_BY_` (C++ *enumerator*), 365
- `i2c_slave_read_write_status_t::I2C_SLAVE_WRITE_BY_` (C++ *enumerator*), 364
- `i2c_slave_receive` (C++ *function*), 356
- `i2c_slave_received_callback_t` (C++ *type*), 360
- `i2c_slave_register_event_callbacks` (C++ *function*), 356
- `i2c_slave_rx_done_event_data_t` (C++ *struct*), 359
- `i2c_slave_rx_done_event_data_t::buffer` (C++ *member*), 359
- `i2c_slave_stretch_cause_t` (C++ *enum*), 364
- `i2c_slave_stretch_cause_t::I2C_SLAVE_STRETCH_CAUS` (C++ *enumerator*), 364
- `i2c_slave_stretch_cause_t::I2C_SLAVE_STRETCH_CAUS` (C++ *enumerator*), 364
- `i2c_slave_stretch_cause_t::I2C_SLAVE_STRETCH_CAUS` (C++ *enumerator*), 364
- `i2c_slave_stretch_cause_t::I2C_SLAVE_STRETCH_CAUS` (C++ *enumerator*), 364
- `i2c_slave_stretch_cause_t::I2C_SLAVE_STRETCH_CAUS` (C++ *enumerator*), 364
- `i2c_slave_transmit` (C++ *function*), 356
- `i2c_slave_write_ram` (C++ *function*), 357
- `i2c_trans_mode_t` (C++ *enum*), 363
- `i2c_trans_mode_t::I2C_DATA_MODE_LSB_FIRST` (C++ *enumerator*), 363
- `i2c_trans_mode_t::I2C_DATA_MODE_MAX` (C++ *enumerator*), 364
- `i2c_trans_mode_t::I2C_DATA_MODE_MSB_FIRST` (C++ *enumerator*), 363
- `i2s_chan_config_t` (C++ *struct*), 406
- `i2s_chan_config_t::auto_clear` (C++ *member*), 406
- `i2s_chan_config_t::dma_desc_num` (C++ *member*), 406
- `i2s_chan_config_t::dma_frame_num` (C++ *member*), 406
- `i2s_chan_config_t::id` (C++ *member*), 406
- `i2s_chan_config_t::intr_priority` (C++ *member*), 406
- `i2s_chan_config_t::role` (C++ *member*), 406
- `i2s_chan_handle_t` (C++ *type*), 408
- `i2s_chan_info_t` (C++ *struct*), 406
- `i2s_chan_info_t::dir` (C++ *member*), 406
- `i2s_chan_info_t::id` (C++ *member*), 406
- `i2s_chan_info_t::mode` (C++ *member*), 407
- `i2s_chan_info_t::pair_chan` (C++ *member*),

- 407
- `i2s_chan_info_t::role` (C++ member), 406
- `I2S_CHANNEL_DEFAULT_CONFIG` (C macro), 407
- `i2s_channel_disable` (C++ function), 403
- `i2s_channel_enable` (C++ function), 403
- `i2s_channel_get_info` (C++ function), 402
- `i2s_channel_init_pdm_rx_mode` (C++ function), 387
- `i2s_channel_init_pdm_tx_mode` (C++ function), 389
- `i2s_channel_init_std_mode` (C++ function), 382
- `i2s_channel_init_tdm_mode` (C++ function), 396
- `i2s_channel_preload_data` (C++ function), 403
- `i2s_channel_read` (C++ function), 404
- `i2s_channel_reconfig_pdm_rx_clock` (C++ function), 388
- `i2s_channel_reconfig_pdm_rx_gpio` (C++ function), 389
- `i2s_channel_reconfig_pdm_rx_slot` (C++ function), 388
- `i2s_channel_reconfig_pdm_tx_clock` (C++ function), 389
- `i2s_channel_reconfig_pdm_tx_gpio` (C++ function), 390
- `i2s_channel_reconfig_pdm_tx_slot` (C++ function), 390
- `i2s_channel_reconfig_std_clock` (C++ function), 383
- `i2s_channel_reconfig_std_gpio` (C++ function), 384
- `i2s_channel_reconfig_std_slot` (C++ function), 383
- `i2s_channel_reconfig_tdm_clock` (C++ function), 397
- `i2s_channel_reconfig_tdm_gpio` (C++ function), 397
- `i2s_channel_reconfig_tdm_slot` (C++ function), 397
- `i2s_channel_register_event_callback` (C++ function), 405
- `i2s_channel_write` (C++ function), 404
- `i2s_clock_src_t` (C++ type), 409
- `i2s_comm_mode_t` (C++ enum), 408
- `i2s_comm_mode_t::I2S_COMM_MODE_NONE` (C++ enumerator), 408
- `i2s_comm_mode_t::I2S_COMM_MODE_PDM` (C++ enumerator), 408
- `i2s_comm_mode_t::I2S_COMM_MODE_STD` (C++ enumerator), 408
- `i2s_comm_mode_t::I2S_COMM_MODE_TDM` (C++ enumerator), 408
- `i2s_data_bit_width_t` (C++ enum), 410
- `i2s_data_bit_width_t::I2S_DATA_BIT_WIDTH_16BIT` (C++ member), 391
- `i2s_data_bit_width_t::I2S_DATA_BIT_WIDTH_24BIT` (C++ member), 391
- `i2s_data_bit_width_t::I2S_DATA_BIT_WIDTH_32BIT` (C++ member), 391
- `i2s_data_bit_width_t::I2S_DATA_BIT_WIDTH_8BIT` (C++ member), 391
- `i2s_del_channel` (C++ function), 402
- `i2s_dir_t` (C++ enum), 409
- `i2s_dir_t::I2S_DIR_RX` (C++ enumerator), 409
- `i2s_dir_t::I2S_DIR_TX` (C++ enumerator), 409
- `i2s_event_callbacks_t` (C++ struct), 405
- `i2s_event_callbacks_t::on_recv` (C++ member), 405
- `i2s_event_callbacks_t::on_recv_q_ovf` (C++ member), 405
- `i2s_event_callbacks_t::on_send_q_ovf` (C++ member), 406
- `i2s_event_callbacks_t::on_sent` (C++ member), 406
- `i2s_event_data_t` (C++ struct), 407
- `i2s_event_data_t::data` (C++ member), 407
- `i2s_event_data_t::size` (C++ member), 407
- `I2S_GPIO_UNUSED` (C macro), 407
- `i2s_isr_callback_t` (C++ type), 408
- `i2s_mclk_multiple_t` (C++ enum), 408
- `i2s_mclk_multiple_t::I2S_MCLK_MULTIPLE_128` (C++ enumerator), 408
- `i2s_mclk_multiple_t::I2S_MCLK_MULTIPLE_256` (C++ enumerator), 408
- `i2s_mclk_multiple_t::I2S_MCLK_MULTIPLE_384` (C++ enumerator), 409
- `i2s_mclk_multiple_t::I2S_MCLK_MULTIPLE_512` (C++ enumerator), 409
- `i2s_new_channel` (C++ function), 402
- `i2s_pcm_compress_t` (C++ enum), 410
- `i2s_pcm_compress_t::I2S_PCM_A_COMPRESS` (C++ enumerator), 411
- `i2s_pcm_compress_t::I2S_PCM_A_DECOMPRESS` (C++ enumerator), 410
- `i2s_pcm_compress_t::I2S_PCM_DISABLE` (C++ enumerator), 410
- `i2s_pcm_compress_t::I2S_PCM_U_COMPRESS` (C++ enumerator), 411
- `i2s_pcm_compress_t::I2S_PCM_U_DECOMPRESS` (C++ enumerator), 411
- `i2s_pdm_dsr_t` (C++ enum), 411
- `i2s_pdm_dsr_t::I2S_PDM_DSR_16S` (C++ enumerator), 411
- `i2s_pdm_dsr_t::I2S_PDM_DSR_8S` (C++ enumerator), 411
- `i2s_pdm_dsr_t::I2S_PDM_DSR_MAX` (C++ enumerator), 411
- `i2s_pdm_rx_clk_config_t` (C++ struct), 391
- `i2s_pdm_rx_clk_config_t::bclk_div` (C++ member), 392
- `i2s_pdm_rx_clk_config_t::clk_src` (C++ member), 391
- `i2s_pdm_rx_clk_config_t::dn_sample_mode` (C++ member), 391

i2s_pdm_rx_clk_config_t::mclk_multiple (C++ enumerator), 412
(C++ member), 391
i2s_pdm_rx_clk_config_t::sample_rate_hz (C++ enumerator), 412
(C++ member), 391
I2S_PDM_RX_CLK_DEFAULT_CONFIG (C macro), 395
i2s_pdm_rx_config_t (C++ struct), 392
i2s_pdm_rx_config_t::clk_cfg (C++ member), 392
i2s_pdm_rx_config_t::gpio_cfg (C++ member), 392
i2s_pdm_rx_config_t::slot_cfg (C++ member), 392
i2s_pdm_rx_gpio_config_t (C++ struct), 392
i2s_pdm_rx_gpio_config_t::clk (C++ member), 392
i2s_pdm_rx_gpio_config_t::clk_inv (C++ member), 392
i2s_pdm_rx_gpio_config_t::din (C++ member), 392
i2s_pdm_rx_gpio_config_t::dins (C++ member), 392
i2s_pdm_rx_gpio_config_t::invert_flags (C++ member), 392
i2s_pdm_rx_slot_config_t (C++ struct), 391
i2s_pdm_rx_slot_config_t::amplify_num (C++ member), 391
i2s_pdm_rx_slot_config_t::data_bit_width (C++ member), 391
i2s_pdm_rx_slot_config_t::hp_cut_off_freq_hz (C++ member), 391
(C++ member), 391
i2s_pdm_rx_slot_config_t::hp_en (C++ member), 391
i2s_pdm_rx_slot_config_t::slot_bit_width (C++ member), 391
i2s_pdm_rx_slot_config_t::slot_mask (C++ member), 391
i2s_pdm_rx_slot_config_t::slot_mode (C++ member), 391
I2S_PDM_RX_SLOT_DEFAULT_CONFIG (C macro), 395
i2s_pdm_sig_scale_t (C++ enum), 411
i2s_pdm_sig_scale_t::I2S_PDM_SIG_SCALING_0_DB (C++ enumerator), 411
i2s_pdm_sig_scale_t::I2S_PDM_SIG_SCALING_2_5_DB (C++ enumerator), 411
i2s_pdm_sig_scale_t::I2S_PDM_SIG_SCALING_5_DB (C++ enumerator), 411
i2s_pdm_sig_scale_t::I2S_PDM_SIG_SCALING_7_5_DB (C++ enumerator), 411
i2s_pdm_slot_mask_t (C++ enum), 412
i2s_pdm_slot_mask_t::I2S_PDM_LINE_SLOT_ALL (C++ enumerator), 413
(C++ enumerator), 413
i2s_pdm_slot_mask_t::I2S_PDM_RX_LINE0_SLOT_LEFT (C++ enumerator), 412
(C++ enumerator), 412
i2s_pdm_slot_mask_t::I2S_PDM_RX_LINE0_SLOT_RIGHT (C++ enumerator), 412
(C++ enumerator), 412
i2s_pdm_slot_mask_t::I2S_PDM_RX_LINE1_SLOT_LEFT (C++ enumerator), 412
(C++ enumerator), 412
i2s_pdm_slot_mask_t::I2S_PDM_RX_LINE1_SLOT_RIGHT (C++ enumerator), 412
(C++ enumerator), 412
i2s_pdm_slot_mask_t::I2S_PDM_RX_LINE2_SLOT_LEFT (C++ enumerator), 413
(C++ enumerator), 413
i2s_pdm_slot_mask_t::I2S_PDM_RX_LINE2_SLOT_RIGHT (C++ enumerator), 413
(C++ enumerator), 413
i2s_pdm_slot_mask_t::I2S_PDM_RX_LINE3_SLOT_LEFT (C++ enumerator), 413
(C++ enumerator), 413
i2s_pdm_slot_mask_t::I2S_PDM_RX_LINE3_SLOT_RIGHT (C++ enumerator), 413
(C++ enumerator), 413
i2s_pdm_slot_mask_t::I2S_PDM_SLOT_BOTH (C++ enumerator), 412
(C++ enumerator), 412
i2s_pdm_slot_mask_t::I2S_PDM_SLOT_LEFT (C++ enumerator), 412
(C++ enumerator), 412
i2s_pdm_slot_mask_t::I2S_PDM_SLOT_RIGHT (C++ enumerator), 412
(C++ enumerator), 412
i2s_pdm_tx_clk_config_t (C++ struct), 393
i2s_pdm_tx_clk_config_t::bclk_div (C++ member), 394
i2s_pdm_tx_clk_config_t::clk_src (C++ member), 394
i2s_pdm_tx_clk_config_t::mclk_multiple (C++ member), 394
(C++ member), 394
i2s_pdm_tx_clk_config_t::sample_rate_hz (C++ member), 393
(C++ member), 393
i2s_pdm_tx_clk_config_t::up_sample_fp (C++ member), 394
(C++ member), 394
i2s_pdm_tx_clk_config_t::up_sample_fs (C++ member), 394
(C++ member), 394
I2S_PDM_TX_CLK_DAC_DEFAULT_CONFIG (C macro), 395
I2S_PDM_TX_CLK_DEFAULT_CONFIG (C macro), 395
i2s_pdm_tx_config_t (C++ struct), 394
i2s_pdm_tx_config_t::clk_cfg (C++ member), 394
(C++ member), 394
i2s_pdm_tx_config_t::gpio_cfg (C++ member), 395
(C++ member), 395
i2s_pdm_tx_config_t::slot_cfg (C++ member), 394
(C++ member), 394
i2s_pdm_tx_gpio_config_t (C++ struct), 394
i2s_pdm_tx_gpio_config_t::clk (C++ member), 394
(C++ member), 394
i2s_pdm_tx_gpio_config_t::clk_inv (C++ member), 394
(C++ member), 394
i2s_pdm_tx_gpio_config_t::dout (C++ member), 394
(C++ member), 394
i2s_pdm_tx_gpio_config_t::dout2 (C++ member), 394
(C++ member), 394
i2s_pdm_tx_gpio_config_t::invert_flags (C++ member), 394
(C++ member), 394
i2s_pdm_tx_line_mode_t (C++ enum), 411
(C++ enumerator), 411
i2s_pdm_tx_line_mode_t::I2S_PDM_TX_ONE_LINE_CODECS (C++ enumerator), 411
(C++ enumerator), 411
i2s_pdm_tx_line_mode_t::I2S_PDM_TX_ONE_LINE_DAC (C++ enumerator), 412
(C++ enumerator), 412
i2s_pdm_tx_line_mode_t::I2S_PDM_TX_TWO_LINE_DAC (C++ enumerator), 412
(C++ enumerator), 412

- (C++ enumerator), 412
- i2s_pdm_tx_slot_config_t (C++ struct), 392
- i2s_pdm_tx_slot_config_t::data_bit_width (C++ member), 393
- i2s_pdm_tx_slot_config_t::hp_cut_off_freq_hz (C++ member), 385
- (C++ member), 393
- i2s_pdm_tx_slot_config_t::hp_en (C++ member), 393
- i2s_pdm_tx_slot_config_t::hp_scale (C++ member), 393
- i2s_pdm_tx_slot_config_t::line_mode (C++ member), 393
- i2s_pdm_tx_slot_config_t::lp_scale (C++ member), 393
- i2s_pdm_tx_slot_config_t::sd_dither (C++ member), 393
- i2s_pdm_tx_slot_config_t::sd_dither2 (C++ member), 393
- i2s_pdm_tx_slot_config_t::sd_prescale (C++ member), 393
- i2s_pdm_tx_slot_config_t::sd_scale (C++ member), 393
- i2s_pdm_tx_slot_config_t::sinc_scale (C++ member), 393
- i2s_pdm_tx_slot_config_t::slot_bit_width (C++ member), 393
- i2s_pdm_tx_slot_config_t::slot_mode (C++ member), 393
- I2S_PDM_TX_SLOT_DAC_DEFAULT_CONFIG (C macro), 395
- I2S_PDM_TX_SLOT_DEFAULT_CONFIG (C macro), 395
- i2s_port_t (C++ enum), 408
- i2s_port_t::I2S_NUM_0 (C++ enumerator), 408
- i2s_port_t::I2S_NUM_1 (C++ enumerator), 408
- i2s_port_t::I2S_NUM_AUTO (C++ enumerator), 408
- i2s_role_t (C++ enum), 409
- i2s_role_t::I2S_ROLE_MASTER (C++ enumerator), 409
- i2s_role_t::I2S_ROLE_SLAVE (C++ enumerator), 410
- i2s_slot_bit_width_t (C++ enum), 410
- i2s_slot_bit_width_t::I2S_SLOT_BIT_WIDTH_16BIT (C++ enumerator), 410
- i2s_slot_bit_width_t::I2S_SLOT_BIT_WIDTH_24BIT (C++ enumerator), 410
- i2s_slot_bit_width_t::I2S_SLOT_BIT_WIDTH_32BIT (C++ enumerator), 410
- i2s_slot_bit_width_t::I2S_SLOT_BIT_WIDTH_AUTO (C++ enumerator), 410
- i2s_slot_mode_t (C++ enum), 409
- i2s_slot_mode_t::I2S_SLOT_MODE_MONO (C++ enumerator), 409
- i2s_slot_mode_t::I2S_SLOT_MODE_STEREO (C++ enumerator), 409
- i2s_std_clk_config_t (C++ struct), 385
- i2s_std_clk_config_t::clk_src (C++ member), 385
- i2s_std_clk_config_t::ext_clk_freq_hz (C++ member), 385
- i2s_std_clk_config_t::mclk_multiple (C++ member), 385
- i2s_std_clk_config_t::sample_rate_hz (C++ member), 385
- I2S_STD_CLK_DEFAULT_CONFIG (C macro), 387
- i2s_std_config_t (C++ struct), 386
- i2s_std_config_t::clk_cfg (C++ member), 386
- i2s_std_config_t::gpio_cfg (C++ member), 386
- i2s_std_config_t::slot_cfg (C++ member), 386
- i2s_std_gpio_config_t (C++ struct), 385
- i2s_std_gpio_config_t::bclk (C++ member), 386
- i2s_std_gpio_config_t::bclk_inv (C++ member), 386
- i2s_std_gpio_config_t::din (C++ member), 386
- i2s_std_gpio_config_t::dout (C++ member), 386
- i2s_std_gpio_config_t::invert_flags (C++ member), 386
- i2s_std_gpio_config_t::mclk (C++ member), 385
- i2s_std_gpio_config_t::mclk_inv (C++ member), 386
- i2s_std_gpio_config_t::ws (C++ member), 386
- i2s_std_gpio_config_t::ws_inv (C++ member), 386
- I2S_STD_MSB_SLOT_DEFAULT_CONFIG (C macro), 387
- I2S_STD_PCM_SLOT_DEFAULT_CONFIG (C macro), 387
- I2S_STD_PHILIPS_SLOT_DEFAULT_CONFIG (C macro), 386
- i2s_std_slot_config_t (C++ struct), 384
- i2s_std_slot_config_t::big_endian (C++ member), 385
- i2s_std_slot_config_t::bit_order_lsb (C++ member), 385
- i2s_std_slot_config_t::bit_shift (C++ member), 385
- i2s_std_slot_config_t::data_bit_width (C++ member), 384
- i2s_std_slot_config_t::left_align (C++ member), 385
- i2s_std_slot_config_t::slot_bit_width (C++ member), 384
- i2s_std_slot_config_t::slot_mask (C++ member), 384
- i2s_std_slot_config_t::slot_mode (C++ member), 384

- member*), 384
- `i2s_std_slot_config_t::ws_pol` (C++ *member*), 385
- `i2s_std_slot_config_t::ws_width` (C++ *member*), 385
- `i2s_std_slot_mask_t` (C++ *enum*), 412
- `i2s_std_slot_mask_t::I2S_STD_SLOT_BOTH` (C++ *enumerator*), 412
- `i2s_std_slot_mask_t::I2S_STD_SLOT_LEFT` (C++ *enumerator*), 412
- `i2s_std_slot_mask_t::I2S_STD_SLOT_RIGHT` (C++ *enumerator*), 412
- `I2S_TDM_AUTO_SLOT_NUM` (C *macro*), 400
- `I2S_TDM_AUTO_WS_WIDTH` (C *macro*), 400
- `i2s_tdm_clk_config_t` (C++ *struct*), 399
- `i2s_tdm_clk_config_t::bclk_div` (C++ *member*), 399
- `i2s_tdm_clk_config_t::clk_src` (C++ *member*), 399
- `i2s_tdm_clk_config_t::ext_clk_freq_hz` (C++ *member*), 399
- `i2s_tdm_clk_config_t::mclk_multiple` (C++ *member*), 399
- `i2s_tdm_clk_config_t::sample_rate_hz` (C++ *member*), 399
- `I2S_TDM_CLK_DEFAULT_CONFIG` (C *macro*), 401
- `i2s_tdm_config_t` (C++ *struct*), 400
- `i2s_tdm_config_t::clk_cfg` (C++ *member*), 400
- `i2s_tdm_config_t::gpio_cfg` (C++ *member*), 400
- `i2s_tdm_config_t::slot_cfg` (C++ *member*), 400
- `i2s_tdm_gpio_config_t` (C++ *struct*), 399
- `i2s_tdm_gpio_config_t::bclk` (C++ *member*), 400
- `i2s_tdm_gpio_config_t::bclk_inv` (C++ *member*), 400
- `i2s_tdm_gpio_config_t::din` (C++ *member*), 400
- `i2s_tdm_gpio_config_t::dout` (C++ *member*), 400
- `i2s_tdm_gpio_config_t::invert_flags` (C++ *member*), 400
- `i2s_tdm_gpio_config_t::mclk` (C++ *member*), 399
- `i2s_tdm_gpio_config_t::mclk_inv` (C++ *member*), 400
- `i2s_tdm_gpio_config_t::ws` (C++ *member*), 400
- `i2s_tdm_gpio_config_t::ws_inv` (C++ *member*), 400
- `I2S_TDM_MSB_SLOT_DEFAULT_CONFIG` (C *macro*), 401
- `I2S_TDM_PCM_LONG_SLOT_DEFAULT_CONFIG` (C *macro*), 401
- `I2S_TDM_PCM_SHORT_SLOT_DEFAULT_CONFIG` (C *macro*), 401
- `I2S_TDM_PHILIPS_SLOT_DEFAULT_CONFIG` (C *macro*), 401
- `i2s_tdm_slot_config_t` (C++ *struct*), 398
- `i2s_tdm_slot_config_t::big_endian` (C++ *member*), 399
- `i2s_tdm_slot_config_t::bit_order_lsb` (C++ *member*), 399
- `i2s_tdm_slot_config_t::bit_shift` (C++ *member*), 398
- `i2s_tdm_slot_config_t::data_bit_width` (C++ *member*), 398
- `i2s_tdm_slot_config_t::left_align` (C++ *member*), 398
- `i2s_tdm_slot_config_t::skip_mask` (C++ *member*), 399
- `i2s_tdm_slot_config_t::slot_bit_width` (C++ *member*), 398
- `i2s_tdm_slot_config_t::slot_mask` (C++ *member*), 398
- `i2s_tdm_slot_config_t::slot_mode` (C++ *member*), 398
- `i2s_tdm_slot_config_t::total_slot` (C++ *member*), 399
- `i2s_tdm_slot_config_t::ws_pol` (C++ *member*), 398
- `i2s_tdm_slot_config_t::ws_width` (C++ *member*), 398
- `i2s_tdm_slot_mask_t` (C++ *enum*), 413
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT0` (C++ *enumerator*), 413
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT1` (C++ *enumerator*), 413
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT10` (C++ *enumerator*), 414
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT11` (C++ *enumerator*), 414
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT12` (C++ *enumerator*), 414
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT13` (C++ *enumerator*), 414
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT14` (C++ *enumerator*), 414
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT15` (C++ *enumerator*), 414
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT2` (C++ *enumerator*), 413
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT3` (C++ *enumerator*), 413
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT4` (C++ *enumerator*), 413
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT5` (C++ *enumerator*), 413
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT6` (C++ *enumerator*), 413
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT7` (C++ *enumerator*), 413
- `i2s_tdm_slot_mask_t::I2S_TDM_SLOT8` (C++ *enumerator*), 414

- i2s_tdm_slot_mask_t::I2S_TDM_SLOT9
 (C++ *enumerator*), 414
 intr_handle_t (C++ *type*), 1484
 intr_handler_t (C++ *type*), 1484
 IP2STR (C *macro*), 249
 IP4ADDR_STRLEN_MAX (C *macro*), 249
 ip_event_add_ip6_t (C++ *struct*), 240
 ip_event_add_ip6_t::addr (C++ *member*),
 240
 ip_event_add_ip6_t::preferred (C++
 member), 240
 ip_event_ap_staipassigned_t (C++ *struct*),
 241
 ip_event_ap_staipassigned_t::esp_netif
 (C++ *member*), 241
 ip_event_ap_staipassigned_t::ip (C++
 member), 241
 ip_event_ap_staipassigned_t::mac (C++
 member), 241
 ip_event_got_ip6_t (C++ *struct*), 240
 ip_event_got_ip6_t::esp_netif (C++
 member), 240
 ip_event_got_ip6_t::ip6_info (C++ *mem-*
 ber), 240
 ip_event_got_ip6_t::ip_index (C++ *mem-*
 ber), 240
 ip_event_got_ip_t (C++ *struct*), 240
 ip_event_got_ip_t::esp_netif (C++ *mem-*
 ber), 240
 ip_event_got_ip_t::ip_changed (C++
 member), 240
 ip_event_got_ip_t::ip_info (C++ *member*),
 240
 ip_event_t (C++ *enum*), 246
 ip_event_t::IP_EVENT_AP_STAIPASSIGNED
 (C++ *enumerator*), 246
 ip_event_t::IP_EVENT_ETH_GOT_IP (C++
 enumerator), 246
 ip_event_t::IP_EVENT_ETH_LOST_IP (C++
 enumerator), 246
 ip_event_t::IP_EVENT_GOT_IP6 (C++ *enu-*
 merator), 246
 ip_event_t::IP_EVENT_PPP_GOT_IP (C++
 enumerator), 246
 ip_event_t::IP_EVENT_PPP_LOST_IP (C++
 enumerator), 246
 ip_event_t::IP_EVENT_STA_GOT_IP (C++
 enumerator), 246
 ip_event_t::IP_EVENT_STA_LOST_IP (C++
 enumerator), 246
 IPSTR (C *macro*), 249
 IPV62STR (C *macro*), 249
 IPV6STR (C *macro*), 249
 is_openthread_internal_mesh_local_addr
 (C++ *function*), 215
- L**
- l2tap_ioctl_opt_t (C++ *enum*), 251
 l2tap_ioctl_opt_t::L2TAP_G_DEVICE_DRV_HNDL
 (C++ *enumerator*), 251
 l2tap_ioctl_opt_t::L2TAP_G_INTF_DEVICE
 (C++ *enumerator*), 251
 l2tap_ioctl_opt_t::L2TAP_G_RCV_FILTER
 (C++ *enumerator*), 251
 l2tap_ioctl_opt_t::L2TAP_S_DEVICE_DRV_HNDL
 (C++ *enumerator*), 251
 l2tap_ioctl_opt_t::L2TAP_S_INTF_DEVICE
 (C++ *enumerator*), 251
 l2tap_ioctl_opt_t::L2TAP_S_RCV_FILTER
 (C++ *enumerator*), 251
 l2tap_iodriver_handle (C++ *type*), 251
 L2TAP_VFS_CONFIG_DEFAULT (C *macro*), 251
 l2tap_vfs_config_t (C++ *struct*), 251
 l2tap_vfs_config_t::base_path (C++
 member), 251
 L2TAP_VFS_DEFAULT_PATH (C *macro*), 251
 lcd_color_range_t (C++ *enum*), 418
 lcd_color_range_t::LCD_COLOR_RANGE_FULL
 (C++ *enumerator*), 419
 lcd_color_range_t::LCD_COLOR_RANGE_LIMIT
 (C++ *enumerator*), 419
 lcd_color_space_t (C++ *enum*), 418
 lcd_color_space_t::LCD_COLOR_SPACE_RGB
 (C++ *enumerator*), 418
 lcd_color_space_t::LCD_COLOR_SPACE_YUV
 (C++ *enumerator*), 418
 lcd_rgb_data_endian_t (C++ *enum*), 418
 lcd_rgb_data_endian_t::LCD_RGB_DATA_ENDIAN_BIG
 (C++ *enumerator*), 418
 lcd_rgb_data_endian_t::LCD_RGB_DATA_ENDIAN_LITTLE
 (C++ *enumerator*), 418
 lcd_rgb_element_order_t (C++ *enum*), 418
 lcd_rgb_element_order_t::LCD_RGB_ELEMENT_ORDER_BG
 (C++ *enumerator*), 418
 lcd_rgb_element_order_t::LCD_RGB_ELEMENT_ORDER_RG
 (C++ *enumerator*), 418
 lcd_yuv_conv_std_t (C++ *enum*), 419
 lcd_yuv_conv_std_t::LCD_YUV_CONV_STD_BT601
 (C++ *enumerator*), 419
 lcd_yuv_conv_std_t::LCD_YUV_CONV_STD_BT709
 (C++ *enumerator*), 419
 lcd_yuv_sample_t (C++ *enum*), 419
 lcd_yuv_sample_t::LCD_YUV_SAMPLE_411
 (C++ *enumerator*), 419
 lcd_yuv_sample_t::LCD_YUV_SAMPLE_420
 (C++ *enumerator*), 419
 lcd_yuv_sample_t::LCD_YUV_SAMPLE_422
 (C++ *enumerator*), 419
 ledc_bind_channel_timer (C++ *function*), 439
 ledc_cb_event_t (C++ *enum*), 450
 ledc_cb_event_t::LEDC_FADE_END_EVT
 (C++ *enumerator*), 450
 ledc_cb_param_t (C++ *struct*), 447
 ledc_cb_param_t::channel (C++ *member*),
 448
 ledc_cb_param_t::duty (C++ *member*), 448

- `ledc_cb_param_t::event` (C++ member), 448
- `ledc_cb_param_t::speed_mode` (C++ member), 448
- `ledc_cb_register` (C++ function), 443
- `ledc_cb_t` (C++ type), 450
- `ledc_cbs_t` (C++ struct), 448
- `ledc_cbs_t::fade_cb` (C++ member), 448
- `ledc_channel_config` (C++ function), 434
- `ledc_channel_config_t` (C++ struct), 446
- `ledc_channel_config_t::channel` (C++ member), 446
- `ledc_channel_config_t::duty` (C++ member), 447
- `ledc_channel_config_t::flags` (C++ member), 447
- `ledc_channel_config_t::gpio_num` (C++ member), 446
- `ledc_channel_config_t::hpoint` (C++ member), 447
- `ledc_channel_config_t::intr_type` (C++ member), 446
- `ledc_channel_config_t::output_invert` (C++ member), 447
- `ledc_channel_config_t::speed_mode` (C++ member), 446
- `ledc_channel_config_t::timer_sel` (C++ member), 446
- `ledc_channel_t` (C++ enum), 452
- `ledc_channel_t::LEDC_CHANNEL_0` (C++ enumerator), 452
- `ledc_channel_t::LEDC_CHANNEL_1` (C++ enumerator), 452
- `ledc_channel_t::LEDC_CHANNEL_2` (C++ enumerator), 452
- `ledc_channel_t::LEDC_CHANNEL_3` (C++ enumerator), 452
- `ledc_channel_t::LEDC_CHANNEL_4` (C++ enumerator), 452
- `ledc_channel_t::LEDC_CHANNEL_5` (C++ enumerator), 452
- `ledc_channel_t::LEDC_CHANNEL_6` (C++ enumerator), 452
- `ledc_channel_t::LEDC_CHANNEL_7` (C++ enumerator), 452
- `ledc_channel_t::LEDC_CHANNEL_MAX` (C++ enumerator), 452
- `ledc_clk_cfg_t` (C++ type), 450
- `ledc_clk_src_t` (C++ enum), 451
- `ledc_clk_src_t::LEDC_SCLK` (C++ enumerator), 451
- `ledc_duty_direction_t` (C++ enum), 451
- `ledc_duty_direction_t::LEDC_DUTY_DIR_DECREASE` (C++ enumerator), 451
- `ledc_duty_direction_t::LEDC_DUTY_DIR_INCREASE` (C++ enumerator), 451
- `ledc_duty_direction_t::LEDC_DUTY_DIR_MAX` (C++ enumerator), 451
- `LEDC_ERR_DUTY` (C macro), 450
- `LEDC_ERR_VAL` (C macro), 450
- `ledc_fade_func_install` (C++ function), 441
- `ledc_fade_func_uninstall` (C++ function), 441
- `ledc_fade_mode_t` (C++ enum), 454
- `ledc_fade_mode_t::LEDC_FADE_MAX` (C++ enumerator), 454
- `ledc_fade_mode_t::LEDC_FADE_NO_WAIT` (C++ enumerator), 454
- `ledc_fade_mode_t::LEDC_FADE_WAIT_DONE` (C++ enumerator), 454
- `ledc_fade_param_config_t` (C++ struct), 448
- `ledc_fade_param_config_t::cycle_num` (C++ member), 449
- `ledc_fade_param_config_t::dir` (C++ member), 449
- `ledc_fade_param_config_t::scale` (C++ member), 449
- `ledc_fade_param_config_t::step_num` (C++ member), 449
- `ledc_fade_start` (C++ function), 441
- `ledc_fade_stop` (C++ function), 441
- `ledc_fill_multi_fade_param_list` (C++ function), 445
- `ledc_find_suitable_duty_resolution` (C++ function), 434
- `ledc_get_duty` (C++ function), 437
- `ledc_get_freq` (C++ function), 436
- `ledc_get_hpoint` (C++ function), 437
- `ledc_intr_type_t` (C++ enum), 450
- `ledc_intr_type_t::LEDC_INTR_DISABLE` (C++ enumerator), 451
- `ledc_intr_type_t::LEDC_INTR_FADE_END` (C++ enumerator), 451
- `ledc_intr_type_t::LEDC_INTR_MAX` (C++ enumerator), 451
- `ledc_isr_handle_t` (C++ type), 450
- `ledc_isr_register` (C++ function), 438
- `ledc_mode_t` (C++ enum), 450
- `ledc_mode_t::LEDC_LOW_SPEED_MODE` (C++ enumerator), 450
- `ledc_mode_t::LEDC_SPEED_MODE_MAX` (C++ enumerator), 450
- `ledc_read_fade_param` (C++ function), 446
- `ledc_set_duty` (C++ function), 437
- `ledc_set_duty_and_update` (C++ function), 442
- `ledc_set_duty_with_hpoint` (C++ function), 436
- `ledc_set_fade` (C++ function), 437
- `ledc_set_fade_step_and_start` (C++ function), 443
- `ledc_set_fade_time_and_start` (C++ function), 442
- `ledc_set_fade_with_step` (C++ function), 439
- `ledc_set_fade_with_time` (C++ function), 440
- `ledc_set_freq` (C++ function), 436
- `ledc_set_multi_fade` (C++ function), 444

ledc_set_multi_fade_and_start (C++ function), 444
 ledc_set_pin (C++ function), 435
 ledc_slow_clk_sel_t (C++ enum), 451
 ledc_slow_clk_sel_t::LEDC_SLOW_CLK_PLL_DIV1 (C++ enumerator), 451
 ledc_slow_clk_sel_t::LEDC_SLOW_CLK_RC_FAST (C++ enumerator), 451
 ledc_slow_clk_sel_t::LEDC_SLOW_CLK_RTC8M (C++ enumerator), 451
 ledc_slow_clk_sel_t::LEDC_SLOW_CLK_XTAL (C++ enumerator), 451
 ledc_stop (C++ function), 435
 ledc_timer_bit_t (C++ enum), 452
 ledc_timer_bit_t::LEDC_TIMER_10_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_11_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_12_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_13_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_14_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_15_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_16_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_17_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_18_BIT (C++ enumerator), 454
 ledc_timer_bit_t::LEDC_TIMER_19_BIT (C++ enumerator), 454
 ledc_timer_bit_t::LEDC_TIMER_1_BIT (C++ enumerator), 452
 ledc_timer_bit_t::LEDC_TIMER_20_BIT (C++ enumerator), 454
 ledc_timer_bit_t::LEDC_TIMER_2_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_3_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_4_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_5_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_6_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_7_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_8_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_9_BIT (C++ enumerator), 453
 ledc_timer_bit_t::LEDC_TIMER_BIT_MAX (C++ enumerator), 454
 ledc_timer_config (C++ function), 434
 ledc_timer_config_t (C++ struct), 447
 ledc_timer_config_t::clk_cfg (C++ member), 447
 ledc_timer_config_t::deconfigure (C++ member), 447
 ledc_timer_config_t::duty_resolution (C++ member), 447
 ledc_timer_config_t::freq_hz (C++ member), 447
 ledc_timer_config_t::speed_mode (C++ member), 447
 ledc_timer_config_t::timer_num (C++ member), 447
 ledc_timer_pause (C++ function), 439
 ledc_timer_resume (C++ function), 439
 ledc_timer_rst (C++ function), 438
 ledc_timer_set (C++ function), 438
 ledc_timer_t (C++ enum), 451
 ledc_timer_t::LEDC_TIMER_0 (C++ enumerator), 452
 ledc_timer_t::LEDC_TIMER_1 (C++ enumerator), 452
 ledc_timer_t::LEDC_TIMER_2 (C++ enumerator), 452
 ledc_timer_t::LEDC_TIMER_3 (C++ enumerator), 452
 ledc_timer_t::LEDC_TIMER_MAX (C++ enumerator), 452
 ledc_update_duty (C++ function), 435
 linenoiseCompletions (C++ type), 1231
 lp_gpio_connect_in_signal (C++ function), 310
 lp_gpio_connect_out_signal (C++ function), 310
 lp_uart_sclk_t (C++ type), 657

M

MAC2STR (C macro), 1506
 MACSTR (C macro), 1506
 MALLOC_CAP_32BIT (C macro), 1438
 MALLOC_CAP_8BIT (C macro), 1438
 MALLOC_CAP_DEFAULT (C macro), 1439
 MALLOC_CAP_DMA (C macro), 1438
 MALLOC_CAP_EXEC (C macro), 1438
 MALLOC_CAP_INTERNAL (C macro), 1439
 MALLOC_CAP_INVALID (C macro), 1439
 MALLOC_CAP_IRAM_8BIT (C macro), 1439
 MALLOC_CAP_PID2 (C macro), 1439
 MALLOC_CAP_PID3 (C macro), 1439
 MALLOC_CAP_PID4 (C macro), 1439
 MALLOC_CAP_PID5 (C macro), 1439
 MALLOC_CAP_PID6 (C macro), 1439
 MALLOC_CAP_PID7 (C macro), 1439
 MALLOC_CAP_RETENTION (C macro), 1439
 MALLOC_CAP_RTCRAM (C macro), 1439
 MALLOC_CAP_SPIRAM (C macro), 1439
 MALLOC_CAP_TCM (C macro), 1439
 MAX_BLE_DEVNAME_LEN (C macro), 1095

MAX_BLE_MANUFACTURER_DATA_LEN (*C macro*), 1095
 MAX_FDS (*C macro*), 1197
 mcpwm_brake_config_t (*C++ struct*), 483
 mcpwm_brake_config_t::brake_mode (*C++ member*), 483
 mcpwm_brake_config_t::cbc_recover_on_time (*C++ member*), 483
 mcpwm_brake_config_t::cbc_recover_on_time_val (*C++ member*), 483
 mcpwm_brake_config_t::fault (*C++ member*), 483
 mcpwm_brake_config_t::flags (*C++ member*), 483
 mcpwm_brake_event_cb_t (*C++ type*), 507
 mcpwm_brake_event_data_t (*C++ struct*), 505
 mcpwm_cap_channel_handle_t (*C++ type*), 506
 mcpwm_cap_timer_handle_t (*C++ type*), 506
 mcpwm_capture_channel_config_t (*C++ struct*), 503
 mcpwm_capture_channel_config_t::flags (*C++ member*), 504
 mcpwm_capture_channel_config_t::gpio_num (*C++ member*), 503
 mcpwm_capture_channel_config_t::intr_priority (*C++ member*), 503
 mcpwm_capture_channel_config_t::invert_input_signal (*C++ member*), 503
 mcpwm_capture_channel_config_t::io_loopback (*C++ member*), 503
 mcpwm_capture_channel_config_t::keep_io_explicit (*C++ member*), 503
 mcpwm_capture_channel_config_t::neg_edge (*C++ member*), 503
 mcpwm_capture_channel_config_t::pos_edge (*C++ member*), 503
 mcpwm_capture_channel_config_t::prescaler (*C++ member*), 503
 mcpwm_capture_channel_config_t::pull_down (*C++ member*), 503
 mcpwm_capture_channel_config_t::pull_up (*C++ member*), 503
 mcpwm_capture_channel_disable (*C++ function*), 501
 mcpwm_capture_channel_enable (*C++ function*), 501
 mcpwm_capture_channel_register_event_callbacks (*C++ function*), 501
 mcpwm_capture_channel_trigger_soft_catch (*C++ function*), 502
 mcpwm_capture_clock_source_t (*C++ type*), 507
 mcpwm_capture_edge_t (*C++ enum*), 509
 mcpwm_capture_edge_t::MCPWM_CAP_EDGE_NEG (*C++ enumerator*), 509
 mcpwm_capture_edge_t::MCPWM_CAP_EDGE_POS (*C++ enumerator*), 509
 mcpwm_capture_event_callbacks_t (*C++ struct*), 504
 mcpwm_capture_event_callbacks_t::on_cap (*C++ member*), 504
 mcpwm_capture_event_cb_t (*C++ type*), 507
 mcpwm_capture_event_data_t (*C++ struct*), 506
 mcpwm_capture_event_data_t::cap_edge (*C++ member*), 506
 mcpwm_capture_event_data_t::cap_value (*C++ member*), 506
 mcpwm_capture_timer_config_t (*C++ struct*), 502
 mcpwm_capture_timer_config_t::clk_src (*C++ member*), 502
 mcpwm_capture_timer_config_t::group_id (*C++ member*), 502
 mcpwm_capture_timer_config_t::resolution_hz (*C++ member*), 502
 mcpwm_capture_timer_disable (*C++ function*), 499
 mcpwm_capture_timer_enable (*C++ function*), 499
 mcpwm_capture_timer_get_resolution (*C++ function*), 500
 mcpwm_capture_timer_set_phase_on_sync (*C++ function*), 500
 mcpwm_capture_timer_start (*C++ function*), 499
 mcpwm_capture_timer_stop (*C++ function*), 500
 mcpwm_capture_timer_sync_phase_config_t (*C++ struct*), 502
 mcpwm_capture_timer_sync_phase_config_t::count_value (*C++ member*), 503
 mcpwm_capture_timer_sync_phase_config_t::direction (*C++ member*), 503
 mcpwm_capture_timer_sync_phase_config_t::sync_src (*C++ member*), 503
 mcpwm_carrier_clock_source_t (*C++ type*), 507
 mcpwm_carrier_config_t (*C++ struct*), 484
 mcpwm_carrier_config_t::clk_src (*C++ member*), 484
 mcpwm_carrier_config_t::duty_cycle (*C++ member*), 484
 mcpwm_carrier_config_t::first_pulse_duration_us (*C++ member*), 484
 mcpwm_carrier_config_t::flags (*C++ member*), 484
 mcpwm_carrier_config_t::frequency_hz (*C++ member*), 484
 mcpwm_carrier_config_t::invert_after_modulate (*C++ member*), 484
 mcpwm_carrier_config_t::invert_before_modulate (*C++ member*), 484
 mcpwm_cmpr_etm_event_config_t (*C++ struct*), 505
 mcpwm_cmpr_etm_event_config_t::event_type

- (C++ member), 505
- mcpwm_cmpr_handle_t (C++ type), 506
- mcpwm_comparator_config_t (C++ struct), 486
- mcpwm_comparator_config_t::flags (C++ member), 486
- mcpwm_comparator_config_t::intr_priority (C++ member), 486
- mcpwm_comparator_config_t::update_cmp_mcpwm (C++ member), 486
- mcpwm_comparator_config_t::update_cmp_mcpwm (C++ member), 486
- mcpwm_comparator_config_t::update_cmp_mcpwm (C++ member), 486
- mcpwm_comparator_etm_event_type_t (C++ enum), 510
- mcpwm_comparator_etm_event_type_t::MCPWM_GEN_ETM_EVENT_TYPE_EQUAL (C++ enumerator), 510
- mcpwm_comparator_etm_event_type_t::MCPWM_GEN_ETM_EVENT_TYPE_WX (C++ enumerator), 510
- mcpwm_comparator_event_callbacks_t (C++ struct), 486
- mcpwm_comparator_event_callbacks_t::on_fault_enter (C++ member), 486
- mcpwm_comparator_new_etm_event (C++ function), 504
- mcpwm_comparator_register_event_callbacks_t (C++ function), 485
- mcpwm_comparator_set_compare_value (C++ function), 486
- mcpwm_compare_event_cb_t (C++ type), 507
- mcpwm_compare_event_data_t (C++ struct), 505
- mcpwm_compare_event_data_t::compare_ticks (C++ member), 505
- mcpwm_compare_event_data_t::direction (C++ member), 506
- mcpwm_dead_time_config_t (C++ struct), 493
- mcpwm_dead_time_config_t::flags (C++ member), 493
- mcpwm_dead_time_config_t::invert_output (C++ member), 493
- mcpwm_dead_time_config_t::negedge_delay_ticks (C++ member), 493
- mcpwm_dead_time_config_t::posedge_delay_ticks (C++ member), 493
- mcpwm_del_capture_channel (C++ function), 501
- mcpwm_del_capture_timer (C++ function), 499
- mcpwm_del_comparator (C++ function), 485
- mcpwm_del_fault (C++ function), 494
- mcpwm_del_generator (C++ function), 487
- mcpwm_del_operator (C++ function), 481
- mcpwm_del_sync_src (C++ function), 497
- mcpwm_del_timer (C++ function), 477
- mcpwm_event_comparator_config_t (C++ struct), 486
- mcpwm_fault_event_callbacks_t (C++ struct), 496
- mcpwm_fault_event_callbacks_t::on_fault_enter (C++ member), 496
- mcpwm_fault_event_callbacks_t::on_fault_exit (C++ member), 496
- mcpwm_fault_event_cb_t (C++ type), 507
- mcpwm_fault_event_data_t (C++ struct), 505
- mcpwm_fault_handle_t (C++ type), 506
- mcpwm_fault_register_event_callbacks_t (C++ function), 495
- MCPWM_GEN_BRAKE_EVENT_ACTION (C macro), 493
- MCPWM_GEN_BRAKE_EVENT_ACTION_END (C macro), 493
- mcpwm_gen_brake_event_action_t (C++ struct), 492
- mcpwm_gen_brake_event_action_t::action (C++ member), 492
- mcpwm_gen_brake_event_action_t::brake_mode (C++ member), 492
- mcpwm_gen_brake_event_action_t::direction (C++ member), 492
- MCPWM_GEN_COMPARE_EVENT_ACTION (C macro), 493
- MCPWM_GEN_COMPARE_EVENT_ACTION_END (C macro), 493
- mcpwm_gen_compare_event_action_t (C++ struct), 491
- mcpwm_gen_compare_event_action_t::action (C++ member), 492
- mcpwm_gen_compare_event_action_t::comparator (C++ member), 492
- mcpwm_gen_compare_event_action_t::direction (C++ member), 492
- MCPWM_GEN_FAULT_EVENT_ACTION (C macro), 493
- mcpwm_gen_fault_event_action_t (C++ struct), 492
- mcpwm_gen_fault_event_action_t::action (C++ member), 492
- mcpwm_gen_fault_event_action_t::direction (C++ member), 492
- mcpwm_gen_fault_event_action_t::fault_ticks (C++ member), 492
- mcpwm_gen_handle_t (C++ type), 506
- MCPWM_GEN_SYNC_EVENT_ACTION (C macro), 493
- mcpwm_gen_sync_event_action_t (C++ struct), 492
- mcpwm_gen_sync_event_action_t::action (C++ member), 493
- mcpwm_gen_sync_event_action_t::direction (C++ member), 493
- mcpwm_gen_sync_event_action_t::sync (C++ member), 493
- MCPWM_GEN_TIMER_EVENT_ACTION (C macro), 493
- MCPWM_GEN_TIMER_EVENT_ACTION_END (C macro), 493

mcpwm_gen_timer_event_action_t (C++ struct), 491
 mcpwm_gen_timer_event_action_t::action (C++ member), 491
 mcpwm_gen_timer_event_action_t::direction (C++ member), 491
 mcpwm_gen_timer_event_action_t::event (C++ member), 491
 mcpwm_generator_action_t (C++ enum), 509
 mcpwm_generator_action_t::MCPWM_GEN_ACTION_KEEP (C++ enumerator), 509
 mcpwm_generator_action_t::MCPWM_GEN_ACTION_KEEP_TOGGLE (C++ enumerator), 509
 mcpwm_generator_action_t::MCPWM_GEN_ACTION_TOGGLE (C++ enumerator), 509
 mcpwm_generator_config_t (C++ struct), 491
 mcpwm_generator_config_t::flags (C++ member), 491
 mcpwm_generator_config_t::gen_gpio_num (C++ member), 491
 mcpwm_generator_config_t::invert_pwm (C++ member), 491
 mcpwm_generator_config_t::io_loop_back (C++ member), 491
 mcpwm_generator_config_t::io_od_mode (C++ member), 491
 mcpwm_generator_config_t::pull_down (C++ member), 491
 mcpwm_generator_config_t::pull_up (C++ member), 491
 mcpwm_generator_set_action_on_brake_event (C++ function), 489
 mcpwm_generator_set_action_on_compare_event (C++ function), 488
 mcpwm_generator_set_action_on_fault_event (C++ function), 490
 mcpwm_generator_set_action_on_sync_event (C++ function), 490
 mcpwm_generator_set_action_on_timer_event (C++ function), 488
 mcpwm_generator_set_actions_on_brake_event (C++ function), 489
 mcpwm_generator_set_actions_on_compare_event (C++ function), 489
 mcpwm_generator_set_actions_on_timer_event (C++ function), 488
 mcpwm_generator_set_dead_time (C++ function), 490
 mcpwm_generator_set_force_level (C++ function), 487
 mcpwm_gpio_fault_config_t (C++ struct), 495
 mcpwm_gpio_fault_config_t::active_level (C++ member), 495
 mcpwm_gpio_fault_config_t::flags (C++ member), 495
 mcpwm_gpio_fault_config_t::gpio_num (C++ member), 495
 mcpwm_gpio_fault_config_t::group_id (C++ member), 495
 mcpwm_gpio_fault_config_t::intr_priority (C++ member), 495
 mcpwm_gpio_fault_config_t::io_loop_back (C++ member), 495
 mcpwm_gpio_fault_config_t::pull_down (C++ member), 495
 mcpwm_gpio_fault_config_t::pull_up (C++ member), 495
 mcpwm_gpio_sync_src_config_t (C++ struct), 498
 mcpwm_gpio_sync_src_config_t::active_neg (C++ member), 498
 mcpwm_gpio_sync_src_config_t::flags (C++ member), 498
 mcpwm_gpio_sync_src_config_t::gpio_num (C++ member), 498
 mcpwm_gpio_sync_src_config_t::group_id (C++ member), 498
 mcpwm_gpio_sync_src_config_t::io_loop_back (C++ member), 498
 mcpwm_gpio_sync_src_config_t::pull_down (C++ member), 498
 mcpwm_gpio_sync_src_config_t::pull_up (C++ member), 498
 mcpwm_new_capture_channel (C++ function), 500
 mcpwm_new_capture_timer (C++ function), 499
 mcpwm_new_comparator (C++ function), 484
 mcpwm_new_event_comparator (C++ function), 485
 mcpwm_new_generator (C++ function), 487
 mcpwm_new_gpio_fault (C++ function), 494
 mcpwm_new_gpio_sync_src (C++ function), 496
 mcpwm_new_operator (C++ function), 481
 mcpwm_new_soft_fault (C++ function), 494
 mcpwm_new_soft_sync_src (C++ function), 497
 mcpwm_new_timer (C++ function), 477
 mcpwm_new_timer_sync_src (C++ function), 496
 mcpwm_oper_handle_t (C++ type), 506
 mcpwm_operator_apply_carrier (C++ function), 482
 mcpwm_operator_brake_mode_t (C++ enum), 509
 mcpwm_operator_brake_mode_t::MCPWM_OPER_BRAKE_MODE_KEEP (C++ enumerator), 509
 mcpwm_operator_brake_mode_t::MCPWM_OPER_BRAKE_MODE_TOGGLE (C++ enumerator), 509
 mcpwm_operator_brake_mode_t::MCPWM_OPER_BRAKE_MODE_TOGGLE_KEEP (C++ enumerator), 509
 mcpwm_operator_config_t (C++ struct), 482
 mcpwm_operator_config_t::flags (C++ member), 483
 mcpwm_operator_config_t::group_id (C++ member), 482

mcpwm_operator_config_t::intr_priority mcpwm_timer_count_mode_t::MCPWM_TIMER_COUNT_MODE_ (C++ member), 482 (C++ enumerator), 508
 mcpwm_operator_config_t::update_dead_time mcpwm_timer_count_mode_t::MCPWM_TIMER_COUNT_MODE_ (C++ member), 483 (C++ enumerator), 508
 mcpwm_operator_config_t::update_dead_time mcpwm_timer_count_mode_t::MCPWM_TIMER_COUNT_MODE_ (C++ member), 483 (C++ enumerator), 508
 mcpwm_operator_config_t::update_dead_time mcpwm_timer_direction_t (C++ enum), 508 (C++ member), 483 mcpwm_timer_direction_t::MCPWM_TIMER_DIRECTION_DO (C++ enumerator), 508
 mcpwm_operator_config_t::update_gen_action_on mcpwm_timer_direction_t::MCPWM_TIMER_DIRECTION_UP (C++ member), 483 (C++ enumerator), 508
 mcpwm_operator_config_t::update_gen_action_on mcpwm_timer_disable (C++ function), 478 (C++ member), 482
 mcpwm_operator_config_t::update_gen_action_on mcpwm_timer_enable (C++ function), 478 (C++ member), 482
 mcpwm_operator_connect_timer (C++ function), 481 mcpwm_timer_event_callbacks_t (C++ struct), 479
 mcpwm_operator_event_callbacks_t (C++ struct), 483 mcpwm_timer_event_callbacks_t::on_empty (C++ member), 479
 mcpwm_operator_event_callbacks_t::on_brake_cb (C++ member), 483 mcpwm_timer_event_callbacks_t::on_full (C++ member), 479
 mcpwm_operator_event_callbacks_t::on_brake_ost (C++ member), 483 mcpwm_timer_event_callbacks_t::on_stop (C++ member), 479
 mcpwm_operator_recover_from_fault mcpwm_timer_event_cb_t (C++ type), 506 (C++ function), 481 mcpwm_timer_event_data_t (C++ struct), 505
 mcpwm_operator_register_event_callbacks mcpwm_timer_event_data_t::count_value (C++ function), 482 (C++ member), 505
 mcpwm_operator_set_brake_on_fault mcpwm_timer_event_data_t::direction (C++ function), 481 (C++ member), 505
 mcpwm_soft_fault_activate (C++ function), 494 mcpwm_timer_event_t (C++ enum), 508
 mcpwm_soft_fault_config_t (C++ struct), 496 mcpwm_timer_event_t::MCPWM_TIMER_EVENT_EMPTY (C++ enumerator), 508
 mcpwm_soft_sync_activate (C++ function), 497 mcpwm_timer_event_t::MCPWM_TIMER_EVENT_FULL (C++ enumerator), 508
 mcpwm_soft_sync_config_t (C++ struct), 498 mcpwm_timer_event_t::MCPWM_TIMER_EVENT_INVALID (C++ enumerator), 508
 mcpwm_sync_handle_t (C++ type), 506 mcpwm_timer_handle_t (C++ type), 506
 mcpwm_timer_clock_source_t (C++ type), 507 mcpwm_timer_register_event_callbacks (C++ function), 478
 mcpwm_timer_config_t (C++ struct), 479 mcpwm_timer_set_period (C++ function), 477
 mcpwm_timer_config_t::clk_src (C++ member), 479 mcpwm_timer_set_phase_on_sync (C++ function), 479
 mcpwm_timer_config_t::count_mode (C++ member), 480 mcpwm_timer_start_stop (C++ function), 478
 mcpwm_timer_config_t::flags (C++ member), 480 mcpwm_timer_start_stop_cmd_t (C++ enum), 508
 mcpwm_timer_config_t::group_id (C++ member), 479 mcpwm_timer_start_stop_cmd_t::MCPWM_TIMER_START_N (C++ enumerator), 509
 mcpwm_timer_config_t::intr_priority (C++ member), 480 mcpwm_timer_start_stop_cmd_t::MCPWM_TIMER_START_S (C++ enumerator), 509
 mcpwm_timer_config_t::period_ticks (C++ member), 480 mcpwm_timer_start_stop_cmd_t::MCPWM_TIMER_START_S (C++ enumerator), 509
 mcpwm_timer_config_t::resolution_hz (C++ member), 479 mcpwm_timer_start_stop_cmd_t::MCPWM_TIMER_STOP_EM (C++ enumerator), 508
 mcpwm_timer_config_t::update_period_on_empty mcpwm_timer_start_stop_cmd_t::MCPWM_TIMER_STOP_FU (C++ member), 480 (C++ enumerator), 508
 mcpwm_timer_config_t::update_period_on_sync mcpwm_timer_sync_phase_config_t (C++ struct), 480 (C++ member), 480
 mcpwm_timer_count_mode_t (C++ enum), 508 mcpwm_timer_sync_phase_config_t::count_value (C++ member), 480
 mcpwm_timer_count_mode_t::MCPWM_TIMER_COUNT_MODE_ (C++ enumerator), 508 mcpwm_timer_sync_phase_config_t::direction (C++ member), 480

- (C++ member), 480
- mcpwm_timer_sync_phase_config_t::sync_src (C++ member), 480
- mcpwm_timer_sync_src_config_t (C++ struct), 497
- mcpwm_timer_sync_src_config_t::flags (C++ member), 498
- mcpwm_timer_sync_src_config_t::propagated (C++ member), 498
- mcpwm_timer_sync_src_config_t::timer_event (C++ member), 498
- MessageBufferHandle_t (C++ type), 1403
- MQTT_ERROR_TYPE_ESP_TLS (C macro), 56
- multi_heap_aligned_alloc (C++ function), 1441
- multi_heap_aligned_alloc_offs (C++ function), 1444
- multi_heap_aligned_free (C++ function), 1442
- multi_heap_check (C++ function), 1443
- multi_heap_dump (C++ function), 1443
- multi_heap_free (C++ function), 1442
- multi_heap_free_size (C++ function), 1443
- multi_heap_get_allocated_size (C++ function), 1442
- multi_heap_get_info (C++ function), 1444
- multi_heap_handle_t (C++ type), 1444
- multi_heap_info_t (C++ struct), 1444
- multi_heap_info_t::allocated_blocks (C++ member), 1444
- multi_heap_info_t::free_blocks (C++ member), 1444
- multi_heap_info_t::largest_free_block (C++ member), 1444
- multi_heap_info_t::minimum_free_bytes (C++ member), 1444
- multi_heap_info_t::total_allocated_bytes (C++ member), 1444
- multi_heap_info_t::total_blocks (C++ member), 1444
- multi_heap_info_t::total_free_bytes (C++ member), 1444
- multi_heap_malloc (C++ function), 1442
- multi_heap_minimum_free_size (C++ function), 1443
- multi_heap_realloc (C++ function), 1442
- multi_heap_register (C++ function), 1442
- multi_heap_set_lock (C++ function), 1443
- nvs_entry_find_in_handle (C++ function), 1137
- nvs_entry_info (C++ function), 1138
- nvs_entry_info_t (C++ struct), 1138
- nvs_entry_info_t::key (C++ member), 1138
- nvs_entry_info_t::namespace_name (C++ member), 1138
- nvs_entry_info_t::type (C++ member), 1138
- nvs_entry_next (C++ function), 1137
- nvs_erase_all (C++ function), 1134
- nvs_erase_key (C++ function), 1134
- nvs_find_key (C++ function), 1134
- nvs_flash_deinit (C++ function), 1125
- nvs_flash_deinit_partition (C++ function), 1125
- nvs_flash_erase (C++ function), 1125
- nvs_flash_erase_partition (C++ function), 1125
- nvs_flash_erase_partition_ptr (C++ function), 1126
- nvs_flash_generate_keys (C++ function), 1127
- nvs_flash_generate_keys_t (C++ type), 1129
- nvs_flash_generate_keys_v2 (C++ function), 1127
- nvs_flash_get_default_security_scheme (C++ function), 1127
- nvs_flash_init (C++ function), 1124
- nvs_flash_init_partition (C++ function), 1125
- nvs_flash_init_partition_ptr (C++ function), 1125
- nvs_flash_read_cfg_t (C++ type), 1129
- nvs_flash_read_security_cfg (C++ function), 1127
- nvs_flash_read_security_cfg_v2 (C++ function), 1128
- nvs_flash_register_security_scheme (C++ function), 1127
- nvs_flash_secure_init (C++ function), 1126
- nvs_flash_secure_init_partition (C++ function), 1126
- nvs_get_blob (C++ function), 1132
- nvs_get_i16 (C++ function), 1131
- nvs_get_i32 (C++ function), 1131
- nvs_get_i64 (C++ function), 1131
- nvs_get_i8 (C++ function), 1130
- nvs_get_stats (C++ function), 1135
- nvs_get_str (C++ function), 1131
- nvs_get_u16 (C++ function), 1131
- nvs_get_u32 (C++ function), 1131
- nvs_get_u64 (C++ function), 1131
- nvs_get_u8 (C++ function), 1131
- nvs_get_used_entry_count (C++ function), 1136
- nvs_handle (C++ type), 1141
- nvs_handle_t (C++ type), 1141
- nvs_iterator_t (C++ type), 1141

N

- name_uuid (C++ struct), 1094
- name_uuid::name (C++ member), 1094
- name_uuid::uuid (C++ member), 1094
- nvs_close (C++ function), 1135
- nvs_commit (C++ function), 1135
- NVS_DEFAULT_PART_NAME (C macro), 1140
- nvs_entry_find (C++ function), 1136

- [NVS_KEY_NAME_MAX_SIZE \(C macro\), 1140](#)
[NVS_KEY_SIZE \(C macro\), 1128](#)
[NVS_NS_NAME_MAX_SIZE \(C macro\), 1140](#)
[nvs_open \(C++ function\), 1132](#)
[nvs_open_from_partition \(C++ function\), 1133](#)
[nvs_open_mode \(C++ type\), 1141](#)
[nvs_open_mode_t \(C++ enum\), 1141](#)
[nvs_open_mode_t::NVS_READONLY \(C++ enumerator\), 1141](#)
[nvs_open_mode_t::NVS_READWRITE \(C++ enumerator\), 1141](#)
[NVS_PART_NAME_MAX_SIZE \(C macro\), 1140](#)
[nvs_release_iterator \(C++ function\), 1138](#)
[nvs_sec_cfg_t \(C++ struct\), 1128](#)
[nvs_sec_cfg_t::eky \(C++ member\), 1128](#)
[nvs_sec_cfg_t::tky \(C++ member\), 1128](#)
[nvs_sec_config_flash_enc_t \(C++ struct\), 1147](#)
[nvs_sec_config_flash_enc_t::nvs_keys_part \(C++ member\), 1147](#)
[nvs_sec_config_hmac_t \(C++ struct\), 1147](#)
[nvs_sec_config_hmac_t::hmac_key_id \(C++ member\), 1147](#)
[NVS_SEC_PROVIDER_CFG_FLASH_ENC_DEFAULT \(C macro\), 1148](#)
[NVS_SEC_PROVIDER_CFG_HMAC_DEFAULT \(C macro\), 1148](#)
[nvs_sec_provider_deregister \(C++ function\), 1147](#)
[nvs_sec_provider_register_flash_enc \(C++ function\), 1147](#)
[nvs_sec_provider_register_hmac \(C++ function\), 1147](#)
[nvs_sec_scheme_id_t \(C++ enum\), 1148](#)
[nvs_sec_scheme_id_t::NVS_SEC_SCHEME_FLASH_ENC \(C++ enumerator\), 1148](#)
[nvs_sec_scheme_id_t::NVS_SEC_SCHEME_HMAC \(C++ enumerator\), 1148](#)
[nvs_sec_scheme_id_t::NVS_SEC_SCHEME_MAX \(C++ enumerator\), 1148](#)
[nvs_sec_scheme_t \(C++ struct\), 1128](#)
[nvs_sec_scheme_t::nvs_flash_key_gen \(C++ member\), 1128](#)
[nvs_sec_scheme_t::nvs_flash_read_cfg \(C++ member\), 1128](#)
[nvs_sec_scheme_t::scheme_data \(C++ member\), 1128](#)
[nvs_sec_scheme_t::scheme_id \(C++ member\), 1128](#)
[nvs_set_blob \(C++ function\), 1133](#)
[nvs_set_i16 \(C++ function\), 1129](#)
[nvs_set_i32 \(C++ function\), 1130](#)
[nvs_set_i64 \(C++ function\), 1130](#)
[nvs_set_i8 \(C++ function\), 1129](#)
[nvs_set_str \(C++ function\), 1130](#)
[nvs_set_u16 \(C++ function\), 1129](#)
[nvs_set_u32 \(C++ function\), 1130](#)
[nvs_set_u64 \(C++ function\), 1130](#)
[nvs_set_u8 \(C++ function\), 1129](#)
[nvs_stats_t \(C++ struct\), 1138](#)
[nvs_stats_t::available_entries \(C++ member\), 1138](#)
[nvs_stats_t::free_entries \(C++ member\), 1138](#)
[nvs_stats_t::namespace_count \(C++ member\), 1139](#)
[nvs_stats_t::total_entries \(C++ member\), 1139](#)
[nvs_stats_t::used_entries \(C++ member\), 1138](#)
[nvs_type_t \(C++ enum\), 1141](#)
[nvs_type_t::NVS_TYPE_ANY \(C++ enumerator\), 1142](#)
[nvs_type_t::NVS_TYPE_BLOB \(C++ enumerator\), 1142](#)
[nvs_type_t::NVS_TYPE_I16 \(C++ enumerator\), 1141](#)
[nvs_type_t::NVS_TYPE_I32 \(C++ enumerator\), 1141](#)
[nvs_type_t::NVS_TYPE_I64 \(C++ enumerator\), 1142](#)
[nvs_type_t::NVS_TYPE_I8 \(C++ enumerator\), 1141](#)
[nvs_type_t::NVS_TYPE_STR \(C++ enumerator\), 1142](#)
[nvs_type_t::NVS_TYPE_U16 \(C++ enumerator\), 1141](#)
[nvs_type_t::NVS_TYPE_U32 \(C++ enumerator\), 1141](#)
[nvs_type_t::NVS_TYPE_U64 \(C++ enumerator\), 1141](#)
[nvs_type_t::NVS_TYPE_U8 \(C++ enumerator\), 1141](#)
- ## O
- [OTA_SIZE_UNKNOWN \(C macro\), 1526](#)
[OTA_WITH_SEQUENTIAL_WRITES \(C macro\), 1526](#)
- ## P
- [parlio_bit_pack_order_t \(C++ enum\), 516](#)
[parlio_bit_pack_order_t::PARLIO_BIT_PACK_ORDER_LSB \(C++ enumerator\), 516](#)
[parlio_bit_pack_order_t::PARLIO_BIT_PACK_ORDER_MSB \(C++ enumerator\), 516](#)
[parlio_clock_source_t \(C++ type\), 516](#)
[parlio_del_tx_unit \(C++ function\), 511](#)
[parlio_new_tx_unit \(C++ function\), 511](#)
[parlio_sample_edge_t \(C++ enum\), 516](#)
[parlio_sample_edge_t::PARLIO_SAMPLE_EDGE_NEG \(C++ enumerator\), 516](#)
[parlio_sample_edge_t::PARLIO_SAMPLE_EDGE_POS \(C++ enumerator\), 516](#)
[parlio_shift_edge_t \(C++ enum\), 516](#)
[parlio_shift_edge_t::PARLIO_SHIFT_EDGE_NEG \(C++ enumerator\), 516](#)

parlio_shift_edge_t::PARLIO_SHIFT_EDGE_POS *(C++ enumerator)*, 516
 parlio_transmit_config_t (C++ struct), 514
 parlio_transmit_config_t::flags (C++ member), 515
 parlio_transmit_config_t::idle_value (C++ member), 515
 parlio_transmit_config_t::queue_nonblocking (C++ member), 515
 parlio_tx_done_callback_t (C++ type), 515
 parlio_tx_done_event_data_t (C++ struct), 514
 parlio_tx_event_callbacks_t (C++ struct), 514
 parlio_tx_event_callbacks_t::on_trans_done (C++ member), 514
 parlio_tx_unit_config_t (C++ struct), 513
 parlio_tx_unit_config_t::bit_pack_order (C++ member), 514
 parlio_tx_unit_config_t::clk_gate_en (C++ member), 514
 parlio_tx_unit_config_t::clk_in_gpio_num (C++ member), 513
 parlio_tx_unit_config_t::clk_out_gpio_num (C++ member), 513
 parlio_tx_unit_config_t::clk_src (C++ member), 513
 parlio_tx_unit_config_t::data_gpio_nums (C++ member), 513
 parlio_tx_unit_config_t::data_width (C++ member), 513
 parlio_tx_unit_config_t::flags (C++ member), 514
 parlio_tx_unit_config_t::input_clk_src_freq_div (C++ member), 513
 parlio_tx_unit_config_t::io_loop_back (C++ member), 514
 parlio_tx_unit_config_t::max_transfer_size (C++ member), 514
 parlio_tx_unit_config_t::output_clk_freq_div (C++ member), 513
 parlio_tx_unit_config_t::sample_edge (C++ member), 514
 parlio_tx_unit_config_t::shift_edge (C++ member), 514
 parlio_tx_unit_config_t::trans_queue_depth (C++ member), 514
 parlio_tx_unit_config_t::valid_gpio_nums (C++ member), 513
 parlio_tx_unit_disable (C++ function), 511
 parlio_tx_unit_enable (C++ function), 511
 parlio_tx_unit_handle_t (C++ type), 515
 PARLIO_TX_UNIT_MAX_DATA_WIDTH (C macro), 515
 parlio_tx_unit_register_event_callbacks (C++ function), 512
 parlio_tx_unit_transmit (C++ function), 512
 parlio_tx_unit_wait_all_done (C++ function), 513
 pcnt_chan_config_t (C++ struct), 530
 pcnt_chan_config_t::edge_gpio_num (C++ member), 530
 pcnt_chan_config_t::flags (C++ member), 531
 pcnt_chan_config_t::invert_edge_input (C++ member), 530
 pcnt_chan_config_t::invert_level_input (C++ member), 530
 pcnt_chan_config_t::io_loop_back (C++ member), 530
 pcnt_chan_config_t::level_gpio_num (C++ member), 530
 pcnt_chan_config_t::virt_edge_io_level (C++ member), 530
 pcnt_chan_config_t::virt_level_io_level (C++ member), 530
 pcnt_channel_edge_action_t (C++ enum), 532
 pcnt_channel_edge_action_t::PCNT_CHANNEL_EDGE_ACTION_NONE (C++ enumerator), 532
 pcnt_channel_edge_action_t::PCNT_CHANNEL_EDGE_ACTION_RISING (C++ enumerator), 532
 pcnt_channel_edge_action_t::PCNT_CHANNEL_EDGE_ACTION_FALLING (C++ enumerator), 532
 pcnt_channel_handle_t (C++ type), 531
 pcnt_channel_level_action_t (C++ enum), 532
 pcnt_channel_level_action_t::PCNT_CHANNEL_LEVEL_ACTION_NONE (C++ enumerator), 532
 pcnt_channel_level_action_t::PCNT_CHANNEL_LEVEL_ACTION_RISING (C++ enumerator), 532
 pcnt_channel_level_action_t::PCNT_CHANNEL_LEVEL_ACTION_FALLING (C++ enumerator), 532
 pcnt_channel_set_edge_action (C++ function), 528
 pcnt_channel_set_level_action (C++ function), 529
 pcnt_clear_signal_config_t (C++ struct), 531
 pcnt_clear_signal_config_t::clear_signal_gpio_num (C++ member), 531
 pcnt_clear_signal_config_t::flags (C++ member), 531
 pcnt_clear_signal_config_t::invert_clear_signal (C++ member), 531
 pcnt_clear_signal_config_t::io_loop_back (C++ member), 531
 pcnt_del_channel (C++ function), 528
 pcnt_del_unit (C++ function), 523
 pcnt_event_callbacks_t (C++ struct), 529
 pcnt_event_callbacks_t::on_reach (C++ member), 529
 pcnt_glitch_filter_config_t (C++ struct), 531
 pcnt_glitch_filter_config_t::max_glitch_ns (C++ member), 531

- pcnt_new_channel (C++ function), 528
 pcnt_new_unit (C++ function), 523
 pcnt_unit_add_watch_point (C++ function), 527
 pcnt_unit_clear_count (C++ function), 526
 pcnt_unit_config_t (C++ struct), 530
 pcnt_unit_config_t::accum_count (C++ member), 530
 pcnt_unit_config_t::flags (C++ member), 530
 pcnt_unit_config_t::high_limit (C++ member), 530
 pcnt_unit_config_t::intr_priority (C++ member), 530
 pcnt_unit_config_t::low_limit (C++ member), 530
 pcnt_unit_disable (C++ function), 525
 pcnt_unit_enable (C++ function), 525
 pcnt_unit_get_count (C++ function), 527
 pcnt_unit_handle_t (C++ type), 531
 pcnt_unit_register_event_callbacks (C++ function), 527
 pcnt_unit_remove_watch_point (C++ function), 528
 pcnt_unit_set_clear_signal (C++ function), 524
 pcnt_unit_set_glitch_filter (C++ function), 524
 pcnt_unit_start (C++ function), 525
 pcnt_unit_stop (C++ function), 526
 pcnt_unit_zero_cross_mode_t (C++ enum), 532
 pcnt_unit_zero_cross_mode_t::PCNT_UNIT_ZERO_CROSS_MODE_NEG_POS (C++ enumerator), 532
 pcnt_unit_zero_cross_mode_t::PCNT_UNIT_ZERO_CROSS_MODE_NEG_ZERO (C++ enumerator), 532
 pcnt_unit_zero_cross_mode_t::PCNT_UNIT_ZERO_CROSS_MODE_POS_ZERO (C++ enumerator), 532
 pcnt_unit_zero_cross_mode_t::PCNT_UNIT_ZERO_CROSS_MODE_POS_POS (C++ enumerator), 532
 pcnt_watch_cb_t (C++ type), 531
 pcnt_watch_event_data_t (C++ struct), 529
 pcnt_watch_event_data_t::watch_point_value (C++ member), 529
 pcnt_watch_event_data_t::zero_cross_mode (C++ member), 529
 pcQueueGetName (C++ function), 1332
 pcTaskGetName (C++ function), 1307
 pcTimerGetName (C++ function), 1366
 PendedFunction_t (C++ type), 1376
 phy_802_3_t (C++ struct), 203
 phy_802_3_t::addr (C++ member), 203
 phy_802_3_t::autonego_timeout_ms (C++ member), 203
 phy_802_3_t::eth (C++ member), 203
 phy_802_3_t::link_status (C++ member), 203
 phy_802_3_t::parent (C++ member), 203
 phy_802_3_t::reset_gpio_num (C++ member), 203
 phy_802_3_t::reset_timeout_ms (C++ member), 203
 protocomm_add_endpoint (C++ function), 1082
 protocomm_ble_config (C++ struct), 1094
 protocomm_ble_config::ble_addr (C++ member), 1095
 protocomm_ble_config::ble_bonding (C++ member), 1095
 protocomm_ble_config::ble_link_encryption (C++ member), 1095
 protocomm_ble_config::ble_notify (C++ member), 1095
 protocomm_ble_config::ble_sm_sc (C++ member), 1095
 protocomm_ble_config::device_name (C++ member), 1095
 protocomm_ble_config::keep_ble_on (C++ member), 1095
 protocomm_ble_config::manufacturer_data (C++ member), 1095
 protocomm_ble_config::manufacturer_data_len (C++ member), 1095
 protocomm_ble_config::nu_lookup (C++ member), 1095
 protocomm_ble_config::nu_lookup_count (C++ member), 1095
 protocomm_ble_config::service_uuid (C++ member), 1095
 protocomm_ble_config_t (C++ type), 1096
 protocomm_ble_event_t (C++ struct), 1094
 protocomm_ble_event_t::conn_handle (C++ member), 1094
 protocomm_ble_event_t::conn_status (C++ member), 1094
 protocomm_ble_event_t::disconnect_reason (C++ member), 1094
 protocomm_ble_event_t::evt_type (C++ member), 1094
 protocomm_ble_name_uuid_t (C++ type), 1096
 protocomm_ble_start (C++ function), 1093
 protocomm_ble_stop (C++ function), 1094
 protocomm_close_session (C++ function), 1083
 protocomm_delete (C++ function), 1082
 protocomm_get_sec_version (C++ function), 1085
 protocomm_http_server_config_t (C++ struct), 1092
 protocomm_http_server_config_t::port (C++ member), 1093
 protocomm_http_server_config_t::stack_size (C++ member), 1093
 protocomm_http_server_config_t::task_priority (C++ member), 1093
 protocomm_httpd_config_data_t (C++ union), 1092

protocomm_httpd_config_data_t::config (C++ member), 1092
 protocomm_httpd_config_data_t::handle (C++ member), 1092
 protocomm_httpd_config_t (C++ struct), 1093
 protocomm_httpd_config_t::data (C++ member), 1093
 protocomm_httpd_config_t::ext_handle_provided (C++ member), 1093
 PROTOCOL_HTTPD_DEFAULT_CONFIG (C macro), 1093
 protocomm_httpd_start (C++ function), 1092
 protocomm_httpd_stop (C++ function), 1092
 protocomm_new (C++ function), 1082
 protocomm_open_session (C++ function), 1083
 protocomm_remove_endpoint (C++ function), 1082
 protocomm_req_handle (C++ function), 1083
 protocomm_req_handler_t (C++ type), 1085
 protocomm_security (C++ struct), 1086
 protocomm_security1_params (C++ struct), 1086
 protocomm_security1_params::data (C++ member), 1086
 protocomm_security1_params::len (C++ member), 1086
 protocomm_security1_params_t (C++ type), 1087
 protocomm_security2_params (C++ struct), 1086
 protocomm_security2_params::salt (C++ member), 1086
 protocomm_security2_params::salt_len (C++ member), 1086
 protocomm_security2_params::verifier (C++ member), 1086
 protocomm_security2_params::verifier_len (C++ member), 1086
 protocomm_security2_params_t (C++ type), 1087
 protocomm_security::cleanup (C++ member), 1087
 protocomm_security::close_transport_session (C++ member), 1087
 protocomm_security::decrypt (C++ member), 1087
 protocomm_security::encrypt (C++ member), 1087
 protocomm_security::init (C++ member), 1086
 protocomm_security::new_transport_session (C++ member), 1087
 protocomm_security::patch_ver (C++ member), 1086
 protocomm_security::security_req_handler (C++ member), 1087
 protocomm_security::ver (C++ member), 1086
 protocomm_security_handle_t (C++ type), 1087
 protocomm_security_pop_t (C++ type), 1087
 protocomm_security_session_event_t (C++ enum), 1087
 protocomm_security_session_event_t::PROTOCOL_SECURITY_SESSION_EVENT (C++ enumerator), 1088
 protocomm_security_session_event_t::PROTOCOL_SECURITY_SESSION_EVENT_INVALID (C++ enumerator), 1088
 protocomm_security_session_event_t::PROTOCOL_SECURITY_SESSION_EVENT_MAX (C++ enumerator), 1087
 protocomm_security_t (C++ type), 1087
 protocomm_set_security (C++ function), 1083
 protocomm_set_version (C++ function), 1084
 protocomm_t (C++ type), 1085
 protocomm_transport_ble_event_t (C++ enum), 1096
 protocomm_transport_ble_event_t::PROTOCOL_TRANSPORT_BLE_EVENT (C++ enumerator), 1096
 protocomm_transport_ble_event_t::PROTOCOL_TRANSPORT_BLE_EVENT_INVALID (C++ enumerator), 1096
 protocomm_unset_security (C++ function), 1084
 protocomm_unset_version (C++ function), 1085
 psk_hint_key_t (C++ type), 72
 psk_key_hint (C++ struct), 69
 psk_key_hint::hint (C++ member), 69
 psk_key_hint::key (C++ member), 69
 psk_key_hint::key_size (C++ member), 69
 PTHREAD_STACK_MIN (C macro), 1539
 pvTaskGetThreadLocalStoragePointer (C++ function), 1308
 pvTimerGetTimerID (C++ function), 1363
 pxTaskGetStackStart (C++ function), 1426

Q

QueueHandle_t (C++ type), 1344
 QueueSetHandle_t (C++ type), 1344
 QueueSetMemberHandle_t (C++ type), 1344

R

RingbufferType_t (C++ enum), 1422
 RingbufferType_t::RINGBUF_TYPE_ALLOWSPLIT (C++ enumerator), 1422
 RingbufferType_t::RINGBUF_TYPE_BYTEBUF (C++ enumerator), 1423
 RingbufferType_t::RINGBUF_TYPE_MAX (C++ enumerator), 1423
 RingbufferType_t::RINGBUF_TYPE_NOSPLIT (C++ enumerator), 1422
 RingbufHandle_t (C++ type), 1422
 rmt_alloc_encoder_mem (C++ function), 557
 rmt_apply_carrier (C++ function), 554
 rmt_bytes_encoder_config_t (C++ struct), 558
 rmt_bytes_encoder_config_t::bit0 (C++ member), 558

[rmt_bytes_encoder_config_t::bit1 \(C++ member\), 558](#)
[rmt_bytes_encoder_config_t::flags \(C++ member\), 558](#)
[rmt_bytes_encoder_config_t::msb_first \(C++ member\), 558](#)
[rmt_bytes_encoder_update_config \(C++ function\), 556](#)
[rmt_carrier_config_t \(C++ struct\), 555](#)
[rmt_carrier_config_t::always_on \(C++ member\), 555](#)
[rmt_carrier_config_t::duty_cycle \(C++ member\), 555](#)
[rmt_carrier_config_t::flags \(C++ member\), 555](#)
[rmt_carrier_config_t::frequency_hz \(C++ member\), 555](#)
[rmt_carrier_config_t::polarity_active_low \(C++ member\), 555](#)
[rmt_channel_handle_t \(C++ type\), 559](#)
[rmt_clock_source_t \(C++ type\), 560](#)
[rmt_copy_encoder_config_t \(C++ struct\), 558](#)
[rmt_del_channel \(C++ function\), 554](#)
[rmt_del_encoder \(C++ function\), 556](#)
[rmt_del_sync_manager \(C++ function\), 548](#)
[rmt_disable \(C++ function\), 555](#)
[rmt_enable \(C++ function\), 554](#)
[rmt_encode_state_t \(C++ enum\), 558](#)
[rmt_encode_state_t::RMT_ENCODING_COMPLETED \(C++ enumerator\), 558](#)
[rmt_encode_state_t::RMT_ENCODING_MEMORY_FULL \(C++ enumerator\), 558](#)
[rmt_encode_state_t::RMT_ENCODING_RESET \(C++ enumerator\), 558](#)
[rmt_encoder_handle_t \(C++ type\), 559](#)
[rmt_encoder_reset \(C++ function\), 557](#)
[rmt_encoder_t \(C++ struct\), 557](#)
[rmt_encoder_t::del \(C++ member\), 557](#)
[rmt_encoder_t::encode \(C++ member\), 557](#)
[rmt_encoder_t::reset \(C++ member\), 557](#)
[rmt_new_bytes_encoder \(C++ function\), 556](#)
[rmt_new_copy_encoder \(C++ function\), 556](#)
[rmt_new_rx_channel \(C++ function\), 551](#)
[rmt_new_sync_manager \(C++ function\), 548](#)
[rmt_new_tx_channel \(C++ function\), 546](#)
[rmt_receive \(C++ function\), 551](#)
[rmt_receive_config_t \(C++ struct\), 553](#)
[rmt_receive_config_t::signal_range_max_ns \(C++ member\), 554](#)
[rmt_receive_config_t::signal_range_min_ns \(C++ member\), 554](#)
[rmt_rx_channel_config_t \(C++ struct\), 553](#)
[rmt_rx_channel_config_t::clk_src \(C++ member\), 553](#)
[rmt_rx_channel_config_t::flags \(C++ member\), 553](#)
[rmt_rx_channel_config_t::gpio_num \(C++ member\), 553](#)
[rmt_rx_channel_config_t::intr_priority \(C++ member\), 553](#)
[rmt_rx_channel_config_t::invert_in \(C++ member\), 553](#)
[rmt_rx_channel_config_t::io_loop_back \(C++ member\), 553](#)
[rmt_rx_channel_config_t::mem_block_symbols \(C++ member\), 553](#)
[rmt_rx_channel_config_t::resolution_hz \(C++ member\), 553](#)
[rmt_rx_channel_config_t::with_dma \(C++ member\), 553](#)
[rmt_rx_done_callback_t \(C++ type\), 559](#)
[rmt_rx_done_event_data_t \(C++ struct\), 559](#)
[rmt_rx_done_event_data_t::num_symbols \(C++ member\), 559](#)
[rmt_rx_done_event_data_t::received_symbols \(C++ member\), 559](#)
[rmt_rx_event_callbacks_t \(C++ struct\), 552](#)
[rmt_rx_event_callbacks_t::on_recv_done \(C++ member\), 553](#)
[rmt_rx_register_event_callbacks \(C++ function\), 552](#)
[rmt_symbol_word_t \(C++ union\), 560](#)
[rmt_symbol_word_t::duration0 \(C++ member\), 560](#)
[rmt_symbol_word_t::duration1 \(C++ member\), 560](#)
[rmt_symbol_word_t::level0 \(C++ member\), 560](#)
[rmt_symbol_word_t::level1 \(C++ member\), 560](#)
[rmt_symbol_word_t::val \(C++ member\), 560](#)
[rmt_symbol_word_t::\[anonymous\] \(C++ member\), 560](#)
[rmt_sync_manager_config_t \(C++ struct\), 551](#)
[rmt_sync_manager_config_t::array_size \(C++ member\), 551](#)
[rmt_sync_manager_config_t::tx_channel_array \(C++ member\), 551](#)
[rmt_sync_manager_handle_t \(C++ type\), 559](#)
[rmt_sync_reset \(C++ function\), 549](#)
[rmt_transmit \(C++ function\), 547](#)
[rmt_transmit_config_t \(C++ struct\), 550](#)
[rmt_transmit_config_t::eot_level \(C++ member\), 550](#)
[rmt_transmit_config_t::flags \(C++ member\), 551](#)
[rmt_transmit_config_t::loop_count \(C++ member\), 550](#)
[rmt_transmit_config_t::queue_nonblocking \(C++ member\), 550](#)
[rmt_tx_channel_config_t \(C++ struct\), 549](#)
[rmt_tx_channel_config_t::clk_src \(C++ member\), 549](#)
[rmt_tx_channel_config_t::flags \(C++ member\), 550](#)
[rmt_tx_channel_config_t::gpio_num](#)

- (C++ member), 549
 - rmt_tx_channel_config_t::init_level (C++ member), 550
 - rmt_tx_channel_config_t::intr_priority (C++ member), 550
 - rmt_tx_channel_config_t::invert_out (C++ member), 550
 - rmt_tx_channel_config_t::io_loop_back (C++ member), 550
 - rmt_tx_channel_config_t::io_od_mode (C++ member), 550
 - rmt_tx_channel_config_t::mem_block_symbols (C++ member), 550
 - rmt_tx_channel_config_t::resolution_hz (C++ member), 550
 - rmt_tx_channel_config_t::trans_queue_depth (C++ member), 550
 - rmt_tx_channel_config_t::with_dma (C++ member), 550
 - rmt_tx_done_callback_t (C++ type), 559
 - rmt_tx_done_event_data_t (C++ struct), 559
 - rmt_tx_done_event_data_t::num_symbols (C++ member), 559
 - rmt_tx_event_callbacks_t (C++ struct), 549
 - rmt_tx_event_callbacks_t::on_trans_done (C++ member), 549
 - rmt_tx_register_event_callbacks (C++ function), 548
 - rmt_tx_switch_gpio (C++ function), 549
 - rmt_tx_wait_all_done (C++ function), 547
 - rtc_gpio_deinit (C++ function), 307
 - rtc_gpio_force_hold_dis_all (C++ function), 309
 - rtc_gpio_force_hold_en_all (C++ function), 309
 - rtc_gpio_get_drive_capability (C++ function), 308
 - rtc_gpio_get_level (C++ function), 307
 - rtc_gpio_hold_dis (C++ function), 309
 - rtc_gpio_hold_en (C++ function), 309
 - rtc_gpio_init (C++ function), 307
 - rtc_gpio_iomux_func_sel (C++ function), 309
 - RTC_GPIO_IS_VALID_GPIO (C macro), 310
 - rtc_gpio_is_valid_gpio (C++ function), 306
 - rtc_gpio_mode_t (C++ enum), 311
 - rtc_gpio_mode_t::RTC_GPIO_MODE_DISABLED (C++ enumerator), 311
 - rtc_gpio_mode_t::RTC_GPIO_MODE_INPUT_ONLY (C++ enumerator), 311
 - rtc_gpio_mode_t::RTC_GPIO_MODE_INPUT_OUTPUT (C++ enumerator), 311
 - rtc_gpio_mode_t::RTC_GPIO_MODE_INPUT_OUTPUT_ONLY (C++ enumerator), 311
 - rtc_gpio_mode_t::RTC_GPIO_MODE_OUTPUT_ONLY (C++ enumerator), 311
 - rtc_gpio_pulldown_dis (C++ function), 308
 - rtc_gpio_pulldown_en (C++ function), 308
 - rtc_gpio_pullup_dis (C++ function), 308
 - rtc_gpio_pullup_en (C++ function), 308
 - rtc_gpio_set_direction (C++ function), 307
 - rtc_gpio_set_direction_in_sleep (C++ function), 307
 - rtc_gpio_set_drive_capability (C++ function), 308
 - rtc_gpio_set_level (C++ function), 307
 - rtc_gpio_wakeup_disable (C++ function), 309
 - rtc_gpio_wakeup_enable (C++ function), 309
 - rtc_gpio_number_get (C++ function), 306
- ## S
- sdmmc_can_discard (C++ function), 1159
 - sdmmc_can_trim (C++ function), 1159
 - sdmmc_card_init (C++ function), 1157
 - sdmmc_card_print_info (C++ function), 1158
 - sdmmc_card_t (C++ struct), 1167
 - sdmmc_card_t::cid (C++ member), 1168
 - sdmmc_card_t::csd (C++ member), 1168
 - sdmmc_card_t::ext_csd (C++ member), 1168
 - sdmmc_card_t::host (C++ member), 1167
 - sdmmc_card_t::is_ddr (C++ member), 1168
 - sdmmc_card_t::is_mem (C++ member), 1168
 - sdmmc_card_t::is_mmc (C++ member), 1168
 - sdmmc_card_t::is_sdio (C++ member), 1168
 - sdmmc_card_t::log_bus_width (C++ member), 1168
 - sdmmc_card_t::max_freq_khz (C++ member), 1168
 - sdmmc_card_t::num_io_functions (C++ member), 1168
 - sdmmc_card_t::ocr (C++ member), 1167
 - sdmmc_card_t::raw_cid (C++ member), 1168
 - sdmmc_card_t::rca (C++ member), 1168
 - sdmmc_card_t::real_freq_khz (C++ member), 1168
 - sdmmc_card_t::reserved (C++ member), 1168
 - sdmmc_card_t::scr (C++ member), 1168
 - sdmmc_card_t::ssr (C++ member), 1168
 - sdmmc_cid_t (C++ struct), 1163
 - sdmmc_cid_t::date (C++ member), 1164
 - sdmmc_cid_t::mfg_id (C++ member), 1163
 - sdmmc_cid_t::name (C++ member), 1163
 - sdmmc_cid_t::oem_id (C++ member), 1163
 - sdmmc_cid_t::revision (C++ member), 1163
 - sdmmc_cid_t::serial (C++ member), 1163
 - sdmmc_command_t (C++ struct), 1165
 - sdmmc_command_t::arg (C++ member), 1165
 - sdmmc_command_t::blklen (C++ member), 1166
 - sdmmc_command_t::buflen (C++ member), 1166
 - sdmmc_command_t::data (C++ member), 1166
 - sdmmc_command_t::datalen (C++ member), 1166
 - sdmmc_command_t::error (C++ member), 1166

- sdmmc_command_t::flags (C++ member), 1166
- sdmmc_command_t::opcode (C++ member), 1165
- sdmmc_command_t::response (C++ member), 1165
- sdmmc_command_t::timeout_ms (C++ member), 1166
- sdmmc_csd_t (C++ struct), 1163
- sdmmc_csd_t::capacity (C++ member), 1163
- sdmmc_csd_t::card_command_class (C++ member), 1163
- sdmmc_csd_t::csd_ver (C++ member), 1163
- sdmmc_csd_t::mmc_ver (C++ member), 1163
- sdmmc_csd_t::read_block_len (C++ member), 1163
- sdmmc_csd_t::sector_size (C++ member), 1163
- sdmmc_csd_t::tr_speed (C++ member), 1163
- sdmmc_delay_phase_t (C++ enum), 1169
- sdmmc_delay_phase_t::SDMMC_DELAY_PHASE_0 (C++ enumerator), 1169
- sdmmc_delay_phase_t::SDMMC_DELAY_PHASE_1 (C++ enumerator), 1169
- sdmmc_delay_phase_t::SDMMC_DELAY_PHASE_2 (C++ enumerator), 1170
- sdmmc_delay_phase_t::SDMMC_DELAY_PHASE_3 (C++ enumerator), 1170
- sdmmc_erase_arg_t (C++ enum), 1170
- sdmmc_erase_arg_t::SDMMC_DISCARD_ARG (C++ enumerator), 1170
- sdmmc_erase_arg_t::SDMMC_ERASE_ARG (C++ enumerator), 1170
- sdmmc_erase_sectors (C++ function), 1158
- sdmmc_ext_csd_t (C++ struct), 1165
- sdmmc_ext_csd_t::erase_mem_state (C++ member), 1165
- sdmmc_ext_csd_t::power_class (C++ member), 1165
- sdmmc_ext_csd_t::rev (C++ member), 1165
- sdmmc_ext_csd_t::sec_feature (C++ member), 1165
- SDMMC_FREQ_26M (C macro), 1169
- SDMMC_FREQ_52M (C macro), 1169
- SDMMC_FREQ_DEFAULT (C macro), 1169
- SDMMC_FREQ_HIGHSPEED (C macro), 1169
- SDMMC_FREQ_PROBING (C macro), 1169
- sdmmc_full_erase (C++ function), 1159
- sdmmc_get_status (C++ function), 1158
- sdmmc_host_deinit (C++ function), 566
- sdmmc_host_do_transaction (C++ function), 565
- SDMMC_HOST_FLAG_1BIT (C macro), 1169
- SDMMC_HOST_FLAG_4BIT (C macro), 1169
- SDMMC_HOST_FLAG_8BIT (C macro), 1169
- SDMMC_HOST_FLAG_DDR (C macro), 1169
- SDMMC_HOST_FLAG_DEINIT_ARG (C macro), 1169
- SDMMC_HOST_FLAG_SPI (C macro), 1169
- sdmmc_host_get_real_freq (C++ function), 566
- sdmmc_host_get_slot_width (C++ function), 565
- sdmmc_host_init (C++ function), 564
- sdmmc_host_init_slot (C++ function), 564
- sdmmc_host_io_int_enable (C++ function), 566
- sdmmc_host_io_int_wait (C++ function), 566
- sdmmc_host_set_bus_ddr_mode (C++ function), 565
- sdmmc_host_set_bus_width (C++ function), 564
- sdmmc_host_set_card_clk (C++ function), 565
- sdmmc_host_set_cclk_always_on (C++ function), 565
- sdmmc_host_set_input_delay (C++ function), 566
- sdmmc_host_t (C++ struct), 1166
- sdmmc_host_t::command_timeout_ms (C++ member), 1167
- sdmmc_host_t::deinit (C++ member), 1167
- sdmmc_host_t::deinit_p (C++ member), 1167
- sdmmc_host_t::do_transaction (C++ member), 1167
- sdmmc_host_t::flags (C++ member), 1166
- sdmmc_host_t::get_bus_width (C++ member), 1166
- sdmmc_host_t::get_real_freq (C++ member), 1167
- sdmmc_host_t::init (C++ member), 1166
- sdmmc_host_t::input_delay_phase (C++ member), 1167
- sdmmc_host_t::io_int_enable (C++ member), 1167
- sdmmc_host_t::io_int_wait (C++ member), 1167
- sdmmc_host_t::io_voltage (C++ member), 1166
- sdmmc_host_t::max_freq_khz (C++ member), 1166
- sdmmc_host_t::set_bus_ddr_mode (C++ member), 1167
- sdmmc_host_t::set_bus_width (C++ member), 1166
- sdmmc_host_t::set_card_clk (C++ member), 1167
- sdmmc_host_t::set_cclk_always_on (C++ member), 1167
- sdmmc_host_t::set_input_delay (C++ member), 1167
- sdmmc_host_t::slot (C++ member), 1166
- sdmmc_io_enable_int (C++ function), 1161
- sdmmc_io_get_cis_data (C++ function), 1162
- sdmmc_io_print_cis_info (C++ function), 1162
- sdmmc_io_read_blocks (C++ function), 1161
- sdmmc_io_read_byte (C++ function), 1159
- sdmmc_io_read_bytes (C++ function), 1160

- `sdmmc_io_wait_int` (C++ function), 1161
- `sdmmc_io_write_blocks` (C++ function), 1161
- `sdmmc_io_write_byte` (C++ function), 1160
- `sdmmc_io_write_bytes` (C++ function), 1160
- `sdmmc_mmc_can_sanitiz` (C++ function), 1159
- `sdmmc_mmc_sanitiz` (C++ function), 1159
- `sdmmc_read_sectors` (C++ function), 1158
- `sdmmc_response_t` (C++ type), 1169
- `sdmmc_scr_t` (C++ struct), 1164
- `sdmmc_scr_t::bus_width` (C++ member), 1164
- `sdmmc_scr_t::erase_mem_state` (C++ member), 1164
- `sdmmc_scr_t::reserved` (C++ member), 1164
- `sdmmc_scr_t::rsvd_mnf` (C++ member), 1164
- `sdmmc_scr_t::sd_spec` (C++ member), 1164
- `sdmmc_slot_config_t` (C++ struct), 567
- `sdmmc_slot_config_t::cd` (C++ member), 568
- `sdmmc_slot_config_t::clk` (C++ member), 567
- `sdmmc_slot_config_t::cmd` (C++ member), 567
- `sdmmc_slot_config_t::d0` (C++ member), 567
- `sdmmc_slot_config_t::d1` (C++ member), 567
- `sdmmc_slot_config_t::d2` (C++ member), 567
- `sdmmc_slot_config_t::d3` (C++ member), 567
- `sdmmc_slot_config_t::d4` (C++ member), 567
- `sdmmc_slot_config_t::d5` (C++ member), 567
- `sdmmc_slot_config_t::d6` (C++ member), 567
- `sdmmc_slot_config_t::d7` (C++ member), 568
- `sdmmc_slot_config_t::flags` (C++ member), 568
- `sdmmc_slot_config_t::gpio_cd` (C++ member), 568
- `sdmmc_slot_config_t::gpio_wp` (C++ member), 568
- `sdmmc_slot_config_t::width` (C++ member), 568
- `sdmmc_slot_config_t::wp` (C++ member), 568
- `SDMMC_SLOT_FLAG_INTERNAL_PULLUP` (C macro), 568
- `SDMMC_SLOT_FLAG_WP_ACTIVE_HIGH` (C macro), 568
- `sdmmc_ssr_t` (C++ struct), 1164
- `sdmmc_ssr_t::alloc_unit_kb` (C++ member), 1164
- `sdmmc_ssr_t::cur_bus_width` (C++ member), 1164
- `sdmmc_ssr_t::discard_support` (C++ member), 1164
- `sdmmc_ssr_t::erase_offset` (C++ member), 1165
- `sdmmc_ssr_t::erase_size_au` (C++ member), 1164
- `sdmmc_ssr_t::erase_timeout` (C++ member), 1164
- `sdmmc_ssr_t::fule_support` (C++ member), 1164
- `sdmmc_ssr_t::reserved` (C++ member), 1165
- `sdmmc_switch_func_rsp_t` (C++ struct), 1165
- `sdmmc_switch_func_rsp_t::data` (C++ member), 1165
- `sdmmc_write_sectors` (C++ function), 1158
- `SDSPI_DEFAULT_DMA` (C macro), 573
- `SDSPI_DEFAULT_HOST` (C macro), 573
- `sdspi_dev_handle_t` (C++ type), 574
- `SDSPI_DEVICE_CONFIG_DEFAULT` (C macro), 574
- `sdspi_device_config_t` (C++ struct), 573
- `sdspi_device_config_t::gpio_cd` (C++ member), 573
- `sdspi_device_config_t::gpio_cs` (C++ member), 573
- `sdspi_device_config_t::gpio_int` (C++ member), 573
- `sdspi_device_config_t::gpio_wp` (C++ member), 573
- `sdspi_device_config_t::gpio_wp_polarity` (C++ member), 573
- `sdspi_device_config_t::host_id` (C++ member), 573
- `SDSPI_HOST_DEFAULT` (C macro), 573
- `sdspi_host_deinit` (C++ function), 572
- `sdspi_host_do_transaction` (C++ function), 571
- `sdspi_host_get_real_freq` (C++ function), 572
- `sdspi_host_init` (C++ function), 571
- `sdspi_host_init_device` (C++ function), 571
- `sdspi_host_io_int_enable` (C++ function), 572
- `sdspi_host_io_int_wait` (C++ function), 572
- `sdspi_host_remove_device` (C++ function), 571
- `sdspi_host_set_card_clk` (C++ function), 572
- `SDSPI_IO_ACTIVE_LOW` (C macro), 574
- `SDSPI_SLOT_NO_CD` (C macro), 573
- `SDSPI_SLOT_NO_CS` (C macro), 573
- `SDSPI_SLOT_NO_INT` (C macro), 574
- `SDSPI_SLOT_NO_WP` (C macro), 573
- `SemaphoreHandle_t` (C++ type), 1359
- `semBINARY_SEMAPHORE_QUEUE_LENGTH` (C macro), 1344
- `semGIVE_BLOCK_TIME` (C macro), 1345
- `semSEMAPHORE_QUEUE_ITEM_LENGTH` (C macro), 1344
- `shared_stack_function` (C++ type), 1219
- `shutdown_handler_t` (C++ type), 1502
- `slave_transaction_cb_t` (C++ type), 635
- `sntp_get_sync_interval` (C++ function), 1574
- `sntp_get_sync_mode` (C++ function), 1573
- `sntp_get_sync_status` (C++ function), 1573
- `sntp_getoperatingmode` (C++ function), 1575
- `sntp_getreachability` (C++ function), 1575
- `sntp_getserver` (C++ function), 1575
- `sntp_getservername` (C++ function), 1575
- `sntp_init` (C++ function), 1575

- SNTP_OPMODE_POLL (*C macro*), 1575
 sntp_restart (*C++ function*), 1574
 sntp_servermode_dhcp (*C++ function*), 1575
 sntp_set_sync_interval (*C++ function*), 1574
 sntp_set_sync_mode (*C++ function*), 1573
 sntp_set_sync_status (*C++ function*), 1573
 sntp_set_time_sync_notification_cb (*C++ function*), 1574
 sntp_setoperatingmode (*C++ function*), 1575
 sntp_setservername (*C++ function*), 1575
 sntp_sync_mode_t (*C++ enum*), 1576
 sntp_sync_mode_t::SNTP_SYNC_MODE_IMMED (*C++ enumerator*), 1576
 sntp_sync_mode_t::SNTP_SYNC_MODE_SMOOTH (*C++ enumerator*), 1576
 sntp_sync_status_t (*C++ enum*), 1576
 sntp_sync_status_t::SNTP_SYNC_STATUS_COMPLETED (*C++ enumerator*), 1576
 sntp_sync_status_t::SNTP_SYNC_STATUS_IN_PROGRESS (*C++ enumerator*), 1576
 sntp_sync_status_t::SNTP_SYNC_STATUS_RESET (*C++ enumerator*), 1576
 sntp_sync_time (*C++ function*), 1573
 sntp_sync_time_cb_t (*C++ type*), 1576
 SOC_ADC_ATTEN_NUM (*C macro*), 1557
 SOC_ADC_CALIBRATION_V1_SUPPORTED (*C macro*), 1557
 SOC_ADC_CHANNEL_NUM (*C macro*), 1556
 SOC_ADC_DIG_SUPPORTED_UNIT (*C macro*), 1556
 SOC_ADC_DIGI_CONTROLLER_NUM (*C macro*), 1557
 SOC_ADC_DIGI_DATA_BYTES_PER_CONV (*C macro*), 1557
 SOC_ADC_DIGI_IIR_FILTER_NUM (*C macro*), 1557
 SOC_ADC_DIGI_MAX_BITWIDTH (*C macro*), 1557
 SOC_ADC_DIGI_MIN_BITWIDTH (*C macro*), 1557
 SOC_ADC_DIGI_MONITOR_NUM (*C macro*), 1557
 SOC_ADC_DIGI_RESULT_BYTES (*C macro*), 1557
 SOC_ADC_MAX_CHANNEL_NUM (*C macro*), 1557
 SOC_ADC_PATT_LEN_MAX (*C macro*), 1557
 SOC_ADC_PERIPH_NUM (*C macro*), 1556
 SOC_ADC_RTC_MAX_BITWIDTH (*C macro*), 1557
 SOC_ADC_RTC_MIN_BITWIDTH (*C macro*), 1557
 SOC_ADC_SAMPLE_FREQ_THRES_HIGH (*C macro*), 1557
 SOC_ADC_SAMPLE_FREQ_THRES_LOW (*C macro*), 1557
 SOC_AES_GDMA (*C macro*), 1556
 SOC_AES_SUPPORT_AES_128 (*C macro*), 1556
 SOC_AES_SUPPORT_AES_256 (*C macro*), 1556
 SOC_AES_SUPPORT_DMA (*C macro*), 1556
 SOC_AHB_GDMA_SUPPORTED (*C macro*), 1555
 SOC_AHB_GDMA_VERSION (*C macro*), 1558
 SOC_ANA_CMPR_CAN_DISTINGUISH_EDGE (*C macro*), 1560
 SOC_ANA_CMPR_CLKS (*C macro*), 267
 SOC_ANA_CMPR_NUM (*C macro*), 1560
 SOC_ANA_CMPR_SUPPORT_ETM (*C macro*), 1560
 SOC_ANA_CMPR_SUPPORTED (*C macro*), 1555
 SOC_APB_BACKUP_DMA (*C macro*), 1557
 SOC_ASYNC_MEMCPY_SUPPORTED (*C macro*), 1555
 SOC_AXI_GDMA_SUPPORT_PSRAM (*C macro*), 1559
 SOC_AXI_GDMA_SUPPORTED (*C macro*), 1555
 SOC_BRANCH_PREDICTOR_SUPPORTED (*C macro*), 1558
 SOC_BROWNOUT_RESET_SUPPORTED (*C macro*), 1557
 SOC_CACHE_FREEZE_SUPPORTED (*C macro*), 1557
 SOC_CACHE_INTERNAL_MEM_VIA_L1CACHE (*C macro*), 1557
 SOC_CACHE_WRITEBACK_SUPPORTED (*C macro*), 1557
 SOC_CLK_APLL_SUPPORTED (*C macro*), 1569
 SOC_CLK_FRC_OSC_SLOW_FREQ_APPROX (*C macro*), 266
 SOC_CLK_FRC_SLOW_SUPPORTED (*C macro*), 1569
 SOC_CLK_RC32K_FREQ_APPROX (*C macro*), 266
 SOC_CLK_RC32K_SUPPORTED (*C macro*), 1569
 SOC_CLK_RC_FAST_FREQ_APPROX (*C macro*), 266
 SOC_CLK_RC_FAST_SUPPORT_CALIBRATION (*C macro*), 1569
 SOC_CLK_RC_SLOW_FREQ_APPROX (*C macro*), 266
 SOC_CLK_XTAL32K_FREQ_APPROX (*C macro*), 266
 SOC_CLK_XTAL32K_SUPPORTED (*C macro*), 1569
 SOC_COEX_HW_PTII (*C macro*), 1568
 SOC_CPU_BREAKPOINTS_NUM (*C macro*), 1558
 soc_cpu_clk_src_t (*C++ enum*), 268
 soc_cpu_clk_src_t::SOC_CPU_CLK_SRC_INVALID (*C++ enumerator*), 268
 soc_cpu_clk_src_t::SOC_CPU_CLK_SRC_PLL (*C++ enumerator*), 268
 soc_cpu_clk_src_t::SOC_CPU_CLK_SRC_RC_FAST (*C++ enumerator*), 268
 soc_cpu_clk_src_t::SOC_CPU_CLK_SRC_XTAL (*C++ enumerator*), 268
 SOC_CPU_COPROC_NUM (*C macro*), 1558
 SOC_CPU_CORES_NUM (*C macro*), 1558
 SOC_CPU_HAS_FLEXIBLE_INTC (*C macro*), 1558
 SOC_CPU_HAS_FPU (*C macro*), 1558
 SOC_CPU_HAS_FPU_EXT_ILL_BUG (*C macro*), 1558
 SOC_CPU_HAS_PMA (*C macro*), 1558
 SOC_CPU_IDRAM_SPLIT_USING_PMP (*C macro*), 1558
 SOC_CPU_INTR_NUM (*C macro*), 1558
 SOC_CPU_WATCHPOINT_MAX_REGION_SIZE (*C macro*), 1558
 SOC_CPU_WATCHPOINTS_NUM (*C macro*), 1558
 SOC_DEDIC_GPIO_IN_CHANNELS_NUM (*C macro*), 1560
 SOC_DEDIC_GPIO_OUT_CHANNELS_NUM (*C macro*), 1560

- SOC_DEDIC_PERIPH_ALWAYS_ENABLE (C macro), 1560
- SOC_DIG_SIGN_SUPPORTED (C macro), 1556
- SOC_DS_KEY_CHECK_MAX_WAIT_US (C macro), 1558
- SOC_DS_KEY_PARAM_MD_IV_LENGTH (C macro), 1558
- SOC_DS_SIGNATURE_MAX_BIT_LEN (C macro), 1558
- SOC_ECC_EXTENDED_MODES_SUPPORTED (C macro), 1556
- SOC_ECC_SUPPORTED (C macro), 1556
- SOC_ECDSA_SUPPORT_EXPORT_PUBKEY (C macro), 1565
- SOC_ECDSA_SUPPORTED (C macro), 1556
- SOC_ECDSA_USES_MPI (C macro), 1565
- SOC_EFUSE_DIS_DIRECT_BOOT (C macro), 1567
- SOC_EFUSE_DIS_DOWNLOAD_ICACHE (C macro), 1567
- SOC_EFUSE_DIS_PAD_JTAG (C macro), 1567
- SOC_EFUSE_DIS_USB_JTAG (C macro), 1567
- SOC_EFUSE_KEY_PURPOSE_FIELD (C macro), 1555
- SOC_EFUSE_REVOKE_BOOT_KEY_DIGESTS (C macro), 1568
- SOC_EFUSE_SECURE_BOOT_KEY_DIGESTS (C macro), 1567
- SOC_EFUSE_SOFT_DIS_JTAG (C macro), 1567
- SOC_EFUSE_SUPPORTED (C macro), 1555
- SOC_ETM_CHANNELS_PER_GROUP (C macro), 1559
- SOC_ETM_GROUPS (C macro), 1559
- SOC_ETM_SUPPORTED (C macro), 1555
- SOC_FLASH_CLKS (C macro), 267
- SOC_FLASH_ENC_SUPPORTED (C macro), 1556
- SOC_FLASH_ENCRYPTED_XTS_AES_BLOCK_MAX (C macro), 1568
- SOC_FLASH_ENCRYPTION_XTS_AES (C macro), 1568
- SOC_FLASH_ENCRYPTION_XTS_AES_128 (C macro), 1568
- SOC_GDMA_NUM_GROUPS_MAX (C macro), 1559
- SOC_GDMA_PAIRS_PER_GROUP_MAX (C macro), 1559
- SOC_GDMA_SUPPORT_CRC (C macro), 1558
- SOC_GDMA_SUPPORTED (C macro), 1555
- SOC_GPIO_DEEP_SLEEP_WAKE_VALID_GPIO_MASK (C macro), 1559
- SOC_GPIO_IN_RANGE_MAX (C macro), 1559
- SOC_GPIO_OUT_RANGE_MAX (C macro), 1559
- SOC_GPIO_PIN_COUNT (C macro), 1559
- SOC_GPIO_PORT (C macro), 1559
- SOC_GPIO_SUPPORT_DEEPSLEEP_WAKEUP (C macro), 1559
- SOC_GPIO_SUPPORT_ETM (C macro), 1559
- SOC_GPIO_SUPPORT_FORCE_HOLD (C macro), 1559
- SOC_GPIO_SUPPORT_HOLD_SINGLE_IO_IN_DSLS (C macro), 1559
- SOC_GPIO_SUPPORT_PIN_HYS_FILTER (C macro), 1559
- SOC_GPIO_SUPPORT_RTC_INDEPENDENT (C macro), 1559
- SOC_GPIO_VALID_DIGITAL_IO_PAD_MASK (C macro), 1559
- SOC_GPIO_VALID_GPIO_MASK (C macro), 1559
- SOC_GPIO_VALID_OUTPUT_GPIO_MASK (C macro), 1559
- SOC_GPSPI_SUPPORTED (C macro), 1555
- SOC_GPTIMER_CLKS (C macro), 266
- SOC_GPTIMER_SUPPORTED (C macro), 1555
- SOC_HMAC_SUPPORTED (C macro), 1556
- SOC_HP_CPU_HAS_MULTIPLE_CORES (C macro), 1558
- SOC_I2C_CLKS (C macro), 267
- SOC_I2C_CMD_REG_NUM (C macro), 1560
- SOC_I2C_FIFO_LEN (C macro), 1560
- SOC_I2C_NUM (C macro), 1560
- SOC_I2C_SLAVE_SUPPORT_BROADCAST (C macro), 1560
- SOC_I2C_SLAVE_SUPPORT_I2CRAM_ACCESS (C macro), 1560
- SOC_I2C_SLAVE_SUPPORT_SLAVE_UNMATCH (C macro), 1560
- SOC_I2C_SUPPORT_10BIT_ADDR (C macro), 1560
- SOC_I2C_SUPPORT_HW_CLR_BUS (C macro), 1560
- SOC_I2C_SUPPORT_HW_FSM_RST (C macro), 1560
- SOC_I2C_SUPPORT_RTC (C macro), 1560
- SOC_I2C_SUPPORT_SLAVE (C macro), 1560
- SOC_I2C_SUPPORT_XTAL (C macro), 1560
- SOC_I2C_SUPPORTED (C macro), 1555
- SOC_I2S_CLKS (C macro), 267
- SOC_I2S_HW_VERSION_2 (C macro), 1561
- SOC_I2S_NUM (C macro), 1560
- SOC_I2S_PDM_MAX_RX_LINES (C macro), 1561
- SOC_I2S_PDM_MAX_TX_LINES (C macro), 1561
- SOC_I2S_SUPPORTED (C macro), 1555
- SOC_I2S_SUPPORTS_APLL (C macro), 1561
- SOC_I2S_SUPPORTS_PCM (C macro), 1561
- SOC_I2S_SUPPORTS_PDM (C macro), 1561
- SOC_I2S_SUPPORTS_PDM_RX (C macro), 1561
- SOC_I2S_SUPPORTS_PDM_RX_HP_FILTER (C macro), 1561
- SOC_I2S_SUPPORTS_PDM_TX (C macro), 1561
- SOC_I2S_SUPPORTS_TDM (C macro), 1561
- SOC_I2S_SUPPORTS_TX_SYNC_CNT (C macro), 1561
- SOC_I2S_SUPPORTS_XTAL (C macro), 1561
- SOC_I2S_TDM_FULL_DATA_WIDTH (C macro), 1561
- SOC_INT_CLIC_SUPPORTED (C macro), 1558
- SOC_INT_HW_NESTED_SUPPORTED (C macro), 1558
- SOC_INT_PLIC_SUPPORTED (C macro), 1558
- SOC_LEDC_CHANNEL_NUM (C macro), 1561
- SOC_LEDC_CLKS (C macro), 267

- SOC_LEDC_FADE_PARAMS_BIT_WIDTH (C macro), 1561
- SOC_LEDC_GAMMA_CURVE_FADE_RANGE_MAX (C macro), 1561
- SOC_LEDC_GAMMA_CURVE_FADE_SUPPORTED (C macro), 1561
- SOC_LEDC_SUPPORT_FADE_STOP (C macro), 1561
- SOC_LEDC_SUPPORT_PLL_DIV_CLOCK (C macro), 1561
- SOC_LEDC_SUPPORT_XTAL_CLOCK (C macro), 1561
- SOC_LEDC_SUPPORTED (C macro), 1555
- SOC_LEDC_TIMER_BIT_WIDTH (C macro), 1561
- SOC_LP_GPIO_MATRIX_SUPPORTED (C macro), 1556
- SOC_LP_PERIPHERALS_SUPPORTED (C macro), 1556
- soc_lp_pll_clk_src_t (C++ enum), 269
- soc_lp_pll_clk_src_t::SOC_LP_PLL_CLK_SRC_INVALID (C++ enumerator), 270
- soc_lp_pll_clk_src_t::SOC_LP_PLL_CLK_SRC_RC32K (C++ enumerator), 270
- soc_lp_pll_clk_src_t::SOC_LP_PLL_CLK_SRC_XTAL32K (C++ enumerator), 270
- SOC_LP_TIMER_BIT_WIDTH_HI (C macro), 1567
- SOC_LP_TIMER_BIT_WIDTH_LO (C macro), 1567
- SOC_LP_UART_CLKS (C macro), 267
- SOC_LP_UART_FIFO_LEN (C macro), 1568
- SOC_MCPWM_CAPTURE_CHANNELS_PER_TIMER (C macro), 1564
- SOC_MCPWM_CAPTURE_CLK_FROM_GROUP (C macro), 1564
- SOC_MCPWM_CAPTURE_CLKS (C macro), 267
- SOC_MCPWM_CAPTURE_TIMERS_PER_GROUP (C macro), 1564
- SOC_MCPWM_CARRIER_CLKS (C macro), 267
- SOC_MCPWM_COMPARATORS_PER_OPERATOR (C macro), 1563
- SOC_MCPWM_EVENT_COMPARATORS_PER_OPERATOR (C macro), 1563
- SOC_MCPWM_GENERATORS_PER_OPERATOR (C macro), 1563
- SOC_MCPWM_GPIO_FAULTS_PER_GROUP (C macro), 1563
- SOC_MCPWM_GPIO_SYNCHROS_PER_GROUP (C macro), 1564
- SOC_MCPWM_GROUPS (C macro), 1563
- SOC_MCPWM_OPERATORS_PER_GROUP (C macro), 1563
- SOC_MCPWM_SUPPORT_ETM (C macro), 1564
- SOC_MCPWM_SUPPORT_EVENT_COMPARATOR (C macro), 1564
- SOC_MCPWM_SUPPORTED (C macro), 1555
- SOC_MCPWM_SWSYNC_CAN_PROPAGATE (C macro), 1564
- SOC_MCPWM_TIMER_CLKS (C macro), 267
- SOC_MCPWM_TIMERS_PER_GROUP (C macro), 1563
- SOC_MCPWM_TRIGGERS_PER_OPERATOR (C macro), 1563
- SOC_MEM_TCM_SUPPORTED (C macro), 1570
- SOC_MEMSPI_IS_INDEPENDENT (C macro), 1566
- SOC_MEMSPI_SRC_FREQ_20M_SUPPORTED (C macro), 1566
- SOC_MEMSPI_SRC_FREQ_40M_SUPPORTED (C macro), 1566
- SOC_MEMSPI_SRC_FREQ_80M_SUPPORTED (C macro), 1566
- SOC_MMU_DI_VADDR_SHARED (C macro), 1562
- SOC_MMU_LINEAR_ADDRESS_REGION_NUM (C macro), 1562
- SOC_MMU_PAGE_SIZE_CONFIGURABLE (C macro), 1561
- SOC_MMU_PERIPH_NUM (C macro), 1562
- SOC_MODEM_CLOCK_IS_INDEPENDENT (C macro), 1569
- soc_module_clk_t (C++ enum), 270
- soc_module_clk_t::SOC_MOD_CLK_APLL (C++ enumerator), 271
- soc_module_clk_t::SOC_MOD_CLK_CPLL (C++ enumerator), 270
- soc_module_clk_t::SOC_MOD_CLK_CPU (C++ enumerator), 270
- soc_module_clk_t::SOC_MOD_CLK_INVALID (C++ enumerator), 271
- soc_module_clk_t::SOC_MOD_CLK_LP_PLL (C++ enumerator), 271
- soc_module_clk_t::SOC_MOD_CLK_MPLL (C++ enumerator), 271
- soc_module_clk_t::SOC_MOD_CLK_PLL_F160M (C++ enumerator), 270
- soc_module_clk_t::SOC_MOD_CLK_PLL_F200M (C++ enumerator), 270
- soc_module_clk_t::SOC_MOD_CLK_PLL_F240M (C++ enumerator), 270
- soc_module_clk_t::SOC_MOD_CLK_PLL_F80M (C++ enumerator), 270
- soc_module_clk_t::SOC_MOD_CLK_RC_FAST (C++ enumerator), 271
- soc_module_clk_t::SOC_MOD_CLK_RTC_FAST (C++ enumerator), 270
- soc_module_clk_t::SOC_MOD_CLK_RTC_SLOW (C++ enumerator), 270
- soc_module_clk_t::SOC_MOD_CLK_SPLL (C++ enumerator), 270
- soc_module_clk_t::SOC_MOD_CLK_XTAL (C++ enumerator), 271
- soc_module_clk_t::SOC_MOD_CLK_XTAL32K (C++ enumerator), 271
- soc_module_clk_t::SOC_MOD_CLK_XTAL_D2 (C++ enumerator), 271
- SOC_MPI_MEM_BLOCKS_NUM (C macro), 1564
- SOC_MPI_OPERATIONS_NUM (C macro), 1564
- SOC_MPI_SUPPORTED (C macro), 1556
- SOC_MPU_CONFIGURABLE_REGIONS_SUPPORTED (C macro), 1562
- SOC_MPU_MIN_REGION_SIZE (C macro), 1562

- SOC_MPU_REGION_RO_SUPPORTED (C macro), 1562
- SOC_MPU_REGION_WO_SUPPORTED (C macro), 1562
- SOC_MPU_REGIONS_MAX_NUM (C macro), 1562
- SOC_MWDT_CLKS (C macro), 267
- SOC_MWDT_SUPPORT_XTAL (C macro), 1567
- SOC_PARLIO_CLKS (C macro), 267
- SOC_PARLIO_GROUPS (C macro), 1564
- SOC_PARLIO_RX_UNIT_MAX_DATA_WIDTH (C macro), 1564
- SOC_PARLIO_RX_UNITS_PER_GROUP (C macro), 1564
- SOC_PARLIO_SUPPORTED (C macro), 1555
- SOC_PARLIO_TX_SIZE_BY_DMA (C macro), 1564
- SOC_PARLIO_TX_UNIT_MAX_DATA_WIDTH (C macro), 1564
- SOC_PARLIO_TX_UNITS_PER_GROUP (C macro), 1564
- SOC_PCNT_CHANNELS_PER_UNIT (C macro), 1562
- SOC_PCNT_GROUPS (C macro), 1562
- SOC_PCNT_SUPPORT_CLEAR_SIGNAL (C macro), 1562
- SOC_PCNT_SUPPORT_RUNTIME_THRES_UPDATE (C macro), 1562
- SOC_PCNT_SUPPORTED (C macro), 1555
- SOC_PCNT_THRES_POINT_PER_UNIT (C macro), 1562
- SOC_PCNT_UNITS_PER_GROUP (C macro), 1562
- soc_periph_ana_cmpr_clk_src_t (C++ enum), 275
- soc_periph_ana_cmpr_clk_src_t::ANA_CMPR_CLK_SRC_DEFAULT (C++ enumerator), 276
- soc_periph_ana_cmpr_clk_src_t::ANA_CMPR_CLK_SRC_FAST (C++ enumerator), 276
- soc_periph_ana_cmpr_clk_src_t::ANA_CMPR_CLK_SRC_XTAL (C++ enumerator), 276
- SOC_PERIPH_CLK_CTRL_SHARED (C macro), 1569
- soc_periph_flash_clk_src_t (C++ enum), 275
- soc_periph_flash_clk_src_t::FLASH_CLK_SRC_DEFAULT (C++ enumerator), 275
- soc_periph_flash_clk_src_t::FLASH_CLK_SRC_FAST (C++ enumerator), 275
- soc_periph_flash_clk_src_t::FLASH_CLK_SRC_XTAL (C++ enumerator), 275
- soc_periph_gptimer_clk_src_t (C++ enum), 271
- soc_periph_gptimer_clk_src_t::GPTIMER_CLK_SRC_DEFAULT (C++ enumerator), 272
- soc_periph_gptimer_clk_src_t::GPTIMER_CLK_SRC_FAST (C++ enumerator), 271
- soc_periph_gptimer_clk_src_t::GPTIMER_CLK_SRC_XTAL (C++ enumerator), 271
- soc_periph_i2c_clk_src_t (C++ enum), 274
- soc_periph_i2c_clk_src_t::I2C_CLK_SRC_DEFAULT (C++ enumerator), 274
- soc_periph_i2c_clk_src_t::I2C_CLK_SRC_RC_FAST (C++ enumerator), 274
- soc_periph_i2c_clk_src_t::I2C_CLK_SRC_XTAL (C++ enumerator), 274
- soc_periph_i2s_clk_src_t (C++ enum), 274
- soc_periph_i2s_clk_src_t::I2S_CLK_SRC_APLL (C++ enumerator), 274
- soc_periph_i2s_clk_src_t::I2S_CLK_SRC_DEFAULT (C++ enumerator), 274
- soc_periph_i2s_clk_src_t::I2S_CLK_SRC_EXTERNAL (C++ enumerator), 274
- soc_periph_i2s_clk_src_t::I2S_CLK_SRC_XTAL (C++ enumerator), 274
- soc_periph_ledc_clk_src_legacy_t (C++ enum), 276
- soc_periph_ledc_clk_src_legacy_t::LEDC_AUTO_CLK (C++ enumerator), 276
- soc_periph_ledc_clk_src_legacy_t::LEDC_USE_PLL_DIV (C++ enumerator), 276
- soc_periph_ledc_clk_src_legacy_t::LEDC_USE_RC_FAST (C++ enumerator), 276
- soc_periph_ledc_clk_src_legacy_t::LEDC_USE_XTAL_CLK (C++ enumerator), 276
- soc_periph_lp_uart_clk_src_t (C++ enum), 273
- soc_periph_lp_uart_clk_src_t::LP_UART_SCLK_DEFAULT (C++ enumerator), 273
- soc_periph_lp_uart_clk_src_t::LP_UART_SCLK_LP_FAST (C++ enumerator), 273
- soc_periph_lp_uart_clk_src_t::LP_UART_SCLK_XTAL_DIV (C++ enumerator), 273
- soc_periph_mcpwm_capture_clk_src_t (C++ enum), 274
- soc_periph_mcpwm_capture_clk_src_t::MCPWM_CAPTURE_DEFAULT (C++ enumerator), 274
- soc_periph_mcpwm_capture_clk_src_t::MCPWM_CAPTURE_FAST (C++ enumerator), 274
- soc_periph_mcpwm_capture_clk_src_t::MCPWM_CAPTURE_XTAL (C++ enumerator), 274
- soc_periph_mcpwm_carrier_clk_src_t (C++ enum), 274
- soc_periph_mcpwm_carrier_clk_src_t::MCPWM_CARRIER_DEFAULT (C++ enumerator), 274
- soc_periph_mcpwm_carrier_clk_src_t::MCPWM_CARRIER_FAST (C++ enumerator), 274
- soc_periph_mcpwm_carrier_clk_src_t::MCPWM_CARRIER_XTAL (C++ enumerator), 274
- soc_periph_mcpwm_timer_clk_src_t (C++ enum), 273
- soc_periph_mcpwm_timer_clk_src_t::MCPWM_TIMER_CLK_DEFAULT (C++ enumerator), 273
- soc_periph_mcpwm_timer_clk_src_t::MCPWM_TIMER_CLK_FAST (C++ enumerator), 273
- soc_periph_mcpwm_timer_clk_src_t::MCPWM_TIMER_CLK_XTAL (C++ enumerator), 273

Espressif Systems

1990

Release v5.2.7

- 1562
- SOC_RMT_SUPPORT_TX_CARRIER_DATA_ONLY (C macro), 1563
- SOC_RMT_SUPPORT_TX_LOOP_AUTO_STOP (C macro), 1563
- SOC_RMT_SUPPORT_TX_LOOP_COUNT (C macro), 1563
- SOC_RMT_SUPPORT_TX_SYNCHRO (C macro), 1563
- SOC_RMT_SUPPORT_XTAL (C macro), 1563
- SOC_RMT_SUPPORTED (C macro), 1555
- SOC_RMT_TX_CANDIDATES_PER_GROUP (C macro), 1562
- soc_root_clk_t (C++ enum), 268
- soc_root_clk_t::SOC_ROOT_CLK_EXT_OSC_SLOW (C++ enumerator), 268
- soc_root_clk_t::SOC_ROOT_CLK_EXT_XTAL (C++ enumerator), 268
- soc_root_clk_t::SOC_ROOT_CLK_EXT_XTAL32K (C++ enumerator), 268
- soc_root_clk_t::SOC_ROOT_CLK_INT_RC32K (C++ enumerator), 268
- soc_root_clk_t::SOC_ROOT_CLK_INT_RC_FAST (C++ enumerator), 268
- soc_root_clk_t::SOC_ROOT_CLK_INT_RC_SLOW (C++ enumerator), 268
- SOC_RSA_MAX_BIT_LEN (C macro), 1564
- soc_rtc_fast_clk_src_t (C++ enum), 269
- soc_rtc_fast_clk_src_t::SOC_RTC_FAST_CLK_SRC_FAST (C++ enumerator), 269
- soc_rtc_fast_clk_src_t::SOC_RTC_FAST_CLK_SRC_MEM (C++ enumerator), 269
- soc_rtc_fast_clk_src_t::SOC_RTC_FAST_CLK_SRC_ROM (C++ enumerator), 269
- soc_rtc_fast_clk_src_t::SOC_RTC_FAST_CLK_SRC_SLOW (C++ enumerator), 269
- soc_rtc_fast_clk_src_t::SOC_RTC_FAST_CLK_SRC_SLOW_DIV (C++ enumerator), 269
- SOC_RTC_FAST_MEM_SUPPORTED (C macro), 1555
- SOC_RTC_MEM_SUPPORTED (C macro), 1555
- soc_rtc_slow_clk_src_t (C++ enum), 268
- soc_rtc_slow_clk_src_t::SOC_RTC_SLOW_CLK_SRC_FAST (C++ enumerator), 269
- soc_rtc_slow_clk_src_t::SOC_RTC_SLOW_CLK_SRC_SLOW (C++ enumerator), 269
- soc_rtc_slow_clk_src_t::SOC_RTC_SLOW_CLK_SRC_SLOW_DIV (C++ enumerator), 269
- SOC_RTCIO_EDGE_WAKE_SUPPORTED (C macro), 1560
- SOC_RTCIO_HOLD_SUPPORTED (C macro), 1559
- SOC_RTCIO_INPUT_OUTPUT_SUPPORTED (C macro), 1559
- SOC_RTCIO_PIN_COUNT (C macro), 1559
- SOC_RTCIO_WAKE_SUPPORTED (C macro), 1559
- SOC_SDM_CHANNELS_PER_GROUP (C macro), 1565
- SOC_SDM_CLK_SUPPORT_PLL_F80M (C macro), 1565
- SOC_SDM_CLK_SUPPORT_XTAL (C macro), 1565
- SOC_SDM_GROUPS (C macro), 1565
- SOC_SDMMC_CLKS (C macro), 268
- SOC_SDMMC_DELAY_PHASE_NUM (C macro), 1565
- SOC_SDMMC_HOST_SUPPORTED (C macro), 1556
- SOC_SDMMC_NUM_SLOTS (C macro), 1565
- SOC_SDMMC_USE_GPIO_MATRIX (C macro), 1565
- SOC_SDMMC_USE_IOMUX (C macro), 1564
- SOC_SECURE_BOOT_SUPPORTED (C macro), 1556
- SOC_SECURE_BOOT_V2_ECC (C macro), 1567
- SOC_SECURE_BOOT_V2_RSA (C macro), 1567
- SOC_SHA_DMA_MAX_BUFFER_SIZE (C macro), 1565
- SOC_SHA_GDMA (C macro), 1565
- SOC_SHA_SUPPORT_DMA (C macro), 1565
- SOC_SHA_SUPPORT_RESUME (C macro), 1565
- SOC_SHA_SUPPORT_SHA1 (C macro), 1565
- SOC_SHA_SUPPORT_SHA224 (C macro), 1565
- SOC_SHA_SUPPORT_SHA256 (C macro), 1565
- SOC_SHARED_IDCACHE_SUPPORTED (C macro), 1557
- SOC_SPI_CLKS (C macro), 267
- SOC_SPI_FLASH_SUPPORTED (C macro), 1556
- SOC_SPI_MAX_CS_NUM (C macro), 1565
- SOC_SPI_MAX_PRE_DIVIDER (C macro), 1566
- SOC_SPI_MAXIMUM_BUFFER_SIZE (C macro), 1565
- SOC_SPI_MEM_SUPPORT_AUTO_RESUME (C macro), 1566
- SOC_SPI_MEM_SUPPORT_AUTO_WAIT_IDLE (C macro), 1566
- SOC_SPI_MEM_SUPPORT_CHECK_SUS (C macro), 1566
- SOC_SPI_MEM_SUPPORT_IDLE_INTR (C macro), 1566
- SOC_SPI_MEM_SUPPORT_SW_SUSPEND (C macro), 1566
- SOC_SPI_MEM_SUPPORT_WRAP (C macro), 1566
- SOC_SPI_PERIPH_CS_NUM (C macro), 1565
- SOC_SPI_PERIPH_NUM (C macro), 1565
- SOC_SPI_PERIPH_SUPPORT_MULTILINE_MODE (C macro), 1566
- SOC_SPI_ROM_SUPPORT_SEG_TRANS (C macro), 1565
- SOC_SPI_ROM_SUPPORT_CD_SIG (C macro), 1566
- SOC_SPI_SUPPORT_CLK_XTAL (C macro), 1566
- SOC_SPI_SUPPORT_DDRCLK (C macro), 1565
- SOC_SPI_SUPPORT_OCT (C macro), 1566
- SOC_SPIRAM_SUPPORTED (C macro), 1556
- SOC_SUPPORT_SECURE_BOOT_REVOKE_KEY (C macro), 1568
- SOC_SUPPORTS_SECURE_DL_MODE (C macro), 1555
- SOC_SYSTIMER_ALARM_MISS_COMPENSATE (C macro), 1566
- SOC_SYSTIMER_ALARM_NUM (C macro), 1566

- [SOC_SYSTIMER_BIT_WIDTH_HI \(C macro\), 1566](#)
[SOC_SYSTIMER_BIT_WIDTH_LO \(C macro\), 1566](#)
[SOC_SYSTIMER_COUNTER_NUM \(C macro\), 1566](#)
[SOC_SYSTIMER_FIXED_DIVIDER \(C macro\), 1566](#)
[SOC_SYSTIMER_INT_LEVEL \(C macro\), 1566](#)
[SOC_SYSTIMER_SUPPORT_RC_FAST \(C macro\), 1566](#)
[SOC_SYSTIMER_SUPPORTED \(C macro\), 1555](#)
[SOC_TEMPERATURE_SENSOR_SUPPORT_FAST_RC \(C macro\), 1569](#)
[SOC_TEMPERATURE_SENSOR_SUPPORT_XTAL \(C macro\), 1569](#)
[SOC_TIMER_GROUP_COUNTER_BIT_WIDTH \(C macro\), 1567](#)
[SOC_TIMER_GROUP_SUPPORT_RC_FAST \(C macro\), 1567](#)
[SOC_TIMER_GROUP_SUPPORT_XTAL \(C macro\), 1567](#)
[SOC_TIMER_GROUP_TIMERS_PER_GROUP \(C macro\), 1567](#)
[SOC_TIMER_GROUP_TOTAL_TIMERS \(C macro\), 1567](#)
[SOC_TIMER_GROUPS \(C macro\), 1567](#)
[SOC_TIMER_SUPPORT_ETM \(C macro\), 1567](#)
[SOC_TWAI_BRP_MAX \(C macro\), 1567](#)
[SOC_TWAI_BRP_MIN \(C macro\), 1567](#)
[SOC_TWAI_CLK_SUPPORT_XTAL \(C macro\), 1567](#)
[SOC_TWAI_CONTROLLER_NUM \(C macro\), 1567](#)
[SOC_TWAI_SUPPORTS_RX_STATUS \(C macro\), 1567](#)
[SOC_UART_BITRATE_MAX \(C macro\), 1568](#)
[SOC_UART_CLKS \(C macro\), 267](#)
[SOC_UART_FIFO_LEN \(C macro\), 1568](#)
[SOC_UART_HAS_LP_UART \(C macro\), 1568](#)
[SOC_UART_HP_NUM \(C macro\), 1568](#)
[SOC_UART_LP_NUM \(C macro\), 1568](#)
[SOC_UART_NUM \(C macro\), 1568](#)
[SOC_UART_SUPPORT_FSM_TX_WAIT_SEND \(C macro\), 1568](#)
[SOC_UART_SUPPORT_PLL_F80M_CLK \(C macro\), 1568](#)
[SOC_UART_SUPPORT_RTC_CLK \(C macro\), 1568](#)
[SOC_UART_SUPPORT_WAKEUP_INT \(C macro\), 1568](#)
[SOC_UART_SUPPORT_XTAL_CLK \(C macro\), 1568](#)
[SOC_UART_SUPPORTED \(C macro\), 1555](#)
[SOC_WDT_SUPPORTED \(C macro\), 1556](#)
[SOC_WIFI_LIGHT_SLEEP_CLK_WIDTH \(C macro\), 1568](#)
[SOC_XTAL_SUPPORT_40M \(C macro\), 1556](#)
[spi_bus_add_device \(C++ function\), 619](#)
[spi_bus_add_flash_device \(C++ function\), 583](#)
[spi_bus_config_t \(C++ struct\), 616](#)
[spi_bus_config_t::data0_io_num \(C++ member\), 616](#)
[spi_bus_config_t::data1_io_num \(C++ member\), 616](#)
[spi_bus_config_t::data2_io_num \(C++ member\), 616](#)
[spi_bus_config_t::data3_io_num \(C++ member\), 617](#)
[spi_bus_config_t::data4_io_num \(C++ member\), 617](#)
[spi_bus_config_t::data5_io_num \(C++ member\), 617](#)
[spi_bus_config_t::data6_io_num \(C++ member\), 617](#)
[spi_bus_config_t::data7_io_num \(C++ member\), 617](#)
[spi_bus_config_t::flags \(C++ member\), 617](#)
[spi_bus_config_t::intr_flags \(C++ member\), 617](#)
[spi_bus_config_t::isr_cpu_id \(C++ member\), 617](#)
[spi_bus_config_t::max_transfer_sz \(C++ member\), 617](#)
[spi_bus_config_t::miso_io_num \(C++ member\), 616](#)
[spi_bus_config_t::mosi_io_num \(C++ member\), 616](#)
[spi_bus_config_t::quadhd_io_num \(C++ member\), 616](#)
[spi_bus_config_t::quadwp_io_num \(C++ member\), 616](#)
[spi_bus_config_t::sclk_io_num \(C++ member\), 616](#)
[spi_bus_free \(C++ function\), 616](#)
[spi_bus_get_max_transaction_len \(C++ function\), 623](#)
[spi_bus_initialize \(C++ function\), 615](#)
[spi_bus_remove_device \(C++ function\), 620](#)
[spi_bus_remove_flash_device \(C++ function\), 584](#)
[spi_clock_source_t \(C++ type\), 613](#)
[spi_command_t \(C++ enum\), 614](#)
[spi_command_t::SPI_CMD_HD_EN_QPI \(C++ enumerator\), 614](#)
[spi_command_t::SPI_CMD_HD_INT0 \(C++ enumerator\), 614](#)
[spi_command_t::SPI_CMD_HD_INT1 \(C++ enumerator\), 614](#)
[spi_command_t::SPI_CMD_HD_INT2 \(C++ enumerator\), 614](#)
[spi_command_t::SPI_CMD_HD_RDBUF \(C++ enumerator\), 614](#)
[spi_command_t::SPI_CMD_HD_RDDMA \(C++ enumerator\), 614](#)
[spi_command_t::SPI_CMD_HD_SEG_END \(C++ enumerator\), 614](#)
[spi_command_t::SPI_CMD_HD_WR_END \(C++ enumerator\), 614](#)
[spi_command_t::SPI_CMD_HD_WRBUF \(C++ enumerator\), 614](#)
[spi_command_t::SPI_CMD_HD_WRDMA \(C++ enumerator\), 614](#)

[spi_common_dma_t \(C++ enum\), 619](#)
[spi_common_dma_t::SPI_DMA_CH_AUTO \(C++ enumerator\), 619](#)
[spi_common_dma_t::SPI_DMA_DISABLED \(C++ enumerator\), 619](#)
[SPI_DEVICE_3WIRE \(C macro\), 627](#)
[spi_device_acquire_bus \(C++ function\), 622](#)
[SPI_DEVICE_BIT_LSBFIRST \(C macro\), 627](#)
[SPI_DEVICE_CLK_AS_CS \(C macro\), 627](#)
[SPI_DEVICE_DDRCLK \(C macro\), 627](#)
[spi_device_get_actual_freq \(C++ function\), 622](#)
[spi_device_get_trans_result \(C++ function\), 620](#)
[SPI_DEVICE_HALFDUPLEX \(C macro\), 627](#)
[spi_device_handle_t \(C++ type\), 628](#)
[spi_device_interface_config_t \(C++ struct\), 623](#)
[spi_device_interface_config_t::address_bits \(C++ member\), 623](#)
[spi_device_interface_config_t::clock_speed \(C++ member\), 624](#)
[spi_device_interface_config_t::clock_spi_div \(C++ member\), 624](#)
[spi_device_interface_config_t::command_bits \(C++ member\), 623](#)
[spi_device_interface_config_t::cs_enable_pos \(C++ member\), 624](#)
[spi_device_interface_config_t::cs_enable_neg \(C++ member\), 624](#)
[spi_device_interface_config_t::dummy_bits \(C++ member\), 623](#)
[spi_device_interface_config_t::duty_cycle \(C++ member\), 624](#)
[spi_device_interface_config_t::flags \(C++ member\), 624](#)
[spi_device_interface_config_t::input_delay \(C++ member\), 624](#)
[spi_device_interface_config_t::mode \(C++ member\), 623](#)
[spi_device_interface_config_t::post_cb \(C++ member\), 624](#)
[spi_device_interface_config_t::pre_cb \(C++ member\), 624](#)
[spi_device_interface_config_t::queue_size \(C++ member\), 624](#)
[spi_device_interface_config_t::sample_rate \(C++ member\), 624](#)
[spi_device_interface_config_t::spics_io \(C++ member\), 624](#)
[SPI_DEVICE_NO_DUMMY \(C macro\), 627](#)
[SPI_DEVICE_NO_RETURN_RESULT \(C macro\), 627](#)
[spi_device_polling_end \(C++ function\), 621](#)
[spi_device_polling_start \(C++ function\), 621](#)
[spi_device_polling_transmit \(C++ function\), 621](#)
[SPI_DEVICE_POSITIVE_CS \(C macro\), 627](#)
[spi_device_queue_trans \(C++ function\), 620](#)
[spi_device_release_bus \(C++ function\), 622](#)
[SPI_DEVICE_RXBIT_LSBFIRST \(C macro\), 627](#)
[spi_device_transmit \(C++ function\), 620](#)
[SPI_DEVICE_TXBIT_LSBFIRST \(C macro\), 627](#)
[spi_dma_chan_t \(C++ type\), 619](#)
[spi_event_t \(C++ enum\), 613](#)
[spi_event_t::SPI_EV_BUF_RX \(C++ enumerator\), 613](#)
[spi_event_t::SPI_EV_BUF_TX \(C++ enumerator\), 613](#)
[spi_event_t::SPI_EV_CMD9 \(C++ enumerator\), 614](#)
[spi_event_t::SPI_EV_CMDA \(C++ enumerator\), 614](#)
[spi_event_t::SPI_EV_RECV \(C++ enumerator\), 614](#)
[spi_event_t::SPI_EV_RECV_DMA_READY \(C++ enumerator\), 614](#)
[spi_event_t::SPI_EV_SEND \(C++ enumerator\), 614](#)
[spi_event_t::SPI_EV_SEND_DMA_READY \(C++ enumerator\), 614](#)
[spi_event_t::SPI_EV_TRANS \(C++ enumerator\), 614](#)
[spi_flash_cache2phys \(C++ function\), 594](#)
[SPI_FLASH_CACHE2PHYS_FAIL \(C macro\), 595](#)
[spi_flash_chip_t \(C++ type\), 592](#)
[SPI_FLASH_CONFIG_CONF_BITS \(C macro\), 599](#)
[spi_flash_counter_t \(C++ type\), 603](#)
[spi_flash_counters_t \(C++ type\), 603](#)
[spi_flash_dump_counters \(C++ function\), 602](#)
[spi_flash_encryption_t \(C++ struct\), 596](#)
[spi_flash_encryption_t::flash_encryption_check \(C++ member\), 597](#)
[spi_flash_encryption_t::flash_encryption_data_prepare \(C++ member\), 597](#)
[spi_flash_encryption_t::flash_encryption_destroy \(C++ member\), 597](#)
[spi_flash_encryption_t::flash_encryption_disable \(C++ member\), 597](#)
[spi_flash_encryption_t::flash_encryption_done \(C++ member\), 597](#)
[spi_flash_encryption_t::flash_encryption_enable \(C++ member\), 597](#)
[spi_flash_get_counters \(C++ function\), 602](#)
[spi_flash_host_driver_s \(C++ struct\), 597](#)
[spi_flash_host_driver_s::check_suspend \(C++ member\), 599](#)
[spi_flash_host_driver_s::common_command \(C++ member\), 597](#)
[spi_flash_host_driver_s::configure_host_io_mode \(C++ member\), 599](#)
[spi_flash_host_driver_s::dev_config \(C++ member\), 597](#)
[spi_flash_host_driver_s::erase_block \(C++ member\), 598](#)

spi_flash_host_driver_s::erase_chip (C++ member), 598
 spi_flash_host_driver_s::erase_sector (C++ member), 598
 spi_flash_host_driver_s::flush_cache (C++ member), 599
 spi_flash_host_driver_s::host_status (C++ member), 598
 spi_flash_host_driver_s::poll_cmd_done (C++ member), 599
 spi_flash_host_driver_s::program_page (C++ member), 598
 spi_flash_host_driver_s::read (C++ member), 598
 spi_flash_host_driver_s::read_data_slice (C++ member), 598
 spi_flash_host_driver_s::read_id (C++ member), 597
 spi_flash_host_driver_s::read_status (C++ member), 598
 spi_flash_host_driver_s::resume (C++ member), 599
 spi_flash_host_driver_s::set_write_protect (C++ member), 598
 spi_flash_host_driver_s::supports_direct_read (C++ member), 598
 spi_flash_host_driver_s::supports_direct_write (C++ member), 598
 spi_flash_host_driver_s::sus_setup (C++ member), 599
 spi_flash_host_driver_s::suspend (C++ member), 599
 spi_flash_host_driver_s::write_data_slice (C++ member), 598
 spi_flash_host_driver_t (C++ type), 600
 spi_flash_host_inst_t (C++ struct), 597
 spi_flash_host_inst_t::driver (C++ member), 597
 spi_flash_mmap (C++ function), 593
 spi_flash_mmap_dump (C++ function), 594
 spi_flash_mmap_get_free_pages (C++ function), 594
 spi_flash_mmap_handle_t (C++ type), 595
 spi_flash_mmap_memory_t (C++ enum), 595
 spi_flash_mmap_memory_t::SPI_FLASH_MMAPPED_DATA (C++ enumerator), 595
 spi_flash_mmap_memory_t::SPI_FLASH_MMAPPED_EMPTY (C++ enumerator), 595
 spi_flash_mmap_pages (C++ function), 593
 SPI_FLASH_MMU_PAGE_SIZE (C macro), 595
 spi_flash_munmap (C++ function), 593
 SPI_FLASH_OPI_FLAG (C macro), 599
 SPI_FLASH_OS_IS_ERASING_STATUS_FLAG (C macro), 592
 spi_flash_phys2cache (C++ function), 594
 SPI_FLASH_READ_MODE_MIN (C macro), 599
 spi_flash_reset_counters (C++ function), 602
 SPI_FLASH_SEC_SIZE (C macro), 595
 spi_flash_sus_cmd_conf (C++ struct), 596
 spi_flash_sus_cmd_conf::cmd_rdsr (C++ member), 596
 spi_flash_sus_cmd_conf::res_cmd (C++ member), 596
 spi_flash_sus_cmd_conf::reserved (C++ member), 596
 spi_flash_sus_cmd_conf::sus_cmd (C++ member), 596
 spi_flash_sus_cmd_conf::sus_mask (C++ member), 596
 SPI_FLASH_TRANS_FLAG_BYTE_SWAP (C macro), 599
 SPI_FLASH_TRANS_FLAG_CMD16 (C macro), 599
 SPI_FLASH_TRANS_FLAG_IGNORE_BASEIO (C macro), 599
 SPI_FLASH_TRANS_FLAG_PE_CMD (C macro), 599
 spi_flash_trans_t (C++ struct), 595
 spi_flash_trans_t::address (C++ member), 596
 spi_flash_trans_t::address_bitlen (C++ member), 596
 spi_flash_trans_t::command (C++ member), 596
 spi_flash_trans_t::dummy_bitlen (C++ member), 596
 spi_flash_trans_t::flags (C++ member), 596
 spi_flash_trans_t::io_mode (C++ member), 596
 spi_flash_trans_t::miso_data (C++ member), 596
 spi_flash_trans_t::miso_len (C++ member), 595
 spi_flash_trans_t::mosi_data (C++ member), 596
 spi_flash_trans_t::mosi_len (C++ member), 595
 spi_flash_trans_t::reserved (C++ member), 595
 SPI_FLASH_YIELD_REQ_SUSPEND (C macro), 592
 SPI_FLASH_YIELD_REQ_YIELD (C macro), 592
 SPI_FLASH_YIELD_STA_RESUME (C macro), 592
 spi_get_actual_clock (C++ function), 622
 spi_get_freq_limit (C++ function), 623
 spi_get_timing (C++ function), 622
 spi_host_device_t (C++ enum), 613
 spi_host_device_t::SPI1_HOST (C++ enumerator), 613
 spi_host_device_t::SPI2_HOST (C++ enumerator), 613
 spi_host_device_t::SPI3_HOST (C++ enumerator), 613
 spi_host_device_t::SPI_HOST_MAX (C++ enumerator), 613

- `spi_line_mode_t` (C++ struct), 613
- `spi_line_mode_t::addr_lines` (C++ member), 613
- `spi_line_mode_t::cmd_lines` (C++ member), 613
- `spi_line_mode_t::data_lines` (C++ member), 613
- `SPI_MASTER_FREQ_10M` (C macro), 626
- `SPI_MASTER_FREQ_11M` (C macro), 626
- `SPI_MASTER_FREQ_13M` (C macro), 626
- `SPI_MASTER_FREQ_16M` (C macro), 626
- `SPI_MASTER_FREQ_20M` (C macro), 626
- `SPI_MASTER_FREQ_26M` (C macro), 626
- `SPI_MASTER_FREQ_40M` (C macro), 627
- `SPI_MASTER_FREQ_80M` (C macro), 627
- `SPI_MASTER_FREQ_8M` (C macro), 626
- `SPI_MASTER_FREQ_9M` (C macro), 626
- `SPI_MAX_DMA_LEN` (C macro), 617
- `spi_sampling_point_t` (C++ enum), 614
- `spi_sampling_point_t::SPI_SAMPLING_POINT_PHASE_0` (C++ enumerator), 615
- `spi_sampling_point_t::SPI_SAMPLING_POINT_PHASE_1` (C++ enumerator), 615
- `SPI_SLAVE_BIT_LSBFIRST` (C macro), 635
- `spi_slave_free` (C++ function), 632
- `spi_slave_get_trans_result` (C++ function), 633
- `spi_slave_initialize` (C++ function), 632
- `spi_slave_interface_config_t` (C++ struct), 633
- `spi_slave_interface_config_t::flags` (C++ member), 633
- `spi_slave_interface_config_t::mode` (C++ member), 634
- `spi_slave_interface_config_t::post_setup` (C++ member), 634
- `spi_slave_interface_config_t::post_transfer` (C++ member), 634
- `spi_slave_interface_config_t::queue_size` (C++ member), 634
- `spi_slave_interface_config_t::spics_io_pin` (C++ member), 633
- `SPI_SLAVE_NO_RETURN_RESULT` (C macro), 635
- `spi_slave_queue_trans` (C++ function), 632
- `SPI_SLAVE_RXBIT_LSBFIRST` (C macro), 635
- `spi_slave_transaction_t` (C++ struct), 634
- `spi_slave_transaction_t::length` (C++ member), 634
- `spi_slave_transaction_t::rx_buffer` (C++ member), 634
- `spi_slave_transaction_t::trans_len` (C++ member), 634
- `spi_slave_transaction_t::tx_buffer` (C++ member), 634
- `spi_slave_transaction_t::user` (C++ member), 634
- `spi_slave_transmit` (C++ function), 633
- `SPI_SLAVE_TXBIT_LSBFIRST` (C macro), 635
- `SPI_SWAP_DATA_RX` (C macro), 617
- `SPI_SWAP_DATA_TX` (C macro), 617
- `SPI_TRANS_CS_KEEP_ACTIVE` (C macro), 628
- `SPI_TRANS_DMA_BUFFER_ALIGN_MANUAL` (C macro), 628
- `SPI_TRANS_MODE_DIO` (C macro), 627
- `SPI_TRANS_MODE_DIOQIO_ADDR` (C macro), 628
- `SPI_TRANS_MODE_OCT` (C macro), 628
- `SPI_TRANS_MODE_QIO` (C macro), 627
- `SPI_TRANS_MULTILINE_ADDR` (C macro), 628
- `SPI_TRANS_MULTILINE_CMD` (C macro), 628
- `SPI_TRANS_USE_RXDATA` (C macro), 627
- `SPI_TRANS_USE_TXDATA` (C macro), 628
- `SPI_TRANS_VARIABLE_ADDR` (C macro), 628
- `SPI_TRANS_VARIABLE_CMD` (C macro), 628
- `SPI_TRANS_VARIABLE_DUMMY` (C macro), 628
- `spi_transaction_ext_t` (C++ struct), 626
- `spi_transaction_ext_t::address_bits` (C++ member), 626
- `spi_transaction_ext_t::base` (C++ member), 626
- `spi_transaction_ext_t::command_bits` (C++ member), 626
- `spi_transaction_ext_t::dummy_bits` (C++ member), 626
- `spi_transaction_t` (C++ struct), 625
- `spi_transaction_t::addr` (C++ member), 625
- `spi_transaction_t::cmd` (C++ member), 625
- `spi_transaction_t::flags` (C++ member), 625
- `spi_transaction_t::length` (C++ member), 625
- `spi_transaction_t::rx_buffer` (C++ member), 625
- `spi_transaction_t::rx_data` (C++ member), 625
- `spi_transaction_t::rxlength` (C++ member), 625
- `spi_transaction_t::tx_buffer` (C++ member), 625
- `spi_transaction_t::tx_data` (C++ member), 625
- `spi_transaction_t::user` (C++ member), 625
- `SPICOMMON_BUSFLAG_DUAL` (C macro), 618
- `SPICOMMON_BUSFLAG_GPIO_PINS` (C macro), 618
- `SPICOMMON_BUSFLAG_IO4_IO7` (C macro), 618
- `SPICOMMON_BUSFLAG_IOMUX_PINS` (C macro), 618
- `SPICOMMON_BUSFLAG_MASTER` (C macro), 618
- `SPICOMMON_BUSFLAG_MISO` (C macro), 618
- `SPICOMMON_BUSFLAG_MOSI` (C macro), 618
- `SPICOMMON_BUSFLAG_NATIVE_PINS` (C macro), 618
- `SPICOMMON_BUSFLAG_OCTAL` (C macro), 618
- `SPICOMMON_BUSFLAG_QUAD` (C macro), 618
- `SPICOMMON_BUSFLAG_SCLK` (C macro), 618
- `SPICOMMON_BUSFLAG_SLAVE` (C macro), 618

- SPICOMMON_BUSFLAG_SLP_ALLOW_PD (C macro), 618
- SPICOMMON_BUSFLAG_WPHD (C macro), 618
- StaticRingbuffer_t (C++ type), 1422
- StreamBufferCallbackFunction_t (C++ type), 1394
- StreamBufferHandle_t (C++ type), 1394
- ## T
- task_wdt_msg_handler (C++ type), 1586
- taskDISABLE_INTERRUPTS (C macro), 1318
- taskENABLE_INTERRUPTS (C macro), 1318
- taskENTER_CRITICAL (C macro), 1318
- taskENTER_CRITICAL_FROM_ISR (C macro), 1318
- taskENTER_CRITICAL_ISR (C macro), 1318
- taskEXIT_CRITICAL (C macro), 1318
- taskEXIT_CRITICAL_FROM_ISR (C macro), 1318
- taskEXIT_CRITICAL_ISR (C macro), 1318
- TaskHandle_t (C++ type), 1323
- TaskHookFunction_t (C++ type), 1323
- taskSCHEDULER_NOT_STARTED (C macro), 1318
- taskSCHEDULER_RUNNING (C macro), 1318
- taskSCHEDULER_SUSPENDED (C macro), 1318
- TaskStatus_t (C++ type), 1323
- taskVALID_CORE_ID (C macro), 1318
- taskYIELD (C macro), 1318
- TimerCallbackFunction_t (C++ type), 1376
- TimerHandle_t (C++ type), 1376
- tls_keep_alive_cfg (C++ struct), 69
- tls_keep_alive_cfg::keep_alive_count (C++ member), 69
- tls_keep_alive_cfg::keep_alive_enable (C++ member), 69
- tls_keep_alive_cfg::keep_alive_idle (C++ member), 69
- tls_keep_alive_cfg::keep_alive_interval (C++ member), 69
- tls_keep_alive_cfg_t (C++ type), 72
- TlsDeleteCallbackFunction_t (C++ type), 1431
- topic_t (C++ struct), 55
- topic_t::filter (C++ member), 55
- topic_t::qos (C++ member), 55
- transaction_cb_t (C++ type), 628
- tskIDLE_PRIORITY (C macro), 1318
- tskNO_AFFINITY (C macro), 1318
- ## U
- uart_at_cmd_t (C++ struct), 656
- uart_at_cmd_t::char_num (C++ member), 656
- uart_at_cmd_t::cmd_char (C++ member), 656
- uart_at_cmd_t::gap_tout (C++ member), 656
- uart_at_cmd_t::post_idle (C++ member), 656
- uart_at_cmd_t::pre_idle (C++ member), 656
- UART_BITRATE_MAX (C macro), 655
- uart_clear_intr_status (C++ function), 644
- uart_config_t (C++ struct), 653
- uart_config_t::backup_before_sleep (C++ member), 654
- uart_config_t::baud_rate (C++ member), 653
- uart_config_t::data_bits (C++ member), 653
- uart_config_t::flags (C++ member), 654
- uart_config_t::flow_ctrl (C++ member), 653
- uart_config_t::lp_source_clk (C++ member), 654
- uart_config_t::parity (C++ member), 653
- uart_config_t::rx_flow_ctrl_thresh (C++ member), 653
- uart_config_t::source_clk (C++ member), 653
- uart_config_t::stop_bits (C++ member), 653
- UART_CTS_GPIO13_DIRECT_CHANNEL (C macro), 661
- UART_CTS_GPIO9_DIRECT_CHANNEL (C macro), 660
- uart_disable_intr_mask (C++ function), 645
- uart_disable_pattern_det_intr (C++ function), 649
- uart_disable_rx_intr (C++ function), 645
- uart_disable_tx_intr (C++ function), 645
- uart_driver_delete (C++ function), 642
- uart_driver_install (C++ function), 641
- uart_enable_intr_mask (C++ function), 645
- uart_enable_pattern_det_baud_intr (C++ function), 649
- uart_enable_rx_intr (C++ function), 645
- uart_enable_tx_intr (C++ function), 645
- uart_event_t (C++ struct), 654
- uart_event_t::size (C++ member), 654
- uart_event_t::timeout_flag (C++ member), 654
- uart_event_t::type (C++ member), 654
- uart_event_type_t (C++ enum), 655
- uart_event_type_t::UART_BREAK (C++ enumerator), 655
- uart_event_type_t::UART_BUFFER_FULL (C++ enumerator), 655
- uart_event_type_t::UART_DATA (C++ enumerator), 655
- uart_event_type_t::UART_DATA_BREAK (C++ enumerator), 655
- uart_event_type_t::UART_EVENT_MAX (C++ enumerator), 655
- uart_event_type_t::UART_FIFO_OVF (C++ enumerator), 655
- uart_event_type_t::UART_FRAME_ERR (C++ enumerator), 655
- uart_event_type_t::UART_PARITY_ERR (C++ enumerator), 655
- uart_event_type_t::UART_PATTERN_DET

- (C++ enumerator), 655
- uart_event_type_t::UART_WAKEUP (C++ enumerator), 655
- UART_FIFO_LEN (C macro), 655
- uart_flush (C++ function), 648
- uart_flush_input (C++ function), 648
- uart_get_baudrate (C++ function), 643
- uart_get_buffered_data_len (C++ function), 649
- uart_get_collision_flag (C++ function), 651
- uart_get_hw_flow_ctrl (C++ function), 644
- uart_get_parity (C++ function), 643
- uart_get_sclk_freq (C++ function), 643
- uart_get_stop_bits (C++ function), 643
- uart_get_tx_buffer_free_size (C++ function), 649
- uart_get_wakeup_threshold (C++ function), 652
- uart_get_word_length (C++ function), 642
- UART_GPIO10_DIRECT_CHANNEL (C macro), 660
- UART_GPIO11_DIRECT_CHANNEL (C macro), 661
- UART_GPIO12_DIRECT_CHANNEL (C macro), 661
- UART_GPIO13_DIRECT_CHANNEL (C macro), 661
- UART_GPIO37_DIRECT_CHANNEL (C macro), 660
- UART_GPIO38_DIRECT_CHANNEL (C macro), 660
- UART_GPIO8_DIRECT_CHANNEL (C macro), 660
- UART_GPIO9_DIRECT_CHANNEL (C macro), 660
- UART_HW_FIFO_LEN (C macro), 655
- uart_hw_flowcontrol_t (C++ enum), 658
- uart_hw_flowcontrol_t::UART_HW_FLOWCONTROL_CTS (C++ enumerator), 659
- uart_hw_flowcontrol_t::UART_HW_FLOWCONTROL_RTS (C++ enumerator), 659
- uart_hw_flowcontrol_t::UART_HW_FLOWCONTROL_BOTH (C++ enumerator), 659
- uart_hw_flowcontrol_t::UART_HW_FLOWCONTROL_DISABLE (C++ enumerator), 659
- uart_hw_flowcontrol_t::UART_HW_FLOWCONTROL_MAX (C++ enumerator), 659
- uart_hw_flowcontrol_t::UART_HW_FLOWCONTROL_RTS (C++ enumerator), 659
- uart_intr_config (C++ function), 647
- uart_intr_config_t (C++ struct), 654
- uart_intr_config_t::intr_enable_mask (C++ member), 654
- uart_intr_config_t::rx_timeout_thresh (C++ member), 654
- uart_intr_config_t::rxfifo_full_thresh (C++ member), 654
- uart_intr_config_t::txfifo_empty_intr_thresh (C++ member), 654
- uart_is_driver_installed (C++ function), 642
- uart_isr_handle_t (C++ type), 655
- uart_mode_t (C++ enum), 657
- uart_mode_t::UART_MODE_IRDA (C++ enumerator), 657
- uart_mode_t::UART_MODE_RS485_APP_CTRL (C++ enumerator), 657
- uart_mode_t::UART_MODE_RS485_COLLISION_DETECT (C++ enumerator), 657
- uart_mode_t::UART_MODE_RS485_HALF_DUPLEX (C++ enumerator), 657
- uart_mode_t::UART_MODE_UART (C++ enumerator), 657
- UART_NUM_0_CTS_DIRECT_GPIO_NUM (C macro), 660
- UART_NUM_0_RTS_DIRECT_GPIO_NUM (C macro), 660
- UART_NUM_0_RXD_DIRECT_GPIO_NUM (C macro), 660
- UART_NUM_0_TXD_DIRECT_GPIO_NUM (C macro), 660
- UART_NUM_1_CTS_DIRECT_GPIO_NUM (C macro), 661
- UART_NUM_1_RTS_DIRECT_GPIO_NUM (C macro), 661
- UART_NUM_1_RXD_DIRECT_GPIO_NUM (C macro), 661
- UART_NUM_1_TXD_DIRECT_GPIO_NUM (C macro), 661
- uart_param_config (C++ function), 647
- uart_parity_t (C++ enum), 658
- uart_parity_t::UART_PARITY_DISABLE (C++ enumerator), 658
- uart_parity_t::UART_PARITY_EVEN (C++ enumerator), 658
- uart_parity_t::UART_PARITY_ODD (C++ enumerator), 658
- uart_pattern_get_pos (C++ function), 650
- uart_pattern_pop_pos (C++ function), 649
- uart_pattern_queue_reset (C++ function), 650
- UART_SAME_NO_CHANGE (C macro), 655
- uart_port_t (C++ enum), 657
- uart_port_t::LP_UART_NUM_0 (C++ enumerator), 657
- uart_port_t::UART_NUM_0 (C++ enumerator), 657
- uart_port_t::UART_NUM_1 (C++ enumerator), 657
- uart_port_t::UART_NUM_2 (C++ enumerator), 657
- uart_port_t::UART_NUM_3 (C++ enumerator), 657
- uart_port_t::UART_NUM_4 (C++ enumerator), 657
- uart_port_t::UART_NUM_MAX (C++ enumerator), 657
- uart_read_bytes (C++ function), 648
- UART_RTS_GPIO12_DIRECT_CHANNEL (C macro), 661
- UART_RTS_GPIO8_DIRECT_CHANNEL (C macro), 660
- UART_RXD_GPIO11_DIRECT_CHANNEL (C macro), 661
- UART_RXD_GPIO38_DIRECT_CHANNEL (C macro), 660

- uart_sclk_t (C++ type), 657
- uart_set_always_rx_timeout (C++ function), 653
- uart_set_baudrate (C++ function), 643
- uart_set_dtr (C++ function), 646
- uart_set_hw_flow_ctrl (C++ function), 644
- uart_set_line_inverse (C++ function), 644
- uart_set_loop_back (C++ function), 652
- uart_set_mode (C++ function), 650
- uart_set_parity (C++ function), 643
- uart_set_pin (C++ function), 646
- uart_set_rts (C++ function), 646
- uart_set_rx_full_threshold (C++ function), 651
- uart_set_rx_timeout (C++ function), 651
- uart_set_stop_bits (C++ function), 642
- uart_set_sw_flow_ctrl (C++ function), 644
- uart_set_tx_empty_threshold (C++ function), 651
- uart_set_tx_idle_num (C++ function), 646
- uart_set_wakeup_threshold (C++ function), 652
- uart_set_word_length (C++ function), 642
- uart_signal_inv_t (C++ enum), 659
- uart_signal_inv_t::UART_SIGNAL_CTS_INV ulTaskGenericNotifyValueClear (C++ function), 1315
- uart_signal_inv_t::UART_SIGNAL_DSR_INV ulTaskGetIdleRunTimeCounter (C++ function), 1311
- uart_signal_inv_t::UART_SIGNAL_DTR_INV ulTaskGetIdleRunTimePercent (C++ function), 1312
- uart_signal_inv_t::UART_SIGNAL_INV_DISABLE ulTaskNotifyTakeIndexed (C macro), 1322
- uart_signal_inv_t::UART_SIGNAL_IRDA_RX_INV ulTaskNotifyValueClear (C macro), 1323
- uart_signal_inv_t::UART_SIGNAL_IRDA_TX_INV ulTaskNotifyValueClearIndexed (C macro), 1323
- uart_signal_inv_t::UART_SIGNAL_IRDA_TX_INV ulQueueMessagesWaiting (C++ function), 1329
- uart_signal_inv_t::UART_SIGNAL_RTS_INV uxQueueMessagesWaitingFromISR (C++ function), 1331
- uart_signal_inv_t::UART_SIGNAL_RXD_INV uxQueueSpacesAvailable (C++ function), 1329
- uart_signal_inv_t::UART_SIGNAL_TXD_INV uxSemaphoreGetCount (C macro), 1358
- uart_signal_inv_t::UART_SIGNAL_TXD_INV uxSemaphoreGetCountFromISR (C macro), 1358
- uart_signal_inv_t::UART_SIGNAL_TXD_INV uxTaskGetNumberOfTasks (C++ function), 1307
- uart_signal_inv_t::UART_SIGNAL_TXD_INV uxTaskGetStackHighWaterMark (C++ function), 1307
- uart_signal_inv_t::UART_SIGNAL_TXD_INV uxTaskGetStackHighWaterMark2 (C++ function), 1308
- uart_signal_inv_t::UART_SIGNAL_TXD_INV uxTaskGetSystemState (C++ function), 1309
- uart_signal_inv_t::UART_SIGNAL_TXD_INV uxTaskPriorityGet (C++ function), 1302
- uart_signal_inv_t::UART_SIGNAL_TXD_INV uxTaskPriorityGetFromISR (C++ function), 1302
- uart_signal_inv_t::UART_SIGNAL_TXD_INV uxTimerGetReloadMode (C++ function), 1366
- uart_sw_flowctrl_t (C++ struct), 656
- uart_sw_flowctrl_t::xoff_char (C++ member), 656
- uart_sw_flowctrl_t::xoff_thrd (C++ member), 656
- uart_sw_flowctrl_t::xon_char (C++ member), 656
- uart_sw_flowctrl_t::xon_thrd (C++ member), 656
- uart_tx_chars (C++ function), 647
- UART_TXD_GPIO10_DIRECT_CHANNEL (C macro), 661
- UART_TXD_GPIO37_DIRECT_CHANNEL (C macro), 660
- uart_wait_tx_done (C++ function), 647
- uart_wait_tx_idle_polling (C++ function), 652
- uart_word_length_t (C++ enum), 658
- uart_word_length_t::UART_DATA_5_BITS (C++ enumerator), 658
- uart_word_length_t::UART_DATA_6_BITS (C++ enumerator), 658
- uart_word_length_t::UART_DATA_7_BITS (C++ enumerator), 658
- uart_word_length_t::UART_DATA_8_BITS (C++ enumerator), 658
- uart_word_length_t::UART_DATA_BITS_MAX (C++ enumerator), 658
- uart_write_bytes (C++ function), 647
- uart_write_bytes_with_break (C++ function), 648

V

- vApplicationGetIdleTaskMemory (C++ function), 1308
- vApplicationGetTimerTaskMemory (C++ function), 1367
- vEventGroupDelete (C++ function), 1383

[vEventGroupDeleteWithCaps \(C++ function\), 1431](#)
[vMessageBufferDelete \(C macro\), 1401](#)
[vMessageBufferDeleteWithCaps \(C++ function\), 1430](#)
[vprintf_like_t \(C++ type\), 1497](#)
[vQueueAddToRegistry \(C++ function\), 1331](#)
[vQueueDelete \(C++ function\), 1329](#)
[vQueueDeleteWithCaps \(C++ function\), 1428](#)
[vQueueUnregisterQueue \(C++ function\), 1332](#)
[vRingbufferDelete \(C++ function\), 1420](#)
[vRingbufferDeleteWithCaps \(C++ function\), 1422](#)
[vRingbufferGetInfo \(C++ function\), 1421](#)
[vRingbufferReturnItem \(C++ function\), 1419](#)
[vRingbufferReturnItemFromISR \(C++ function\), 1419](#)
[vSemaphoreCreateBinary \(C macro\), 1345](#)
[vSemaphoreDelete \(C macro\), 1358](#)
[vSemaphoreDeleteWithCaps \(C++ function\), 1429](#)
[vStreamBufferDelete \(C++ function\), 1390](#)
[vStreamBufferDeleteWithCaps \(C++ function\), 1430](#)
[vTaskAllocateMPURegions \(C++ function\), 1299](#)
[vTaskDelay \(C++ function\), 1300](#)
[vTaskDelete \(C++ function\), 1299](#)
[vTaskDeleteWithCaps \(C++ function\), 1428](#)
[vTaskGenericNotifyGiveFromISR \(C++ function\), 1313](#)
[vTaskGetInfo \(C++ function\), 1302](#)
[vTaskGetRunTimeStats \(C++ function\), 1311](#)
[vTaskList \(C++ function\), 1310](#)
[vTaskNotifyGiveFromISR \(C macro\), 1322](#)
[vTaskNotifyGiveIndexedFromISR \(C macro\), 1322](#)
[vTaskPrioritySet \(C++ function\), 1303](#)
[vTaskResume \(C++ function\), 1304](#)
[vTaskSetApplicationTaskTag \(C++ function\), 1308](#)
[vTaskSetThreadLocalStoragePointer \(C++ function\), 1308](#)
[vTaskSetThreadLocalStoragePointerAndDeleteCallback \(C++ function\), 1427](#)
[vTaskSetTimeOutState \(C++ function\), 1315](#)
[vTaskSuspend \(C++ function\), 1304](#)
[vTaskSuspendAll \(C++ function\), 1305](#)
[vTimerSetReloadMode \(C++ function\), 1366](#)
[vTimerSetTimerID \(C++ function\), 1363](#)

W

[wl_erase_range \(C++ function\), 1202](#)
[wl_handle_t \(C++ type\), 1203](#)
[WL_INVALID_HANDLE \(C macro\), 1203](#)
[wl_mount \(C++ function\), 1202](#)
[wl_read \(C++ function\), 1203](#)
[wl_sector_size \(C++ function\), 1203](#)

[wl_size \(C++ function\), 1203](#)
[wl_unmount \(C++ function\), 1202](#)
[wl_write \(C++ function\), 1202](#)

X

[xEventGroupClearBits \(C++ function\), 1379](#)
[xEventGroupClearBitsFromISR \(C macro\), 1383](#)
[xEventGroupCreate \(C++ function\), 1377](#)
[xEventGroupCreateStatic \(C++ function\), 1377](#)
[xEventGroupCreateWithCaps \(C++ function\), 1430](#)
[xEventGroupGetBits \(C macro\), 1385](#)
[xEventGroupGetBitsFromISR \(C++ function\), 1383](#)
[xEventGroupGetStaticBuffer \(C++ function\), 1383](#)
[xEventGroupSetBits \(C++ function\), 1380](#)
[xEventGroupSetBitsFromISR \(C macro\), 1384](#)
[xEventGroupSync \(C++ function\), 1381](#)
[xEventGroupWaitBits \(C++ function\), 1378](#)
[xMessageBufferCreateStaticWithCallback \(C macro\), 1395](#)
[xMessageBufferCreateWithCallback \(C macro\), 1395](#)
[xMessageBufferCreateWithCaps \(C++ function\), 1430](#)
[xMessageBufferGetStaticBuffers \(C macro\), 1396](#)
[xMessageBufferIsEmpty \(C macro\), 1402](#)
[xMessageBufferIsFull \(C macro\), 1402](#)
[xMessageBufferNextLengthBytes \(C macro\), 1402](#)
[xMessageBufferReceive \(C macro\), 1399](#)
[xMessageBufferReceiveCompletedFromISR \(C macro\), 1403](#)
[xMessageBufferReceiveFromISR \(C macro\), 1400](#)
[xMessageBufferReset \(C macro\), 1402](#)
[xMessageBufferSend \(C macro\), 1397](#)
[xMessageBufferSendCompletedFromISR \(C macro\), 1402](#)
[xMessageBufferSendFromISR \(C macro\), 1398](#)
[xMessageBufferSpaceAvailable \(C macro\), 1402](#)
[xMessageBufferSpacesAvailable \(C macro\), 1402](#)
[xQueueAddToSet \(C++ function\), 1333](#)
[xQueueCreate \(C macro\), 1334](#)
[xQueueCreateSet \(C++ function\), 1332](#)
[xQueueCreateStatic \(C macro\), 1335](#)
[xQueueCreateWithCaps \(C++ function\), 1428](#)
[xQueueGenericSend \(C++ function\), 1325](#)
[xQueueGenericSendFromISR \(C++ function\), 1329](#)
[xQueueGetStaticBuffers \(C macro\), 1336](#)
[xQueueGiveFromISR \(C++ function\), 1330](#)

- `xQueueIsQueueEmptyFromISR` (C++ *function*), [1331](#)
- `xQueueIsQueueFullFromISR` (C++ *function*), [1331](#)
- `xQueueOverwrite` (C *macro*), [1339](#)
- `xQueueOverwriteFromISR` (C *macro*), [1342](#)
- `xQueuePeek` (C++ *function*), [1326](#)
- `xQueuePeekFromISR` (C++ *function*), [1327](#)
- `xQueueReceive` (C++ *function*), [1328](#)
- `xQueueReceiveFromISR` (C++ *function*), [1330](#)
- `xQueueRemoveFromSet` (C++ *function*), [1333](#)
- `xQueueReset` (C *macro*), [1344](#)
- `xQueueSelectFromSet` (C++ *function*), [1333](#)
- `xQueueSelectFromSetFromISR` (C++ *function*), [1334](#)
- `xQueueSend` (C *macro*), [1338](#)
- `xQueueSendFromISR` (C *macro*), [1343](#)
- `xQueueSendToBack` (C *macro*), [1337](#)
- `xQueueSendToBackFromISR` (C *macro*), [1341](#)
- `xQueueSendToFront` (C *macro*), [1336](#)
- `xQueueSendToFrontFromISR` (C *macro*), [1340](#)
- `xRingbufferAddToQueueSetRead` (C++ *function*), [1420](#)
- `xRingbufferCanRead` (C++ *function*), [1420](#)
- `xRingbufferCreate` (C++ *function*), [1414](#)
- `xRingbufferCreateNoSplit` (C++ *function*), [1414](#)
- `xRingbufferCreateStatic` (C++ *function*), [1414](#)
- `xRingbufferCreateWithCaps` (C++ *function*), [1421](#)
- `xRingbufferGetCurFreeSize` (C++ *function*), [1420](#)
- `xRingbufferGetMaxItemSize` (C++ *function*), [1420](#)
- `xRingbufferGetStaticBuffer` (C++ *function*), [1421](#)
- `xRingbufferPrintInfo` (C++ *function*), [1421](#)
- `xRingbufferReceive` (C++ *function*), [1416](#)
- `xRingbufferReceiveFromISR` (C++ *function*), [1416](#)
- `xRingbufferReceiveSplit` (C++ *function*), [1417](#)
- `xRingbufferReceiveSplitFromISR` (C++ *function*), [1418](#)
- `xRingbufferReceiveUpTo` (C++ *function*), [1418](#)
- `xRingbufferReceiveUpToFromISR` (C++ *function*), [1419](#)
- `xRingbufferRemoveFromQueueSetRead` (C++ *function*), [1421](#)
- `xRingbufferSend` (C++ *function*), [1414](#)
- `xRingbufferSendAcquire` (C++ *function*), [1415](#)
- `xRingbufferSendComplete` (C++ *function*), [1416](#)
- `xRingbufferSendFromISR` (C++ *function*), [1415](#)
- `xSemaphoreCreateBinary` (C *macro*), [1345](#)
- `xSemaphoreCreateBinaryStatic` (C *macro*), [1346](#)
- `xSemaphoreCreateBinaryWithCaps` (C++ *function*), [1428](#)
- `xSemaphoreCreateCounting` (C *macro*), [1356](#)
- `xSemaphoreCreateCountingStatic` (C *macro*), [1357](#)
- `xSemaphoreCreateCountingWithCaps` (C++ *function*), [1429](#)
- `xSemaphoreCreateMutex` (C *macro*), [1353](#)
- `xSemaphoreCreateMutexStatic` (C *macro*), [1353](#)
- `xSemaphoreCreateMutexWithCaps` (C++ *function*), [1429](#)
- `xSemaphoreCreateRecursiveMutex` (C *macro*), [1354](#)
- `xSemaphoreCreateRecursiveMutexStatic` (C *macro*), [1355](#)
- `xSemaphoreCreateRecursiveMutexWithCaps` (C++ *function*), [1429](#)
- `xSemaphoreGetMutexHolder` (C *macro*), [1358](#)
- `xSemaphoreGetMutexHolderFromISR` (C *macro*), [1358](#)
- `xSemaphoreGetStaticBuffer` (C *macro*), [1358](#)
- `xSemaphoreGive` (C *macro*), [1349](#)
- `xSemaphoreGiveFromISR` (C *macro*), [1351](#)
- `xSemaphoreGiveRecursive` (C *macro*), [1350](#)
- `xSemaphoreTake` (C *macro*), [1347](#)
- `xSemaphoreTakeFromISR` (C *macro*), [1352](#)
- `xSemaphoreTakeRecursive` (C *macro*), [1348](#)
- `xSTATIC_RINGBUFFER` (C++ *struct*), [1422](#)
- `xStreamBufferBytesAvailable` (C++ *function*), [1391](#)
- `xStreamBufferCreateStaticWithCallback` (C *macro*), [1393](#)
- `xStreamBufferCreateWithCallback` (C *macro*), [1392](#)
- `xStreamBufferCreateWithCaps` (C++ *function*), [1430](#)
- `xStreamBufferGetStaticBuffers` (C++ *function*), [1385](#)
- `xStreamBufferIsEmpty` (C++ *function*), [1390](#)
- `xStreamBufferIsFull` (C++ *function*), [1390](#)
- `xStreamBufferReceive` (C++ *function*), [1388](#)
- `xStreamBufferReceiveCompletedFromISR` (C++ *function*), [1392](#)
- `xStreamBufferReceiveFromISR` (C++ *function*), [1389](#)
- `xStreamBufferReset` (C++ *function*), [1390](#)
- `xStreamBufferSend` (C++ *function*), [1386](#)
- `xStreamBufferSendCompletedFromISR` (C++ *function*), [1391](#)
- `xStreamBufferSendFromISR` (C++ *function*), [1387](#)
- `xStreamBufferSetTriggerLevel` (C++ *function*), [1391](#)
- `xStreamBufferSpacesAvailable` (C++ *function*), [1391](#)
- `xTASK_STATUS` (C++ *struct*), [1317](#)
- `xTASK_STATUS::eCurrentState` (C++ *mem-*

- ber*), [1317](#)
- `xTASK_STATUS::pcTaskName` (C++ *member*), [1317](#)
- `xTASK_STATUS::pxStackBase` (C++ *member*), [1317](#)
- `xTASK_STATUS::ulRunTimeCounter` (C++ *member*), [1317](#)
- `xTASK_STATUS::usStackHighWaterMark` (C++ *member*), [1317](#)
- `xTASK_STATUS::uxBasePriority` (C++ *member*), [1317](#)
- `xTASK_STATUS::uxCurrentPriority` (C++ *member*), [1317](#)
- `xTASK_STATUS::xCoreID` (C++ *member*), [1318](#)
- `xTASK_STATUS::xHandle` (C++ *member*), [1317](#)
- `xTASK_STATUS::xTaskNumber` (C++ *member*), [1317](#)
- `xTaskAbortDelay` (C++ *function*), [1301](#)
- `xTaskCallApplicationTaskHook` (C++ *function*), [1309](#)
- `xTaskCatchUpTicks` (C++ *function*), [1316](#)
- `xTaskCheckForTimeOut` (C++ *function*), [1315](#)
- `xTaskCreate` (C++ *function*), [1296](#)
- `xTaskCreatePinnedToCore` (C++ *function*), [1425](#)
- `xTaskCreatePinnedToCoreWithCaps` (C++ *function*), [1427](#)
- `xTaskCreateStatic` (C++ *function*), [1297](#)
- `xTaskCreateStaticPinnedToCore` (C++ *function*), [1425](#)
- `xTaskCreateWithCaps` (C++ *function*), [1427](#)
- `xTaskDelayUntil` (C++ *function*), [1300](#)
- `xTaskGenericNotifyStateClear` (C++ *function*), [1314](#)
- `xTaskGenericNotifyWait` (C++ *function*), [1312](#)
- `xTaskGetApplicationTaskTag` (C++ *function*), [1308](#)
- `xTaskGetApplicationTaskTagFromISR` (C++ *function*), [1308](#)
- `xTaskGetCoreID` (C++ *function*), [1426](#)
- `xTaskGetCurrentTaskHandleForCore` (C++ *function*), [1426](#)
- `xTaskGetHandle` (C++ *function*), [1307](#)
- `xTaskGetIdleTaskHandle` (C++ *function*), [1309](#)
- `xTaskGetIdleTaskHandleForCore` (C++ *function*), [1426](#)
- `xTaskGetStaticBuffers` (C++ *function*), [1307](#)
- `xTaskGetTickCount` (C++ *function*), [1307](#)
- `xTaskGetTickCountFromISR` (C++ *function*), [1307](#)
- `xTaskNotifyAndQueryIndexed` (C macro), [1320](#)
- `xTaskNotifyAndQueryIndexedFromISR` (C macro), [1321](#)
- `xTaskNotifyGiveIndexed` (C macro), [1321](#)
- `xTaskNotifyIndexed` (C macro), [1318](#)
- `xTaskNotifyIndexedFromISR` (C macro), [1320](#)
- `xTaskNotifyStateClear` (C macro), [1323](#)
- `xTaskNotifyStateClearIndexed` (C macro), [1323](#)
- `xTaskNotifyWait` (C macro), [1321](#)
- `xTaskNotifyWaitIndexed` (C macro), [1321](#)
- `xTaskResumeAll` (C++ *function*), [1306](#)
- `xTaskResumeFromISR` (C++ *function*), [1305](#)
- `xTimerChangePeriod` (C macro), [1368](#)
- `xTimerChangePeriodFromISR` (C macro), [1374](#)
- `xTimerCreate` (C++ *function*), [1359](#)
- `xTimerCreateStatic` (C++ *function*), [1361](#)
- `xTimerDelete` (C macro), [1369](#)
- `xTimerGetExpiryTime` (C++ *function*), [1366](#)
- `xTimerGetPeriod` (C++ *function*), [1366](#)
- `xTimerGetReloadMode` (C++ *function*), [1366](#)
- `xTimerGetStaticBuffer` (C++ *function*), [1367](#)
- `xTimerGetTimerDaemonTaskHandle` (C++ *function*), [1364](#)
- `xTimerIsTimerActive` (C++ *function*), [1364](#)
- `xTimerPendFunctionCall` (C++ *function*), [1365](#)
- `xTimerPendFunctionCallFromISR` (C++ *function*), [1364](#)
- `xTimerReset` (C macro), [1370](#)
- `xTimerResetFromISR` (C macro), [1375](#)
- `xTimerStart` (C macro), [1367](#)
- `xTimerStartFromISR` (C macro), [1372](#)
- `xTimerStop` (C macro), [1368](#)
- `xTimerStopFromISR` (C macro), [1373](#)