
Trusted Firmware-A Tests

unknown

May 23, 2023

CONTENTS

1	About	1
2	Getting Started	5
3	Processes & Policies	17
4	Framework Design	19
5	Implementing Tests	23
6	Porting	25
7	Change Log & Release Notes	33
8	License	59
9	Getting started	61

1.1 Feature Overview

This page provides an overview of the current TF-A Tests feature set. The *Change Log & Release Notes* document provides details of changes made since the last release.

1.1.1 Current Features

The following TF-A features are currently tested to some extent (this list is not exhaustive):

- SMC Calling Convention
- Power State Coordination Interface (PSCI)
- Software Delegated Exception Interface (SDEI)
- Performance Measurement Framework (PMF)
- Communication and interaction with the Test Secure Payload (TSP)
- Firmware update (or recovery mode)
- EL3 payload boot flow
- Secure partition support
- True Random Number Generator Firmware Interface (TRNG_FW)

These tests are not a compliance test suite for the Arm interface standards used in TF-A (such as PSCI).

They do not cover 100% of the TF-A code. The fact that all tests pass does not mean that TF-A is free of bugs.

They are not reference code. They should not be considered as the official way to test hardware/firmware features. Instead, they are provided as example code to experiment with and improve on.

1.1.2 Still to come

- Additional tests
- Support for new platforms
- Design improvements
- Stability improvements
- Enhancements to the test framework to make it easier to implement tests
- Fixes for known issues (detailed in *Change Log & Release Notes*)

Copyright (c) 2019-2023, Arm Limited. All rights reserved.

1.2 Platform Support

1.2.1 Juno Arm Development Platform

The AArch64 build of this release has been tested on variants r0, r1 and r2 of the *Juno Arm Development Platform*. The AArch32 build has only been tested on variant r0.

1.2.2 Armv8 Architecture Fixed Virtual Platforms

The AArch64 build has been tested on the following Armv8 Architecture Fixed Virtual Platforms (FVP):

- FVP_Base_AEMv8A-AEMv8A
- FVP_Base_Cortex-A35x4
- FVP_Base_Cortex-A57x4-A53x4
- FVP_Base_RevC-2xAEMv8A
- Foundation_Platform

The AArch32 build has been tested on the following FVPs:

- FVP_Base_Cortex-A32x4
- FVP_Base_RevC-2xAEMv8A

NOTE: Unless otherwise stated, the model version is version 11.6, build 45.

1.2.3 System Guidance for Infrastructure (SGI) Fixed Virtual Platforms

The AArch64 build has been tested on the following Fixed Virtual Platforms (FVP):

- FVP_CSS_SGI-575
- FVP_RD_N1Edge

NOTE:

- For FVP_CSS_SGI-575 and FVP_RD_N1Edge, internal version of the models were used.

Copyright (c) 2019, Arm Limited. All rights reserved.

1.3 Maintainers

Trusted Firmware-A Tests (TF-A Tests) is a community maintained project. All contributions are ultimately merged by the maintainers listed below.

Please note the maintainers' bandwidth is limited and contributions to this project will be reviewed and handled on a best-effort basis.

1.3.1 Maintainers List

- Alexei Fedorov <alexei.fedorov@arm.com>
 - Bipin Ravi <bipin.ravi@arm.com>
 - Dan Handley <dan.handley@arm.com>
 - Joanna Farley <joanna.farley@arm.com>
 - Manish Pandey <manish.pandey2@arm.com>
 - Mark Dykes <mark.dykes@arm.com>
 - Olivier Deprez <olivier.deprez@arm.com>
 - Sandrine Bailleux <sandrine.bailleux@arm.com>
 - Soby Mathew <soby.mathew@arm.com>
-

Copyright (c) 2018-2020, Arm Limited. All rights reserved.

1.4 Support & Contact

We welcome any feedback on TF-A Tests and there are several methods for providing it or for obtaining support.

1.4.1 Mailing Lists

Public mailing lists for TF-A Tests and the wider Trusted Firmware project are hosted on TrustedFirmware.org. The mailing lists can be used for general enquiries, enhancement requests and issue reports, or to follow and participate in technical or organizational discussions around the project. These discussions include design proposals, advance notice of changes and upcoming events.

The relevant list for the TF-A Tests project is [TF-A-Tests development](#)

You can see a [summary of all the lists](#) on the TrustedFirmware.org website.

1.4.2 Issue Tracker

Specific issues may be raised using the [issue tracker](#) on the TrustedFirmware.org website. Using this tracker makes it easy for the maintainers to prioritise and respond to your ticket.

Copyright (c) 2019-2022, Arm Limited. All rights reserved.

Copyright (c) 2019, Arm Limited. All rights reserved.

GETTING STARTED

2.1 Prerequisites & Requirements

This document describes the software and hardware requirements for building TF-A Tests for AArch32 and AArch64 target platforms.

It may be possible to build TF-A Tests with combinations of software and hardware that are different from those listed below. The software and hardware described in this document are officially supported.

2.1.1 Build Host

TF-A Tests may be built using a Linux build host machine with a recent Linux distribution. We have performed tests using Ubuntu 22.04 LTS (64-bit), but other distributions should also work fine, provided that the tools and libraries can be installed.

2.1.2 Toolchain

Install the required packages to build TF-A Tests with the following command:

```
sudo apt-get install device-tree-compiler build-essential git perl libxml-libxml-perl
```

Download and install the GNU cross-toolchain from Arm. The TF-A Tests have been tested with version 12.2.Rel1 (gcc 12.2):

- [GCC cross-toolchain](#)

In addition, the following optional packages and tools may be needed:

- For debugging, Arm [Development Studio \(Arm-DS\)](#).

Copyright (c) 2019-2022, Arm Limited. All rights reserved.

2.2 Building Documentation

To create a rendered copy of this documentation locally you can use the [Sphinx](#) tool to build and package the plain-text documents into HTML-formatted pages.

If you are building the documentation for the first time then you will need to check that you have the required software packages, as described in the *Prerequisites* section that follows.

Note: An online copy of the documentation is available at <https://trustedfirmware-a-tests.readthedocs.io>, if you want to view a rendered copy without doing a local build.

2.2.1 Prerequisites

For building a local copy of the documentation you will need, at minimum:

- Python 3 (3.5 or later)
- PlantUML (1.2017.15 or later)

You must also install the Python modules that are specified in the `requirements.txt` file in the root of the `docs` directory. These modules can be installed using `pip3` (the Python Package Installer). Passing this requirements file as an argument to `pip3` automatically installs the specific module versions required by TF-A Tests.

An example set of installation commands for Ubuntu 18.04 LTS follows, assuming that the working directory is `docs`:

```
sudo apt install python3 python3-pip plantuml
pip3 install [--user] -r requirements.txt
```

Note: Several other modules will be installed as dependencies. Please review the list to ensure that there will be no conflicts with other modules already installed in your environment.

Passing the optional `--user` argument to `pip3` will install the Python packages only for the current user. Omitting this argument will attempt to install the packages globally and this will likely require the command to be run as root or using `sudo`.

Note: More advanced usage instructions for *pip* are beyond the scope of this document but you can refer to the [pip homepage](#) for detailed guides.

2.2.2 Building rendered documentation

The documentation can be built into HTML-formatted pages from the project's root directory by running the following command.

```
make doc
```

Output from the build process will be placed in:

```
docs/build/html
```

We also support building documentation in other formats. From the docs directory of the project, run the following command to see the supported formats. It is important to note that you will not get the correct result if the command is run from the project's root directory, as that would invoke the top-level Makefile for TF-A Tests themselves.

```
make help
```

2.2.3 Building rendered documentation from a container

There may be cases where you can not either install or upgrade required dependencies to generate the documents, so in this case, one way to create the documentation is through a docker container. The first step is to check if `docker` is installed in your host, otherwise check main docker page for installation instructions. Once installed, run the following script from project root directory

```
docker run --rm -v $PWD:/TF sphinxdoc/sphinx \
    bash -c 'cd /TF && \
    pip3 install plantuml -r ./docs/requirements.txt && make doc'
```

The above command fetches the `sphinxdoc/sphinx` container from [docker hub](https://hub.docker.com/r/sphinxdoc/sphinx), launches the container, installs documentation requirements and finally creates the documentation. Once done, exit the container and output from the build process will be placed in:

```
docs/build/html
```

Copyright (c) 2020, Arm Limited. All rights reserved.

2.3 Obtaining Source Code

Download the TF-A Tests source code using the following command:

```
git clone https://git.trustedfirmware.org/TF-A/tf-a-tests.git
```

Copyright (c) 2019, Arm Limited. All rights reserved.

2.4 Building TF-A Tests

- Before building TF-A Tests, the environment variable `CROSS_COMPILE` must point to the cross compiler.

For AArch64:

```
export CROSS_COMPILE=<path-to-aarch64-gcc>/bin/aarch64-none-elf-
```

For AArch32:

```
export CROSS_COMPILE=<path-to-aarch32-gcc>/bin/arm-eabi-
```

- Change to the root directory of the TF-A Tests source tree and build.

For AArch64:

```
make PLAT=<platform>
```

For AArch32:

```
make PLAT=<platform> ARCH=aarch32
```

Notes:

- If PLAT is not specified, fvp is assumed by default. See the [TF-A documentation](#) for more information on available build options.
- By default this produces a release version of the build. To produce a debug version instead, build the code with DEBUG=1.
- The build process creates products in a build/ directory tree, building the objects and binaries for each test image in separate sub-directories. The following binary files are created from the corresponding ELF files:

- * build/<platform>/<build-type>/tftf.bin
- * build/<platform>/<build-type>/ns_bl1u.bin
- * build/<platform>/<build-type>/ns_bl2u.bin
- * build/<platform>/<build-type>/el3_payload.bin
- * build/<platform>/<build-type>/cactus_mm.bin
- * build/<platform>/<build-type>/cactus.bin
- * build/<platform>/<build-type>/ivy.bin

where <platform> is the name of the chosen platform and <build-type> is either debug or release. The actual number of images might differ depending on the platform.

Refer to the sections below for more information about each image.

- Build products for a specific build variant can be removed using:

```
make DEBUG=<D> PLAT=<platform> clean
```

... where <D> is 0 or 1, as specified when building.

The build tree can be removed completely using:

```
make realclean
```

- Use the following command to list all supported build commands:

```
make help
```

2.4.1 TFTF test image

`tftf.bin` is the main test image to exercise the TF-A features. The other test images provided in this repository are optional dependencies that TFTF needs to test some specific features.

`tftf.bin` may be built independently of the other test images using the following command:

```
make PLAT=<platform> tftf
```

In TF-A boot flow, `tftf.bin` replaces the BL33 image and should be injected in the FIP image. This might be achieved by running the following command from the TF-A root directory:

```
BL33=<path/to/tftf.bin> make PLAT=<platform> fip
```

Please refer to the [TF-A documentation](#) for further details.

2.4.2 Realm payload test image

`realm.bin` is the realm payload test image and is packaged along with `tftf` for Realm Management Extension (RME) testing. This can be built using the following command:

```
make PLAT=<platform> realm
```

The generated `realm.bin` needs to be packaged as part of `tftf.bin` to be used as a single BL33 image and can be done using the following command:

```
make PLAT=<platform> pack_realm
```

Please refer to the [TF-A RME documentation](#) for build and run instructions.

2.4.3 NS_BL1U and NS_BL2U test images

`ns_bl1u.bin` and `ns_bl2u.bin` are test images that exercise the *Firmware Update (FWU)* feature of TF-A¹. Throughout this document, they will be referred as the *FWU test images*.

In addition to updating the firmware, the FWU test images also embed some tests that exercise the FWU state machine implemented in the TF-A. They send valid and invalid SMC requests to the TF-A BL1 image in order to test its robustness.

NS_BL1U test image

The NS_BL1U image acts as the *Application Processor (AP) Firmware Update Boot ROM*. This typically is the first software agent executing on the AP in the Normal World during a firmware update operation. Its primary purpose is to load subsequent firmware update images from an external interface, such as NOR Flash, and communicate with BL1 to authenticate those images.

The NS_BL1U test image provided in this repository performs the following tasks:

- Load FWU images from external non-volatile storage (typically flash memory) to Non-Secure RAM.
- Request TF-A BL1 to copy these images in Secure RAM and authenticate them.
- Jump to NS_BL2U which carries out the next steps in the firmware update process.

¹ Therefore, the Trusted Board Boot feature must be enabled in TF-A for the FWU test images to work. Please refer the [TF-A documentation](#) for further details.

This image may be built independently of the other test images using the following command:

```
make PLAT=<platform> ns_bl1u
```

NS_BL2U test image

The NS_BL2U image acts as the *AP Firmware Updater*. Its primary responsibility is to load a new set of firmware images from an external interface and write them into non-volatile storage.

The NS_BL2U test image provided in this repository overrides the original FIP image stored in flash with the backup FIP image (see below).

This image may be built independently of the other test images using the following command:

```
make PLAT=<platform> ns_bl2u
```

Putting it all together

The FWU test images should be used in conjunction with the TFTP image, as the latter initiates the FWU process by corrupting the FIP image and resetting the target. Once the FWU process is complete, TFTP takes over again and checks that the firmware was successfully updated.

To sum up, 3 images must be built out of the TF-A Tests repository in order to test the TF-A Firmware Update feature:

- ns_bl1u.bin
- ns_bl2u.bin
- tftf.bin

Once that's done, they must be combined in the right way.

- ns_bl1u.bin is a standalone image and does not require any further processing.
- ns_bl2u.bin must be injected into the FWU_FIP image. This might be achieved by setting NS_BL2U=ns_bl2u.bin when building the FWU_FIP image out of the TF-A repository. Please refer to the section Building FIP images with support for Trusted Board Boot in the [TF-A documentation](#).
- tftf.bin must be injected in the standard FIP image, as explained in section [TFTP test image](#).

Additionally, on Juno platform, the FWU FIP must contain a SCP_BL2U image. This image can simply be a copy of the standard SCP_BL2 image if no specific firmware update operations need to be carried on the SCP side.

Finally, the backup FIP image must be created. This can simply be a copy of the standard FIP image, which means that the Firmware Update process will restore the original, uncorrupted FIP image.

2.4.4 EL3 test payload

el3_payload.bin is a test image exercising the alternative EL3 payload boot flow in TF-A. Refer to the [EL3 test payload README file](#) for more details about its behaviour and how to build and run it.

2.4.5 SPM test images

This repository contains three sample Secure Partitions (SP) meant to be used with one implementation of a Secure Partition Manager (SPM):

- Cactus-MM
- Cactus and Ivy

They are only supported on AArch64 FVP. They can be built independently of the other test images using the following command:

```
make PLAT=fvp cactus ivy cactus_mm
```

To run the full set of tests in the Secure Partitions, they should be used in conjunction with the TFTP image.

Please refer to the [TF-A documentation](#) for further details.

Cactus-MM

Cactus-MM is designed to test the TF-A EL3 SPM implementation ([TF-A Secure Partition Manager \(MM\)](#)) based on the [Arm Management Mode Interface \(MM\)](#)

This SP runs in Secure-EL0 and performs the following tasks:

- Test that TF-A has correctly setup the secure partition environment: it should be allowed to perform cache maintenance operations, access floating point registers, etc.
- Test that TF-A accepts to change data access permissions and instruction permissions on behalf of the Secure Partition for memory regions the latter owns.
- Test communication with SPM through MM interface.

In the TF-A boot flow, the partition replaces the BL32 image and should be injected in the FIP image. To test SPM-MM with Cactus-MM, it is enough to use `cactus_mm.bin` as BL32 image.

For SPM-MM, build TF-A following [Building TF-A Secure Partition Manager \(MM\)](#) and the following commands can be used to build the tests:

```
# TF-A-Tests repository:
make PLAT=fvp TESTS=spm-mm tftf cactus_mm
```

Cactus and Ivy

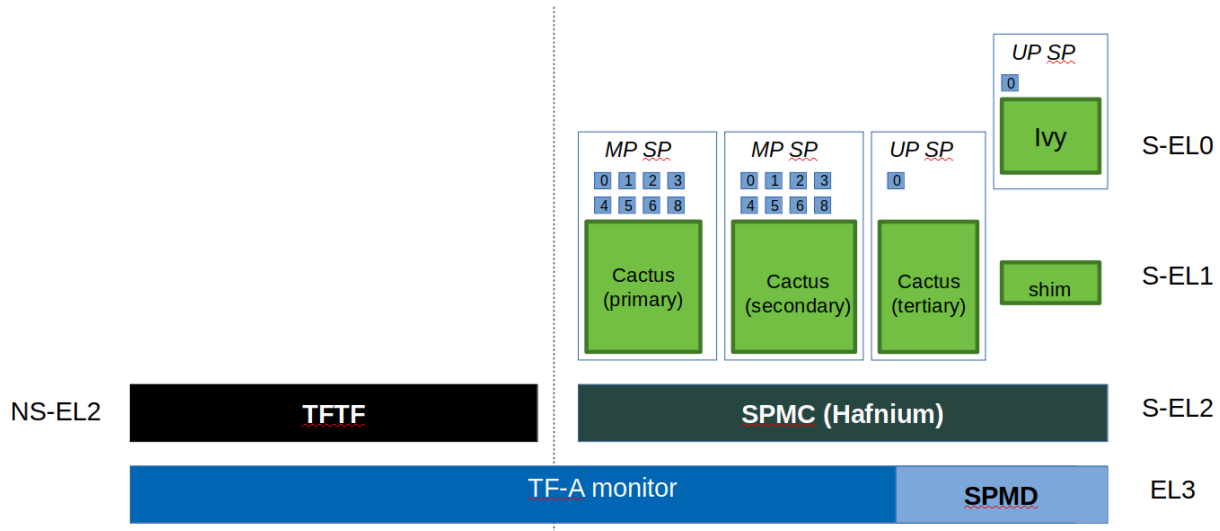
Cactus and Ivy are designed to test the FF-A based SPM implementation with secure virtualization enabled. Refer to [Arm Firmware Framework for Armv8-A](#)

In the TF-A reference code base, BL31 implements the SPMD and BL32 the SPMC. The SPMC runs at S-EL2 and acts as a partition manager for multiple secure partitions ([TF-A Secure Partition Manager \(FF-A\)](#)):

- Cactus is a sample FF-A compliant S-EL1 partition. As a matter of providing a realistic test harness, three instances of the same partition binary are launched as separate SPs (hence assigned three different FF-A IDs corresponding each to a different secure partition). Each secure partition instance has a separate manifest ([Cactus sample manifest](#), [Cactus secondary manifest](#), [Cactus tertiary manifest](#)). First two instances are MP SPs. Third instance is a UP SP. Each instance runs a set of built-in tests at boot time. They exercise SP to SPMC FF-A interfaces contained in the secure world. The partition interacts with the SPMC through SMC. Once the NWd and TFTP are started, another set of run-time tests exercise the normal world to secure world primitives.

- Ivy is a specific kind of S-EL1 UP partition, where the S-EL1 exception level consists of a thin shim layer. The applicative part of the partition is held at S-EL0. The shim provides early bootstrap code, MMU configuration and a vector table trapping S-EL0 requests. The application interacts with the shim through FF-A protocol by the use of SVC instruction. The shim relays the request to the SPMC by an SMC. The S-EL0 application doesn't require knowledge of the shim, and can be self contained.

This picture illustrates the test setup:



To build TFTF with SPM tests, Cactus and Ivy use:

```
# TF-A-Tests repository:
make PLAT=fvp TESTS=spm tftf cactus ivy
```

Copyright (c) 2019-2021, Arm Limited. All rights reserved.

2.5 Build Options Summary

As far as possible, TF-A Tests dynamically detects the platform hardware components and available features. There are a few build options to select specific features where the dynamic detection falls short.

Unless mentioned otherwise, these options are expected to be specified at the build command line and are not to be modified in any component makefiles.

Note: The build system doesn't track dependencies for build options. Therefore, if any of the build options are changed from a previous build, a clean build must be performed.

2.5.1 Common (Shared) Build Options

Most of the build options listed in this section apply to TFTP, the FWU test images and Cactus, unless otherwise specified. These do not influence the EL3 payload, whose simplistic build system is mostly independent.

- **ARCH:** Choose the target build architecture for TF-A Tests. It can take either `aarch64` or `aarch32` as values. By default, it is defined to `aarch64`. Not all test images support this build option.
- **ARM_ARCH_FEATURE:** Optional Arm Architecture build option which specifies one or more feature modifiers. This option has the form `[no]feature+...` and defaults to `none`. It translates into compiler option `-march=armvX[.Y]-a+[no]feature+...`. See compiler's documentation for the list of supported feature modifiers.
- **ARM_ARCH_MAJOR:** The major version of Arm Architecture to target when compiling TF-A Tests. Its value must be numeric, and defaults to 8.
- **ARM_ARCH_MINOR:** The minor version of Arm Architecture to target when compiling TF-A Tests. Its value must be a numeric, and defaults to 0.
- **BRANCH_PROTECTION:** Numeric value to enable ARMv8.3 Pointer Authentication (ARMv8.3-PAuth) and ARMv8.5 Branch Target Identification (ARMv8.5-BTI) support in the Trusted Firmware-A Test Framework itself. If enabled, it is needed to use a compiler that supports the option `-mbranch-protection` (GCC 9 and later). Selects the branch protection features to use:
 - 0: Default value turns off all types of branch protection
 - 1: Enables all types of branch protection features
 - 2: Return address signing to its standard level
 - 3: Extend the signing to include leaf functions
 - 4: Turn on branch target identification mechanism

The table below summarizes **BRANCH_PROTECTION** values, GCC compilation options and resulting PAuth/BTI features.

Value	GCC option	PAuth	BTI
0	<code>none</code>	N	N
1	<code>standard</code>	Y	Y
2	<code>pac-ret</code>	Y	N
3	<code>pac-ret+leaf</code>	Y	N
4	<code>bti</code>	N	Y

This option defaults to 0 and this is an experimental feature.

- **DEBUG:** Chooses between a debug and a release build. A debug build typically embeds assertions checking the validity of some assumptions and its output is more verbose. The option can take either 0 (release) or 1 (debug) as values. 0 is the default.
- **ENABLE_ASSERTIONS:** This option controls whether calls to `assert()` are compiled out.
 - For debug builds, this option defaults to 1, and calls to `assert()` are compiled in.
 - For release builds, this option defaults to 0 and calls to `assert()` are compiled out.

This option can be set independently of **DEBUG**. It can also be used to hide any auxiliary code that is only required for the assertion and does not fit in the assertion itself.

- **LOG_LEVEL:** Chooses the log level, which controls the amount of console log output compiled into the build. This should be one of the following:

```
0 (LOG_LEVEL_NONE)
10 (LOG_LEVEL_ERROR)
20 (LOG_LEVEL_NOTICE)
30 (LOG_LEVEL_WARNING)
40 (LOG_LEVEL_INFO)
50 (LOG_LEVEL_VERBOSE)
```

All log output up to and including the selected log level is compiled into the build. The default value is 40 in debug builds and 20 in release builds.

- **PLAT:** Choose a platform to build TF-A Tests for. The chosen platform name must be a subdirectory of any depth under `plat/`, and must contain a platform makefile named `platform.mk`. For example, to build TF-A Tests for the Arm Juno board, select `PLAT=juno`.
- **V:** Verbose build. If assigned anything other than 0, the build commands are printed. Default is 0.

2.5.2 Arm FVP Platform Specific Build Options

- **FVP_CLUSTER_COUNT :** Configures the cluster count to be used to build the topology tree within TFTP. By default TFTP is configured for dual cluster for CPUs with single thread (ST) and single cluster for SMT CPUs. For ST CPUs this option can be used to override the default number of clusters with a value in the range 1-4.
- **FVP_MAX_CPUS_PER_CLUSTER:** Sets the maximum number of CPUs implemented in a single cluster. This option defaults to the maximum value of 4 for ST CPUs and maximum value of 8 for SMT CPUs.
- **FVP_MAX_PE_PER_CPU:** Sets the maximum number of PEs implemented on any CPU in the system. This option defaults to 1 to select ST CPUs. For platforms with SMT CPUs this value must be set to 2.

2.5.3 TFTP-specific Build Options

- **NEW_TEST_SESSION:** Choose whether a new test session should be started every time or whether the framework should determine whether a previous session was interrupted and resume it. It can take either 1 (always start new session) or 0 (resume session as appropriate). 1 is the default.
- **TESTS:** Set of tests to run. Use the following command to list all possible sets of tests:

```
make help_tests
```

If no set of tests is specified, the standard tests will be selected (see `tftp/tests/tests-standard.xml`).

- **USE_NVM:** Used to select the location of test results. It can take either 0 (RAM) or 1 (non-volatile memory like flash) as test results storage. Default value is 0, as writing to the flash significantly slows tests down.

2.5.4 Realm payload specific Build Options

- **TFTP_MAX_IMAGE_SIZE:** The option needs to be either set by the user or by the platform makefile to specify the maximum size of TFTP binary. This is needed so that the Realm payload binary can be appended to TFTP binary via `make pack_realm` build command.

2.5.5 FWU-specific Build Options

- **FIRMWARE_UPDATE**: Whether the Firmware Update test images (i.e. NS_BL1U and NS_BL2U) should be built. The default value is 0. The platform makefile is free to override this value if Firmware Update is supported on this platform.

Copyright (c) 2019-2020, Arm Limited. All rights reserved.

2.6 Running Tests

Refer to the [Juno and FVP platform documentation](#) in the *TF-A documentation*. The same instructions mostly apply to running the TF-A Tests on those two platforms. The difference is that the following images are not needed here:

- Normal World bootloader. The TFTP replaces it in the boot flow;
- Linux Kernel;
- Device tree;
- Filesystem.

In other words, only the following software images are needed:

- BL1 firmware image;
- FIP image containing the following images:
 - BL2;
 - SCP_BL2 if required by the platform (e.g. Juno);
 - BL31;
 - BL32 (optional);
 - `tftf.bin` (standing as the BL33 image).

2.6.1 Running Manual Tests on FVP

The manual tests rely on storing state in non-volatile memory (NVM) across reboot. On FVP the NVM is not persistent across reboots, so the following flag must be used to write the NVM to a file when the model exits.

```
-C bp.flashloader0.fnameWrite=[filename]
```

To ensure the model exits on shutdown the following flag must be used:

```
-C bp.ve_sysregs.exit_on_shutdown=1
```

After the model has been shutdown, this file must be fed back in to continue the test. Note this flash file includes the FIP image, so the original `fip.bin` does not need to be passed in. The following flag is used:

```
-C bp.flashloader0.fname=[filename]
```

2.6.2 Running Firmware Update (FWU) Tests

As previously mentioned in *Putting it all together*, there are a couple of extra images involved when running the FWU tests. They need to be loaded at the right addresses, which depend on the platform.

On FVP

In addition to the usual BL1 and FIP images, the following extra images must be loaded:

- NS_BL1U image at address `0x0BEB8000` (i.e. `NS_BL1U_BASE` macro in TF-A)
- FWU_FIP image at address `0x08400000` (i.e. `NS_BL2U_BASE` macro in TF-A)
- Backup FIP image at address `0x09000000` (i.e. `FIP_BKP_ADDRESS` macro in TF-A tests).

An example script is provided in `scripts/run_fw_fvp.sh`.

On Juno

The same set of extra images and load addresses apply for Juno as for FVP.

The new images must be programmed in flash memory by adding some entries in the `SITE1/HBI0262x/images.txt` configuration file on the Juno SD card (where `x` depends on the revision of the Juno board). Refer to the [Juno Getting Started Guide](#), section 2.3 “Flash memory programming” for more information. Users should ensure these do not overlap with any other entries in the file.

Addresses in this file are expressed as an offset from the base address of the flash (that is, `0x08000000`).

```
NOR10UPDATE: AUTO                ; Image Update:NONE/AUTO/FORCE
NOR10ADDRESS: 0x00400000          ; Image Flash Address
NOR10FILE: \SOFTWARE\fwu_fip.bin  ; Image File Name
NOR10LOAD: 00000000              ; Image Load Address
NOR10ENTRY: 00000000             ; Image Entry Point

NOR11UPDATE: AUTO                ; Image Update:NONE/AUTO/FORCE
NOR11ADDRESS: 0x03EB8000          ; Image Flash Address
NOR11FILE: \SOFTWARE\ns_bluu.bin  ; Image File Name
NOR11LOAD: 00000000              ; Image Load Address
NOR11ENTRY: 00000000             ; Image Load Address

NOR12UPDATE: AUTO                ; Image Update:NONE/AUTO/FORCE
NOR12ADDRESS: 0x01000000          ; Image Flash Address
NOR12FILE: \SOFTWARE\backup_fip.bin ; Image File Name
NOR12LOAD: 00000000              ; Image Load Address
NOR12ENTRY: 00000000             ; Image Entry Point
```

Copyright (c) 2019, Arm Limited. All rights reserved.

This document describes how to build the Trusted Firmware-A Tests (TF-A Tests) and run them on a set of platforms. It assumes that the reader has previous experience building and running [Trusted Firmware-A \(TF-A\)](#).

Copyright (c) 2019-2020, Arm Limited. All rights reserved.

PROCESSES & POLICIES

3.1 Checking source code style

When making changes to the source for submission to the project, the source must be in compliance with the Linux style guide. To assist with this, the project Makefile provides two targets, which both utilise the `checkpatch.pl` script that ships with the Linux source tree.

To check the entire source tree, you must first download copies of `checkpatch.pl`, `spelling.txt` and `const_structs.checkpatch` available in the [Linux master tree](#) scripts directory, then set the `CHECKPATCH` environment variable to point to `checkpatch.pl` (with the other 2 files in the same directory).

Then use the following command:

```
make CHECKPATCH=<path-to-linux>/linux/scripts/checkpatch.pl checkcodebase
```

To limit the coding style checks to your local changes, use:

```
make CHECKPATCH=<path-to-linux>/linux/scripts/checkpatch.pl checkpatch
```

By default, this will check all patches between `origin/master` and your local branch. If you wish to use a different reference commit, this can be specified using the `BASE_COMMIT` variable.

Copyright (c) 2019, Arm Limited. All rights reserved.

Copyright (c) 2019, Arm Limited. All rights reserved.

FRAMEWORK DESIGN

This document provides some details about the internals of the TF-A Tests design. It is incomplete at the moment.

4.1 High-Level Behaviour

The EL3 firmware is expected to hand over to the TF-A tests with all secondary cores powered down, i.e. only the primary core should enter the TF-A tests.

The primary CPU initialises the platform and the TF-A tests internal data structures.

Then the test session begins. The TF-A tests are executed one after the other. Tests results are saved in non-volatile memory as we go along.

Once all tests have completed, a report is printed over the serial console.

4.2 Global Code Structure

The code is organised into the following categories (present as directories at the top level or under the `tftf/` directory):

- **Drivers.**

Some examples follow, this list might not be exhaustive.

- Generic GIC driver.

`arm_gic.h` contains the public APIs that tests might use. Both GIC architecture versions 2 and 3 are supported.

- PL011 UART driver.
- VExpress NOR flash driver.

Note that tests are not expected to use this driver in most cases. Instead, they should use the `tftf_nvm_read()` and `tftf_nvm_write()` wrapper APIs. See definitions in `tftf/framework/include/nvm.h`. See also the NVM validation test cases (`tftf/tests/framework_validation_tests/test_validation_nvm.c`) for an example of usage of these functions.

- SP805 watchdog.

Used solely to generate an interrupt that will reset the system on purpose (used in `tftf_plat_reset()`).

- SP804 timer.

This is used as the system timer on Juno. It is configured such that an interrupt is generated when it reaches 0. It is programmed in one-shot mode, i.e. it must be rearmed every time it reaches 0.

- **Framework.**

Core features of the test framework.

- **Library code.**

Firstly, there is `include/libc/` which provides standard C library functions like `memcpy()`, `printf()` and so on. Additionally, various other APIs are provided under `include/lib/`. The below list gives some examples but might not be exhaustive.

- `aarch64/`

Architecture helper functions for e.g. system registers access, cache maintenance operations, MMU configuration, ...

- `events.h`

Events API. Used to create synchronisation points between CPUs in tests.

- `irq.h`

IRQ handling support. Used to configure IRQs and register/unregister handlers called upon reception of a specific IRQ.

- `power_management.h`

Power management operations (CPU ON/OFF, CPU suspend, etc.).

- `sgi.h`

Software Generated Interrupt support. Used as an inter-CPU communication mechanism.

- `spinlock.h`

Lightweight implementation of synchronisation locks. Used to prevent concurrent accesses to shared data structures.

- `timer.h`

Support for programming the timer. Any timer which is in the *always-on* power domain can be used to exit CPUs from suspend state.

- `tftf_lib.h`

Miscellaneous helper functions/macros: MP-safe `printf()`, low-level PSCI wrappers, insertion of delays, raw SMC interface, support for writing a string in the test report, macros to skip tests on platforms that do not meet topology requirements, etc.

- `io_storage.h`

Low-level IO operations. Tests are not expected to use these APIs directly. They should use higher-level APIs like `tftf_nvmm_read()` and `tftf_nvmm_write()`.

- **Platform specific.**

Note that `include/plat/common/plat_topology.h` provides the interfaces that a platform must implement to support topology discovery (i.e. how many CPUs and clusters there are).

- **Tests.**

The tests are divided into the following categories (present as directories in the `tftf/tests/` directory):

- **Framework validation tests.**

Tests that exercise the core features of the framework. Verify that the test framework itself works properly.

- **Runtime services tests.**

Tests that exercise the runtime services offered by the EL3 Firmware to the Normal World software. For example, this includes tests for the Standard Service (to which PSCI belongs to), the Trusted OS service or the SiP service.

- **CPU extensions tests.**

Tests some CPU extensions features. For example, the AMU tests ensure that the counters provided by the Activity Monitor Unit are behaving correctly.

- **Firmware Update tests.**

Tests that exercise the [Firmware Update](#) feature of TF-A.

- **Template tests.**

Sample test code showing how to write tests in practice. Serves as documentation.

- **Performance tests.**

Simple tests measuring the latency of an SMC call.

- **Miscellaneous tests.**

Tests for RAS support, correct system setup, ...

All assembler files have the `.S` extension. The linker source file has the extension `.ld.S`. This is processed by GCC to create the linker script which has the extension `.ld`.

4.3 Detailed Code Structure

The cold boot entry point is `tftf_entrpoint` (see `tftf/framework/aarch64/entrpoint.S`). As explained in [High-Level Behaviour](#), only the primary CPU is expected to execute this code.

Tests can power on other CPUs using the function `tftf_cpu_on()`. This uses the PSCI CPU_ON API of the EL3 Firmware. When entering the Normal World, execution starts at the warm boot entry point, which is `tftf_hotplug_entry()` (see `tftf/framework/aarch64/entrpoint.S`).

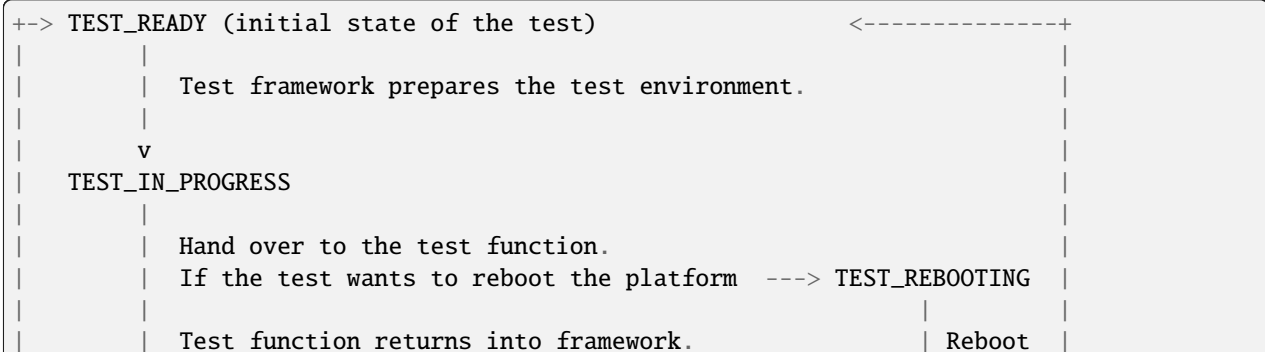
Information about the progression of the test session and tests results are written into Non-Volatile Memory as we go along. This consists of the following data (see struct `tftf_state_t` typedef in `tftf/framework/include/nvm.h`):

- `test_to_run`

Reference to the test to run.

- `test_progress`

Progress in the execution of `test_to_run`. This is used to implement the following state machine:



(continues on next page)

(continued from previous page)

```
| |                                     | |
```

```
| |                                     +-----+
```

```
| v
```

```
| TEST_COMPLETE
```

```
| |
```

```
| |   Do some framework management.
```

```
| |   Move to next test.
```

```
+-----+
```

- testcase_buffer

A buffer that the test can use as a scratch area for whatever it is doing.

- testcase_results

- result_buffer_size

- result_buffer

Buffer holding the tests output. Tests output are concatenated.

4.4 Interrupt Management

The TF-A tests expect SGIs #0 to #7 to be available for their own usage. In particular, this means that Trusted World software must configure them as non-secure interrupts.

SGI #7 has a special status. It is the SGI that the timer management framework sends to all CPUs when the system timer fires off (see the definition of the constant `IRQ_WAKE_SGI` in the header file `include/lib/irq.h`). Although test cases can use this specific SGI - e.g. they can register an IRQ handler for it and use it as an inter-CPU communication mechanism - they have to be aware of the underlying consequences. Some tests, like the PSCI CPU_SUSPEND tests, rely on this SGI to be enabled in order to wake up CPUs from their suspend state. If it is disabled, these tests will leave the system in an unresponsive state.

Copyright (c) 2018-2019, Arm Limited. All rights reserved.

IMPLEMENTING TESTS

This document aims at providing some pointers to help implementing new tests in the TFTF image.

5.1 Test Structure

A test might be divided into 3 logical parts, detailed in the following sections.

5.1.1 Prologue

A test has a main entry point function, whose type is:

```
typedef test_result_t (*test_function_t)(void);
```

See `tftf/framework/include/tftf.h`.

Only the primary CPU enters this function, while other CPUs are powered down.

First of all, the test function should check whether this test is applicable to this platform and environment. Some tests rely on specific hardware features or firmware capabilities to be present. If these are not available, the test should be skipped. For example, a multi-core test requires at least 2 CPUs to run. Macros and functions are provided in `include/common/test_helpers.h` to help test code verify that their requirements are met.

5.1.2 Core

This is completely dependent on the purpose of the test. The paragraphs below just provide some useful, general information.

The primary CPU may power on other CPUs by calling the function `tftf_cpu_on()`. It provides an address to which secondary CPUs should jump to once they have been initialized by the test framework. This address should be different from the primary CPU test function.

Synchronization primitives are provided in `include/lib/events.h` in case CPUs' execution threads need to be synchronized. Most multi-processing tests will need some synchronisation points that all/some CPUs need to reach before test execution may continue.

Any CPU that is involved in a test must return from its test function. Failure to do so will put the framework in an unrecoverable state, see the [Change Log & Release Notes](#) for details on this and other known limitations. The return code indicates the test result from the point of view of this CPU. At the end of the test, individual CPU results are aggregated and the overall test result is derived from that. A test is considered as passed if all involved CPUs reported a success status code.

5.1.3 Epilogue

Each test is responsible for releasing any allocated resources and putting the system back in a clean state when it finishes. Any change to the system configuration (e.g. MMU setup, GIC configuration, system registers, ...) must be undone and the original configuration must be restored. This guarantees that the next test is not affected by the actions of the previous one.

One exception to this rule is that CPUs powered on as part of a test must not be powered down. As already stated above, as soon as a CPU enters the test, the framework expects it to return from the test.

5.2 Template Test Code

Some template test code is provided in `tftf/tests/template_tests`. It can be used as a starting point for developing new tests. Template code for both single-core and multi-core tests is provided.

5.3 Build System Integration

All test code is located under the `tftf/tests` directory. Tests are usually divided into categories represented as sub-directories under `tftf/tests/`.

The source file implementing the new test code should be added to the appropriate tests makefile, see `.*mk` files under `tftf/tests`.

The new test code should also appear in a tests manifest, see `*.xml` files under `tftf/tests`. A unique name and test function must be provided. An optional description may be provided as well.

For example, to create a test case named “Foo test case”, whose test function is `foo()`, add the following line in the tests manifest:

```
<testcase name="Foo test case" function="foo" />
```

A testcase must be part of a testsuite. The testcase XML node above must be inside a testsuite XML node. A unique name and a description must be provided for the testsuite.

For example, to create a test suite named “Bar test suite”, whose description is: “An example test suite”, add the following 2 lines:

```
<testsuite name="Bar test suite" description="An example test suite">
</testsuite>
```

See the template test manifest for reference: `tftf/tests/tests-template.xml`.

Copyright (c) 2018-2020, Arm Limited. All rights reserved.

6.1 Platform Requirements

The TF-A Tests rely on the following features to be present on the platform and accessible from Normal World.

- Watchdog
- Non-Volatile Memory
- System Timer

This also means that a platform port of the TF-A Tests must include software drivers for those features.

Copyright (c) 2019, Arm Limited. All rights reserved.

6.2 Storage Abstraction Layer

In order to improve platform independence and portability a storage abstraction layer is used to store test results to non-volatile platform storage.

Each platform should register devices and their drivers via the storage layer. These drivers then need to be initialized using the `tftf_platform_setup()` function.

Warning: It is mandatory to implement at least one storage driver.

For the FVP and Juno platforms the NOR Flash driver is provided as the default means to store test results to storage. The storage layer is described in the header file `include/lib/io_storage.h`. The implementation of the common library is in `drivers/io/io_storage.c` and the driver files are located in `drivers/io/`.

Copyright (c) 2019, Arm Limited. All rights reserved.

6.3 Build Flags

The platform may define build flags to control inclusion or exclusion of certain tests. These flags must be defined in the platform makefile which is included by the build system.

- **PLAT_TESTS_SKIP_LIST**

This build flag can be defined by the platform to control exclusion of some testcases from the default test plan for a platform. If used this needs to point to a text file which follows the following criteria:

- Contain a list of tests to skip for this platform.
- Specify 1 test per line, using the following format:

`testsuite_name/testcase_name`

where `testsuite_name` and `testcase_name` are the names that appear in the XML tests file.

- Alternatively, it is possible to disable a test suite entirely, which will disable all test cases part of this test suite. To do so, only specify the test suite name, omitting the `/testcase_name` part.

Copyright (c) 2019, Arm Limited. All rights reserved.

6.4 Mandatory Modifications

6.4.1 File : platform_def.h

Each platform must ensure that a header file of this name is in the system include path with the following constants defined. This may require updating the list of `PLAT_INCLUDES` in the `platform.mk` file. In the ARM FVP port, this file is found in `plat/arm/board/fvp/include/platform_def.h`.

- **#define : PLATFORM_LINKER_FORMAT**

Defines the linker format used by the platform, for example `elf64-littleaarch64` used by the FVP.

- **#define : PLATFORM_LINKER_ARCH**

Defines the processor architecture for the linker by the platform, for example `aarch64` used by the FVP.

- **#define : PLATFORM_STACK_SIZE**

Defines the stack memory available to each CPU. This constant is used by `plat/common/aarch64/platform_mp_stack.S`.

- **#define : PLATFORM_CLUSTER_COUNT**

Defines the total number of clusters implemented by the platform in the system.

- **#define : PLATFORM_CORE_COUNT**

Defines the total number of CPUs implemented by the platform across all clusters in the system.

- **#define : PLATFORM_NUM_AFFS**

Defines the total number of nodes in the affinity hierarchy at all affinity levels used by the platform.

- **#define : PLATFORM_MAX_AFFLVL**

Defines the maximum number of affinity levels in the system that the platform implements. ARMv8-A has support for 4 affinity levels. It is likely that hardware will implement fewer affinity levels. For example, the Base

AEM FVP implements two clusters with a configurable number of CPUs. It reports the maximum affinity level as 1.

- **#define : PLAT_MAX_SPI_OFFSET_ID**

Defines the offset of the last Shared Peripheral Interrupt supported by the TF-A Tests on this platform. SPI numbers are mapped onto GIC interrupt IDs, starting from interrupt ID 32. In other words, this offset ID corresponds to the last SPI number, to which 32 must be added to get the corresponding last GIC IRQ ID.

E.g. If PLAT_MAX_SPI_OFFSET_ID is 10, this means that IRQ #42 is the last SPI.

- **#define : PLAT_LOCAL_PSTATE_WIDTH**

Defines the bit-field width of the local state in State-ID field of the power-state parameter. This macro will be used to compose the State-ID field given the local power state at different affinity levels.

- **#define : PLAT_MAX_PWR_STATES_PER_LVL**

Defines the maximum number of power states at a power domain level for the platform. This macro will be used by the PSCI_STAT_COUNT/RESIDENCY tests to determine the size of the array to allocate for storing the statistics.

- **#define : TFTF_BASE**

Defines the base address of the TFTF binary in DRAM. Used by the linker script to link the image at the right address. Must be aligned on a page-size boundary.

- **#define : IRQ_PCPU_NS_TIMER**

Defines the IRQ ID of the per-CPU Non-Secure timer of the platform.

- **#define : IRQ_CNTPSIRQ1**

Defines the IRQ ID of the System timer of the platform.

- **#define : TFTF_NVM_OFFSET**

The TFTF needs some Non-Volatile Memory to store persistent data. This defines the offset from the beginning of this memory that the TFTF can use.

- **#define : TFTF_NVM_SIZE**

Defines the size of the Non-Volatile Memory allocated for TFTF usage.

If the platform port uses the ARM Watchdog Module (SP805) peripheral, the following constant needs to be defined:

- **#define : SP805_WDOG_BASE**

Defines the base address of the SP805 watchdog peripheral.

If the platform port uses the IO storage framework, the following constants must also be defined:

- **#define : MAX_IO_DEVICES**

Defines the maximum number of registered IO devices. Attempting to register more devices than this value using `io_register_device()` will fail with `IO_RESOURCES_EXHAUSTED`.

- **#define : MAX_IO_HANDLES**

Defines the maximum number of open IO handles. Attempting to open more IO entities than this value using `io_open()` will fail with `IO_RESOURCES_EXHAUSTED`.

If the platform port uses the VExpress NOR flash driver (see `drivers/io/vexpress_nor/`), the following constants must also be defined:

- **#define : NOR_FLASH_BLOCK_SIZE**

Defines the largest block size as seen by the software while writing to NOR flash.

6.4.2 Function : `tftf_plat_arch_setup()`

Argument : void
Return : void

This function performs any platform-specific and architectural setup that the platform requires.

In both the ARM FVP and Juno ports, this function configures and enables the MMU.

6.4.3 Function : `tftf_early_platform_setup()`

Argument : void
Return : void

This function executes with the MMU and data caches disabled. It is only called by the primary CPU. It is used to perform platform-specific actions very early in the boot.

In both the ARM FVP and Juno ports, this function configures the console.

6.4.4 Function : `tftf_platform_setup()`

Argument : void
Return : void

This function executes with the MMU and data caches enabled. It is responsible for performing any remaining platform-specific setup that can occur after the MMU and data cache have been enabled.

This function is also responsible for initializing the storage abstraction layer used to access non-volatile memory for permanent storage of test results. It also initialises the GIC and detects the platform topology using platform-specific means.

6.4.5 Function : `plat_get_nvm_handle()`

Argument : `uintptr_t *`
Return : void

It is needed if the platform port uses IO storage framework. This function is responsible for getting the pointer to the initialised non-volatile memory entity.

6.4.6 Function : `tftf_plat_get_pwr_domain_tree_desc()`

Argument : void
Return : `const unsigned char *`

This function returns the platform topology description array in a suitable format as expected by TFTP. The size of the array is expected to be `PLATFORM_NUM_AFFS - PLATFORM_CORE_COUNT + 1`. The format used to describe this array is :

1. The first entry in the array specifies the number of power domains at the highest power level implemented in the platform. This caters for platforms where the power domain tree does not have a single root node e.g. the FVP which has two cluster power domains at the highest level (that is, 1).

2. Each subsequent entry corresponds to a power domain and contains the number of power domains that are its direct children.

The array format is the same as the one used by Trusted Firmware-A and more details of its description can be found in the [Trusted Firmware-A documentation](#).

6.4.7 Function : `tftf_plat_get_mpidr()`

Argument : unsigned `int`
Return : `uint64_t`

This function converts a given *core_pos* into a valid MPIDR if the CPU is present in the platform. The *core_pos* is a unique number less than the `PLATFORM_CORE_COUNT` returned by `platform_get_core_pos()` for a given CPU. This API is used by the topology framework in TFTP to query the presence of a CPU and, if present, returns the corresponding MPIDR for it. If the CPU referred to by the *core_pos* is absent, then this function returns `INVALID_MPIDR`.

6.4.8 Function : `plat_get_state_prop()`

Argument : unsigned `int`
Return : `const plat_state_prop_t *`

This functions returns the `plat_state_prop_t` array for all the valid low power states from platform for a specified affinity level and returns `NULL` for an invalid affinity level. The array is expected to be `NULL`-terminated. This function is expected to be used by tests that need to compose the power state parameter for use in `PSCI_CPU_SUSPEND` API or `PSCI_STAT/RESIDENCY` API.

6.4.9 Function : `plat_fwu_io_setup()`

Argument : `void`
Return : `void`

This function initializes the IO system used by the firmware update.

6.4.10 Function : `plat_arm_gic_init()`

Argument : `void`
Return : `void`

This function initializes the ARM Generic Interrupt Controller (GIC).

6.4.11 Function : `platform_get_core_pos()`

```
Argument : u_register_t  
Return   : unsigned int
```

This function returns a linear core ID from a MPID.

6.4.12 Function : `plat_crash_console_init()`

```
Argument : void  
Return   : int
```

This function initializes a platform-specific console for crash reporting.

6.4.13 Function : `plat_crash_console_putc()`

```
Argument : int  
Return   : int
```

This function prints a character on the platform-specific crash console.

6.4.14 Function : `plat_crash_console_flush()`

```
Argument : void  
Return   : int
```

This function waits until all the characters of the platform-specific crash console have been actually printed.

Copyright (c) 2019, Arm Limited. All rights reserved.

6.5 Optional Modifications

The following are helper functions implemented by the test framework that perform common platform-specific tasks. A platform may choose to override these definitions.

6.5.1 Function : `platform_get_stack()`

```
Argument : unsigned long  
Return   : unsigned long
```

This function returns the base address of the memory stack that has been allocated for the CPU specified by MPIDR. The size of the stack allocated to each CPU is specified by the platform defined constant `PLATFORM_STACK_SIZE`.

Common implementation of this function is provided in `plat/common/aarch64/platform_mp_stack.S`.

6.5.2 Function : `tftf_platform_end()`

Argument : void
Return : void

This function performs any operation required by the platform to properly finish the test session.

The default implementation sends an EOT (End Of Transmission) character on the UART. This can be used to automatically shutdown the FVP models. When running on real hardware, the UART output may be parsed by an external tool looking for this character and rebooting the platform for example.

6.5.3 Function : `tftf_plat_reset()`

Argument : void
Return : void

This function resets the platform.

The default implementation uses the ARM watchdog peripheral ([SP805](#)) to generate a watchdog timeout interrupt. This interrupt remains deliberately unserviced, which eventually asserts the reset signal.

Copyright (c) 2019, Arm Limited. All rights reserved.

Porting the TF-A Tests to a new platform involves making some mandatory and optional modifications for both the cold and warm boot paths. Modifications consist of:

- Implementing a platform-specific function or variable,
- Setting up the execution context in a certain way, or
- Defining certain constants (for example `#defines`).

The platform-specific functions and variables are all declared in `include/plat/common/platform.h`. The framework provides a default implementation of variables and functions to fulfill the optional requirements. These implementations are all weakly defined; they are provided to ease the porting effort. Each platform port can override them with its own implementation if the default implementation is inadequate.

Copyright (c) 2019, Arm Limited. All rights reserved.

CHANGE LOG & RELEASE NOTES

Please note that the Trusted Firmware-A Tests version follows the Trusted Firmware-A version for simplicity. At any point in time, TF-A Tests version *x.y* aims at testing TF-A version *x.y*. Different versions of TF-A and TF-A Tests are not guaranteed to be compatible. This also means that a version upgrade on the TF-A-Tests side might not necessarily introduce any new feature.

7.1 Version 2.9

7.1.1 New features

- More tests are made available in this release to help validate the functionalities in the following areas:
 - FF-A Features
 - Realm Management Extension
 - New Architecture Specific features related to v8.8
 - New platform ports

TFTF

- FF-A testing:
 - Reordered logs in the memory sharing tests.
 - Memory share bumped to v1.1 EAC0.
 - Updated tests for FFA_FEATURES(FFA_MEM_RETRIEVE_REQ).
 - Fixed issues with RXTX buffer unmapping and dependencies on tests.
 - Added check for execution state property of partitions.
- New tests:
 - Tests for Errata management firmware interface.
 - Ensure FPU state registers context is preserved in RL/SE/NS.
 - Modified FEAT_HCX test to also check for HCRX_EL2 init value.
 - Added basic SME2 tests.
 - PSCI tests for OS-initiated mode.
 - Added “nop” test to be used in conjunction with TFX.

- Introduced capability to generate Sync External Aborts (EA) in TFTP.
- New test to generate an SError.
- Tests to check whether the PMU is functional and if the state is preserved when switching worlds. PMEV-TYPER.evtCount width extended.
- Added support for more SPE versions.
- Platforms:
 - RD-N2-Cfg3:
 - * Added TFTP support.
- Miscellaneous:
 - SIMD/FPU save/restore routine moved to common lib.
 - Updated toolchain requirements documentation.
 - Update SME/Mortlach tests.
 - Unified Firmware First handling of lower EL EA.
 - Moved trusted wdog API to spm_common.
 - Added the ability to skip tests for AArch32.
 - Added config file to allow doc defaults be changed.
 - Modified tests for FEAT_RNG_TRAP.
 - Moved 'Stress test timer framework' to a new test suite 'tests-timer-stress'.
 - Support for new binutils versions.
 - Removed deprecated SPM libs and test code.

Realm Management Extension (RME)

- Added helper macro for RME tests.
- Test Secure interrupt can preempt Realm EL1.
- Added PMU Realm tests.
- Added BP_OPTION to REALM_CFLAGS to allow build realm payload with BTI/Pauth support.
- Fixed build issues introduced by the TFTP Realm extension enhancement tests.
- Test case return codes updated according to RMM Bet0 specification.
- Fixed build problem related to rmi_rec_enter verbose log.
- Added randomization of SMC RMI commands parameters and checking of X4-X7 return values as per SMCCC v1.2.

Cactus (Secure-EL1 test partition)

- Use of FFA_CONSOLE_LOG for debug logs.
- Test for consecutive same memory donation to other partitions.
- Now validating NWd can't share forbidden addresses.
- Support for registering irq handlers.
- Fixed attributes for NS memory region.
- Removal of memory regions not page-aligned.
- Added check for core linear id matching id passed by SPMC.

7.1.2 Issues resolved since last release

- Build issue for older toolchains and other failures resolved.
- Dropped invalid access test from CI.
- Now checking that the PMU is supported before using any of it.
- Use of write instead of read to generate an SError to avoid sync exceptions instead.
- Fixed broken link to TRNG_FW documentation.
- INIT_TFTF_MAILBOX() is called first for the invalid access test.

7.2 Version 2.8

7.2.1 New features

- More tests are made available in this release to help validate the functionalities in the following areas:
 - FF-A Features
 - Realm Management Extension
 - New Architecture Specific features related to v8.8
 - New platform ports

TFTF

- FF-A testing:
 - UUID included in partition information descriptors.
 - Checks for size of partition information descriptors.
 - Renamed FFA_MSG_RUN ABI function to FFA_RUN and allowed it to return from Waiting state.
 - Made ffa_tests available for Ivy.
 - Updated verbose message log structure.
 - Prevented generate_json.sh from being called more than once by requiring a list of partitions to be supplied.
 - Added a temporary workaround for unexpected affinity info state to prevent a system panic.

- Added test to exercise FFA_CONSOLE_LOG ABI.
- FF-A v1.1 Secure interrupts
 - * Added managed exit to first and second SP in call chain.
 - * Added test to exercise managed exit by two SPs in a call chain.
 - * Added tests to exercise NS interrupt being queued and signaled to SP.
- New tests:
 - Tests for SVE operations in Normal World and discover SVE vector length.
 - Added cleanup TRNG service tests.
 - Added test for SMCCC_ARCH_WORKAROUND_3.
 - Updated PAuth helpers to support QARMA3 algorithm.
 - Added tests for RNG_TRAP.
- Platforms:
 - SGI:
 - * Introduced platform variant build option.
 - * Re-organized header files.
 - * Migrated to secure uart port for routing tftf logs.
 - N1SDP:
 - * Added TFTP support for N1SDP.
 - RD-N2:
 - * Added TFTP support for RD-N2.
 - RD-N2-Cfg1:
 - * Added TFTP support for RD-N2-Cfg1.
 - RD-V1:
 - * Added TFTP support for RD-V1.
- Miscellaneous:
 - Added a missing ISB instruction in SME test.
 - Refactor to make some helper functions re-usable.
 - Updated build command to clean EL3 payload image.
 - Move renaming of the primary dts file for ivy partitions.
 - Added check that verifies if a platform supports el3_payload before building it.
 - Updated memory share test to meet Hafnium specification.
 - Updated toolchain requirements documentation.

Realm Management Extension (RME)

- Added Realm payload management capabilities to TFTP to act as a NS Host.
- Added test to verify that RMM and SPM can co-exist and work properly.
- Added function to reset delegated buffers to non-delegated state.
- Re-used existing `wait_for_non_lead_cpus()` function helper.
- Refactored RMI FID macros to simplify usage.
- Added userguide for realm payload testing.

Cactus (Secure-EL1 test partition)

- Corrected some tests message types from ERROR to VERBOSE.
- Increased the cactus number of xlat to allow the use of 48b PA size for memory sharing between SPs.
- Introduced a new direct request message command to resume after managed exit.
- Skip enabling virtual maintenance interrupts explicitly.
- Allowed sender to resume interrupted target vCPU.
- Added support for handling managed exit through vIRQ.
- Added support for discovering interrupt IDs of managed exit signals.
- Specified action in response to NS interrupt in manifest.

Ivy (Secure-EL0 test partition)

- Allowed testing using VHE.
- Allowed Ivy partitions to use `ffa_helpers` functions.
- Requirement of common name for Ivy partitions for consistency.
- Specified action in response to NS interrupt in manifest.

7.2.2 Issues resolved since last release

- Fixed SME header guard name.
- Fixed response for incorrect direct message request for FF-A.

7.3 Version 2.7

7.3.1 New features

- More tests are made available in this release to help validate the functionalities in the following areas:
 - FF-A Features
 - New Architecture Specific features related to v8.7
 - New platform port

TFTF

- FF-A testing:
 - FF-A partition information structure is updated to include UUIDs.
 - Memory Management helper functions are refactored to fetch the details of smc call failures in tftf and cactus.
 - Added test to validate memory sharing operations from SP to NS-endpoint are denied by SPMC.
 - Added test to ensure an endpoint that sets its version to v1.0 receives v1.0 partition information descriptors as defined in v1.0 FF-A specification.
 - Added test to validate that memory is cleared on memory sharing operations between normal world and secure world.
 - FF-A v1.1 Secure interrupts
 - * Added support to enhance the secure interrupt handling test.
 - * Support for registering and unregistering custom handler that is invoked by SP at the tail end of the virtual interrupt processing.
 - * Added support for querying the ID of the last serviced virtual interrupt.
- New tests:
 - Added test to validate that realm region access is being prevented from normal world.
 - Added test to validate that secure region access is being prevented from normal world.
 - Added test to validate that secure region access is being prevented from realm world.
 - Added test to validate that root region access is being prevented from realm world.
 - Added a test for v8.7 Advanced floating-point behavior (FEAT_AFP).
 - Added a SPE test that reads static profiling system registers of available SPE version i.e. FEAT_SPE/FEAT_SPEv1p1/FEAT_SPEv1p2.
 - Added a test to validate functionality of WFET and WFIT instructions introduced by v8.7 FEAT_WFxT.
 - Added basic SME tests to ensure feature enablement by EL3 is proper for its usage at lower non-secure ELs.
 - Added test to check Data Independent timing (DIT) field of PSTATE is retained on exception.
 - Added test to ensure that EL3 has properly enabled access to FEAT_BRBE from non-secure ELs.
- Platforms:
 - Add initial platform support for corstone1000.
 - TC:
 - * Support for notification in tertiary SP manifest.
 - FVP:
 - * Support to provide test memory addresses to validate the invalid memory access test from tftf(ns-el2).
- Miscellaneous:
 - Added support to configure the physical/virtual address space for FVP.
 - Added common header file for defining macros with size to support all the platforms.
 - Introduced handler for synchronous exceptions (AArch64).

- Added macros to extract the ISS portion of an ELx ESR exception syndrome register.
- Support to dynamically map/unmap test region to validate invalid memory access tests.
- Added support to receive boot information through secure partitions, according to the FF-A v1.1 EAC0 specification.
- Added an helper API function from SPM test suite to initialize FFA-mailbox and enable FF-A based message with SP.
- Updated the build string to display the rc-tagged version.

Cactus (Secure-EL1 test partition)

- Added test for nonsecure memory sharing between Secure Partitions(SPs).
- Added test to validate that a realm region cannot be accessed from secure world.
- Added test to permit checking a root region cannot be accessed from secure world.
- Extended the test command CACTUS_MEM_SEND_CMD to add support for memory sharing flags.
- Added support to save the state of general purpose registers x0-x4 at the entry to cold boot and restore them before jumping to entypoint of cactus.

7.3.2 Issues resolved since last release

- Fixed a bug to align RMI FIDs with SMCCC.
- Fixed encoding of vCPU and receiver IDs in the FFA_NOTIFICATION_GET interface to comply with the FF-A v1.1 beta0 specification.
- Fixed memory retrieve request attributes by enforcing them to be inner shareable rather than outer.
- Fixed static memory mapping of EL3 in EL2.
- Fixed a spurious error log message with memory share test.
- Aligning RMI FIDs with SMCCC.
- Fixed PSCI system suspend test suite execution in a four world system.
- Configured the build system to use DWARF 4 standard for debug builds with ArmDS.
- Introduced macro IRQ_TWDOG_INTID for the Tegra210, Tegra186 and Tegra194 platforms to fix the compilation failures.

7.4 Version 2.6

7.4.1 New features

- More tests are made available in this release to help validate the functionalities in the following areas:
 - Firmware Framework for Arm A-profile(FF-A)
 - Realm Management Extensions(RME)
 - Embedded Trace Extension and Trace Buffer Extension (ETE and TRBE)

TFTF

- FF-A testing:
 - Update FF-A version to v1.1
 - Added helpers for SPM tests to check partition info of SPs from normal world.
 - Added tests to check for ffa_features supported.
 - Added test for FFA_RXTX_UNMAP ABI.
 - Added test for FFA_SPM_ID_GET.
 - FF-A v1.1 Notifications
 - * Added test for notifications bitmap create and destroy ABIs.
 - * Added test for notifications set and get ABIs.
 - * Added test for notification INFO_GET ABI.
 - * Added test to check notifications pending interrupt is injected into and handled by the expected vCPU in a MP setup.
 - * Added test for signaling from MP SP to UP SP.
 - * Added test to check notifications interrupt IDs retrieved with FFA_FEATURES ABI.
 - * Added test to check functionality of notifications scheduled receiver interrupt.
 - FF-A v1.1 Secure interrupts
 - * Added support for handling secure interrupts in Cactus SP.
 - * Added several tests to exercise secure interrupt handling while SP is in WAITING/RUNNING/BLOCKED state.
- New tests:
 - Enabled SVE tests
 - Added test for trace system registers access.
 - Added test for trace filter control registers access.
 - Added test for trace buffer control registers access.
 - Added test to check PSTATE in SDEI handler.
 - Added test to check if HCRX_EL2 is accessible.
- Platforms:
 - TC0:
 - * Support for direct messaging with managed exit.
 - * Support for building S-EL0 Ivy partition.
 - FVP:
 - * Update Cactus secure partitions to indicate Managed exit support.
- Miscellaneous
 - Added random seed generation capability and ability to specify build parameters for SMC Fuzzer tool.

Cactus (Secure-EL1 test partition)

- Added helper for Cactus SP sleep.
- Added test commands to request use of notifications interfaces.
- Added several commands that generate direct message requests to assist in testing secure interrupt handling and notifications features in FF-A v1.1
- Added support for SP805 Trusted Watchdog module.

Ivy (Secure-EL1 test partition)

- Add shim layer to Ivy partition and enable PIE.
- Define Ivy partition manifest and use FF-A for message handling.
- Prepare S-EL1/0 environment for enabling S-EL0 application.

Realm Management Extension(RME)

- Added tests to run RMI and SPM on multiple CPUs concurrently.
- Added tests for multi CPU delegation and fail conditions.
- Added tests to query RMI version on multiple CPUs.

7.4.2 Issues resolved since last release

- Fixed Ivy partition start address for TC0.
- Fixed SP manifests to use little endian format UUID.
- Fixed a bug in memory sharing test for Cactus SP.
- Invalidate data cache for NS_BL1U and NS_BL2U images.
- Fixed attributes to Read-Write only for memory regions described in partition manifests.

7.5 Version 2.5

7.5.1 New features

- More tests are made available in this release to help validate the functionalities in the following areas:
 - True Random Number Generator (TRNG) test scenarios.
 - Multicore / Power State Controller Interface (PSCI) tests.
 - v8.6 Activity Monitors Unit (AMU) enhancements test scenarios.
 - **Secure Partition Manager (SPM) / Firmware Framework (FF-A) v1.0 testing.**
 - * Interrupt Handling between Non-secure and Secure world.
 - * Direct messages and memory sharing between Secure Partitions(SP).
 - * Many tests to exercise FF-A v1.0 ABIs.

- * SPM saving/restoring the NS SIMD context enabling a normal world FF-A endpoint (TFTF) and a secure partition to use SIMD vectors and instructions independently.

TFTF

- **SPM / FF-A v1.0 testing.**
 - **Refactor FF-A memory sharing tests**
 - * Created helper functions to initialize `ffa_memory_region` and to send the respective memory region to the SP, making it possible to reuse the logic in SP-to-SP memory share tests.
 - * Added comments to document relevant aspects about memory sharing.
 - **Trigger direct messaging between SPs.**
 - * Use cactus command `‘CACTUS_REQ_ECHO_SEND_CMD’` to make cactus SPs communicate with each other using direct message interfaces.
 - **Added helpers for SPM tests.**
 - * Checking SPMC has expected `FFA_VERSION`.
 - * Checking that expected FF-A endpoints are deployed in the system.
 - * Getting global TFTF mailbox.
- Replace `‘.inst’` AArch64 machine directives with CPU Memory Tagging Extension instructions in `‘test_mte_instructions’` function.
- **Add build option for Arm Feature Modifiers.**
 - This patch adds a new `ARM_ARCH_FEATURE` build option to add support for compiler’s feature modifiers.
- Enable 8 cores support for Theodul DSU(DynamiQ Shared Unit) for the Total Compute (TC0) platform.
- New tests:
 - **Remove redundant code and add better tests for TRNG SMCs.**
 - * Tests that the Version, Features, and RND calls conform to the spec.
 - **New tests for v8.6 AMU enhancements (FEAT_AMUv1p1)**
 - * Make sure AMU offsets are being saved and restored properly.
 - Tests to request SP-to-SP memory share.
 - **SP-to-SP direct messaging deadlock test.**
 - * TFTF sends `CACTUS_REQ_DEADLOCK_CMD` to cactus SP.

Cactus(Secure-EL1 test partition)

- Enable managed exit for primary cactus secure partition.
- Helper commands needed for interrupt testing.
- Add handler from managed exit FIQ interrupt.
- Make `ffa_id` global.
- Implement `HF_INTERRUPT_ENABLE` Hafnium hypervisor call wrapper. With this service, a secure partition calls into the SPMC to enable/disable a particular virtual interrupt.

- Invalidate the data cache for the cactus image.
- **Helper commands needed for interrupt testing.**
 - CACTUS_SLEEP_CMD & CACTUS_INTERRUPT_CMD added.
- **Decouple exception handling from tftf framework.**
 - With new interrupt related tests coming up in Cactus, added separate exception handler code for irq/fiq in Cactus.
- Hypervisor calls moved to a separate module.
- Add secondary entry point register function.
- Declare third SP instance as UP SP.
- Provision a cold boot path for secondary cores (or secondary pinned execution contexts).
- Tidy message loop, commands definitions, direct messaging API definitions.
- Helpers for error logging after FF-A calls.
- Properly placing Cactus test files.
- Tidying FF-A Memory Sharing tests.
- Use CACTUS_ECHO_CMD in direct message tests.
- **Refactor handling of commands.**
 - Added helper macros to define a command handler, build a command table in which each element is a pair of the handler and respective command ID. Available tests have been moved to their own command handler.
- **Extend arguments in commands responses.**
 - In the test commands framework, added template to extend number of values to include in a command response.
- **Check FF-A return is a valid direct response.**
 - Added a helper function to check if return of FFA_MSG_SEND_DIRECT_REQ is FFA_MSG_SEND_DIRECT_RESP.
- FFA_MSG_DIRECT_RESP call extended to use 5 registers.
- **Added accessors for arguments from FF-A calls.**
 - Some accessors for arguments from FF-A calls, namely for func id, error code, and direct message destination/source.
- **Use virtual counter for sp_sleep.**
 - Changes sp_sleep() to use virtual counter instead of physical counter.
- Checks if SIMD vectors are preserved in the normal world while transitioning from normal world to secure world and back to normal world.
- Tidying common code to tftf and cactus.
- Refactor cactus_test_cmds.h to incorporate static inline functions instead of macros to enforce type checking.
- Removed reference to Hafnium in name from helper function and macro to make them generic.
- For consistency added the cmd id 'CACTUS_MEM_SEND_CMD'.
- Add command to request memory sharing between SPs.

- Add & handle commands ‘CACTUS_REQ_ECHO_CMD’ and ‘CACTUS_ECHO_CMD’.
- Update README with list of sample partitions.
- Remove reference to PSA from xml test file.
- **Reduce tests verbosity in release mode.**
 - Update few NOTICE messages to VERBOSE/INFO.
- Fix conversion issues on cactus responses.
- Create RXTX map/configure helper macros and use them.
- **Update OP-TEE version used for testing to 3.10.**
 - SPMC as S-EL1 tests using OP-TEE depend on a static binary stored as a CI file. This binary corresponds to a build of OP-TEE v3.10.
- **Add uart2 to device-regions node.**
 - First SP no longer has an open access to the full system peripheral range and devices must be explicitly declared in the SP manifest.
- New tests:
 - Test for exercising SMMUv3 driver to perform stage2 translation.
 - Test handling of non-secure interrupt while running SP.
 - Add secondary cores direct messaging test for SPM.
 - **Testing deadlock by FF-A direct message.**
 - * Added command CACTUS_DEADLOCK_CMD to file cactus_test_cmds.h to create a deadlock scenario using FF-A direct message interfaces.
 - **Test SP-to-SP memory share operations**
 - * Handle ‘CACTUS_REQ_MEM_SEND_CMD’ by sending memory to the receiver SP.
 - Implemented test to validate FFA_RXTX_MAP ABI.

7.6 Version 2.4

7.6.1 New features

- More tests are made available in this release to help validate the functionalities in the following areas: - SMCCC. - New architecture specific features. - FF-A features. - New platform ports.
- Various improvements to test framework and test suite such as documentation, removing un-necessary dependencies, etc.

TFTF

- Remove dependencies from FVP to generic code by converting some FVP platform specific macros to the common macros.
- Remove make as a package dependency to compile TF-A test code.
- Move defaults values and macro defs in a separate folder from Makefile.
- Allow alternate stdout to be used apart from pl011 UART.
- Get FVP platform's topology from build options to make FVP platform configuration more flexible and eliminate test errors when the platform is configured with number of CPUs less than default values in the makefile.
- Update the FIP corrupt address which is used to corrupt BL2 image that helps to trigger firmware update process.
- Add explicit barrier before sev() in tftf_send_event_common API to avoid core hang.
- Align output properly on issuing make help_tests by removing dashes and sort tests.
- Moved a few FVP and Juno specific defined from common header files to platform specific header files.
- Replace SPCI with PSA FF-A in code as SPCI is now called as FF-A.
- Add owner field to sp_layout generation to differentiate owner of SP which could either be Silicon Provider or Platform provider.
- Add v8.5 Branch Target Identifier(BTI) support in TFTF.
- Remove dependency on SYS_CNT_BASE1 to read the memory mapped timers.
- Enables SError aborts for all CPUs, during their power on sequence.
- Documentation:
 - Use conditional assignment on sphinx variables so that they can be overwritten by environment and/or command line.
 - Add support for documentation build as a target in Makefile.
 - Update list of maintainers.
 - Update documentation to explain how to locally build the documentation.
 - Add .editorconfig from TF-A to define the coding style.
 - Fix documentation to include 'path/to' prefix when specifying tftf.bin on make fip cmd.
 - Use docker to build documentation.
 - Replace SPCI with PSA FF-A in documentation as SPCI is now called as FF-A.
- NVIDIA Tegra194:
 - Skip CPU suspend tests requiring SGI as wake source as Tegra194 platforms do not support CPU suspend power down and cannot be woken up with an SGI.
 - Disable some system suspend test cases.
 - Create dummy SMMU context for system resume to allow the System Resume Firmware to complete without any errors or warnings.
 - Increase RTC step value to 5ms as RTC consumes 250us for each register read/write. Increase the step value to 5ms to cover all the register read/write in program_timer().
 - Skip some timer framework validation tests as CPUs on Tegra194 platforms cannot be woken up with the RTC timer interrupt after power off.

- Introduce per-CPU Hypervisor Timer Interrupt ID.
- Skip PSCI STAT tests requiring `PSTATE_TYPE_POWERDOWN` as Tegra194 platforms do not support CPU suspend with state type as `PSTATE_TYPE_POWERDOWN`.
- Disable boot requirement tests as Tegra194 platforms do not support memory mapped timers.
- Skips the test “Create all power states and validate EL3 power state parsing” from the “EL3 power state parser validation” test suite as it is not in sync with this expectation.
- Moved reset, timers, wake, watchdog drivers from Tegra194 specific folder to common driver folder so that these drivers can be used for other NVIDIA platforms.
- New tests:
 - Add test for SDEI RM_ANY routing mode.
 - Add initial platform support for TC0.
 - Add SMC fuzzing module test.
 - Add test case for `SMCCC_ARCH_SOC_ID` feature.
 - Add test that supports ARMv8.6-FGT in TF-A.
 - Add test that supports ARMv8.6-ECV in TF-A.
 - Add test for `FFA_VERSION` interface.
 - Add test for `FFA_FEATURES` interface.
 - Add console driver for the TI UART 16550.
 - Add tests for FF-A memory sharing interfaces between tftf and cactus secure partitions.
 - NVIDIA Tegra194:
 - * Introduce platform port for Tegra194 to initialize the tftf framework and execute tests on the CPUs.
 - * Introduce power management support.
 - * Introduce support for RTC as wake source.
 - * Introduce system reset functionality test.
 - * Introduce watchdog timer test.
 - * Introduce support for NVIDIA Denver CPUs.
 - * Introduce RAS uncorrectable error injection test.
 - * Introduce tests to verify the Video Memory resize interface.
 - * Introduce test to inject RAS corrected errors for all supported nodes from all CPUs.
 - * Introduce a test to get return value from SMC SiP function `TEGRA_SIP_GET_SMMU_PER`.
 - NVIDIA Tegra196:
 - * Introduce initial support for Tegra186 platforms.
 - NVIDIA Tegra210:
 - * Introduce initial support for Tegra210 platforms.

Secure partition - Cactus

- TFTP doesn't need to boot Secondary Cactus as Hafnium now boots all partitions according to "boot-order" field value in the partition manifests.
- Remove test files related to deprecated SPCI Alpha specification and SPRT interface.
- Select different stdout device at runtime as primary VM can access to UART while secondary VM's use hypervisor call to SPM for debug logging.
- An SP maps its RX/TX buffers in its EL1&0 Stage-1 translation regime. The same RX/TX buffers are mapped by the SPMC in the SP's EL1&0 Stage-2 translation regime during boot time.
- Update memory/device region nodes in manifest. Memory region has 3 entries such as RX buffer, TX buffer and dummy. These memory region entries are mapped with attributes as "RX buffer: read-only", "TX buffer: read-write" and "dummy: read-write-execute". Device region mapped with read-write attribute.
- Create tertiary partition without RX_TX region specified to test the RXTX_MAP API.
- Add third partition to ffa_partition_info_get test to test that a partition can successfully get information about the third cactus partition.
- Map RXTX region to third partition to point the mailbox to this RXTX region.
- Adjust the number of EC context to max number of PEs as per the FF-A specification mandating that a SP must either "Implement as many ECs as the number of PEs (in case of a "multi-processor" SP with pinned contexts)" or "Implement a single EC (in case of a migratable "uni-processor" SP).
- Updated cactus test payload and TFTP ids as it is decided to have secure partition FF-A ids in the range from 0x8001 to 0xfffe, 0x8000 and 0xffff FF-A ids are reserved for the SPMC and the SPMD respectively and in the non-secure worlds, FF-A id 0 is reserved for the hypervisor and 1 to 0x7fff FF-A ids are reserved for VMs.
- Break the message loop on bad message request instead of replying with the FF-A error ABI to the SPMC.
- Remove deprecated hypervisor calls spm_vm_get_count and spm_vcpu_get_count. Instead use FFA_PARTITION_INFO_GET discovery ABI.
- Implement hvc call 'SPM_INTERRUPT_GET' to get interrupt id.
- Re-structure platform dependent files by moving platform dependent files and macros to platform specific folder.
- Adjust partition info get properties to support receipt of direct message request.
- New tests:
 - Add FFA Version Test.
 - Add FFA_FEATURES test.
 - Add FFA_MEM_SHARE test
 - Add FFA_MEM_LEND test.
 - Add FFA_MEM_DONATE test.
 - Add FFA_PARTITION_INFO_GET test.
 - Add exception/interrupt framework.
 - Add cactus support for TC0 platform.

7.6.2 Issues resolved since last release

- Update link to SMCCC specification.
- Trim down the top-level readme file to give brief overview of the project and also fix/update a number of broken/out-dated links in it.
- Bug fix in Multicore IRQ spurious test.
- Fix memory regions mapping with no NS bit set.
- Reenable PSCI NODE_HW_STATE test which was disabled earlier due to outdated SCP firmware.
- Fix AArch32 zeromem() function by avoiding infinite loop in 'zeromem' function and optimizing 'memcpy4' function.
- Add missing help_tests info on help target in the top-level Makefile.
- Trim down the readme file as it does not need to provide detailed information, instead it can simply be a landing page providing a brief overview of the project and redirecting the reader to RTD for further information.
- Fix maximum number of CPUs in DSU cluster by setting maximum number of CPUs in DSU cluster to 8.

7.7 Version 2.3

7.7.1 New features

- More tests are made available in this release to help validate the functionality of TF-A.
- CI upgraded to use GCC 9.2-2019.12 toolchain for tf-a-tests.
- Various improvements to test framework and test suite.

TFTF

- Support for extended register usage as per SMCCC v1.2 specification.
- Support for FVP platforms with SMT capabilities.
- Improved support for documentation through addition of basic Sphinx configuration and Makefile similar to TF-A repository.
- Enhancement to libc library synchronous to TF-A code base.
- ARMv8.3-PAuth enabled for all FWU tests in TFTP.
- TFTP made RFC 4122 compliant by converting UUIDs to network order format.
- Build improvement by deprecating custom AARCH64/AARCH32 macros in favor of __arch64__ macro provided by compiler.
- Support for HVC as a SMCCC conduit in TFTP.
- New tests:
 - AArch32 tests for checking if PMU counters leak in secure world.
 - Add new debug filesystem (debugfs) test.
 - Add a SPCI direct messaging test targeting bare-metal cactus SP.

Secure partitions

Cactus

- Several build improvements and symbol relocation fixup to make it position independent executable.
- Update of sample manifest to SPCI Beta1 format.
- Support for generating JSON file as required by TF-A.

7.7.2 Issues resolved since last release

- Makefile bug fix for performing parallel builds.
- Add missing D-cache invalidation of RW memory in tftf_entrpoint to safeguard against possible corruption.
- Fixes in GIC drivers to support base addresses beyond 4G range.
- Fix build with XML::LibXML 2.0202 Perl module

7.7.3 Known issues and limitations

The sections below list the known issues and limitations of each test image provided in this repository. Unless and otherwise stated, issues and limitations stated in previous release continue to exist in this release.

TFTF

- NODE_HW_STATE test has been temporarily disabled for sgi575 platform due to a dependency on SCP binaries version 2.5

7.8 Version 2.2

7.8.1 New features

- A wide range of tests are made available in this release to help validate the functionality of TF-A.
- Various improvements to test framework and test suite.

TFTF

- Enhancement to xlat table library synchronous to TF-A code base.
- Enabled strict alignment checks (SCTLR.A & SCTLR.SA) in all images.
- Support for a simple console driver. Currently it serves as a placeholder with empty functions.
- A topology helper API is added in the framework to get parent node info.
- Support for FVP with clusters having upto 8 CPUs.
- Enhanced linker script to separate code and RO data sections.
- Relax SMC calls tests. The SMCCC specification recommends Trusted OSES to mitigate the risk of leaking information by either preserving the register state over the call, or returning a constant value, such as zero, in each register. Tests only allowed the former behaviour and have been extended to allow the latter as well.

- Pointer Authentication enabled on warm boot path with individual APIAKey generation for each CPU.
- New tests:
 - Basic unit tests for xlat table library v2.
 - Tests for validating SVE support in TF-A.
 - Stress tests for dynamic xlat table library.
 - PSCI test to measure latencies when turning ON a cluster.
 - Series of AArch64 tests that stress the secure world to leak sensitive counter values.
 - Test to validate PSCI SYSTEM_RESET call.
 - Basic tests to validate Memory Tagging Extensions are being enabled and ensuring no undesired leak of sensitive data occurs.
- Enhanced tests:
 - Improved tests for Pointer Authentication support. Checks are performed to see if pointer authentication keys are accessible as well as validate if secure keys are being leaked after a PSCI version call or TSP call.
 - Improved AMU test to remove unexecuted code iterating over Group1 counters and fix the conditional check of AMU Group0 counter value.

Secure partitions

A new Secure Partition Quark is introduced in this release.

Quark

The Quark test secure partition provided is a simple service which returns a magic number. Further, a simple test is added to test if Quark is functional.

7.8.2 Issues resolved since last release

- Bug fix in libc memchr implementation.
- Bug fix in calculation of number of CPUs.
- Streamlined SMC WORKAROUND_2 test and fixed a false fail on Cortex-A76 CPU.
- Pointer Authentication support is now available for secondary CPUs and the corresponding tests are stable in this release.

7.8.3 Known issues and limitations

The sections below list the known issues and limitations of each test image provided in this repository. Unless and otherwise stated, issues and limitations stated in previous release continue to exist in this release.

TFTF

- Multicore spurious interrupt test is observed to have unstable behavior. As a temporary solution, this test is skipped for AArch64 Juno configurations.
- Generating SVE instructions requires *O3* compilation optimization. Since the current build structure does not allow compilation flag modification for specific files, the function which tests support for SVE has been pre-compiled and added as an assembly file.

7.9 Version 2.1

7.9.1 New features

- Add initial support for testing Secure Partition Client Interface (SPCI) and Secure Partition Run-Time (SPRT) standards.

Exercise the full communication flow throughout the software stack, involving:

- A Secure-EL0 test partition as the Trusted World agent.
 - TFTF as the Normal World agent.
 - The Secure Partition Manager (SPM) in TF-A.
- Various stability improvements, code refactoring and clean ups.

TFTF

- Reorganize tests build infrastructure to allow the selection of a subset of tests.
- Reorganize the platform layer for improved clarity and simplicity.
- Sanitise inclusion of drivers header files.
- Enhance the test report format for improved clarity and conciseness.
- Dump CPU registers when hitting an unexpected exception. Previously, this would silently loop forever.
- Import libc from TF-A to better align the two code bases.
- New tests:
 - SPM tests for exercising communication through either the MM or SPCI/SPRT interfaces.
 - SMC calling convention tests.
 - Initial tests for Armv8.3 Pointer Authentication support (experimental).
- New platform ports:
 - [Arm SGI-575 FVP](#).
 - Hikey960 board (experimental).
 - [Arm Neoverse Reference Design N1 Edge \(RD-N1-Edge\) FVP](#) (experimental).

Secure partitions

We now have 3 Secure Partitions to test the SPM implementation in TF-A.

Cactus-MM

The Cactus test secure partition provided in version 2.0 has been renamed into “*Cactus-MM*”. It is still responsible for testing the SPM implementation based on the Arm Management Mode Interface.

Cactus

This is a new test secure partition (as the former “*Cactus*” has been renamed into “*Cactus-MM*”, see above).

Unlike *Cactus-MM*, this image tests the SPM implementation based on the SPCI and SPRT draft specifications.

It runs in Secure-EL0 and performs the following tasks:

- Test that TF-A has correctly setup the secure partition environment (access to cache maintenance operations, to floating point registers, etc.)
- Test that TF-A accepts to change data access permissions and instruction permissions on behalf of Cactus for memory regions the latter owns.
- Test communication with SPM through SPCI/SPRT interfaces.

Ivy

This is also a new test secure partition. It is provided in order to test multiple partitions support in TF-A. It is derived from Cactus and essentially provides the same services but with different identifiers at the moment.

EL3 payload

- New platform ports:
 - [Arm SGI-575 FVP](#).
 - [Arm Neoverse Reference Design N1 Edge \(RD-N1-Edge\) FVP](#) (experimental).

7.9.2 Issues resolved since last release

- The GICv2 spurious IRQ test is no longer Juno-specific. It is now only GICv2-specific.
- The manual tests in AArch32 state now work properly. After investigation, we identified that this issue was not AArch32 specific but concerned any test relying on state information persisting across reboots. It was due to an incorrect build configuration.
- Cactus-MM now successfully links with GNU toolchain 7.3.1.

7.9.3 Known issues and limitations

The sections below lists the known issues and limitations of each test image provided in this repository.

TFTF

The TFTF test image might be conceptually sub-divided further in 2 parts: the tests themselves, and the test framework they are based upon.

Test framework

- Some stability issues.
- No mechanism to abort tests when they time out (e.g. this could be implemented using a watchdog).
- No convenient way to include or exclude tests on a per-platform basis.
- Power domains and affinity levels are considered equivalent but they may not necessarily be.
- Need to provide better support to alleviate duplication of test code. There are some recurrent test patterns for which helper functions should be provided. For example, bringing up all CPUs on the platform and executing the same function on all of them, or programming an interrupt and waiting for it to trigger.
- Every CPU that participates in a test must return from the test function. If it does not - e.g. because it powered itself off for testing purposes - then the test framework will wait forever for this CPU. This limitation is too restrictive for some tests.
- No protection against interrupted flash operations. If the target is reset while some data is written to flash, the test framework might behave incorrectly on reset.
- When compiling the code, if the generation of the `tests_list.c` and/or `tests_list.h` files fails, the build process is not aborted immediately and will only fail later on.
- The directory layout requires further improvements. Most of the test framework code has been moved under the `tftf/` directory to better isolate it but this effort is not complete. As a result, there are still some TFTF files scattered around.
- Pointer Authentication testing is experimental and incomplete at this stage. It is only enabled on the primary CPU on the cold boot.

Tests

- Some tests are implemented for AArch64 only and are skipped on AArch32.
- Some tests are not robust enough:
 - Some tests might hang in some circumstances. For example, they might wait forever for a condition to become true.
 - Some tests rely on arbitrary time delays instead of proper synchronization when executing order-sensitive steps.
 - Some tests have been implemented in a practical manner: they seem to work on actual hardware but they make assumptions that are not guaranteed by the Arm architecture. Therefore, they might fail on some other platforms.
- PSCI stress tests are very unreliable and will often hang. The root cause is not known for sure but this might be due to bad synchronization between CPUs.

- The GICv2 spurious IRQ test sometimes fails with the following error message:

SMC @ lead CPU returned 0xFFFFFFFF 0x8 0xC

The root cause is unknown.

- The FWU tests take a long time to complete. This is because they wait for the watchdog to reset the system. On FVP, TF-A configures the watchdog period to about 4 min. This limit is excessive for an automated testing context and leaves the user without feedback and unable to determine if the tests are proceeding properly.
- The test “Target timer to a power down cpu” sometimes fails with the following error message:

Expected timer switch: 4 Actual: 3

The root cause is unknown.

FWU images

- The FWU tests do not work on the revC of the Base AEM FVP. They only work on the revB.
- NS-BL1U and NS-BL2U images reuse TFTP-specific code for legacy reasons. This is not a clean design and may cause confusion.

Test secure partitions (Cactus, Cactus-MM, Ivy)

- This is experimental code. It’s likely to change a lot as the secure partition software architecture evolves.
- Supported on AArch64 FVP platform only.

All test images

- TF-A Tests are derived from a fork of TF-A so:
 - they’ve got some code in common but lag behind on some features.
 - there might still be some irrelevant references to TF-A.
- Some design issues. E.g. TF-A Tests inherited from the I/O layer of TF-A, which still needs a major rework.
- Cannot build TF-A Tests with Clang. Only GCC is supported.
- The build system does not cope well with parallel building. The user should not attempt to run multiple jobs in parallel with the `-j` option of *GNU make*.
- The build system does not properly track build options. A clean build must be performed every time a build option changes.
- UUIDs are not compliant to RFC 4122.
- No floating point support. The code is compiled with GCC flag `-mgeneral-regs-only`, which prevents the compiler from generating code that accesses floating point registers. This might limit some test scenarios.
- The documentation is too lightweight.
- Missing instruction barriers in some places before reading the system counter value. As a result, the CPU could speculatively read it and any delay loop calculations might be off (because based on stale values). We need to examine all such direct reads of the `CNTPCT_EL0` register and replace them with a call to `syscounter_read()` where appropriate.

7.10 Version 2.0

7.10.1 New features

This is the first public release of the Trusted Firmware-A Tests source code.

TFTF

- Provides a baremetal test framework to exercise TF-A features through its SMC interface.
- Integrates easily with TF-A: the TFTF binary is packaged in the FIP image as a BL33 component.
- Standalone binary that runs on the target without human intervention (except for some specific tests that require a manual target reset).
- Designed for multi-core testing. The various sub-frameworks allow maximum parallelism in order to stress the firmware.
- Displays test results on the UART output. This may then be parsed by an external tool and integrated in a continuous integration system.
- Supports running in AArch64 (NS-EL2 or NS-EL1) and AArch32 states.
- Supports parsing a tests manifest (XML file) listing the tests to include in the binary.
- Detects most platform features at run time (e.g. topology, GIC version, ...).
- Provides a topology enumeration framework. Allows tests to easily go through affinity levels and power domain nodes.
- Provides an event framework to synchronize CPU operations in a multi-core context.
- Provides a timer framework. Relies on a single global timer to generate interrupts for all CPUs in the system. This allows tests to easily program interrupts on demand to use as a wake-up event source to come out of CPU suspend state for example.
- Provides a power-state enumeration framework. Abstracts the valid power states supported on the platform.
- Provides helper functions for power management operations (CPU hotplug, CPU suspend, system suspend, ...) with proper saving of the hardware state.
- Supports rebooting the platform at the end of each test for greater independence between tests.
- Supports interrupting and resuming a test session. This relies on storing test results in non-volatile memory (e.g. flash).

FWU images

- Provides example code to exercise the Firmware Update feature of TF-A.
- Tests the robustness of the FWU state machine implemented in the TF-A by sending valid and invalid authentication, copy and image execution requests to the TF-A BL1 image.

EL3 test payload

- Tests the ability of TF-A to load an EL3 payload.

Cactus test secure partition

- Tests that TF-A has correctly setup the secure partition environment: it should be allowed to perform cache maintenance operations, access floating point registers, etc.
- Tests the ability of a secure partition to request changing data access permissions and instruction permissions of memory regions it owns.
- Tests the ability of a secure partition to handle StandaloneMM requests.

7.10.2 Known issues and limitations

The sections below lists the known issues and limitations of each test image provided in this repository.

TFTF

The TFTF test image might be conceptually sub-divided further in 2 parts: the tests themselves, and the test framework they are based upon.

Test framework

- Some stability issues.
- No mechanism to abort tests when they time out (e.g. this could be implemented using a watchdog).
- No convenient way to include or exclude tests on a per-platform basis.
- Power domains and affinity levels are considered equivalent but they may not necessarily be.
- Need to provide better support to alleviate duplication of test code. There are some recurrent test patterns for which helper functions should be provided. For example, bringing up all CPUs on the platform and executing the same function on all of them, or programming an interrupt and waiting for it to trigger.
- Every CPU that participates in a test must return from the test function. If it does not - e.g. because it powered itself off for testing purposes - then the test framework will wait forever for this CPU. This limitation is too restrictive for some tests.
- No protection against interrupted flash operations. If the target is reset while some data is written to flash, the test framework might behave incorrectly on reset.
- When compiling the code, if the generation of the `tests_list.c` and/or `tests_list.h` files fails, the build process is not aborted immediately and will only fail later on.
- The directory layout is confusing. Most of the test framework code has been moved under the `tftf/` directory to better isolate it but this effort is not complete. As a result, there are still some TFTF files scattered around.

Tests

- Some tests are implemented for AArch64 only and are skipped on AArch32.
- Some tests are not robust enough:
 - Some tests might hang in some circumstances. For example, they might wait forever for a condition to become true.
 - Some tests rely on arbitrary time delays instead of proper synchronization when executing order-sensitive steps.
 - Some tests have been implemented in a practical manner: they seem to work on actual hardware but they make assumptions that are not guaranteed by the Arm architecture. Therefore, they might fail on some other platforms.
- PSCI stress tests are very unreliable and will often hang. The root cause is not known for sure but this might be due to bad synchronization between CPUs.
- The GICv2 spurious IRQ test is Juno-specific. In reality, it should only be GICv2-specific. It should be reworked to remove any platform-specific assumption.
- The GICv2 spurious IRQ test sometimes fails with the following error message:


```
SMC @ lead CPU returned 0xFFFFFFFF 0x8 0xC
```

The root cause is unknown.
- The manual tests in AArch32 mode do not work properly. They save some state information into non-volatile memory in order to detect the reset reason but this state does not appear to be retained. As a result, these tests keep resetting infinitely.
- The FWU tests take a long time to complete. This is because they wait for the watchdog to reset the system. On FVP, TF-A configures the watchdog period to about 4 min. This is way too long in an automated testing context. Besides, the user gets not feedback, which may let them think that the tests are not working properly.
- The test “Target timer to a power down cpu” sometimes fails with the following error message:


```
Expected timer switch: 4 Actual: 3
```

The root cause is unknown.

FWU images

- The FWU tests do not work on the revC of the Base AEM FVP. They only work on the revB.
- NS-BL1U and NS-BL2U images reuse TFTP-specific code for legacy reasons. This is not a clean design and may cause confusion.

Cactus test secure partition

- Cactus is experimental code. It’s likely to change a lot as the secure partition software architecture evolves.
- Fails to link with GNU toolchain 7.3.1.
- Cactus is supported on AArch64 FVP platform only.

All test images

- TF-A Tests are derived from a fork of TF-A so:
 - they've got some code in common but lag behind on some features.
 - there might still be some irrelevant references to TF-A.
- Some design issues. E.g. TF-A Tests inherited from the I/O layer of TF-A, which still needs a major rework.
- Cannot build TF-A Tests with Clang. Only GCC is supported.
- The build system does not cope well with parallel building. The user should not attempt to run multiple jobs in parallel with the `-j` option of *GNU make*.
- The build system does not properly track build options. A clean build must be performed every time a build option changes.
- SMCCC v2 is not properly supported.
- UUIDs are not compliant to RFC 4122.
- No floating point support. The code is compiled with GCC flag `-mgeneral-regs-only`, which prevents the compiler from generating code that accesses floating point registers. This might limit some test scenarios.
- The documentation is too lightweight.

Copyright (c) 2018-2022, Arm Limited. All rights reserved.

LICENSE

The software is provided under a BSD-3-Clause license.

Copyright (c) <year> <owner>. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

8.1 SPDX Identifiers

Individual files contain the following tag instead of the full license text.

SPDX-License-Identifier: BSD-3-Clause

This enables machine processing of license information based on the SPDX License Identifiers that are here available:
<http://spdx.org/licenses/>

8.2 Other Projects

This project contains code from other projects as listed below. The original license text is included in those source files.

- The libc source code is derived from [FreeBSD](#) and [SCC](#). FreeBSD uses various BSD licenses, including BSD-3-Clause and BSD-2-Clause. The SCC code is used under the BSD-3-Clause license with the author's permission.
- The [LLVM compiler-rt](#) source code is disjunctively dual licensed (NCSA OR MIT). It is used by this project under the terms of the NCSA license (also known as the University of Illinois/NCSA Open Source License), which is a permissive license compatible with BSD-3-Clause. Any contributions to this code must be made under the terms of both licenses.

The Trusted Firmware-A Tests (TF-A-Tests) is a suite of baremetal tests to exercise the [Trusted Firmware-A \(TF-A\)](#) features from the Normal World. It enables strong TF-A functional testing without dependency on a Rich OS. It mainly interacts with TF-A through its SMC interface.

It provides a basis for TF-A developers to validate their own platform ports and add their own test cases.

GETTING STARTED

Get the TF-A Tests source code from trustedfirmware.org.

See content under the *Getting Started* chapter for instructions on how to install, build and use the TF-A Tests.

See *Framework Design* for information on how the TF-A Tests work internally.

See content under the *Porting* chapter for information about how to use this software on another Armv8-A platform.

See content under the *Process* chapter for information on how to contribute to this project.

Copyright (c)2019, Arm Limited. All rights reserved.