



Arm® Firmware Framework for Armv8-A

Document number	DEN0077A
Document version	1.0
Document confidentiality	Non-confidential

Copyright © 2020 Arm Limited or its affiliates. All rights reserved.

Contents

Arm® Firmware Framework for Armv8-A

	Release information	vii
	Arm Non-Confidential Document Licence (“Licence”)	ix
	References	xi
	Feedback	xii
Chapter 1	Introduction	
	1.1 Overview	15
	1.2 Document organization	17
Chapter 2	Concepts	
	2.1 Partition manager	19
	2.2 SPM architecture	20
	2.2.1 SPM architecture with Secure EL2	22
	2.2.2 SPM architecture without Secure EL2	22
	2.3 FF-A instances	26
	2.4 Conduits	27
	2.5 Execution state	28
	2.6 Memory types	29
	2.7 Memory granularity and alignment	30
	2.8 Partition identification and discovery	31
	2.9 Execution context	32
	2.10 System resource management	34
	2.11 Primary scheduler	35
	2.12 Run-time states	37
Chapter 3	Partition setup	
	3.1 Overview	39
	3.2 Partition manifest at virtual FF-A instance	40
	3.3 SP manifest at physical FF-A instance	45
	3.4 Independent peripheral device manifest	46
	3.5 Partition boot protocol	48
	3.5.1 Register state	48
	3.5.2 Protocol for passing data	48
	3.5.3 Protocol to initialize an execution context	49
Chapter 4	Message passing	
	4.1 Overview	51
	4.1.1 Indirect messaging	51
	4.1.2 Direct messaging	53
	4.2 Message transmission	55
	4.2.1 Overview	55
	4.2.2 RX/TX buffers	56
	4.3 Indirect messaging usage	67
	4.3.1 Discovery and setup	67
	4.3.2 Message delivery and scheduler notification	67
	4.3.3 Scheduling the Receiver	68
	4.4 Direct messaging usage	70
	4.4.1 Discovery and setup	71
	4.4.2 Message delivery and Receiver execution	72

4.5	Partition message processing	74
4.5.1	Indirect message processing	74
4.5.2	Direct message processing	74
4.5.3	Preemption during message processing	74
4.5.4	Managed exit	76

Chapter 5

Memory Management

5.1	Overview	82
5.2	Direct memory access	83
5.2.1	Stream endpoint	83
5.3	Address translation regimes	85
5.4	Ownership and access attributes	86
5.4.1	Ownership and access rules	86
5.4.2	Ownership and access states	87
5.5	Memory management transactions	90
5.5.1	Component roles	90
5.5.2	Transaction life cycle	92
5.6	Donate memory transaction	94
5.6.1	Donate memory state machine	94
5.6.2	Donate memory transaction lifecycle	94
5.7	Lend memory transaction	96
5.7.1	Lend memory transaction state machine	96
5.7.2	Lend memory transaction lifecycle	96
5.8	Share memory transaction	98
5.8.1	Share memory transaction state machine	98
5.8.2	Share memory transaction lifecycle	98
5.9	Relinquish memory transaction	100
5.9.1	Relinquish memory access state machine	100
5.9.2	Relinquish memory transaction lifecycle	101
5.10	Memory region description	102
5.10.1	Composite memory region descriptor	102
5.10.2	Memory region handle	105
5.11	Memory region properties	106
5.11.1	ABI-specific flags usage	107
5.11.2	Data access permissions usage	108
5.11.3	Instruction access permissions usage	110
5.11.4	Memory region attributes usage	111
5.12	Lend, donate, and share transaction descriptor	115
5.12.1	Handle usage	116
5.12.2	Tag usage	116
5.12.3	Endpoint memory access descriptor array usage	116
5.12.4	Flags usage	119

Chapter 6

Interface overview

6.1	Compliance requirements	124
6.1.1	Common compliance requirements	124
6.1.2	Compliance requirements for deploying an SP	125
6.1.3	Compliance requirements for deploying a VM	125
6.1.4	Compliance requirements for memory management	126

Chapter 7

Status reporting interfaces

7.1	Overview	128
7.2	FFA_ERROR	129
7.3	FFA_SUCCESS	131
7.4	FFA_INTERRUPT	133

Chapter 8	Setup and discovery interfaces	
8.1	FFA_VERSION	135
8.1.1	Overview	136
8.1.2	Usage	136
8.1.3	SPM usage	137
8.2	FFA_FEATURES	138
8.3	FFA_RX_RELEASE	141
8.4	FFA_RXTX_MAP	142
8.5	FFA_RXTX_UNMAP	145
8.6	FFA_PARTITION_INFO_GET	147
8.7	FFA_ID_GET	150
Chapter 9	CPU cycle management interfaces	
9.1	FFA_MSG_WAIT	153
9.1.1	Component responsibilities for FFA_MSG_WAIT	154
9.2	FFA_YIELD	156
9.2.1	Component responsibilities for FFA_YIELD	157
9.3	FFA_RUN	158
9.3.1	Component responsibilities for FFA_RUN	159
9.4	FFA_NORMAL_WORLD_RESUME	161
9.4.1	Overview	161
Chapter 10	Messaging interfaces	
10.1	FFA_MSG_SEND	164
10.1.1	Target availability notification	165
10.1.2	Component responsibilities for FFA_MSG_SEND	166
10.1.3	Mechanism for scheduler notification	167
10.2	FFA_MSG_SEND_DIRECT_REQ	168
10.2.1	Component responsibilities for FFA_MSG_SEND_DIRECT_REQ	169
10.3	FFA_MSG_SEND_DIRECT_RESP	171
10.3.1	Component responsibilities for FFA_MSG_SEND_DIRECT_RESP	172
10.4	FFA_MSG_POLL	174
Chapter 11	Memory management interfaces	
11.1	FFA_MEM_DONATE	176
11.1.1	Component responsibilities for FFA_MEM_DONATE	177
11.2	FFA_MEM_LEND	180
11.2.1	Component responsibilities for FFA_MEM_LEND	181
11.3	FFA_MEM_SHARE	184
11.3.1	Component responsibilities for FFA_MEM_SHARE	185
11.4	FFA_MEM_RETRIEVE_REQ	188
11.4.1	Component responsibilities for FFA_MEM_RETRIEVE_REQ	189
11.4.2	Support for multiple retrievals by a Borrower	191
11.4.3	Support for retrieval by the Hypervisor	191
11.5	FFA_MEM_RETRIEVE_RESP	193
11.5.1	Component responsibilities for FFA_MEM_RETRIEVE_RESP	194
11.6	FFA_MEM_RELINQUISH	195
11.6.1	Component responsibilities for FFA_MEM_RELINQUISH	197
11.7	FFA_MEM_RECLAIM	199
11.7.1	Component responsibilities for FFA_MEM_RECLAIM	200
Chapter 12	Appendix	
12.1	S-EL0 & User mode partitions	203
12.1.1	UEFI PI Standalone Management Mode partitions	203
12.2	Additional memory management features	208

Contents
Contents

12.2.1	Transmission of transaction descriptor in dynamically allocated buffers .	208
12.2.2	Transmission of transaction descriptor in fragments	210
12.2.3	Time slicing of memory management operations	218
Terms and abbreviations		223

Release information

Date	Version	Changes
2020/Jul/24	REL	<ul style="list-style-type: none">• Language fixes based upon feedback from editorial review• Removed reference to PSA from document title• Converted document to Arm spec format• Converted ffa_init_info C structure into a table• Clarified use of Sender ID field in FFA_FRAG_RX/TX• Fixed clash in FIDs of FFA_NORMAL_WORLD_RESUME and FFA_MEM_FRAG_RX• Clarified use of FFA_MSG_POLL with RX full interrupt• Clarified multi-endpoint memory management is an optional feature• Clarified how a receiver should request retransmission of a fragmented memory region description• Clarified 64-bit registers can be used in direct messaging
2020/Apr/24	EAC	<ul style="list-style-type: none">• Replaced occurrences of SPCI with PSA FF-A• Added flag to identify other borrowers in a memory retrieve operation• Allowed time slicing of memory management operations at Non-secure physical SPCI instance• Replaced Cookie with Handle in fragmented and time-sliced memory management operations• Added separate ABIs for fragmented memory management operations• Allowed multiple retrievals by a Borrower of a memory region• Allowed retrieval by Hypervisor of a memory region on behalf of a VM• Replaced separate memory transaction descriptors with a single one• Removed Write-through attribute to cater for S2FWB• Specified coherency requirements for memory zeroing• Moved to 64-bit memory Handles• Clarifications to existing memory management guidance• Made guidance on power management IMPLEMENTATION DEFINED• Allowed discovery of minimum buffer size through FFA_FEATURES• Changed FFA_VERSION for negotiation of version number between caller and callee• Clarified usage and description of FFA_FEATURES• Added section on compliance requirements• Other errata fixes and language clarifications based on feedback from beta 1
2019/Dec/20	beta 1	<ul style="list-style-type: none">• Added ability to pause and resume memory management transactions• Restricted indirect messaging to Normal world• Reworded guidance on Stream endpoint IDs (SEPIDs)• Added ABI to resume Normal world execution after a Secure interrupt• Reworded guidance on SPCI instances and Split SPM configuration• Added clearer guidance on optional and mandatory interfaces• Other errata fixes and language clarifications based on feedback from beta 0

Date	Version	Changes
2019/Nov/13	beta 0	<ul style="list-style-type: none">• Replaced some occurrences of ARM with Arm• Non-confidential release of beta 0 spec
2019/Sep/17	beta 0	<ul style="list-style-type: none">• Added guidance on partition manifest and setup• Significant rewrite of section on message passing• Added support for multi-component memory management• Added new interfaces for RX/TX management and deprecated old interfaces• Device reassignment has been removed from the scope of this release
2019/Apr/26	alpha 3 Draft 0	<ul style="list-style-type: none">• Significant rewrite of section on message passing• Chapter on scheduling models has been removed• Significant rewrite of section on memory management• Chapter 5 has become Chapter 10. Its scope has been reduced temporarily due to preceding changes.
2018/Dec/21	alpha 2	<ul style="list-style-type: none">• Changed content based on partner feedback since alpha 1• There is a clear separation between message passing and scheduling• Introduced use of RX/TX buffers to enable message passing

Arm Non-Confidential Document Licence (“Licence”)

This Licence is a legal agreement between you and Arm Limited (“**Arm**”) for the use of Arm’s intellectual property (including, without limitation, any copyright) embodied in the document accompanying this Licence (“**Document**”). Arm licenses its intellectual property in the Document to you on condition that you agree to the terms of this Licence. By using or copying the Document you indicate that you agree to be bound by the terms of this Licence.

“**Subsidiary**” means any company the majority of whose voting shares is now or hereafter owner or controlled, directly or indirectly, by you. A company shall be a Subsidiary only for the period during which such control exists.

This Document is **NON-CONFIDENTIAL** and any use by you and your Subsidiaries (“**Licensee**”) is subject to the terms of this Licence between you and Arm.

Subject to the terms and conditions of this Licence, Arm hereby grants to Licensee under the intellectual property in the Document owned or controlled by Arm, a non-exclusive, non-transferable, non-sub-licensable, royalty-free, worldwide licence to:

- (i) use and copy the Document for the purpose of designing and having designed products that comply with the Document;
- (ii) manufacture and have manufactured products which have been created under the licence granted in (i) above; and
- (iii) sell, supply and distribute products which have been created under the licence granted in (i) above.

Licensee hereby agrees that the licences granted above shall not extend to any portion or function of a product that is not itself compliant with part of the Document.

Except as expressly licensed above, Licensee acquires no right, title or interest in any Arm technology or any intellectual property embodied therein.

THE DOCUMENT IS PROVIDED “AS IS”. ARM PROVIDES NO REPRESENTATIONS AND NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT OR FITNESS FOR A PARTICULAR PURPOSE WITH RESPECT TO THE DOCUMENT. Arm may make changes to the Document at any time and without notice. For the avoidance of doubt, Arm makes no representation with respect to, and has undertaken no analysis to identify or understand the scope and content of, third party patents, copyrights, trade secrets, or other rights.

NOTWITHSTANDING ANYTHING TO THE CONTRARY CONTAINED IN THIS LICENCE, TO THE FULLEST EXTENT PERMITTED BY LAW, IN NO EVENT WILL ARM BE LIABLE FOR ANY DAMAGES, IN CONTRACT, TORT OR OTHERWISE, IN CONNECTION WITH THE SUBJECT MATTER OF THIS LICENCE (INCLUDING WITHOUT LIMITATION) (I) LICENSEE’S USE OF THE DOCUMENT; AND (II) THE IMPLEMENTATION OF THE DOCUMENT IN ANY PRODUCT CREATED BY LICENSEE UNDER THIS LICENCE). THE EXISTENCE OF MORE THAN ONE CLAIM OR SUIT WILL NOT ENLARGE OR EXTEND THE LIMIT. LICENSEE RELEASES ARM FROM ALL OBLIGATIONS, LIABILITY, CLAIMS OR DEMANDS IN EXCESS OF THIS LIMITATION.

This Licence shall remain in force until terminated by Licensee or by Arm. Without prejudice to any of its other rights, if Licensee is in breach of any of the terms and conditions of this Licence then Arm may terminate this Licence immediately upon giving written notice to Licensee. Licensee may terminate this Licence at any time. Upon termination of this Licence by Licensee or by Arm, Licensee shall stop using the Document and destroy all copies of the Document in its possession. Upon termination of this Licence, all terms shall survive except for the licence grants.

Contents

Arm Non-Confidential Document Licence ("Licence")

Any breach of this Licence by a Subsidiary shall entitle Arm to terminate this Licence as if you were the party in breach. Any termination of this Licence shall be effective in respect of all Subsidiaries. Any rights granted to any Subsidiary hereunder shall automatically terminate upon such Subsidiary ceasing to be a Subsidiary.

The Document consists solely of commercial items. Licensee shall be responsible for ensuring that any use, duplication or disclosure of the Document complies fully with any relevant export laws and regulations to assure that the Document or any portion thereof is not exported, directly or indirectly, in violation of such export laws.

This Licence may be translated into other languages for convenience, and Licensee agrees that if there is any conflict between the English version of this Licence and any translation, the terms of the English version of this Licence shall prevail.

The Arm corporate logo and words marked with ® or ™ are registered trademarks or trademarks of Arm Limited (or its subsidiaries) in the US and/or elsewhere. All rights reserved. Other brands and names mentioned in this document may be the trademarks of their respective owners. No licence, express, implied or otherwise, is granted to Licensee under this Licence, to use the Arm trade marks in connection with the Document or any products based thereon. Visit Arm's website at <http://www.arm.com/company/policies/trademarks> for more information about Arm's trademarks.

The validity, construction and performance of this Licence shall be governed by English Law.

Copyright © 2020 Arm Limited (or its affiliates). All rights reserved.

Arm Limited. Company 02557590 registered in England.

110 Fulbourn Road, Cambridge, England CB1 9NJ.

LES-PRE-21585 version 4.0

Copyright © 2020 Arm Limited. All rights reserved.

References

This section lists publications by Arm® and by third parties.

See Arm® Developer (<http://developer.arm.com>) for access to Arm® documentation.

- [1] *Arm® System Memory Management Unit Architecture specification versions 3.0, 3.1 and 3.2*. See <https://developer.arm.com/documentation/ih0070/ca>
- [2] *Arm® System Memory Management Unit Architecture specification version 2.0*. See https://static.docs.arm.com/ih0062/dc/IH0062D_c_system_mmu_architecture_specification.pdf
- [3] *Isolation using virtualization in the Secure world*. See <https://developer.arm.com/products/architecture/security-architectures>
- [4] *SMC Calling Convention*. See https://static.docs.arm.com/den0028/c/Q1-ARM-DEN-0028_SMC_Calling_Convention_v1_2_Non_Conf_EAC.pdf
- [5] *Arm® Architecture Reference Manual for the ARMv8-A architecture*. See https://static.docs.arm.com/ddi0487/ea/DDI0487E_a_armv8_arm.pdf
- [6] *Universally Unique IDentifier*. See <https://tools.ietf.org/html/rfc4122>
- [7] *Power State Coordination Interface*. See https://static.docs.arm.com/den0022/d/Power_State_Coordination_Interface_PDD_v1_1_DEN0022D.pdf
- [8] *Arm® GIC architecture specification versions 3.0 and 4.0*. See https://static.docs.arm.com/ih0069/e/Q1-IH0069E_gic_architecture_specification_v3.1_19_01_21.pdf
- [9] *VOLUME 4: Platform Initialization Specification, Management Mode Core Interface*. See http://www.uefi.org/sites/default/files/resources/PI_Spec_1_6.pdf
- [10] *Management Mode Interface Specification*. See http://infocenter.arm.com/help/topic/com.arm.doc.den0060a/DEN0060A_ARM_MM_Interface_Specification.pdf

Feedback

Arm welcomes feedback on its documentation.

If you have comments on the content of this manual, send an e-mail to errata@arm.com. Give:

- The title (Arm® Firmware Framework for Armv8-A).
- The document ID and version (DEN0077A 1.0).
- The page numbers to which your comments apply.
- A concise explanation of your comments.

Arm® also welcomes general suggestions for additions and improvements.

Chapter 1

Introduction

The Armv8.4 architecture introduces the Virtualization extension in the Secure state. The Arm® SMMU v3.2 architecture [1] adds support for stage 2 translations for Secure streams to complement the Secure EL2 translation regime in an Armv8.4 PE. These architectural features enable *isolation* of mutually mistrusting software components in the Secure state from each other. *Isolation* is a mechanism for implementing the principle of least privilege:

A software component must be able to access only regions in the physical address space and system resources for example, interrupts in the GIC that are necessary for its correct operation.

Virtualization in the Secure state enables application of this principle in the following ways:

1. Firmware in EL3 can be isolated from software in S-EL1 for example, a Trusted OS.
2. Firmware components in EL3 can be isolated from each other by migrating vendor-specific components to a sandbox in S-EL1.
3. Normal world software can be isolated from software in S-EL1 to mitigate privilege escalation attacks.

This specification describes a software architecture that achieves the following goals.

1. Uses the Virtualization extension to isolate software images provided by an ecosystem of vendors from each other.
2. Describes interfaces that standardize communication between the various software images. This includes communication between images in the Secure world and Normal world.

This software architecture is the *Firmware Framework*¹ for Arm® A-profile processors. The term *Framework* and abbreviation *FF-A* are used interchangeably with *Firmware Framework* in this specification.

This Framework also goes beyond the preceding goals to ensure that the interfaces can be used to standardize communication:

¹This document was called the *Secure Partition Client Interface (SPCI)* specification until its BETA1 release.

1. In the absence of the Virtualization extension in the Secure state. This provides a migration path for existing Secure world software images to a system that implements the Virtualization extension in the Secure state.
2. Between VMs managed by a Hypervisor in the Normal world. The Virtualization extension in the Secure state mirrors its counterpart in the Non-secure state (see also [2]). Therefore, a Hypervisor could use the Firmware Framework interfaces to enable communication between VMs it manages.

More rationale about the introduction of the Virtualization extension in Secure state and goals of the Firmware Framework is provided in the white-paper titled *Isolation using virtualization in the Secure world* [3].

1.1 Overview

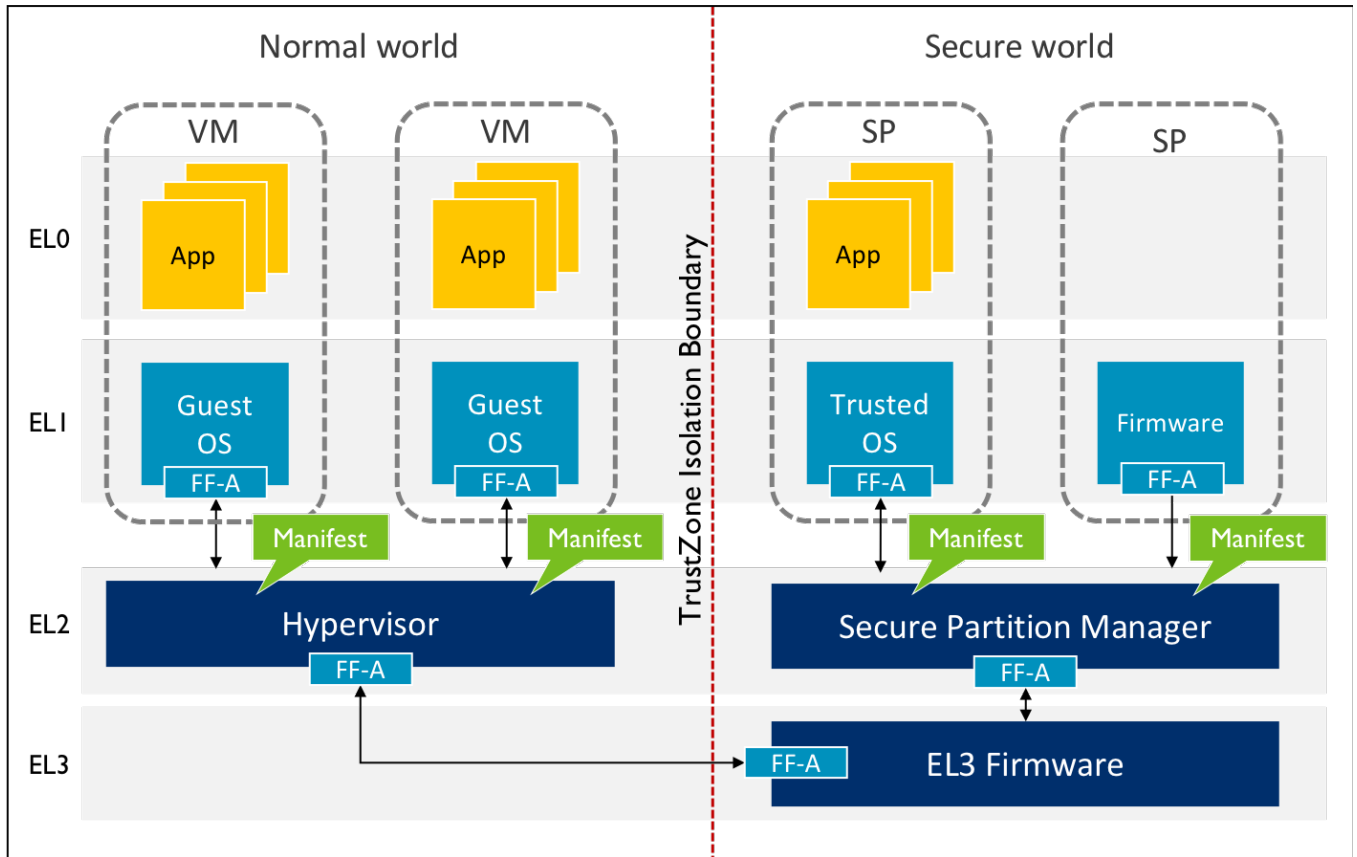


Figure 1.1: Firmware Framework with Secure EL2

The **building blocks** of the Firmware Framework described in this specification are as follows. [Figure 1.1](#) illustrates an implementation of this Framework and its components.

1. **One or more partitions** that provide a sandboxed software execution environment. These could be VMs running in the Normal or Secure world. A Secure world VM is called a *Secure Partitions* (SP) to distinguish it from VMs in the Normal world.

An SP typically encompasses the S-EL1 and S-EL0 Exception levels. The Firmware Framework supports SPs that run only in S-EL0 as well. A S-EL0 SP could be managed by software in S-EL1 or EL3. This is an IMPLEMENTATION DEFINED choice.

The term *endpoint* is used interchangeably with the term *partition*.

- In the Normal world, an endpoint could be a VM when the Virtualization extension is enabled or the OS Kernel when the Virtualization extension is disabled or unavailable. These endpoints are called *NS-Endpoints* in scenarios where it is not necessary to distinguish between them.
- In the Secure world, an endpoint is an SP running in one of the following Exception levels:
 - Secure EL0.
 - Secure User mode.
 - Secure EL1.
 - Secure Supervisor mode.

These endpoints are called *S-Endpoints* in scenarios where it is not necessary to distinguish between them.

2. A **partition manifest** that describes the system resources a partition needs, identity of the partition to enable discovery of services that it implements and other attributes of the partition that govern its run-time behavior.
3. A **partition manager** (PM) that assigns system resources to partitions and manages isolation among them. In the Secure world, this component is called the *Secure Partition Manager* (SPM). In the Normal world it is a Hypervisor². The SPM and Hypervisor are collectively referred to as the *Partition managers* in scenarios where they have the same responsibilities and it is not necessary to distinguish between them. See [2.1 Partition manager](#) for a description of this component.
4. **Application binary interfaces** that partitions can invoke at their Exception level boundaries for the following purposes.
 1. Discover the presence of a partition, its properties and services it implements.
 2. Message passing among partitions and partition managers.
 3. Memory management between partitions.

[Table 1.1](#) summarizes software configurations supported by the Firmware Framework in each Security state with regard to availability of the Virtualization extension. Furthermore, each configuration in one Security state can co-exist with any configuration in the other Security state.

Table 1.1: Firmware Framework configurations

Config. No.	Security state	Virtualization extension enabled	Description
1.	Non-secure	No	Rich OS in EL1 uses the Firmware Framework to communicate with one or more S-Endpoints.
2.	Non-secure	Yes	One or more VMs in EL1 use the Firmware Framework to: <ul style="list-style-type: none"> • Communicate with one or more S-Endpoints. • Communicate with each other. • Isolate themselves from each other.
3.	Secure	No	One or more S-Endpoints use the Firmware Framework as follows: <ul style="list-style-type: none"> • A single SP for example, a Trusted OS in S-EL1 uses the Framework to communicate with one or more NS-Endpoints. • One or more SPs in S-EL0 use the Firmware Framework to: <ul style="list-style-type: none"> – Communicate with one or more NS-Endpoints. – Communicate with each other. – Isolate themselves from each other.
4.	Secure	Yes	One or more SPs use the Firmware Framework to: <ul style="list-style-type: none"> • Communicate with one or more NS-Endpoints. • Communicate with each other. • Isolate themselves from each other.

²A hypervisor implementation could span EL1 and EL2. In this specification, this term refers to the layer of software that runs in EL2 and is responsible for providing isolation guarantees between VMs through use of the Arm® virtualization extension.

1.2 Document organization

The rest of this document is organized as follows.

1. [Chapter 2 Concepts](#) describes some fundamental concepts that are used to define the Firmware Framework architecture.
2. [Chapter 3 Partition setup](#) specifies the information contained in a partition manifest and how it is used to initialize a partition by a partition manager.
3. [Chapter 4 Message passing](#) describes the mechanisms that partitions can use for message passing.
4. [Chapter 5 Memory Management](#) describes the mechanisms that partitions can use for memory management.
5. [Chapter 6 Interface overview](#) provides an overview of the ABIs defined by the Firmware Framework.
6. [Chapter 7 Status reporting interfaces](#), [Chapter 8 Setup and discovery interfaces](#), [Chapter 9 CPU cycle management interfaces](#), [Chapter 10 Messaging interfaces](#) & [Chapter 11 Memory management interfaces](#) specify ABIs used in the Firmware Framework for status reporting, setup and discovery of partitions, messaging, and memory management.
7. [Chapter 12 Appendix](#) describes how an example use case of S-EL0 SPs can be implemented using the Firmware Framework.

Chapter 2

Concepts

2.1 Partition manager

The partition manager is responsible for initializing a partition as per the requirements stated in its manifest (see [Chapter 3 Partition setup](#)). A partition describes the regions in the system physical address space and resources it needs access to through its manifest. The partition manager uses the manifest to validate the resource requests and assign resources to the endpoint if the validation succeeds.

The following trust boundaries are defined by the Firmware Framework vis-a-vis the partition managers and partitions.

- The SPM is a part of the TCB for a system resource or physical address space range assigned to the Secure state.
- Both the Hypervisor and SPM are a part of the TCB for a system resource or physical address space range assigned to the Non-secure state.
- A VM must trust the Hypervisor and SPM to protect its resources from other endpoints.
- An SP must trust the SPM to protect its Secure resources from other SPs.
- An SP must not trust the state of any Non-secure resource it has access to. Therefore, it must not trust the Hypervisor or a NS-Endpoint that could also access the same resource.

A partition manager protects partition resources from other FF-A components by utilizing the following features implemented by the CPU and system architecture.

- The Arm® TrustZone Security extension is used to protect the Secure physical address space ranges and system resources assigned to FF-A components in the Secure state from components in the Non-secure state.
- Virtual memory-based memory protection provided by the Armv8-A VMSA is used to protect the physical address space ranges assigned to a Security state and FF-A component from other FF-A components. Its usage by the SPM is described in [2.2 SPM architecture](#). The Hypervisor uses this feature as follows.
 - If the *EL1&0 stage 2 translation regime, when EL2 is enabled* is implemented by a System Memory Management Unit (SMMU) in the Non-secure state, the Hypervisor uses it to restrict visibility of the Non-secure physical address space from a device upstream of the SMMU to only those regions that have been assigned to the VM that controls the device.
 - If the *Secure EL1&0 stage 2 translation regime, when EL2 is enabled* is implemented by a CPU in the Non-secure state, the Hypervisor uses it to restrict visibility of the Non-secure physical address space from a VM to only those regions that have been assigned to the VM.

The partition manager enables partitions to exchange messages (see [Chapter 4 Message passing](#)). It also enables a partition to manage access to memory regions that are assigned to it from other partitions (see [Chapter 5 Memory Management](#)).

In an implementation of this Framework, the SPM must use the concepts and interfaces described in this specification to fulfill the preceding responsibilities. A Hypervisor could use the Framework only for communication and memory management between the Normal world and Secure world. In this case, the Hypervisor must:

- Initialize VMs and isolate them from other VMs through IMPLEMENTATION DEFINED mechanisms.
- Implement a partition manager component that uses the Firmware Framework to enable communication and memory management between two endpoints. for example, this could be an FF-A driver implemented in the Hypervisor.

In this version of the Firmware Framework, it is assumed that a partition manager is initialized through an IMPLEMENTATION DEFINED mechanism.

2.2 SPM architecture

The responsibilities of the SPM are split between the two components as follows.

1. The *SPM Core* (SPMC) component which is responsible for:
 - Partition initialization and isolation at boot time.
 - Inter-partition isolation at run-time.
 - Inter-partition communication at run-time between:
 - S-Endpoints.
 - S-Endpoints and NS-Endpoints.
2. The *SPM Dispatcher* (SPMD) component which is responsible for:
 - *SPM Core* initialization at boot time.
 - Forwarding FF-A calls from Normal world to the *SPM Core*.
 - Forwarding FF-A calls from the *SPM Core* to the Normal world.

The term SPM is used when it is not necessary to distinguish between these two components. Some properties of the two components are as follows.

- Both components have access to the entire physical address space and are a part of the *Trusted computing base*.
- If the two components reside in separate Exception levels:
 - They must implement and report a mutually compatible version of the Firmware Framework. See [8.1.3 SPM usage](#) for details.
 - They must use the ABIs defined in this specification for communication.
- The mechanism used by the SPMD to initialize the SPMC is IMPLEMENTATION DEFINED.
- The SPMD component must execute in either EL3 in AArch64 or the Monitor mode in AArch32 Execution states.

The Firmware Framework supports SPMC configurations listed in [Table 2.1](#) & [Table 2.2](#).

Table 2.1: SPMC configurations in AArch64 Execution state

SPM config. number	SPMD EL	SPMC EL	Virtualization extension enabled	Name of configuration
1.	EL3	S-EL1	No	S-EL1 SPMC
2.	EL3	S-EL2	Yes	S-EL2 SPMC
3.	EL3	EL3	No	EL3 SPMC

Table 2.2: SPMC configurations in AArch32 Execution state

SPM config. number	SPMD EL	SPMC EL	Virtualization extension enabled	Name of configuration
1.	EL3	Secure Supervisor	No	S-Supervisor SPMC
2.	Monitor	Secure Supervisor	No	S-Supervisor SPMC

[Table 2.3](#) lists combinations of SPMC and SP configurations supported by this version of the Framework.

- The first row lists all possible SP configurations.
- The first column lists all possible SPMC configurations.
- An intersection of a row and a column indicates whether the SP configuration in the row is supported by the SPMC configuration in the column.

Table 2.3: Valid combinations of SPMC and SP configurations

SPMC config. name	S-EL0 SP	Secure User mode SP	S-EL1 SP	Secure Supervisor mode SP	Notes
S-EL1 SPMC	Yes	Yes	Yes	NA	See 2.2.2.1 S-EL1 SPM core component
S-Supervisor SPMC	NA	Yes	NA	Yes	See 2.2.2.2 Secure Supervisor mode SPM core component
EL3 SPMC	Yes	Yes	Yes ¹	No	See 2.2.2.3 EL3 SPM core component
S-EL2 SPMC	Yes	Yes	Yes	Yes	See 2.2.1 SPM architecture with Secure EL2

¹Either a S-EL1 SP or S-EL0/User mode SPs must be supported but not both.

2.2.1 SPM architecture with Secure EL2




Figure 2.1: Example S-EL2 SPM Core and SP configuration

The S-EL2 SPMC (see SPM configuration 2 in [Table 2.1](#)) is fundamental to enforcing the principle of least privilege in the Secure state on Armv8.4 or later systems as described in [Chapter 1 Introduction](#). It must support at least one of the SP configurations as follows.

1. The SPMC uses *Armv8.1 VHE* to manage one or more SPs that run in S-EL0 or Secure User mode. It fulfills all the responsibilities listed in [2.2 SPM architecture](#).

The physical address space assigned to an SP is isolated from other FF-A components through the single stage of address translation implemented by the *Secure EL2&0 translation regime*.

2. The SPMC manages one or more SPs that run in S-EL1 or S-Supervisor mode. It fulfills all the responsibilities listed in [2.2 SPM architecture](#).

The physical address space assigned to an SP is isolated from other FF-A components by the *Secure EL1&0 stage 2 translation regime, when EL2 is enabled*.

An example of these configurations is illustrated in [Figure 2.1](#).

2.2.2 SPM architecture without Secure EL2

In the absence of Secure EL2, SPM could be used in the following scenarios.

- Reduce the size of the *Trusted computing base* on an Armv8.3 or earlier system by migrating EL3 & S-EL1 firmware components to one or more SPs that run in S-EL0 or Secure User mode.
 - The SPMC configurations described in [2.2.2.1 S-EL1 SPM core component](#), [2.2.2.2 Secure Supervisor mode SPM core component](#) and [2.2.2.3 EL3 SPM core component](#) could be used in this scenario.

- See [12.1 S-EL0 & User mode partitions](#) for an example use case of this configuration.
- Migrate a Secure world software stack that runs on Armv8.3 or earlier systems to the Firmware Framework in preparation for Armv8.4 or later systems.
 - The SPMC configurations described in [2.2.2.1 S-EL1 SPM core component](#) and [2.2.2.2 Secure Supervisor mode SPM core component](#) could be used in this scenario.

2.2.2.1 S-EL1 SPM core component

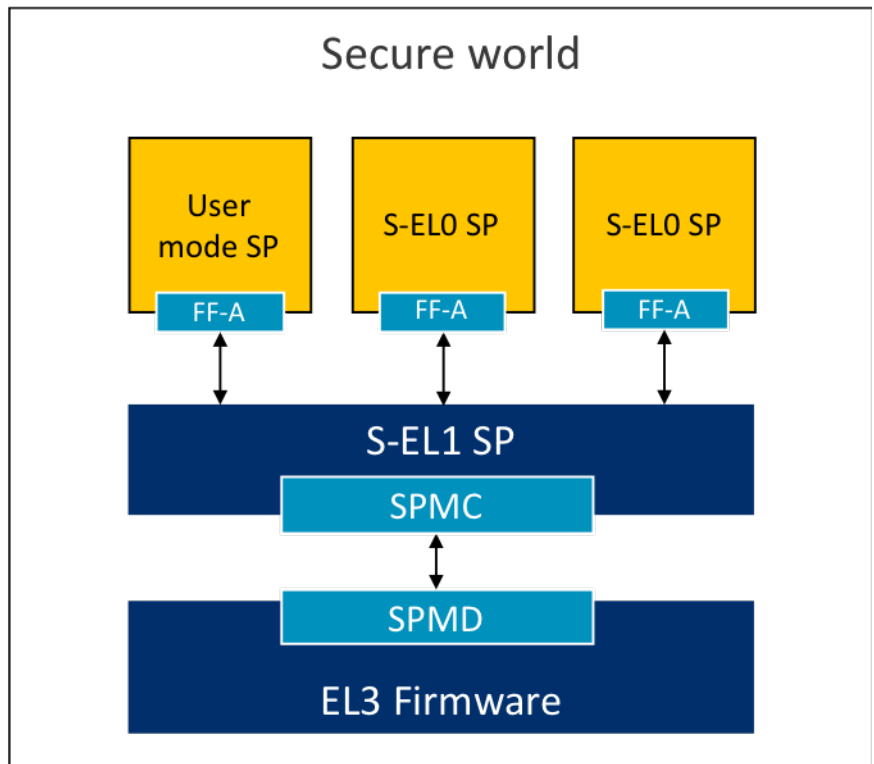


Figure 2.2: Example S-EL1 SPM Core and SP configuration

A S-EL1 SPMC must support at least one of the SP configurations as follows.

1. The SPMC manages one or more SPs that run in S-EL0 or Secure User mode. It fulfills all the responsibilities listed in [2.2 SPM architecture](#).

The physical address space assigned to an SP is isolated from other FF-A components through the single stage of address translation implemented by the:

- *Secure EL1&0 translation regime* for S-EL0 SPs.
- *Secure PL1&0 translation regime* for Secure User mode SPs.

2. The SPMC manages a single SP that also runs in S-EL1. The SPMD, SPMC, and SP components have the same level of access to the physical address space and are a part of the *Trusted computing base*.

In this configuration:

- The Framework assumes that the SPMC is packaged in the SP software image.
- The interface between the SPMC and the SP component is IMPLEMENTATION DEFINED for example, a set of C programming language APIs.
- Any FF-A calls targeted to the SP from the Normal world must be received by the SPMC and forwarded to the appropriate SP component through the IMPLEMENTATION DEFINED interface.

- The SPM and Normal world software cannot be isolated from the SP at boot time. See [Chapter 3 Partition setup](#) for more information on the implications of this constraint on partition setup and boot.
- The SPM and Normal world software cannot be isolated from the SP at run-time. See [5.5.1 Component roles](#) for more information on the implications of this constraint on memory management transactions between the SP and the Normal world.
- The SP must be capable of receiving and sending messages just like the SPM. See [4.1.1 Indirect messaging](#) & [4.4.1 Discovery and setup](#) for more information on the implications of this constraint on communication between the SP and the Normal world.

[Figure 2.2](#) illustrates a combination of these configurations.

2.2.2.2 Secure Supervisor mode SPM core component




Figure 2.3: Example Supervisor mode SPM Core and SP configuration

The S-Supervisor SPMC must support the same SP configurations described in [2.2.2.1 S-EL1 SPM core component](#) with the following caveats.

1. The SPMC manages one or more SPs that run only in Secure User mode.
2. The SPMC coexists with a single SP that also runs in Secure Supervisor mode.
3. The SPM is isolated from the Secure User mode SPs through the single stage of address translation implemented by the *Secure PL1&0 translation regime*.

This configuration is illustrated in [Figure 2.3](#).

2.2.2.3 EL3 SPM core component

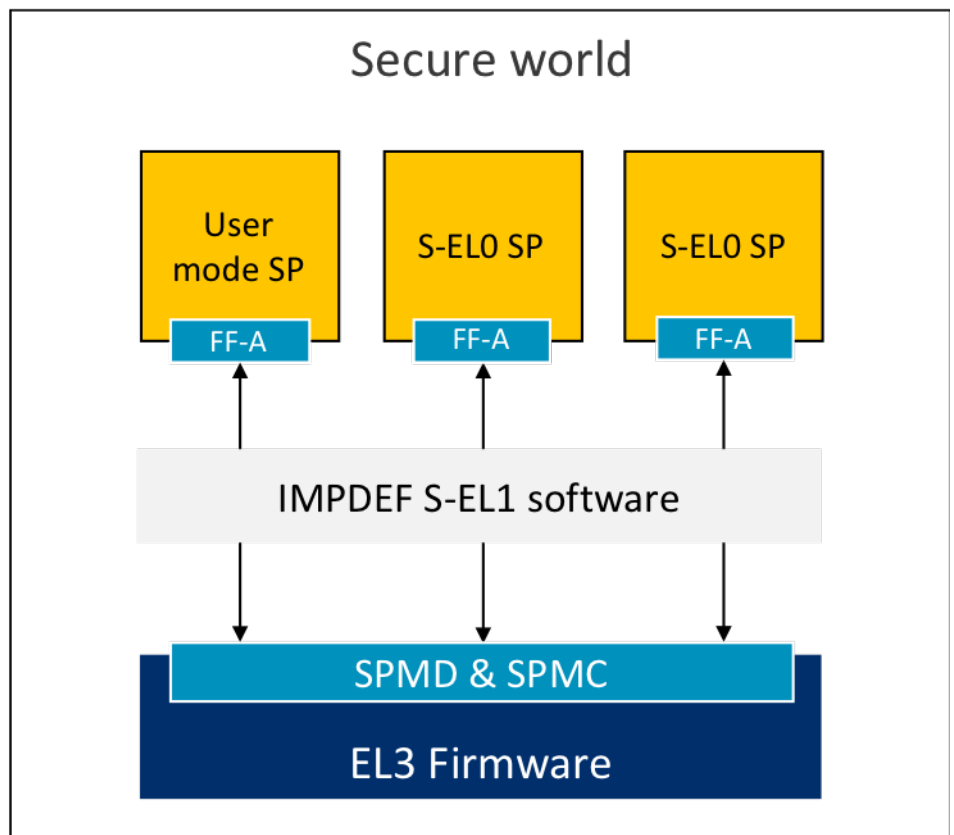


Figure 2.4: Example EL3 SPM Core and SP configuration

The EL3 SPMC co-exists with the SPMD and manages one of the following SP configurations. It fulfills all the responsibilities listed in [2.2 SPM architecture](#).

1. One or more SPs that run in S-EL0 or Secure User mode. This configuration is illustrated in [Figure 2.4](#).

The physical address space assigned to an SP is isolated from other FF-A components through the single stage of address translation implemented by the:

- *Secure EL1&0 translation regime* for S-EL0 SPs.
- *Secure PL1&0 translation regime* for Secure User mode SPs.

2. A single SP that resides in S-EL1. The SPMD, SPMC, and SP components have the same level of access to the physical address space and are a part of the *Trusted computing base*. The roles of the SPMC and SPMD are combined such that they are collectively responsible for:

- SP initialization at boot time.
- Inter-partition communication between the SP and NS-Endpoints at runtime.

2.3 FF-A instances

An *FF-A instance* is a valid combination of two FF-A components at an Exception level boundary. These instances are used to describe the interfaces specified by the Firmware Framework. An interface is accessed at an FF-A instance through a conduit described in 2.4 *Conduits*. The responsibilities of the caller and callee in each interface depend on the FF-A instance at which it is invoked.

- An instance is *physical* if:
 - Each component can independently manage its translation regime.
 - The translation regimes of each component map virtual addresses to physical addresses.
- An instance is *virtual* if it is not physical.
- The instance between the SPMC and SPMD is called the *Secure physical FF-A instance*.
- The instance between the SPMC and an SP is called the *Secure virtual FF-A instance*.
- In the Normal world, the instance between:
 - The Hypervisor and a VM is called the *Non-secure virtual FF-A instance*.
 - The Hypervisor and SPMD is called the *Non-secure physical FF-A instance*.
 - The OS kernel and SPMD, in the absence of a Hypervisor is called the *Non-secure physical FF-A instance*.

Table 2.4 lists the valid Secure FF-A instances. Table 2.5 lists the valid Non-secure FF-A instances.

- Entries in the first row represent the higher Exception level at an Exception level boundary.
- Entries in the first column represent the lower Exception level at an Exception level boundary.
- Combinations of Exception levels that are not architecturally feasible are listed as *Not applicable (NA)*.
- Combinations of Exception levels that are not supported by the Firmware Framework are listed as *Invalid (INV)*.

Table 2.4: Secure FF-A instances

EL boundary	EL3	Monitor	S-EL2	S-EL1	Secure Supervisor
S-EL2	Secure physical	NA	NA	NA	NA
S-EL1	Secure physical	NA	Secure virtual	NA	NA
Secure Supervisor	Secure physical	Secure physical	Secure virtual	NA	NA
S-EL0	Secure virtual	NA	Secure virtual	Secure virtual	NA
User	Secure virtual	INV	Secure virtual	Secure virtual	Secure virtual

Table 2.5: Non-secure FF-A instances

EL boundary	EL3	Monitor	EL2	Hypervisor
EL2	Non-secure physical	NA	NA	NA
Hypervisor	Non-secure physical	Non-secure physical	NA	NA
EL1	Non-secure physical	NA	Non-secure virtual	Non-secure virtual
Supervisor	Non-secure physical	Non-secure physical	Non-secure virtual	Non-secure virtual

2.4 Conduits

The Framework defines interfaces to enable communication between various FF-A components (see [Chapter 6 Interface overview](#)). Each interface is accessible through one or more conduits as follows.

The SMC conduit as described in [4] should be used to invoke an interface by an FF-A component executing in EL1 or S-EL1. When an interface is invoked from EL1, the SMC execution must be trapped by the Hypervisor at EL2. Similarly, when an interface is invoked at S-EL1 and the SPM resides in S-EL2, the SMC execution must be trapped by the SPM. This implies that the SMC conduit provides the flexibility that is required to support implementations with and without a hypervisor in EL2 or SPM in S-EL2.

If an endpoint executing in EL1 or S-EL1 cannot use the SMC conduit, it must use the HVC conduit instead.

A S-EL0 SP must use the SVC (Supervisor Call) instruction as a conduit to call into S-EL1. The SMC32 and SMC64 calling conventions are mirrored as SVC32 and SVC64 calling conventions respectively.

The Firmware Framework enables message exchange between any two FF-A components that might be at the same or a different Exception level relative to each other. A request, its results, or an error status could be sent from:

- A lower EL to a higher EL
- A higher EL to a lower EL.

To fulfill this requirement, this version of the Framework uses the *ERET* instruction as a conduit for transmitting requests and responses from a higher EL to a lower EL.

The parameter register usage in an SMC, HVC, or SVC call is mirrored in an ERET call for example, *w0* contains a *function identifier* parameter in the ERET call. This ensures that messages can be passed at any FF-A instance irrespective of their direction of travel. An invocation through the SMC, HVC, or SVC conduits is completed through the ERET conduit. An invocation through the ERET conduit is completed through the SMC, HVC, or SVC conduits.

This usage of the *ERET* instruction as a conduit along with the SMC, HVC, and SVC conduits enables half-duplex communication between two FF-A components at an EL boundary at any FF-A instance.

The taxonomy of information transmitted through a conduit at an FF-A instance is as follows.

1. An interface invocation described in [Chapter 6 Interface overview](#).
2. Results from the successful completion of the invoked interface.
3. Error code from an unsuccessful completion of the invoked interface.

Based on the preceding taxonomy, an interface invocation through one conduit at an FF-A instance can complete through another conduit in one of the following ways.

- An error code. The *FFA_ERROR* function is used to return the error code (see [7.2 FFA_ERROR](#)).
- Results of the request. The *FFA_SUCCESS* function is used to return the results (see [7.3 FFA_SUCCESS](#)).
- An invocation of another interface described in [Chapter 6 Interface overview](#).

An invocation of a non-FF-A interface from a lower Exception level to a higher Exception level for example, through the SMC, HVC, or SVC conduits must not complete with an invocation of an FF-A function through the ERET conduit unless, the caller implements support to distinguish between the FF-A and non-FF-A register usage on completion. For example, *w0* would contain a status code in the latter case while it will contain a *function identifier* in the former case.

2.5 Execution state

The Armv8-A architecture defines two Execution states, AArch32 and AArch64 as described in [5]. The Execution states that are applicable to each FF-A component are as follows.

- The SPM in S-EL2 or EL3 must run in AArch64.
- The SPM in S-EL1 could run in AArch64 or AArch32.
- The Hypervisor in EL2 must run in AArch64.
- A S-EL0 SP could run in AArch64 or AArch32.
- An endpoint in S-EL1 or EL1 could run in either AArch64 or AArch32.

2.6 Memory types

Each memory region is assigned to either the Secure or Non-secure physical address space at system reset or during system boot. Normal world can only access memory regions in the Non-secure physical address space. Secure world can access memory regions in both address spaces. The Non-secure (NS) attribute bit in the translation table descriptor determines whether an access is to Secure or Non-secure memory. In this version of the Framework:

- Memory that is accessed with the NS bit set in the translation regime of any component is called *Normal memory*.
- Memory that is accessed with the NS bit cleared in the component translation regime is called *Secure memory*.

2.7 Memory granularity and alignment

The Firmware Framework specifies support to map a memory region in the translation regimes of the two FF-A components at an FF-A instance (see [4.2.2.3 Buffer attributes](#) & [Chapter 5 Memory Management](#)). The translation regimes could use the same or a different translation granule size. To map the memory region correctly in both translation regimes, the following constraints must be met:

- If X is the larger translation granule size used by the two translation regimes, then the size of the memory region must be a multiple of X .
- The base address of the memory region must be aligned to X .

For example, at the Non-secure virtual FF-A instance, a VM and the Hypervisor could use translation granule sizes of 4K and 64K respectively. The size of any memory region that must be mapped in both their translation regimes must be a multiple of 64K and aligned to the 64K boundary.

An endpoint could specify its translation granule size in its partition manifest as described in [3.2 Partition manifest at virtual FF-A instance](#). The Hypervisor and SPM could also use an IMPLEMENTATION DEFINED mechanism to determine the translation granule size of an endpoint.

An endpoint must discover the minimum size and alignment boundary (that is, the minimum value of X) to share a memory region with its partition manager through the *FFA_FEATURES* interface (see [8.2 FFA_FEATURES](#)).

2.8 Partition identification and discovery

Partitions are identified by a 16-bit *ID* and a UUID (Unique Universal Identifier) (see [6]). This helps partitions discover the presence of other partitions and their properties.

A partition must use the *FFA_ID_GET* interface (also see 8.7 *FFA_ID_GET*) to discover its ID.

All FF-A components can discover the identities and properties of other partitions through the *FFA_PARTITION_INFO_GET* interface. Once discovered, the IDs must be used in the messaging interfaces to identify the target of a message.

A unique ID must be assigned to each partition in the system. The mechanism used to assign an *ID* to a partition is IMPLEMENTATION DEFINED. The *ID* could be specified in the manifest of the partition or allocated at boot by the partition manager responsible for managing the partition. A Hypervisor could use the *FFA_PARTITION_INFO_GET* interface to determine the *IDs* assigned to SPs to avoid clashes with VM *IDs*

The id value 0 is reserved for the Hypervisor as described in [4]. The id value assigned to the SPM is IMPLEMENTATION DEFINED.

2.9 Execution context

Each endpoint has one or more *execution contexts* depending on its implementation. An execution context comprises of general-purpose, system, and any memory mapped register state that must be maintained by a partition manager.

A partition manager is responsible for allocating, initializing, and running the execution context of an endpoint on a physical or virtual PE in the system. An execution context is identified by using a 16-bit ID. This ID is referred to as the *vCPU* or *execution context ID*. Each execution context must be allocated an ID that is unique among all execution contexts that belong to the endpoint. [Figure 2.5](#) illustrates an example configuration of endpoints and their execution contexts.

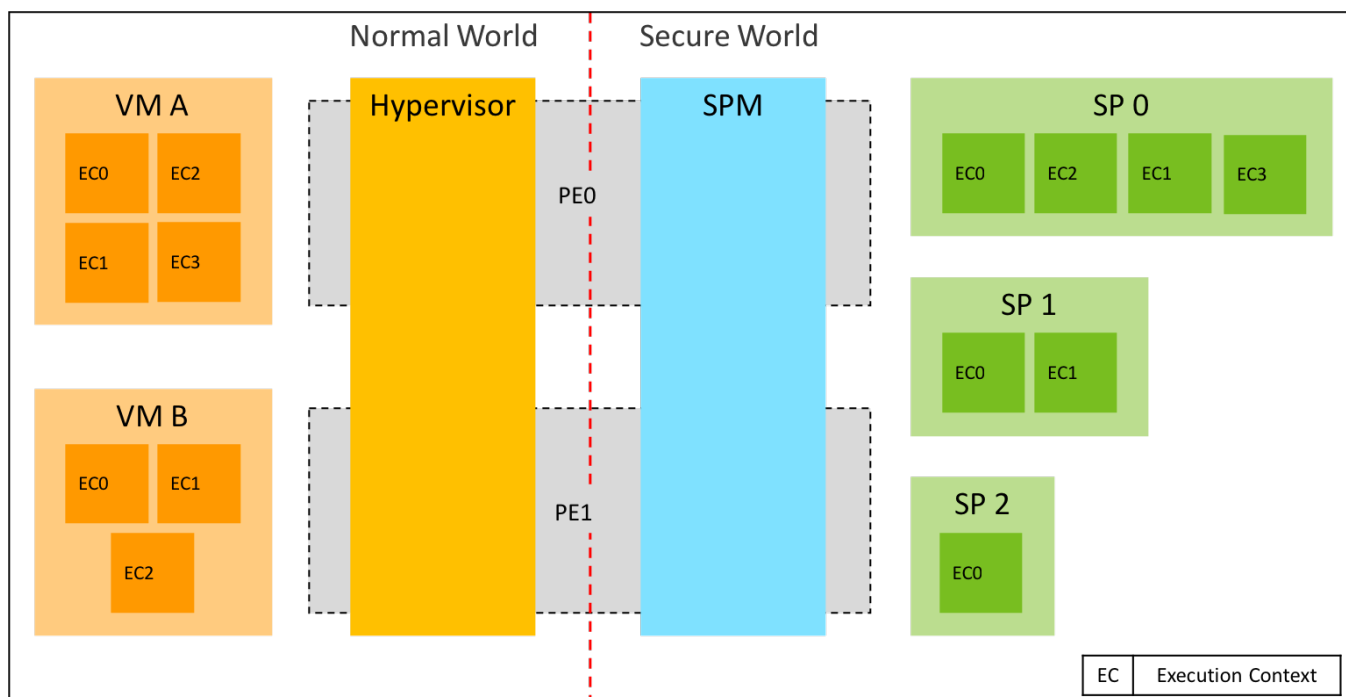


Figure 2.5: SP and VM vCPUs

An execution context of an endpoint represents a logical processor to the partition manager. The partition manager delegates message processing to an execution context of an endpoint. It is independent of threads implemented inside an endpoint to process the messages and logic to schedule these threads (see also [2.11 Primary scheduler](#)). [Figure 2.6](#) illustrates this relationship.

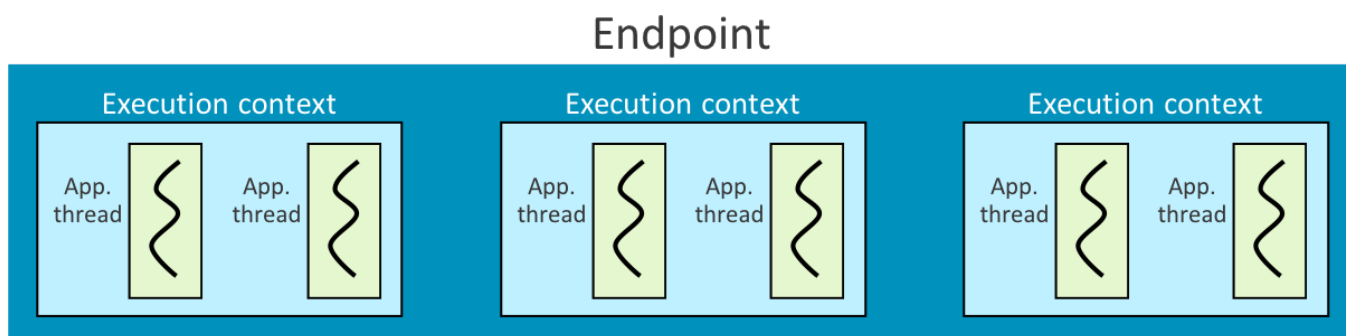


Figure 2.6: Example endpoint with execution contexts and threads

An endpoint must be one of the following types:

- Implements a single execution context and is not capable of Symmetric multi-processing. It runs only on a single PE in the system at any point of time. This type of endpoint is called a **UP** endpoint.
- Implements multiple execution contexts and is capable of Symmetric multi-processing. These contexts run concurrently on separate PEs in the system. These endpoints are called **MP** endpoints.

An execution context of an endpoint could be capable of *migrating*. Migration capability means that the partition manager could save the execution context of an endpoint on one PE. It could then restore the saved execution context on another PE and resume endpoint execution. The endpoint must not make any assumptions about the PE it runs on.

This version of the Framework requires the following:

- UP endpoints must be capable of migrating.
- Execution contexts of MP endpoints could be capable of migrating between PEs or could be fixed to a particular PE. The latter are called *pinned contexts*.
- The migration capability must be specified in the endpoint manifest (see [3.2 Partition manifest at virtual FF-A instance](#)).
- S-EL0 partitions must be UP.

The number of execution contexts an endpoint implements can differ from the number of PEs in the system. This must be specified in the manifest of the endpoint (see [2.10 System resource management](#)). For example, a VM in the Normal world must use the manifest to inform the Hypervisor how many vCPUs it implements. The Hypervisor must maintain an execution context for each vCPU.

2.10 System resource management

Components in the Firmware Framework require access to the following system resources.

- Memory regions.
- Devices.
- CPU cycles.

The Framework associates the attributes of *ownership* and *access* with these resources. The Owner governs the following capabilities of non-Owners for each resource.

- The level of access a non-Owner has for using the resource. This could be exclusive, shared or no-access.
- The ability to grant access to the resource to other non-Owners. This is called access forwarding.

Also, the Owner could relinquish ownership to another component.

The Framework also specifies the transitions that result in a change of ownership and access attributes associated with a resource. A combination of these attributes and transitions determines how a resource is managed among components.

Rules associated with ownership and access of memory regions are described in [Chapter 5 Memory Management](#).

Rules associated with ownership and access of CPU cycles are described in [2.11 Primary scheduler](#).

For a device that is upstream of an SMMU, its access to the physical address space is managed using the rules associated with management of memory regions (also see [5.2 Direct memory access](#)).

For all devices, ownership and access attributes are associated with its *MMIO* region. A partition could request access and/or ownership of a device through its manifest (see [Table 3.3](#)). This is done through one of the following ways.

- A partition requests ownership and exclusive access to the MMIO region of a device during boot time (see [Chapter 3 Partition setup](#)). The corresponding partition manager assigns the MMIO region with these attributes to the partition.
- One or more partitions request access to the MMIO region of a device during boot time. The corresponding partition manager is the Owner of the MMIO region and grants access to all the partitions.

This version of the Framework does not permit:

- Ownership of a device MMIO region to be transferred to another partition during run-time.
- Access to a device MMIO region to be granted to another partition during run-time.
- Access to a device MMIO region to be revoked from a partition during run-time.

2.11 Primary scheduler

The primary scheduler is a module which is responsible for ensuring that all FF-A components are allocated CPU cycles for message processing on any PE in the system.

In terms of ownership and access rules, it is the Owner of CPU cycles and lends them to *threads of execution* it manages. It does this in response to a *scheduling decision*.

In a scheduling decision, the scheduler chooses a thread of execution to run on a particular PE in the system as per the scheduling policy. The thread of execution is chosen from a set of *runnable* threads. Any thread that is *not runnable* is either *blocked* or *sleeping* until some event occurs.

The first type of thread of execution considered in the Framework are *Application threads*. These can run inside any FF-A component for example, an application that provides streaming media services inside the OS Kernel.

The second type of thread of execution considered in the Framework are *Execution context threads*. Each thread corresponds to an execution context of an endpoint in the system for example, For every VM managed by a Hypervisor, it implements an execution context thread for each vCPU of the VM.

The primary scheduler can manage one or both types of threads. An endpoint that is allocated CPU cycles by the primary scheduler could implement a *secondary scheduler* to manage the allocated cycles among its application threads.

Figure 2.7 illustrates an example of a primary scheduler that manages both types of threads of execution and a VM that manages application threads inside itself.

The Framework assumes that execution context threads are managed only by the primary scheduler. See 4.3 *Indirect messaging usage* for details.

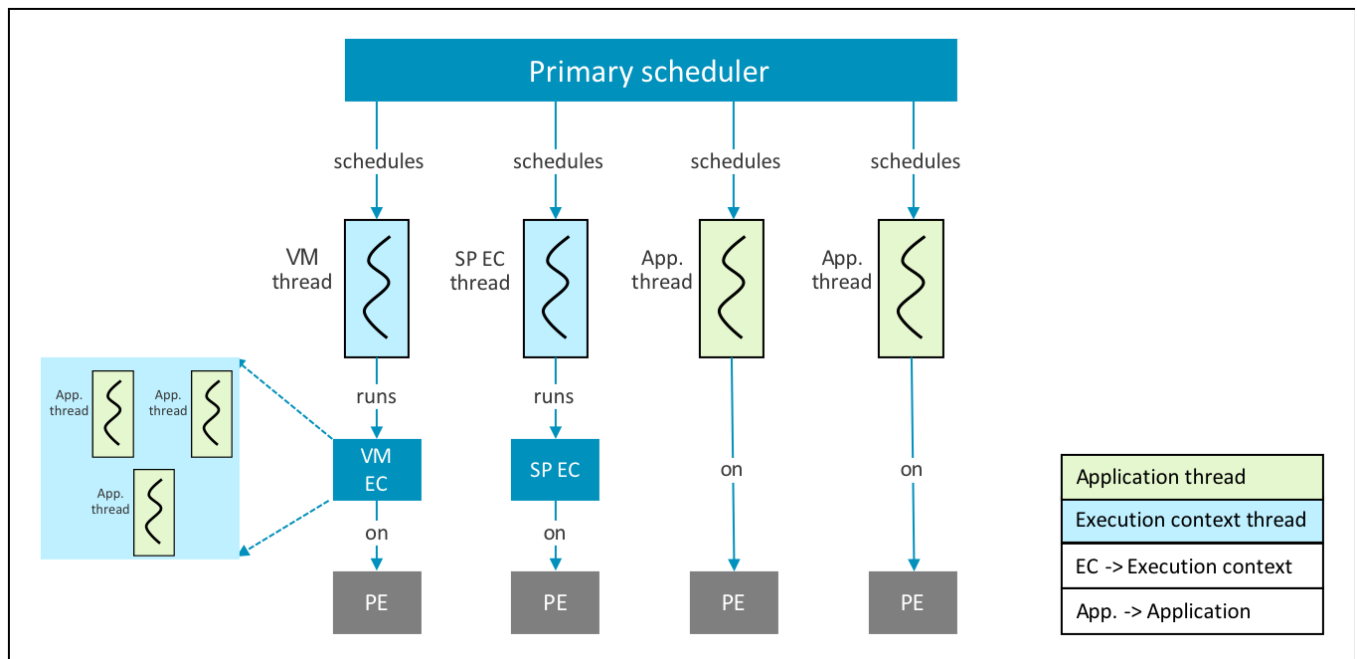


Figure 2.7: Example primary scheduler configuration

The primary scheduler could be a part of:

- The OS kernel running in EL1 if the Virtualization extension is not used in the Normal world
- The Host OS running in EL2 in the case of a Type 2 Hypervisor when the Virtualization host extension is used.

- The Host OS running in EL1 in the case of a Type 2 Hypervisor when the Virtualization host extension is not implemented or used (see [5])
- A separate VM running in EL1 that has been delegated the responsibility of scheduling by the Hypervisor.
- The Hypervisor running in EL2

The following assumptions have been made regarding the role of the primary scheduler.

- There is a single primary scheduler across all PEs in the system.
- An instance of the primary scheduler executes on each PE in the system.

Endpoint services could be accessed before the primary scheduler is initialized. At boot time, one or more software components are responsible for booting the primary scheduler. The Framework assumes that each component executes a boot stage on the primary CPU. Ownership of cycles is relayed from one component to the next across these boot stages. Each component lends cycles to an endpoint if it accesses the services of the endpoint.

2.12 Run-time states

Run-time refers to the stage during system boot when all the endpoints are initialized and application threads in an endpoint can access services implemented in other endpoints or partition managers through FF-A functions.

During run-time, the execution context of an endpoint can be in one of the following states from the perspective of the primary scheduler, SPM, and Hypervisor. These states are used to describe messaging mechanisms in [Chapter 4 Message passing](#).

- *Idle*. The execution context is waiting to be allocated CPU cycles to do work.
- *Busy*. The execution context has been allocated CPU cycles and is doing work for example, running an application thread to process one or more messages.
- *Preempted*. The execution context was preempted by an interrupt while doing work.

Chapter 3

Partition setup

3.1 Overview

The Firmware Framework is responsible for partition setup during boot time. *Boot time* refers to the stage during system boot before an endpoint is able to provide its services to other FF-A components. Given a valid partition manifest and image, a partition manager must be able to set up all the execution contexts of a partition so that it can provide its services by using the Firmware Framework. The list of actions to set up a partition depends on the FF-A instance at which the partition resides.

- A non-exhaustive list of actions that a partition manager must perform to set up an endpoint at a virtual FF-A instance is as follows.
 1. Validate the contents of a partition manifest and use them as follows. The contents of the manifest are described in [3.2 Partition manifest at virtual FF-A instance](#).
 - Configure the partition as per the properties described in the manifest.
 - Assign the requested physical address space ranges and system resources to the partition.
 - Isolate the partition by ensuring it only has visibility of resources that it has requested.
 2. Initialize the partition on the boot CPU and non-boot CPUs (if required). This is described in [3.5 Partition boot protocol](#).
- In the absence of a Hypervisor at the Non-secure physical FF-A instance, the OS kernel is set up through an IMPLEMENTATION DEFINED mechanism. Guidance provided in this section could be used to perform the setup.
- An SP could co-reside with the S-EL1 or S-Supervisor SPMC at the Secure physical FF-A instance as described in [2.2.2.1 S-EL1 SPM core component](#) and [2.2.2.2 Secure Supervisor mode SPM core component](#). A non-exhaustive list of actions in this case is as follows:
 - The SPMD is responsible for initializing the SP and the SPMC on both boot and non-boot CPUs.
 - The partition manifest is used to specify the SP properties. It is described in [3.3 SP manifest at physical FF-A instance](#).

3.2 Partition manifest at virtual FF-A instance

The following information must be specified in the manifest of a partition.

- Partition properties as described in [Table 3.1](#).
- Memory regions as described in [Table 3.2](#).
- Devices as described in [Table 3.3](#).
- Partition boot protocol as described in [Table 3.10](#).

The following aspects of the partition manifest are IMPLEMENTATION DEFINED.

- Format of the manifest.
- Time of creation of manifest. This could be at:
 - Build time.
 - Boot time.
 - Combination of both.
- Mechanism used by the Hypervisor and SPM to obtain the information in the manifest and interpret its contents.

Table 3.1: Partition properties

Information fields	Mandatory	Description
FF-A version	Yes	<ul style="list-style-type: none"> • Version of FF-A expected by the partition at the FF-A instance it will execute.
UUID	Yes	<ul style="list-style-type: none"> • UUID of service implemented by this partition. • UUID can be shared by multiple instances of partitions that offer the same service. • For example, <ul style="list-style-type: none"> – If there are multiple instances of a Trusted OS, then the UUID can be shared by all instances. – The TEE driver in the HLOS can use the UUID with the <code>FFA_PARTITION_INFO_GET</code> interface to determine the: <ul style="list-style-type: none"> * Number of Trusted OSs. * The partition ID of each instance of the Trusted OS.
Partition ID	No	<ul style="list-style-type: none"> • Pre-allocated partition ID.
Auxiliary IDs	No	<ul style="list-style-type: none"> • List of pre-allocated 16-bit IDs that could be used in memory management transactions to allow a partition manager to handle the transaction in an IMPLEMENTATION DEFINED manner.
Name	No	<ul style="list-style-type: none"> • Name of the partition for example, for debugging purposes.
Number of execution contexts	Yes	<ul style="list-style-type: none"> • Number of vCPUs that a VM or SP wants to instantiate. • In the absence of virtualization, this is the number of execution contexts that a partition implements. • If value of this field = 1 and number of PEs > 1 then the partition is treated as UP & migrate capable. • If the value of this field > 1 then the partition is treated as an MP capable partition irrespective of the number of PEs.

Information fields	Mandatory	Description
Run-time EL	Yes	<ul style="list-style-type: none"> • EL1 or Secure EL1. • Secure EL0.
Execution state	Yes	<ul style="list-style-type: none"> • AArch64. • AArch32.
Load address	No	<ul style="list-style-type: none"> • Absence of this field indicates that the partition is position independent and can be loaded at any address chosen at boot time.
Entry point offset	No	<ul style="list-style-type: none"> • Absence of this field indicates that the entry point is at offset 0x0 from the base of the partition binary image. • If present, this field specifies the offset of the entry point from the base of the partition binary image.
Translation Granule	No	<ul style="list-style-type: none"> • 4KB (default value if not specified). • 16KB. • 64KB.
Boot order	No	<ul style="list-style-type: none"> • A unique number among all partitions that specifies if this partition must be booted before others. • For example, a partition could provide a service that other partitions need to initialize themselves. The manifest of this partition can use this field to ensure it is booted before others.
RX/TX information	No	<ul style="list-style-type: none"> • Reference to memory region entries in this manifest that describes the RX/TX buffers expected by the partition. • The memory region entries must specify the base addresses of both buffers. • The size and attributes fields must fulfill the requirements specified in 4.2.2.3 Buffer attributes.
Messaging method	Yes	<ul style="list-style-type: none"> • This field specifies which messaging methods are supported by the partition. This could be one or both of direct and indirect messaging. These methods are described in Chapter 4 Message passing. The following information must be provided in the manifest: • Indirect messaging is supported. This always includes support for both sending and receiving indirect messages. • Direct messaging is supported. 4.4.1 Discovery and setup specifies the information that must be provided. • Managed exits are supported. 4.5.4 Managed exit specifies the information that must be provided.
Primary Scheduler implemented	No	<ul style="list-style-type: none"> • Presence of this field indicates that the partition implements the primary scheduler. • Run-time EL must be EL1 if this field is specified.

Information fields	Mandatory	Description
Run-time model	No	<ul style="list-style-type: none"> If the run-time EL is S-EL0 or User mode then this field specifies the run-time model that the SPM must enforce for this SP. <ul style="list-style-type: none"> <i>Run to completion</i>. SP execution must not be preempted. An execution context of this SP must only transition between the <i>idle</i> and <i>busy</i> states described in 2.12 Run-time states. <i>Preemptible</i>. SP execution can be preempted. An execution context of this SP can transition between all states described in 2.12 Run-time states. This is the default run-time model for a S-EL0/User mode SP if this field is not specified in the partition manifest.
Tuples of (Name, SEPID, SMMU ID, Stream IDs)	No	<ul style="list-style-type: none"> If present, then each tuple specifies the association between its members that the partition manager must create. The members are as follows. <ul style="list-style-type: none"> <i>Stream endpoint ID</i> that this endpoint is a <i>proxy</i> for. The dependent device must not be assigned to this endpoint (see 5.2.1 Stream endpoint). <i>SMMU ID</i> identifies the SMMU instance on a system with multiple SMMUs. One or more <i>Stream IDs</i> associate the device that generates them with the <i>SEPID</i> in the SMMU identified by <i>SMMU ID</i>. An optional <i>Name</i> for the SEPID for debugging purposes.

Table 3.2: Memory regions

Information fields	Mandatory	Description
Base address	No	<ul style="list-style-type: none"> Absence of this field indicates that a memory region of specified size and attributes must be mapped into the partition translation regime. The PM must describe the memory region to the partition through an IMPLEMENTATION DEFINED mechanism. If present, this field could specify a PA, VA (for S-EL0 partitions) or IPA (for S-EL1 and EL1 partitions). This information must be specified using an IMPLEMENTATION DEFINED mechanism. <ul style="list-style-type: none"> If a PA is specified, then the memory region must be identity mapped with the same IPA or VA as the PA. If a VA or IPA is specified, then the memory could be identity or non-identity mapped. If present, the address must be aligned to the Translation granule size.
Page count	Yes	<ul style="list-style-type: none"> Size of memory region expressed as a count of 4K pages. For example, if the memory region size is 16K, value of this field is 4.

Information fields	Mandatory	Description
Attributes	Yes	<ul style="list-style-type: none"> Memory access permissions. <ul style="list-style-type: none"> Instruction access permission. Data access permission. Memory region attributes. <ul style="list-style-type: none"> Memory type. Shareability attributes. Cacheability attributes. Memory Security state. <ul style="list-style-type: none"> Non-secure for a NS-Endpoint. Non-secure or Secure for an S-Endpoint.
Name	No	<ul style="list-style-type: none"> Name of the memory region for example, for debugging purposes.

Table 3.3: Device regions

Information fields	Mandatory	Description
Physical base address	Yes	<ul style="list-style-type: none"> PA of base of a device MMIO region. If the MMIO region is not physically contiguous, then an entry for each physically contiguous constituent region must be specified. Each entry must specify the PA and size of the constituent region. The size must be expressed as a count of 4K pages.
Page count	Yes	<ul style="list-style-type: none"> Total size of MMIO region expressed as a count of 4K pages. For example, if the MMIO region size is 16K, value of this field is 4.
Attributes	Yes	<ul style="list-style-type: none"> Memory attributes must be Device-nGnRnE. Instruction access permission must be not executable. Data access permissions must be one of the following: <ul style="list-style-type: none"> Read/write. Read-only. Security attributes must be: <ul style="list-style-type: none"> Non-secure for a NS-Endpoint. Non-secure or Secure for an S-Endpoint.

Information fields	Mandatory	Description
Interrupts	No	<ul style="list-style-type: none"> List of physical interrupt IDs. Attributes of each interrupt ID. <ul style="list-style-type: none"> Interrupt type. <ul style="list-style-type: none"> SPI. PPI. SGL. Interrupt configuration. <ul style="list-style-type: none"> Edge triggered. Level triggered. Interrupt Security state. <ul style="list-style-type: none"> Secure. Non-secure. Interrupt priority value. Target execution context/vCPU for each SPI. <ul style="list-style-type: none"> This field is optional even if other interrupt properties are specified since interrupt affinity could be managed through an IMPLEMENTATION DEFINED interface between the endpoint and its partition manager.
SMMU ID	No	<ul style="list-style-type: none"> If present, then on a system with multiple SMMUs, this field must help the partition manager determine which SMMU instance is this device upstream of. Absence of this field implies that the device is not upstream of an SMMU.
Stream IDs	No	<ul style="list-style-type: none"> List of Stream IDs assigned to this device. Absence of Stream ID list indicates that the device is not upstream of an SMMU.
Exclusive access and ownership	No	<ul style="list-style-type: none"> If present, this field implies that this endpoint must be granted exclusive access and ownership of the MMIO region of the device. Absence of this field implies that access to the MMIO region of the device could be shared among multiple endpoints.
Name	No	<ul style="list-style-type: none"> Name of the device region for example, for debugging purposes.

3.3 SP manifest at physical FF-A instance

[Table 3.4](#) describes the fields in the manifest that must be used by the SPMD to set up an SP that co-resides with the SPMC and executes in S-EL1 or Secure Supervisor mode at the physical FF-A instance.

The following aspects of the partition manifest are IMPLEMENTATION DEFINED.

- Format of the manifest.
- Time of creation of the manifest. This could be at:
 - Build time.
 - Boot time.
 - Combination of both.
- Mechanism used by the SPMD to obtain information in the manifest and interpret its contents.

Table 3.4: Properties of an SP at physical FF-A instance

Information fields	Mandatory	Description
FF-A version	Yes	<ul style="list-style-type: none"> • Version of Firmware Framework implemented by the SPMC component in the SP. See 8.1 FFA_VERSION for more information about the usage of this field.
UUID	No	<ul style="list-style-type: none"> • UUID to identify the single SP.
Partition ID	No	<ul style="list-style-type: none"> • Pre-allocated partition ID.
Name	No	<ul style="list-style-type: none"> • Name of the partition for example, for debugging purposes.
Execution state	Yes	<ul style="list-style-type: none"> • AArch64. • AArch32.
Load address	No	<ul style="list-style-type: none"> • Absence of this field indicates that the partition is position independent and can be loaded at any address chosen at boot time.
Entry point offset	No	<ul style="list-style-type: none"> • Absence of this field indicates that the entry point is at offset 0x0 from the base of the SP binary image. • If present, this field specifies the offset of the entypoint from the base of the SP binary image.
FF-A boot protocol usage	No	<ul style="list-style-type: none"> • See Table 3.10.

3.4 Independent peripheral device manifest

This manifest must be used by *independent peripheral devices* to describe their properties to a partition manager. See [5.2 Direct memory access](#) for more details.

Table 3.5: Device properties

Information fields	Mandatory	Description
FF-A version	Yes	<ul style="list-style-type: none">• Version of the Firmware Framework expected by the device.
Name	No	<ul style="list-style-type: none">• Name of the partition for example, for debugging purposes.
Translation Granule	Yes	<ul style="list-style-type: none">• 4KB.• 16KB.• 64KB.
SEPID	Yes	<ul style="list-style-type: none">• Pre-allocated Stream endpoint ID.

Table 3.6: Memory regions accessible by the device

Information fields	Mandatory	Description
Base address	Yes	<ul style="list-style-type: none">• This field could specify a PA or IPA. This distinction must be specified using an IMPLEMENTATION DEFINED mechanism.<ul style="list-style-type: none">– If a PA is specified, then the memory region must be identity mapped with the same IPA as the PA.– If a IPA is specified, then the memory could be identity or non-identity mapped.• The address must be aligned to the Translation granule size.
Page count	Yes	<ul style="list-style-type: none">• Size of memory region expressed as a count of 4K pages.• For example, if the memory region size is 16K, value of this field is 4.
Properties	Yes	<ul style="list-style-type: none">• Memory region properties (see 5.11 Memory region properties).• Security attributes.<ul style="list-style-type: none">– Non-secure for a Non-secure device.– Non-secure or Secure for a Secure device.
Name	No	<ul style="list-style-type: none">• Name of the memory region for example, for debugging purposes.

Table 3.7: Device regions

Information fields	Mandatory	Description
Physical base address	Yes	<ul style="list-style-type: none"> • PA of base of a device MMIO region. • If the MMIO region is not physically contiguous, then an entry for each physically contiguous constituent region must be specified. • Each entry must specify the PA and size of the constituent region. The size must be expressed as a count of 4K pages.
Properties	Yes	<ul style="list-style-type: none"> • Memory type must be Device-nGnRnE. • Instruction access permission must be not executable. • Data access permissions must be one of the following: <ul style="list-style-type: none"> – Read/write. – Read-only. • Security attributes must be: <ul style="list-style-type: none"> – Non-secure for a Non-secure device. – Non-secure or Secure for a Secure device.
Page count	Yes	<ul style="list-style-type: none"> • Total size of MMIO region expressed as a count of 4K pages. • For example, if the MMIO region size is 16K, value of this field is 4.
SMMU ID	Yes	<ul style="list-style-type: none"> • On a system with multiple SMMUs, this field must help the partition manager determine which SMMU instance is this device upstream of.
Stream IDs	Yes	<ul style="list-style-type: none"> • List of Stream IDs assigned to this device.
Name	No	<ul style="list-style-type: none"> • Name of the device region for example, for debugging purposes.

3.5 Partition boot protocol

The partition manager must program an entry into each execution context of a partition as follows.

3.5.1 Register state

The partition manager must program system and general-purpose registers that influence partition execution as follows.

- The MMU must be disabled for a partition that does not run in S-EL0 in either Execution state. The MMU must be enabled for S-EL0 partition that runs in either Execution state.
- The partition manager must ensure that all memory regions allocated to a partition are clean to the Point of Coherency. Also, there must be no stale cached copies of executable memory held in any instruction caches visible to a PE on which the execution contexts of the partition may execute.

This could be achieved by executing cache maintenance instructions, after initialising the memory regions for a partition.

- The state of other System registers is IMPLEMENTATION DEFINED. If the partition manager must program a System register to fulfill a specific partition requirement then this must be encoded in its manifest through an IMPLEMENTATION DEFINED mechanism.
 - For example, an S-EL0 partition could want the instruction alignment check to be disabled by setting SCTLR_EL1.A, bit[1] = b'0.
- The state of general-purpose registers is IMPLEMENTATION DEFINED. Also see [Table 3.10](#).

3.5.2 Protocol for passing data

The partition manager could also pass an array of *name-value-size* pairs to a partition execution context when it is entered. This information is encoded in an initialization descriptor specified in [Table 3.9](#).

- The partition must specify the information it expects to be populated in an initialization descriptor in its manifest through an IMPLEMENTATION DEFINED mechanism.
- In this version of the Firmware Framework, it is assumed that information in an initialization descriptor is passed only to the first execution context of a partition that is initialized by the partition manager.
- The partition manager must fulfill the following requirements for the memory region where an initialization descriptor is populated.
 - Size of memory region must be a multiple of the translation granule size used by the partition.
 - Address of memory region must be aligned to the translation granule size used by the partition.
 - The memory region must be mapped in the translation regime of the partition that is managed by the Hypervisor (see [2.1 Partition manager](#)) or SPM (see [2.2 SPM architecture](#)).
 - The memory region must be mapped with the same memory attributes as the RX/TX buffers as described in [4.2.2.3 Buffer attributes](#) in the partition translation regime managed by the Hypervisor or SPM.
 - Boot information must be populated at offset 0 in the memory region
 - The address of boot information must be passed in the general-purpose register specified in the partition manifest (see [Table 3.10](#)).

Table 3.8: Name value size tuple descriptor

Field	Byte length	Byte offset	Description
Name	16	0	• Name of an object passed to the partition.

Field	Byte length	Byte offset	Description
Value	8	16	<ul style="list-style-type: none"> Value of the object identified by the <i>Name</i> field.
Size	8	24	<ul style="list-style-type: none"> Size of the object identified by the <i>Name</i> field.

Table 3.9: Initialization descriptor

Field	Byte length	Byte offset	Description
Signature	4	0	<ul style="list-style-type: none"> ASCII string “FF-A” to identify this descriptor.
NVS count	4	4	<ul style="list-style-type: none"> Count of <i>Name value size</i> tuple descriptors.
NVS array	–	8	<ul style="list-style-type: none"> Array of <i>Name value size</i> tuple descriptors. See Table 3.8.

Table 3.10: Boot protocol information

Information fields	Mandatory	Description
FF-A boot protocol usage	No	<ul style="list-style-type: none"> Presence of this field indicates that the partition expects the address of an initialization descriptor to be passed in a general-purpose register (see Table 3.9). The register in which the address of an initialization descriptor will be passed must be specified. Register must be between <i>w0/x0-w3/x3</i>. The width of the register is derived from its Execution state specified in the partition manifest.

3.5.3 Protocol to initialize an execution context

As a part of initializing a partition, the partition manager must program an entry into the first execution context of the partition. This execution context is hereby called the *boot execution context*. If the partition is an MP endpoint, initialization of other execution contexts must be done through an IMPLEMENTATION DEFINED mechanism.

For example, The Hypervisor is responsible for initializing a VM. It initiates this process by programming an entry into the boot execution context corresponding to a vCPU of the VM. The vCPU could then use the *PSCI_CPU_ON* interface described in the PSCI specification [7] to request the Hypervisor to initialize another vCPU of the VM.

A partition must use the *FFA_MSG_WAIT* (also see [9.1 FFA_MSG_WAIT](#)) interface or an IMPLEMENTATION DEFINED mechanism to indicate completion of initialization of the boot execution context to the partition manager.

A partition must use the *FFA_ERROR* (also see [7.2 FFA_ERROR](#)) interface or an IMPLEMENTATION DEFINED mechanism to report an error during initialization of the boot execution context to the partition manager.

Chapter 4

Message passing

4.1 Overview

The Firmware Framework defines a set of ABIs that enable FF-A components to exchange messages with each other. A message exchange comprises of the following phases.

1. Transmission of the message payload from the Sender to the Receiver. The mechanisms specified by the Framework to do this are described in [4.2 Message transmission](#).
2. Allocation of CPU cycles to the Receiver to process the message on a PE in the system. Cycles could be allocated by the Sender or the primary scheduler. These methods are described in [4.1.1 Indirect messaging](#) & [4.1.2 Direct messaging](#). The partition manager must make at least one method available to an endpoint it manages.
3. Message processing using the allocated cycles. The role of the Framework during message processing is described in [4.5 Partition message processing](#).

FF-A components (also see [2.3 FF-A instances](#)) participate in a message exchange by playing one or more of the following roles.

- *Sender*.
- *Receiver*.
- *Relayer*.
 - Validates and forwards a message to the Receiver.
 - Provide access to CPU cycles to the Receiver to process the message.
- *Primary scheduler*. See [2.11 Primary scheduler](#).

[Table 4.1](#) specifies the roles that each component can play in a message exchange. The responsibilities of a component in a role depend on the type and phase of a message exchange. This is described in detail in subsequent sections. In all messaging scenarios, in the absence of a Hypervisor, the SPM subsumes its responsibilities. In a configuration with the S-EL1 or S-Supervisor SPMC, the role of the SPM as a Relayer is subsumed by the SP.

Table 4.1: FF-A component roles in a message exchange

Config. No.	FF-A Component	Sender	Receiver	Relayer	Scheduler
1.	NS-Endpoint	Yes	Yes	No	Yes
2.	S-Endpoint	Yes	Yes	No	No
3.	Hypervisor	Yes	Yes	Yes	Yes
4.	SPM	Yes	Yes	Yes	No

4.1.1 Indirect messaging

In this method, the message Sender requests the primary scheduler to allocate CPU cycles to the Receiver for processing the message. The Receiver is scheduled by the primary scheduler. The Sender could make progress concurrently with the Receiver either on the same or a different PE. The Sender either polls or is notified when a response from the Receiver is available.

The term **Indirect messaging** is used to describe this method for CPU cycle allocation along with interfaces to transmit the message payload. A detailed description of the usage of this method is provided in [4.3 Indirect](#)

messaging usage. [Figure 4.1](#) illustrates this method. It assumes that both the Sender and Receiver run on the same PE.

In this method, the primary scheduler manages the execution context threads of the Receiver (also see [2.11 Primary scheduler](#)). It chooses an execution context of the Receiver and allocates cycles to it for message processing.

The Framework assumes that this method is only used for passing messages between VMs. All message passing between the Normal and Secure world and within the Secure world must be done through direct messaging see [4.1.2 Direct messaging](#).

The following ABIs are used to implement indirect messaging between endpoints.

- *FFA_MSG_SEND*. This interface is used by a Sender to send a message payload to a Receiver (also see [10.1 FFA_MSG_SEND](#)).
- *FFA_RUN*. This interface is used by the primary scheduler to allocate CPU cycles and hand control to a Receiver for message processing (also see [9.3 FFA_RUN](#) and [4.5 Partition message processing](#)).
- *FFA_MSG_WAIT*. This interface is used by a Receiver to indicate that it is ready to receive a new message (also see [9.1 FFA_MSG_WAIT](#)).
- *FFA_MSG_POLL*. This interface is used by a Receiver to poll for a new message (also see [10.4 FFA_MSG_POLL](#)).

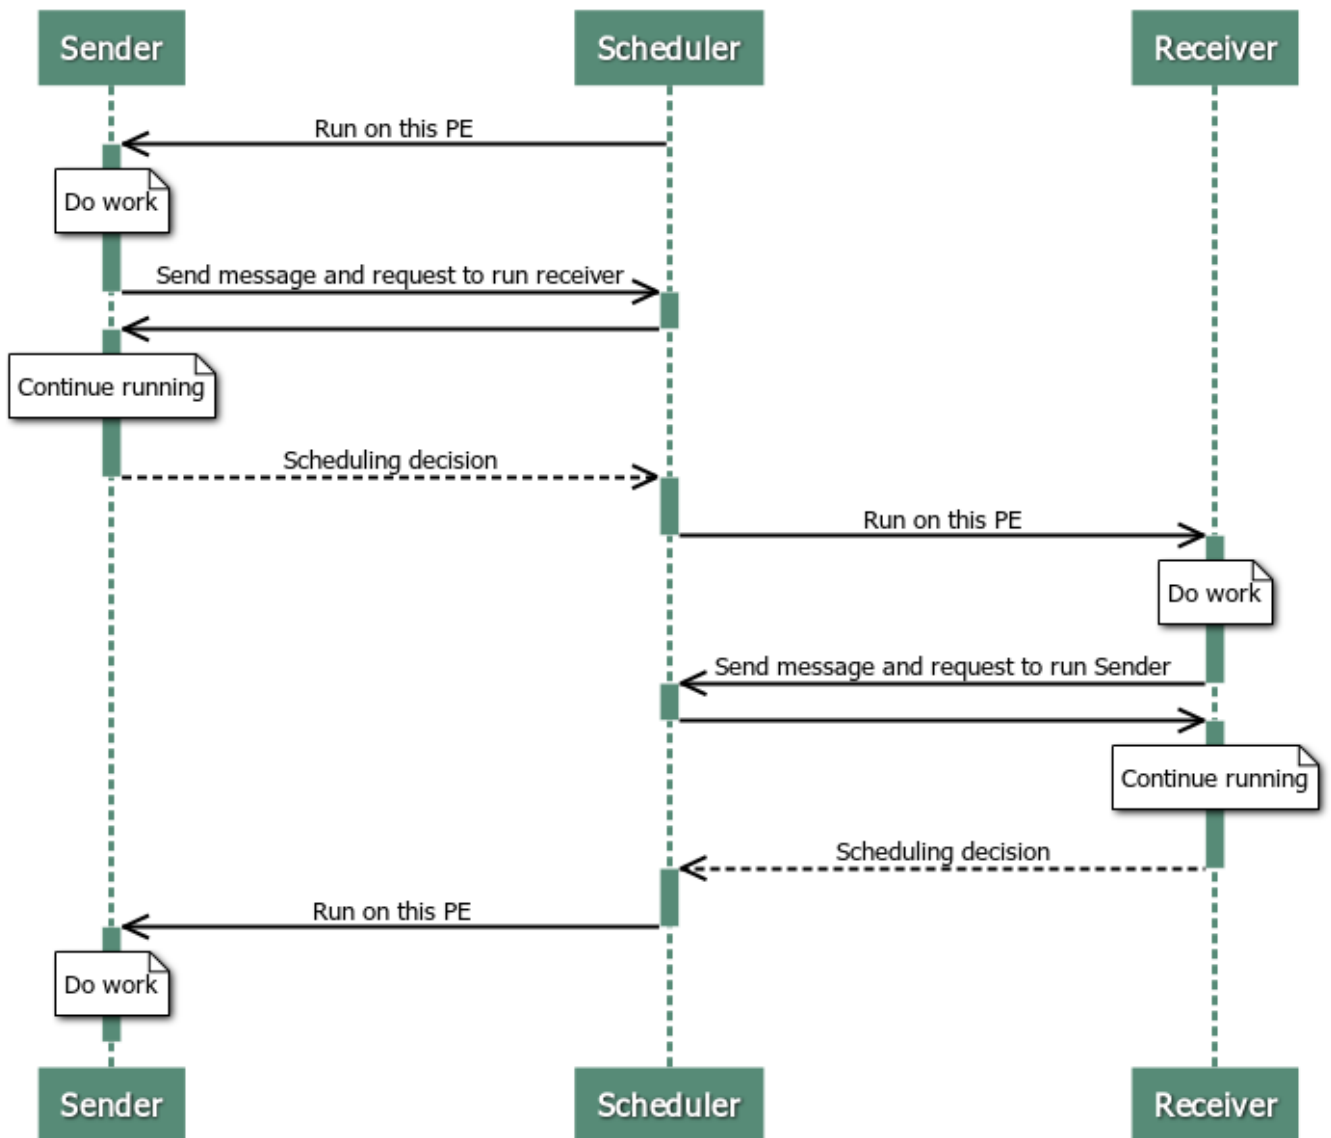


Figure 4.1: Indirect messaging

4.1.2 Direct messaging

In this method, the message Sender lends CPU cycles to the Receiver so that it can make progress. The Receiver is *scheduled* by the Sender. The Sender does not make progress until either a response is returned or execution is returned back to it. Both the Sender and Receiver run on the same PE.

The term **Direct messaging** is used to describe this method for CPU cycle allocation along with interfaces to transmit the message payload. A detailed description of this method is provided in [4.4 Direct messaging usage](#). [Figure 4.2](#) illustrates this method.

This method is used for messaging between an endpoint and the Hypervisor or SPM. It is also used for messaging between endpoints before the primary scheduler is initialized and when the primary scheduler cannot be involved in the message exchange for example, it is not possible to communicate with the scheduler, the latency associated with scheduler communication cannot be tolerated by the use case or the Receiver must be run on the same PE as the Sender.

The Framework allows this method to be used for passing messages in only the following scenarios.

- Between FF-A endpoints under the constraints described in [4.4 Direct messaging usage](#).
- From an endpoint to the Hypervisor or SPM.
- From the Hypervisor or SPM to an endpoint.
- Between the Hypervisor and SPM.

The following ABIs are used to implement direct messaging between endpoints.

- *FFA_MSG_SEND_DIRECT_REQ*. This interface is used by a Sender to send a request message payload to a Receiver, lend CPU cycles to the Receiver and wait for a response to arrive. Also see [10.2 FFA_MSG_SEND_DIRECT_REQ](#).
- *FFA_MSG_SEND_DIRECT_RESP*. This interface is used by a Sender to send a response message payload to a Receiver, return CPU cycles to the Receiver and wait for a new message to arrive. Also see [10.3 FFA_MSG_SEND_DIRECT_RESP](#).
- *FFA_INTERRUPT*. This interface is used by the Relay to inform the Sender that direct message processing in the Receiver was preempted.
- *FFA_RUN*. This interface is used by the Sender to resume a preempted Receiver.

The following is a non-exhaustive list of ABIs used to implement direct messaging between endpoints and Hypervisor or SPM and between the Hypervisor and SPM.

- FFA_VERSION.
- FFA_RXTX_MAP.
- FFA_RXTX_UNMAP.
- FFA_MEM_DONATE.
- FFA_MEM_SHARE.
- FFA_MEM_LEND.
- FFA_MEM_RETRIEVE_REQ.
- FFA_MEM_RETRIEVE_RESP.
- FFA_MEM_RELINQUISH.
- FFA_MEM_RECLAIM.

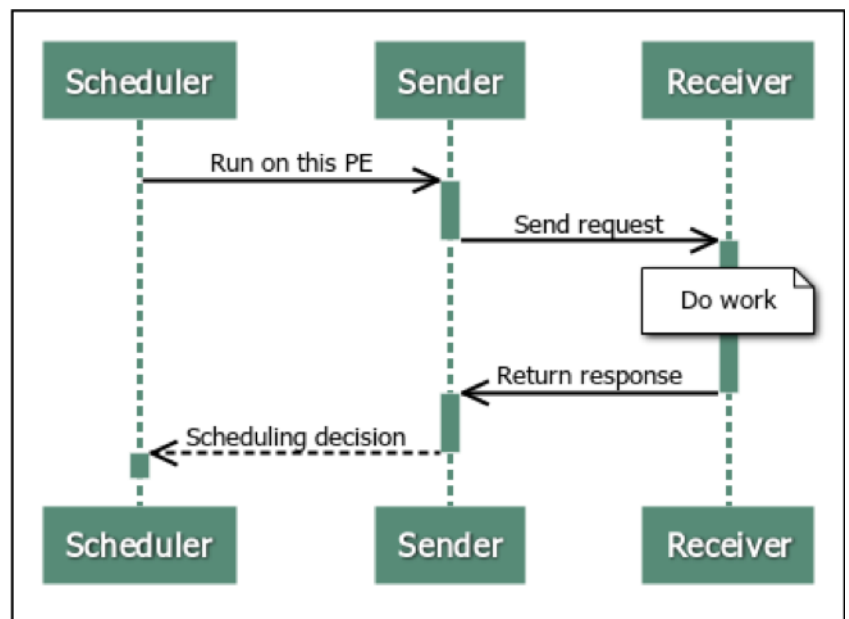


Figure 4.2: Direct messaging

4.2 Message transmission

4.2.1 Overview

Message payloads are exchanged between two FF-A components through general purpose registers or a single pair of shared memory regions to transmit and receive messages called *RX/TX buffers* (see also [4.2.2 RX/TX buffers](#)).

- Direct messaging can use both these mechanisms along with the ABIs described in [4.1.2 Direct messaging](#).
- Indirect messaging must use only the RX/TX message buffers along with the ABIs described in [4.1.1 Indirect messaging](#).

Each message has a header and a payload. The header describes the:

- Source and target of the message.
- Size and type of message payload.

The version of the message header and payloads is the same as the version of the Firmware Framework as returned by **FFA_VERSION** (see [8.1 FFA_VERSION](#)).

The header is encoded in the parameter registers of the ABI used for message transmission. This information is used to validate and route the message and decide if the message payload must be interpreted by a Relayer.

There are two types of message payloads.

1. Payloads that are defined by the communication Framework for example, memory management messages. They have the same definition in any implementation of a particular version of the Firmware Framework. Messages with these payloads are called **Framework messages**.

Framework message payloads can be interpreted by the Relayer, Sender and Receiver. They are used when:

- Relayer participation is required to validate or modify message contents before delivery to the Receiver.
- The Hypervisor or SPM is the destination of the message payload. It processes the message and provides a response.

In this version of the Firmware Framework, Framework messages are exchanged only in the following scenarios.

- Between an endpoint and Hypervisor or SPM.
- Between the Hypervisor or SPM and an endpoint.
- Between the Hypervisor and SPM.
- Between the SPM and Hypervisor.

2. Payloads that are defined by the services implemented inside a partition. The format of these messages is specific to the service or partition implementation. Messages with these payloads are called **Partition messages**.

Partition message payloads are only interpreted by the Sender and Receiver endpoints. A Relayer only uses the header information to route them correctly. Hence, by definition these messages are only exchanged between partitions.

The properties of Framework and Partition messages influence direct and indirect messaging as follows.

- Direct messaging can be used to transmit both Framework and partition messages. Framework messages can be transmitted in both RX/TX buffers and registers. Partition messages can only be transmitted in registers.
- Indirect messaging can be used to only transmit Partition messages in the RX/TX buffers.

[Table 4.2](#) lists examples of ABIs available for transmitting messages of both types in registers or RX/TX buffers using both types of messaging methods.

Table 4.2: Combinations of messaging and message transmission mechanisms

Messaging method	Message type	Message payload location	Message transmission interface
Direct	Partition	Register	<ul style="list-style-type: none"> • FFA_MSG_SEND_DIRECT_REQ. • FFA_MSG_SEND_DIRECT_RESP.
Direct	Partition	RX/TX	<ul style="list-style-type: none"> • Invalid usage.
Direct	Framework	Register	<ul style="list-style-type: none"> • Any interface to send or receive information from the Hypervisor or SPM for example, <ul style="list-style-type: none"> – FFA_VERSION. – FFA_RX_RELEASE. – FFA_YIELD. – FFA_RXTX_MAP. – FFA_RXTX_UNMAP. – FFA_RUN.
Direct	Framework	RX/TX	<ul style="list-style-type: none"> • Any interface to send or receive information from the Hypervisor or SPM for example, <ul style="list-style-type: none"> – FFA_MEM_DONATE. – FFA_MEM_SHARE. – FFA_MEM_LEND. – FFA_MEM_RELINQUISH. – FFA_MEM_RETRIEVE_REQ. – FFA_MEM_RETRIEVE_RESP. – FFA_MEM_RECLAIM.
Indirect	Partition	Register	<ul style="list-style-type: none"> • Invalid usage.
Indirect	Partition	RX/TX	<ul style="list-style-type: none"> • FFA_MSG_SEND.
Indirect	Framework	Register	<ul style="list-style-type: none"> • Invalid usage.
Indirect	Framework	RX/TX	<ul style="list-style-type: none"> • Invalid usage.

4.2.2 RX/TX buffers

A RX/TX buffer pair is shared between two FF-A components at an FF-A instance.

- The FF-A component at the lower EL is the *Consumer* of the RX buffer and *Producer* of the TX buffer.
- The FF-A component at the higher EL is the *Producer* of the RX buffer and the *Consumer* of the TX buffer.
- An FF-A component is permitted to share only a single RX/TX buffer pair with another FF-A component at an FF-A instance.

The endianness of all message payloads populated in the RX/TX buffers is *little-endian*.

In the Normal world,

- Each VM has a Non-secure buffer pair. Also see [4.2.2.3 Buffer attributes](#).
- Each VM shares its buffer pair with the Hypervisor and SPM.
- The OS kernel shares its buffer pair with the SPM.

- The Hypervisor shares its buffer pair with the SPM.

In the Secure world,

- Each SP has a Secure buffer pair. Also see [4.2.2.3 Buffer attributes](#).
- Each SP shares its Secure buffer pair with the SPM.
- The SPM is split into the SPMD and SPMC components as described in [2.2 SPM architecture](#). In configurations where the SPMC resides in a separate Exception level from the SPMD (see [Table 2.1](#) & [Table 2.2](#)), it is IMPLEMENTATION DEFINED whether the two SPM components share an RX/TX buffer pair.

These message buffer configurations are illustrated in [Figure 4.3](#).




Figure 4.3: Configurations of RX/TX buffer pair between FF-A components

Mechanisms for message transmission through RX/TX buffers are described in [4.2.2.1 Buffer-based message transmission](#).

Mechanisms for discovery and setup of a RX/TX buffer pair are described in [4.2.2.2 Buffer discovery and setup](#).

Requirements for correctly mapping a RX/TX buffer pair in the translation regimes of both FF-A components at any FF-A instance are described in [4.2.2.3 Buffer attributes](#).

4.2.2.1 Buffer-based message transmission

A message is transmitted between VMs by copying it from the TX buffer of the Sender VM to the RX buffer of the Receiver VM. The message copy is done by the Hypervisor.

A message is transmitted between the Hypervisor and SPM through the RX/TX buffer pair shared between them. Both the Hypervisor and SPM can be the Sender or Receiver. A TX to RX copy is not required as both components can access the single buffer pair.

A message is transmitted between a VM as a Sender and Hypervisor or SPM as Receiver through the Non-secure RX/TX buffer pairs they share. The following message copies could be required in this case.

- A copy from the TX buffer of the VM to the RX buffer shared between the Hypervisor and SPM.
- A copy from the TX buffer shared between the Hypervisor and SPM to the RX buffer of the VM.

If the message payload is modified by the Hypervisor before being forwarded to the SPM, the preceding message copies are required since the VM can see a modified message payload in its TX buffer.

If the message payload is forwarded unmodified to the SPM from the Hypervisor or to the Hypervisor from the SPM, the preceding message copies are not required since the SPM can access the buffer pair of the VM.

A message is transmitted between an SP as a Sender and SPM as Receiver through the Secure RX/TX buffer they share. A TX to RX copy is not required as the SPM can access the Secure buffer pair of the SP.

A message is transmitted between an SP as a Sender and Hypervisor as Receiver by copying it from the TX buffer of the SP to the TX buffer shared between the Hypervisor and SPM since the Hypervisor cannot access the Secure buffer pair of the SP.

4.2.2.2 Buffer discovery and setup

This version of the Framework enables discovery and setup of RX/TX buffer pairs between FF-A components as follows.

1. The Hypervisor or SPM could allocate the buffer pair on behalf of a VM or SP respectively and convey this information to the endpoint through an IMPLEMENTATION DEFINED mechanism.

The endpoint must specify this requirement in its manifest by specifying the base addresses (as IPAs or VAs) of where it expects its RX/TX buffer pair to be mapped by the SPM or Hypervisor as applicable (see also [3.2 Partition manifest at virtual FF-A instance](#)).

The SPM is not permitted to allocate a buffer pair on behalf of the Hypervisor or a VM.

2. An endpoint could allocate the buffer pair and use the **FFA_RXTX_MAP** interface to map it with the Hypervisor or SPM as applicable.
 3. An endpoint could use the **FFA_RXTX_UNMAP** interface to unmap a buffer pair from the Hypervisor or SPM as applicable.
 4. The Hypervisor must allocate the buffer pair it shares with the SPM and use the **FFA_RXTX_MAP** interface to map it with the SPM.
 5. The Hypervisor could use the **FFA_RXTX_UNMAP** interface to unmap the buffer pair from the SPM.
 6. The Hypervisor must forward an invocation of the **FFA_RXTX_MAP** interface from a VM to map its RX/TX buffer with the SPM.
 7. The Hypervisor must forward an invocation of the **FFA_RXTX_UNMAP** interface from a VM to unmap its RX/TX buffer from the SPM.
- See [8.4 FFA_RXTX_MAP](#) for a description of the **FFA_RXTX_MAP** interface.
 - See [8.5 FFA_RXTX_UNMAP](#) for a description of the **FFA_RXTX_UNMAP** interface.

[Figure 4.4](#) illustrates an example RX/TX buffer setup with a single VM and SP using the preceding interfaces where the:

- SPM allocates the buffer pair on behalf of the SP.
- Hypervisor registers its buffer pair with the SPM.
- VM registers its buffer pair with the Hypervisor and SPM.
- VM unregisters its buffer pair with the Hypervisor and SPM.




Figure 4.4: RX/TX Buffer setup

4.2.2.3 Buffer attributes

The size of the RX and TX buffers in a pair must be the same and a multiple of the larger translation granule size used by the FF-A components at an FF-A instance.

The alignment of the RX and TX buffers in a pair must be equal to the larger translation granule size used by the FF-A components at an FF-A instance (see also [2.7 Memory granularity and alignment](#)).

An endpoint must discover the minimum size and alignment boundary for the RX/TX buffers by passing the function ID of the *FFA_RXTX_MAP* ABI as input in the *FFA_FEATURES* interface (see [8.2 FFA_FEATURES](#)).

The Framework assumes that the memory attributes of all RX and TX buffers in the system are configured as follows.

- If a stage of address translation is enabled in a translation regime from where the buffer is accessed, it must be mapped with the following memory region attributes in that stage.
 - Normal memory.
 - Write-Back Cacheable.
 - Non-transient Read-Allocate.
 - Non-transient Write-Allocate.
 - Inner Shareable.
 - Buffers shared between an SP and SPM must have the NS bit = 0 in their translation table descriptors.
 - Buffers shared between an VM, Hypervisor and SPM must have the NS bit = 1 in their translation table descriptors.
- [Table 4.3](#) describes the minimum permission requirements of RX/TX buffer.

Table 4.3: RX/TX buffer minimum permission requirements

Buffer Type	Producer	Consumer	Description
RX	RW, XN	RO, XN	<ul style="list-style-type: none"> • Producer must have write access to populate message payload. • Consumer must have at least read access to read message payload.
TX	RW, XN	RO, XN	<ul style="list-style-type: none"> • Producer must have Write-access to populate message payload. • Consumer must have at least read access to copy the message payload to the target RX buffer. • Consumer must also have Write- access to modify message payload if required.

- A buffer pair could be accessed with different memory region attributes from the translation regime of the Producer and Consumer, if address translation is disabled in one of them.

To avoid memory coherency issues in this scenario, the FF-A component that has address translation disabled must perform cache maintenance on the buffer in scenarios listed in [Table 4.4](#). The cache maintenance must ensure that the buffer contents at any intermediate cache levels are not out of sync with the buffer contents at the *Point of coherence* (see [\[5\]](#)).

- As a Producer, this must be done before the Consumer reads the buffer (see [4.2.2.4 Buffer synchronization](#)).
- As a Consumer, this must be done before reading the buffer populated by the Producer.

Table 4.4: RX/TX buffer cache maintenance requirements

Config No.	Address translation in Producer	Address translation in Consumer	Cache maintenance required
1.	Disabled	Disabled	No
2.	Disabled	Enabled	Yes
3.	Enabled	Disabled	Yes
4.	Enabled	Enabled	No

4.2.2.4 Buffer synchronization

The RX and TX buffers are written to by a Producer and read by a Consumer as described in [Table 4.5](#). Concurrent accesses to these buffers from both entities on either side of an FF-A instance must be synchronized to preserve the integrity of their contents.

Table 4.5: Producers and Consumers of RX/TX buffers

Buffer Type	Producers	Consumers
VM RX	Hypervisor	VM
VM TX	VM	Hypervisor, SPM
OS Kernel RX	SPM	OS Kernel
OS Kernel TX	OS Kernel	SPM
SP RX	SPM	SP
SP TX	SP	SPM
Hypervisor RX	SPM	Hypervisor
Hypervisor TX	Hypervisor	SPM

The Framework defines states, access, and ownership rules that must be followed by the Producer and Consumer of each buffer.

- Each buffer can be either *empty* or *full* (has a message in it) at any given time. This state must be tracked internally by the Producer and Consumer using an IMPLEMENTATION DEFINED mechanism.
- The Producer must not assume it can read from or write to the buffer when it does not own it.
- The Consumer must not assume it can read from or write to the buffer when it does not own it.
- The Producer of a buffer owns it when it is empty.
- The Consumer of a buffer owns it when it is full.
- The Producer must write to the buffer only when it is empty.
- The Consumer must read from the buffer only when it is not empty.

After a Producer has written to a buffer, it must transfer its ownership to the Consumer for reading the message.

The mechanism used for this depends on the following factors.

1. **Exception level of each entity.** The Producer could reside in a higher Exception level than the Consumer and vice-versa. This in turn governs the conduits available to signal transfer of ownership. These are described in the [Table 4.6](#).
2. **Type of Secure partition.** Architectural constraints on the highest EL in which an SP runs determine which conduit can be used to signal the transfer of ownership for example, a S-EL1 SP can use all conduits described in [Table 4.6](#) while a S-EL0 SP cannot use the interrupt conduit.

Table 4.6: Conduits to signal transfer of buffer ownership

	Producer to Consumer	Consumer to Producer
Producer at lower EL	SMC, SVC, HVC	ERET
Producer at higher EL	Interrupt or ERET	SMC, SVC, HVC

As Producers, the SPM and Hypervisor could use an interrupt to indicate availability of the RX buffer they share with each partition they manage. They must describe this interrupt to each partition during its initialization using an IMPLEMENTATION DEFINED mechanism for example, a device tree.

Transfer of ownership of the TX buffer takes place as follows.

- Through an invocation of the **FFA_MSG_SEND** interface from the Producer to the Consumer.
- Through an invocation of the **ERET** instruction from the Consumer to complete a previous invocation of the **FFA_MSG_SEND** interface from the Producer.
- Through a memory management interface (see [Chapter 5 Memory Management](#)) that could use the TX buffer to transmit information about the memory management operation. This is as follows.
 - An invocation of such an interface transfers ownership of the TX buffer from the Producer to the Consumer. The interfaces are as follows.
 - * **FFA_MEM_DONATE.**
 - * **FFA_MEM_LEND.**
 - * **FFA_MEM_SHARE.**
 - * **FFA_MEM_RETRIEVE_REQ.**
 - * **FFA_MEM_RELINQUISH.**
 - * **FFA_MEM_FRAG_TX.**
 - The completion of a preceding interface transfers ownership of the TX buffer from the Consumer to the Producer.

Transfer of ownership of the RX buffer takes place as follows.

- Through an invocation of the **FFA_MSG_SEND** function through the **ERET** conduit from the Producer to the Consumer.
- Through an interrupt pended by the Producer to indicate to the Consumer that the RX buffer is full.
- Through an invocation of the **FFA_RX_RELEASE** or **FFA_MSG_WAIT** interfaces from the Consumer to the Producer.
- Through an invocation of the **FFA_MEM_RETRIEVE_RESP** interface from the Producer to the Consumer.
- Through the completion of a previous invocation of the **FFA_PARTITION_INFO_GET** function from the Consumer to the Producer.

U

Implementation Note

A buffer could be shared among multiple Producers, Consumers, and multiple instances of the same Producer and Consumer (also see [Table 4.5](#)). Both the Producers and the Consumers must use an IMPLEMENTATION DEFINED

synchronization mechanism to protect the buffer from concurrent accesses that are internal to them. A Producer or Consumer could implement additional states internally to prevent concurrent accesses. Such states are outside the scope of this version of the Firmware Framework.

For example, multiple instances of the SPM will run concurrently on different PEs. As the Producer for an RX buffer or as a Consumer for a TX buffer, the SPM could use a spinlock to protect each buffer from accesses made concurrently by its own instances.

4.2.2.5 Example buffer synchronization flows

This section illustrates examples of how the states and mechanisms to transfer ownership of a buffer can be used in an implementation.

[Figure 4.5](#) illustrates interaction between a VM and the Hypervisor for transferring ownership of the RX buffer through the **ERET** instruction and **FFA_RX_RELEASE** interface.

[Figure 4.6](#) illustrates interaction between a VM and the Hypervisor for transferring ownership of the RX buffer through an interrupt and **FFA_RX_RELEASE** interface.

[Figure 4.7](#) illustrates interaction between a VM and the Hypervisor for transferring ownership of the TX buffer through the **FFA_MSG_SEND** interface and **ERET** instruction.

The following aspects of these interactions have been assumed.

- The same interactions can be applied to a RX/TX buffer pair shared between the Hypervisor and SPM as well.
- The VM and Hypervisor are MP-capable. This means that it has multiple instances that can run concurrently on separate physical PEs.

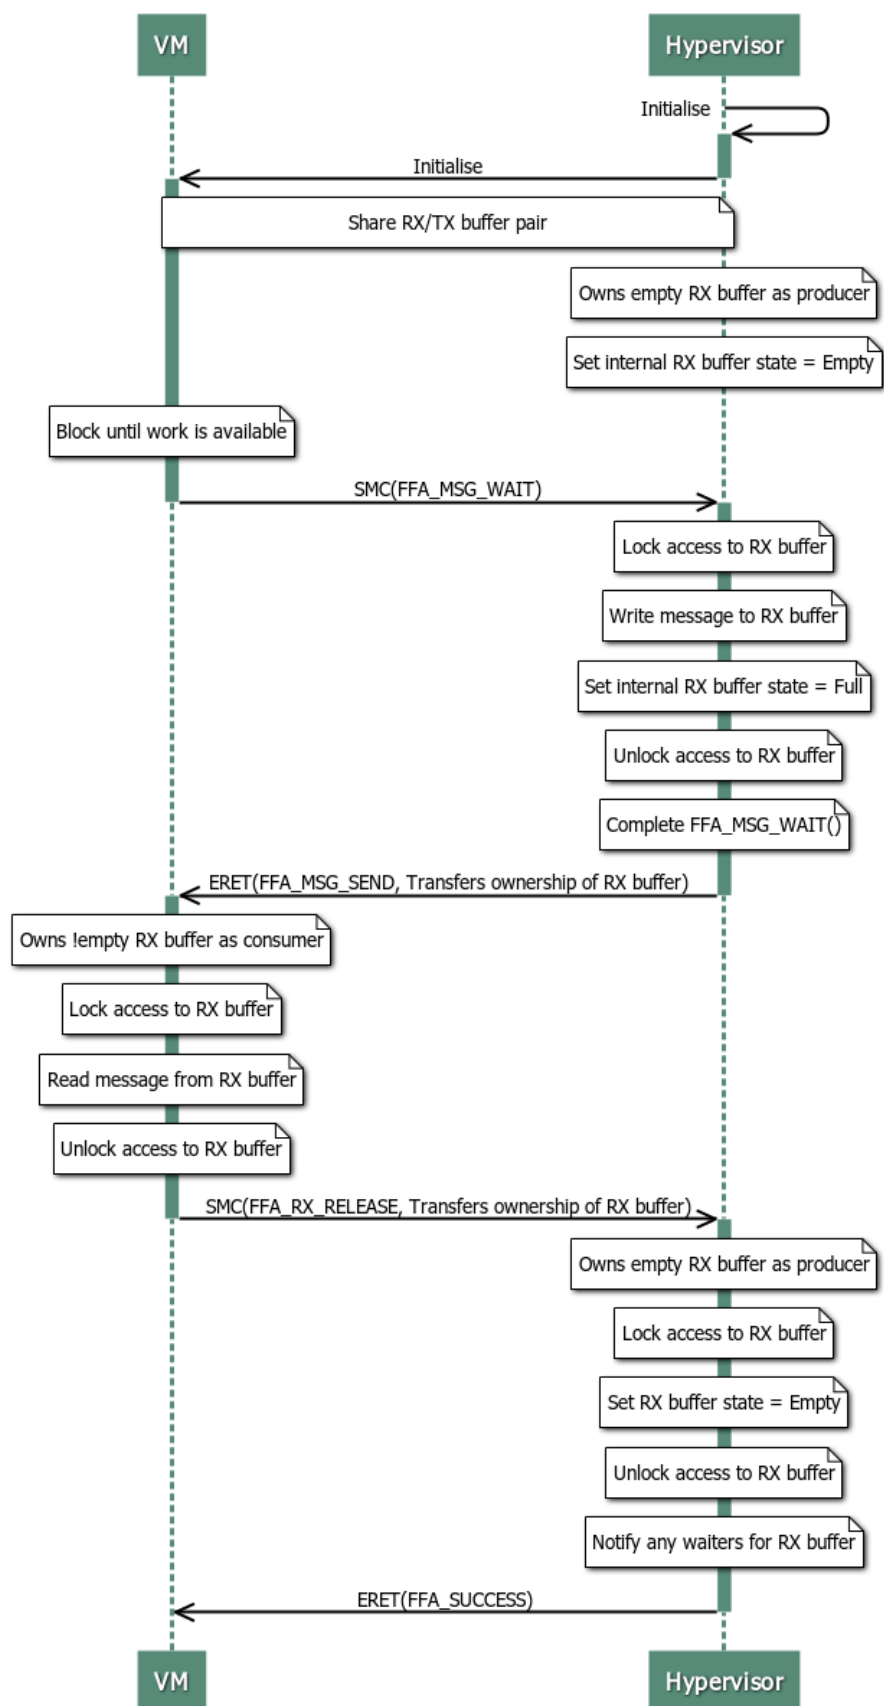


Figure 4.5: Buffer synchronization with Producer at higher EL



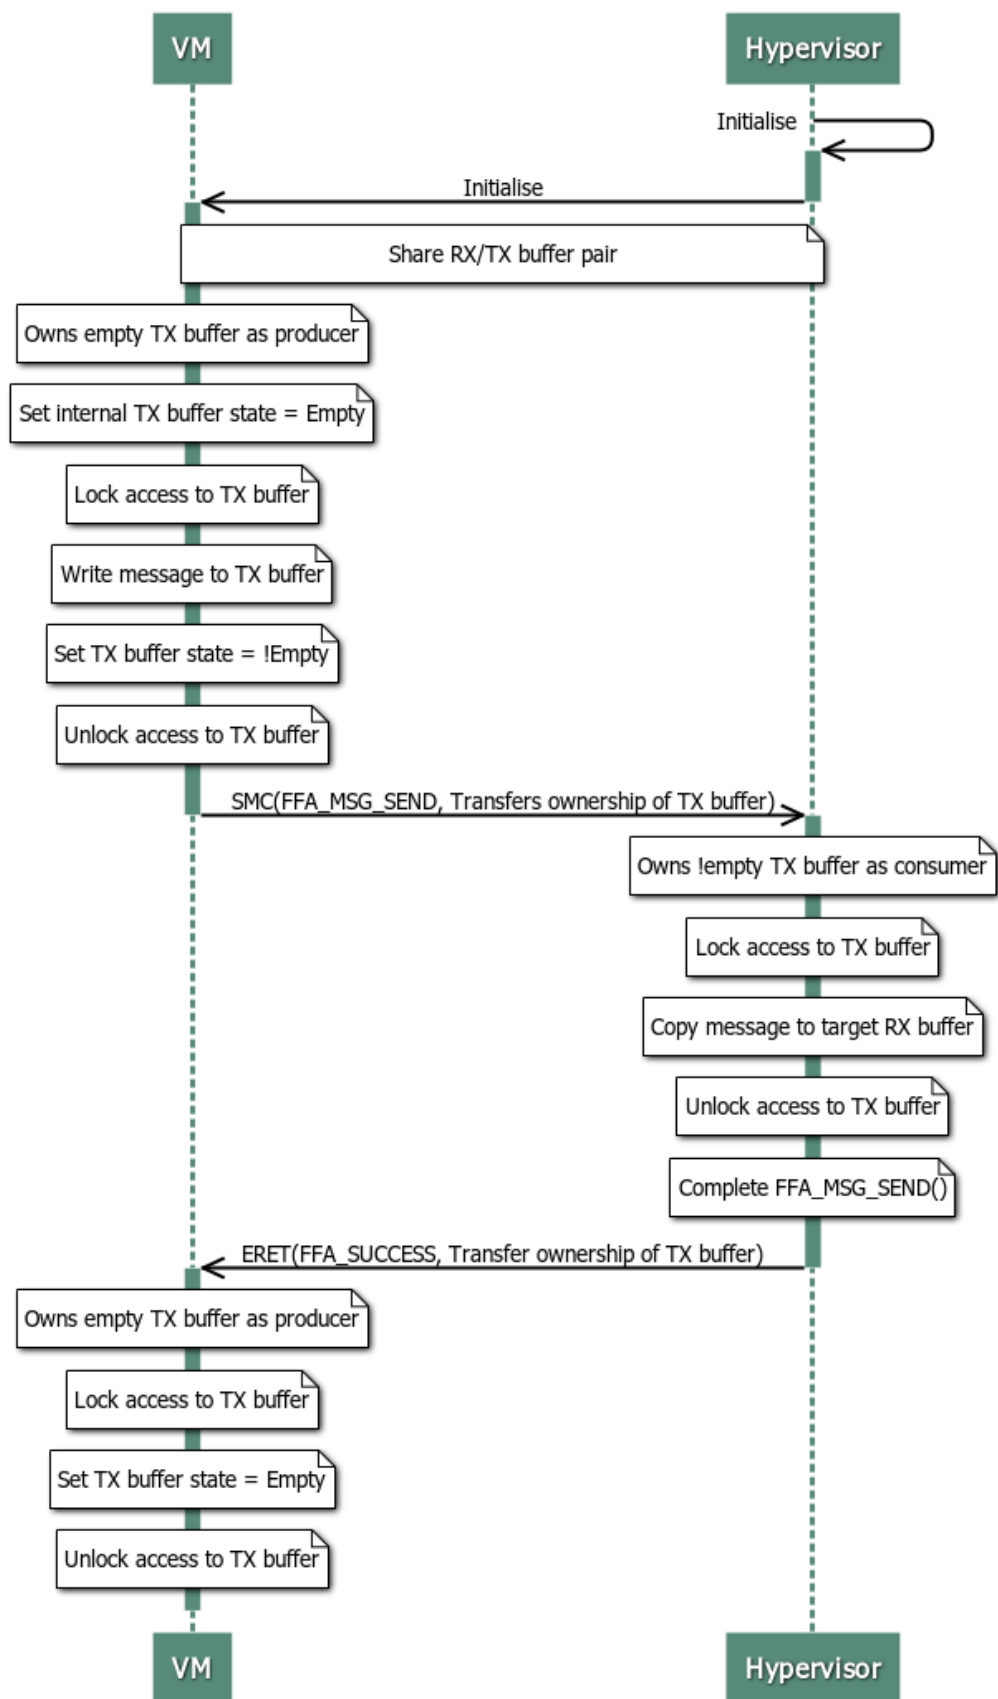


Figure 4.7: Buffer synchronization with Producer at lower EL

4.3 Indirect messaging usage

There are three phases in an indirect message exchange

- Transmission of message from Sender to Receiver.
- Notification to the primary scheduler that the Receiver is runnable.
- Allocation of CPU cycles by the primary scheduler to the Receiver to process the delivered message.

An FF-A component could play one or more of the roles listed in [Table 4.1](#). A list of configurations that result from valid combinations of these roles for indirect messaging is specified in [Table 4.7](#).

In each configuration, the primary scheduler could be resident either in a NS-Endpoint or the Hypervisor. This implies that the primary scheduler could be:

- Co-resident with the Sender if the Sender is in the Normal world.
- Co-resident with the Receiver if the Receiver is in the Normal world.
- Co-resident with the Relay that is, the Hypervisor.
- Resident in a VM different from the Sender and Receiver.

Table 4.7: Valid configurations for indirect messaging

Config no.	Sender	Receiver	Relayer
1.	VM	VM	Hypervisor

4.3.1 Discovery and setup

An endpoint that can receive messages through indirect messaging must specify this property in its manifest (see [3.2 Partition manifest at virtual FF-A instance](#)). The primary scheduler must be aware of the presence of this component and the number of execution contexts it implements. The Hypervisor could provide this information to the primary scheduler through an IMPLEMENTATION DEFINED mechanism for example, Device tree. Alternatively, the FF-A component that implements the primary scheduler could use the **FFA_PARTITION_INFO_GET** interface to obtain this information.

An attempt to send a direct message to a VM that only supports indirect messaging must be rejected by the Hypervisor.

4.3.2 Message delivery and scheduler notification

The Framework defines the **FFA_MSG_SEND** interface to transmit a message from a Sender to a Receiver and notify the primary scheduler that the Receiver is runnable.

[10.1 FFA_MSG_SEND](#) describes the **FFA_MSG_SEND** interface. [10.1.2 Component responsibilities for FFA_MSG_SEND](#) describes how a message is transmitted using this interface and the responsibilities of the participating components in all the configurations listed in [Table 4.7](#).

The Sender must notify the primary scheduler to schedule the Receiver to process the message in its RX buffer. The mechanism to do this depends on the FF-A component where the primary scheduler is implemented relative to the Sender and Relayer. An applicable mechanism in [10.1.3 Mechanism for scheduler notification](#) must be used. In all cases, after the scheduler has been notified, an invocation of the **FFA_MSG_SEND** call must be completed to enable the Sender to make progress.

[Figure 4.8](#) illustrates an example flow in which VM0 sends an indirect message to a VM1. The primary scheduler is resident in a different VM. The Hypervisor performs the TX to RX copy, notifies the scheduler and completes the **FFA_MSG_SEND** call from VM0. The primary scheduler runs VM1 during a scheduling decision to process the message sent by VM0.

4.3.3 Scheduling the Receiver

The primary scheduler must make the following choices after receiving the notification that the Receiver is runnable.

- Choose a PE on which CPU cycles will be allocated to the Receiver.
- Choose an execution context of the Receiver to which the cycles will be allocated on the chosen PE.

The mechanism used by the primary scheduler to run the Receiver depends on the FF-A component where the scheduler is implemented relative to the Receiver and Relayer. A combination of applicable mechanisms as follows must be used.

1. The primary scheduler resides in the Receiver VM. The scheduler must use an IMPLEMENTATION DEFINED mechanism to run the thread responsible for processing the message for example, use an OS primitive to run an application thread.
2. The primary scheduler is co-resident with the Hypervisor. If the Receiver is a VM, the Hypervisor is responsible for programming a return to the Receiver in response to a decision by the primary scheduler to run the Receiver. The method used to do this depends on the state of the Receiver (also see [2.12 Run-time states](#)).
 - If the execution context of the Receiver is in the *idle* state, the Relayer must complete the invocation of the interface used by the Receiver to enter this state for example, *FFA_MSG_WAIT*, *FFA_MSG_SEND_DIRECT_RESP*.
 - If the execution context of the Receiver is in the *preempted* state, the Relayer must resume its execution.
3. In all other scenarios, the primary scheduler must use the *FFA_RUN* interface to run the Receiver. [9.3.1 Component responsibilities for FFA_RUN](#) describes the responsibilities of the participating components in an invocation of this interface.

Once the Receiver starts processing the message after a scheduling decision, it could interact with the primary scheduler in ways described in [4.5 Partition message processing](#).

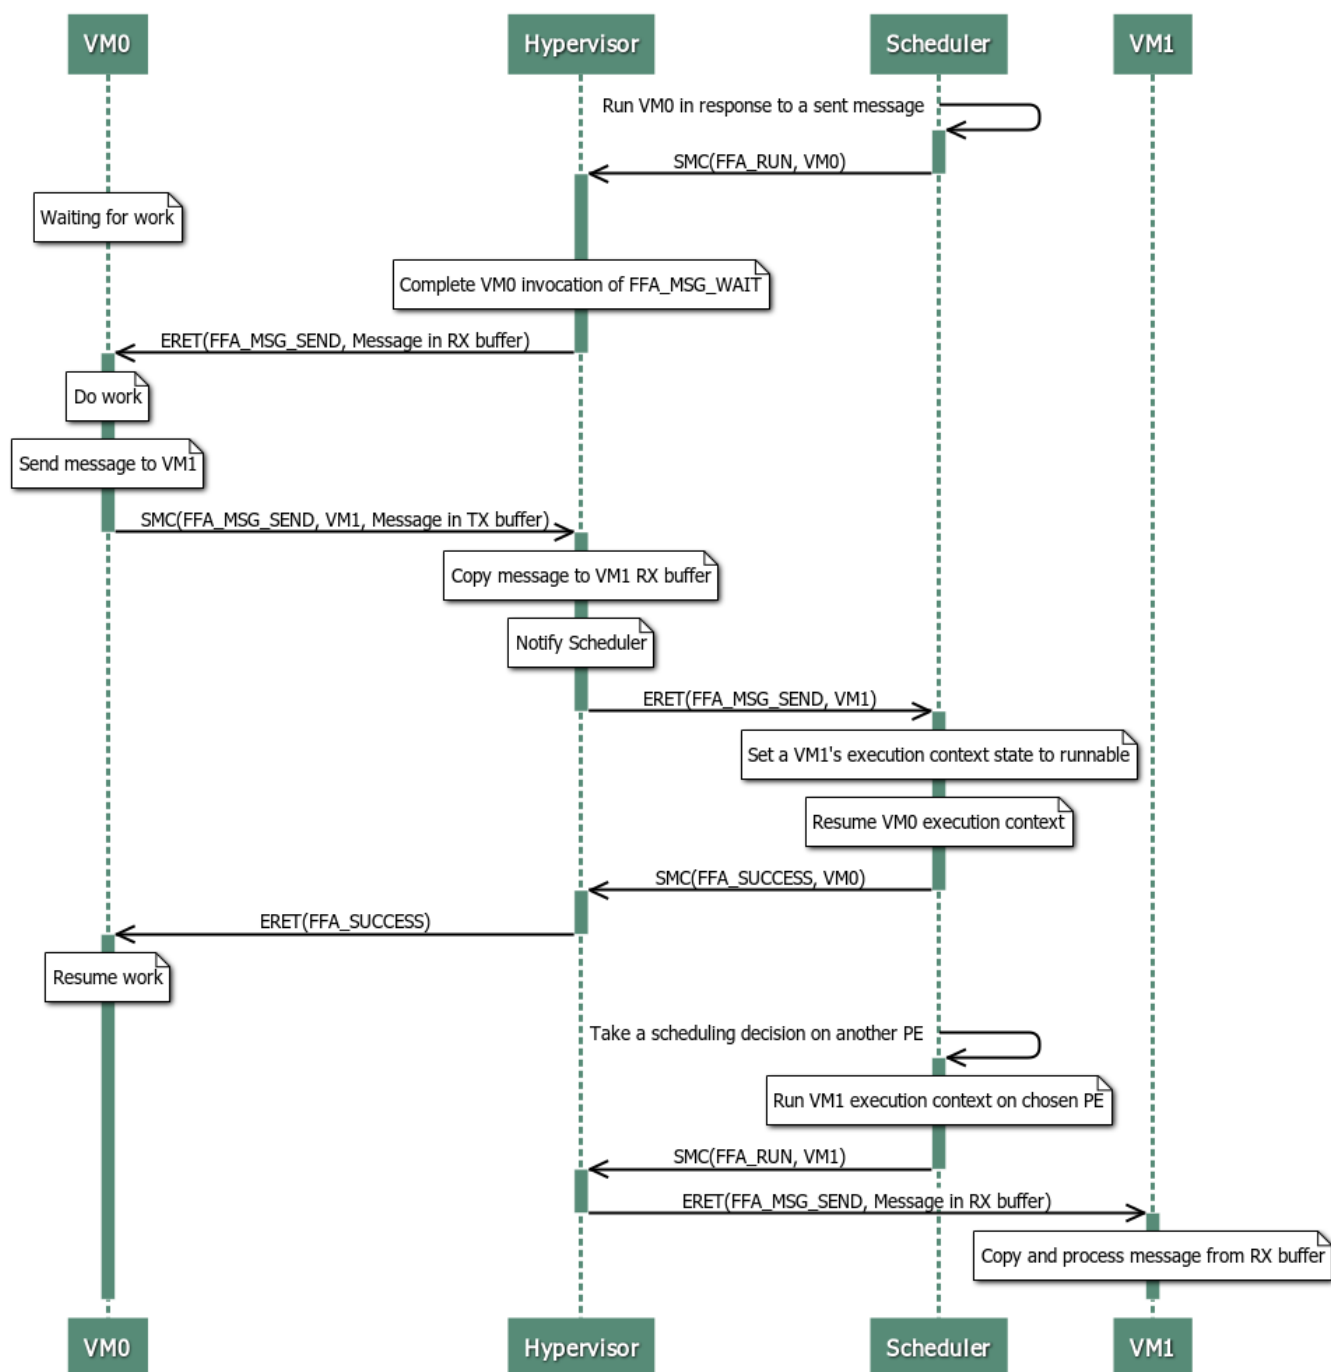


Figure 4.8: Indirect messaging flow between two VMs

4.4 Direct messaging usage

In a direct message exchange, transmission of the message from the Sender to the Receiver takes place in tandem with allocation of CPU cycles to the Receiver to process the message.

The Framework assumes that direct messaging is used by a Sender as an equivalent of invoking a procedure or function in the Receiver. The Receiver executes the function and returns the results through a direct message. For Framework messages, execution of the function in the Hypervisor or SPM runs to completion from the perspective of the Sender. For Partition messages, execution of the function in an endpoint could run to completion or be preempted by interrupts one or more times. In the latter case, the communication framework is responsible for resuming function execution.

An FF-A component could play one or more of the roles listed in [Table 4.1](#) except for the role of the primary scheduler during a direct message exchange. Direct messaging is used for:

1. Exchanging Framework messages with the Hypervisor and SPM in the configurations listed in [Table 4.8](#). These messages can be exchanged in both RX/TX buffers and registers.
2. Exchanging Partition messages between endpoints in the configurations listed in [Table 4.9](#). These messages can be exchanged only in registers.

Table 4.8: Valid configurations for exchanging Framework messages through direct messaging

Config no.	Sender	Receiver	Relayer
1.	VM	Hypervisor	•
2.	NS-Endpoint	SPM	Hypervisor (if present)
3.	SP	SPM	•
4.	SP	Hypervisor	SPM
5.	Hypervisor	VM	•
6.	Hypervisor	SPM	•
7.	Hypervisor	SP	SPM
8.	SPM	NS-Endpoint	Hypervisor (if present)
9.	SPM	Hypervisor	•
10.	SPM	SP	•

Table 4.9: Valid configurations for exchanging Partition messages through direct messaging

Config no.	Sender	Receiver	Relayer
1.	VM	VM	Hypervisor
2.	NS-Endpoint	SP	Hypervisor (if present) and SPM
3.	SP	SP	SPM
4.	SP	NS-Endpoint	SPM and Hypervisor (if present)

4.4.1 Discovery and setup

An endpoint could be capable of receiving direct messages, sending direct messages or both. *This property must be specified in the manifest of the endpoint (see Table 3.1 in 3.2 Partition manifest at virtual FF-A instance).* A Sender of direct requests must be able to receive direct responses. A Receiver of direct requests must be able to send direct responses.

When a direct message is sent to an endpoint capable of receiving it, it is possible that the primary scheduler cannot participate in choosing the execution context of the Receiver and the PE it is run on. To minimize conflict with the scheduling decisions of the primary scheduler and aid its availability on a PE, the Receiver must make one of the following implementation choices.

- The Receiver could be implemented as a **UP** endpoint. This enables the SPM or Hypervisor to migrate the endpoint execution context to the PE on which a direct messaging request is made.
- The Receiver could be implemented as a **MP** endpoint. In this case, the number of execution contexts that the endpoint implements must be equal to the number of PEs in the system. Each execution context must be pinned to a PE at system boot. This enables the SPM or Hypervisor to guarantee availability of an endpoint execution context for direct messages on the same PE as the Sender.

This implementation choice must be specified in the manifest of the endpoint (see Table 3.1 in 3.2 Partition manifest at virtual FF-A instance).

A partition manager can discover the properties of an endpoint it manages through the endpoint manifest. It can discover the properties of endpoints it does not manage through the **FFA_PARTITION_INFO_GET** interface (see 8.6 **FFA_PARTITION_INFO_GET**). An endpoint could use the same interface to determine properties of other endpoints as well.

An attempt to send an indirect message to an endpoint that only supports receipt of direct requests must be rejected as follows.

- By the Hypervisor if the Sender is a VM.
- By the SPM if the Sender is an SP, Hypervisor, or NS-Endpoint.

In this version of the Firmware Framework, a partition manager can only send and receive Framework messages through direct messaging. To support this model, it must be implemented as per the constraints listed as follows for an **MP** endpoint.

- It must have as many execution contexts as PEs in the system.
- Each execution context runs only on the PE where it was initialized during boot. Hence, it can be considered to be *pinned* to that PE.

In the SPM configuration where the SPMC coexists with an SP at S-EL1 or Secure Supervisor mode (see [Table 2.3](#)), the SP must be implemented as per the constraints that apply to the SPM implementation.

4.4.2 Message delivery and Receiver execution

This version of the Firmware Framework defines a number of interfaces for passing Framework messages with and between the Hypervisor and SPM. The description of these interfaces and the responsibilities of the participating components in all the configurations listed in [Table 4.8](#) is provided in [Chapter 6 Interface overview](#). A non-exhaustive list of these interfaces is as follows.

- FFA_VERSION.
- FFA_RXTX_MAP.
- FFA_RXTX_UNMAP.
- FFA_MEM_DONATE.
- FFA_MEM_SHARE.
- FFA_MEM_LEND.
- FFA_MEM_RETRIEVE_REQ.
- FFA_MEM_RETRIEVE_RESP.
- FFA_MEM_RELINQUISH.
- FFA_MEM_RECLAIM.

The rest of this section describes direct messaging between endpoints.

The Framework defines the **FFA_MSG_SEND_DIRECT_REQ** and **FFA_MSG_SEND_DIRECT_RESP** interfaces (also see [10.2 FFA_MSG_SEND_DIRECT_REQ](#) & [10.3 FFA_MSG_SEND_DIRECT_RESP](#)) to transmit a direct message from a Sender endpoint to a Receiver endpoint.

[10.2.1 Component responsibilities for FFA_MSG_SEND_DIRECT_REQ](#) describes how a message is transmitted using the *FFA_MSG_SEND_DIRECT_REQ* interface and the responsibilities of the participating components in all the configurations listed in [Table 4.9](#).

[10.3.1 Component responsibilities for FFA_MSG_SEND_DIRECT_RESP](#) describes how a message is transmitted using the *FFA_MSG_SEND_DIRECT_RESP* interface and the responsibilities of the participating components in all the configurations listed in [Table 4.9](#).

[Figure 4.9](#) illustrates an example flow in which a VM sends a direct message to an SP through the *FFA_MSG_SEND_DIRECT_REQ* interface. The SP processes the messages and returns the results using the *FFA_MSG_SEND_DIRECT_RESP* interface.




Figure 4.9: Example direct message exchange between a VM and SP

4.5 Partition message processing

The run-time model available to an endpoint for processing a Partition message depends on the method used for sending the message.

- [4.5.1 Indirect message processing](#) describes processing of an indirect message.
- [4.5.2 Direct message processing](#) describes processing of a direct message.

In both cases, an endpoint can exchange Framework messages with the Hypervisor and SPM. It is assumed that the processing of these messages runs to completion from the perspective of the endpoint.

4.5.1 Indirect message processing

After a VM starts processing an indirect message, one or more of the following events could occur.

1. It relinquishes control back to the primary scheduler after completing message processing. This is done using the `FFA_MSG_WAIT` interface and described in [9.1 FFA_MSG_WAIT](#).
2. It sends an indirect message to another VM and requests the primary scheduler to schedule the message target. This is described in [4.3 Indirect messaging usage](#).
3. It gets preempted by an interrupt targeted to the primary scheduler or another FF-A component. This is described in [4.5.3 Preemption during message processing](#).
4. It yields control back to the primary scheduler because it cannot continue execution due to an internal dependency. This is done using the `FFA_YIELD` interface and described in [9.2 FFA_YIELD](#).
5. It sends a direct Partition message request to an endpoint and blocks until it does receives a response. This is described in [4.4 Direct messaging usage](#).
6. It sends a direct Framework message to a partition manager and blocks until it receives a response.

4.5.2 Direct message processing

After a Receiver starts processing a message, one or more of the following events could occur. The Hypervisor or SPM must treat the invocation of any other events as invalid.

1. It sends a direct Partition message response to the endpoint that sent it the request. It blocks until it receives another direct or indirect message. This is described in [4.4 Direct messaging usage](#).
2. It sends a direct Partition message request to an endpoint and blocks until it receives a response. This is described in [4.4 Direct messaging usage](#).
3. It gets preempted by an interrupt targeted to the primary scheduler or another FF-A component. This is described in [4.5.3 Preemption during message processing](#).

4.5.3 Preemption during message processing

An endpoint execution context could be interrupted during message processing by a physical interrupt targeted to any FF-A component. The endpoint execution context must be resumed after the interrupt has been handled so that it can continue processing the message.

The endpoint execution context could be resumed on the same PE where it was preempted or migrated to a different PE if it is not pinned to the original PE.

The method to do this depends on the location of the endpoint relative to the FF-A component that implements the interrupt handler. Responsibilities of participating components are as follows.

- An endpoint execution context must enter the *preempted* state on being interrupted.
- A VM's execution context must be saved on interruption and restored on resumption by the Hypervisor.
- A SP's execution context must be saved on interruption and restored on resumption by the SPM.
- The OS kernel's execution context must be saved on interruption and restored on resumption by the SPM.

- If the Hypervisor or SPM must pass control to another FF-A component for handling the physical interrupt, the *FFA_INTERRUPT* interface must be used with the appropriate conduit (also see [7.4 FFA_INTERRUPT](#)) if the component is:
 - The primary scheduler that resides in EL1.
 - An endpoint waiting for direct message response.
 - An endpoint in the *idle* state.
- If the physical interrupt is handled in the FF-A component that implements the primary scheduler, then it must use the *FFA_RUN* interface to resume the preempted endpoint.

Figure 4.10 illustrates an example flow where *Client 0* in a NS-Endpoint sends a direct message to the single execution context EC0 on CPU0 of a UP-Migrate capable SP. Message processing in SP EC0 is preempted by a Non-secure interrupt. It is later resumed on CPU1 by the NS-Endpoint.

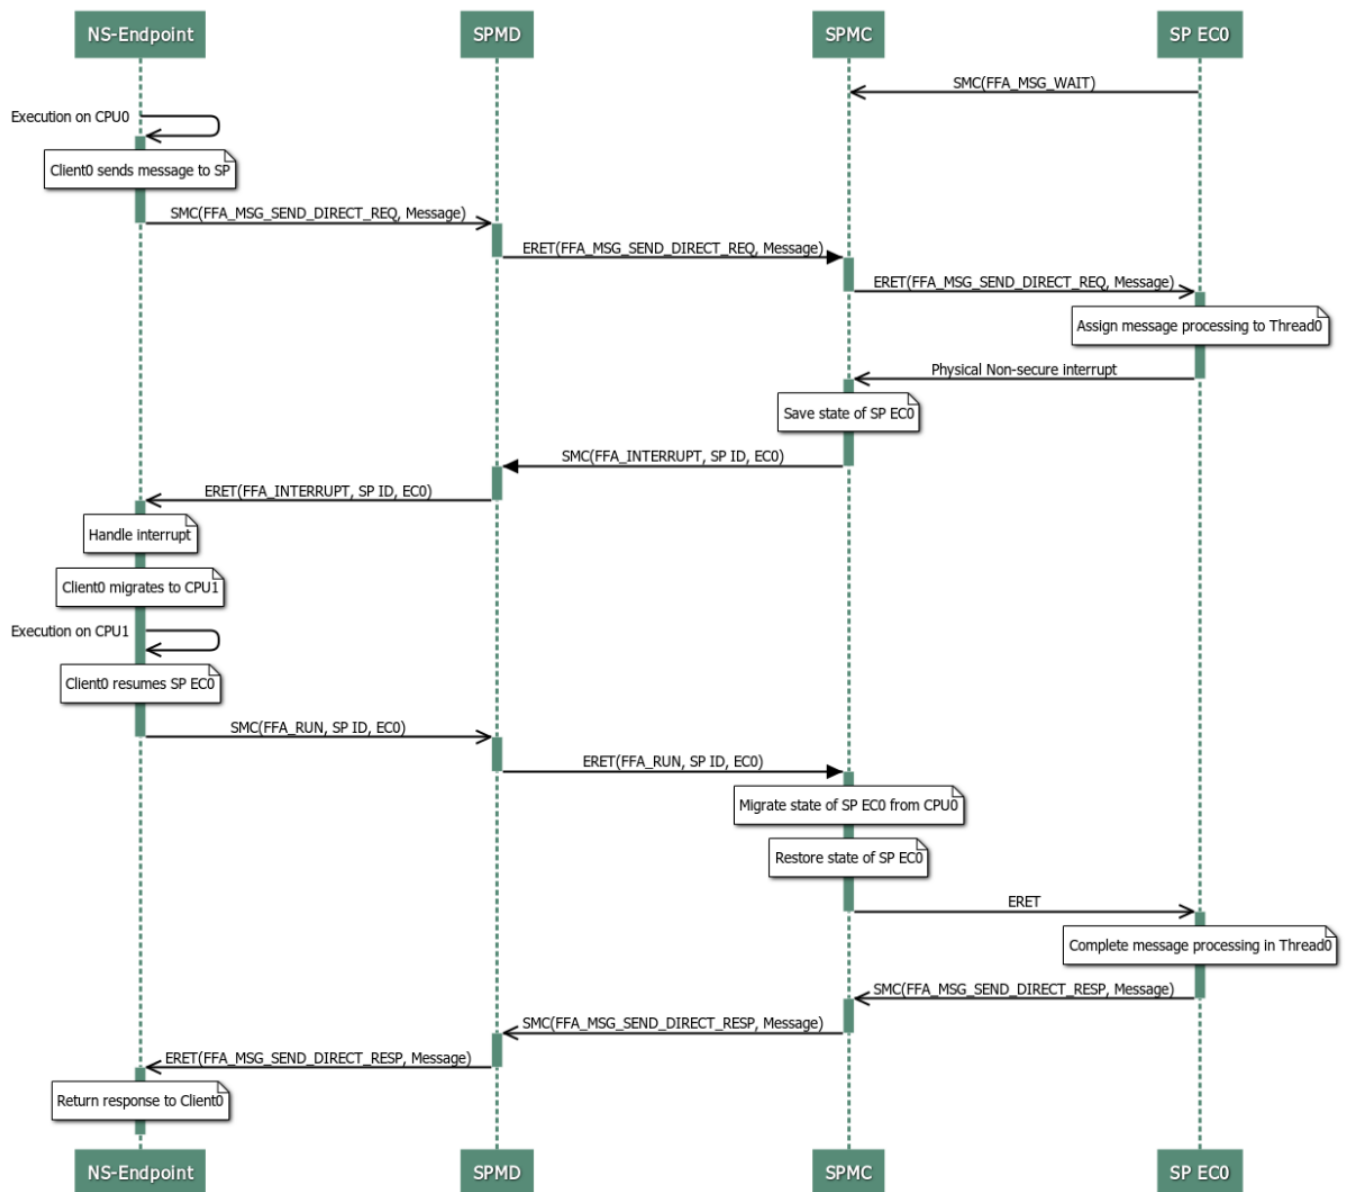


Figure 4.10: Example endpoint preemption flow

4.5.4 Managed exit

A managed exit is a mechanism in which an endpoint execution context that is processing a direct message is notified about the occurrence of an interrupt targeted to another FF-A component. This allows the endpoint to manage the state of its application threads before relinquishing control to the component where the interrupt must be handled.

A managed exit stands in contrast to preemption of an endpoint execution context during message processing (see [10.3 FFA_MSG_SEND_DIRECT_RESP](#)). In the latter case, the endpoint does not get an opportunity to manage its internal state before control is handed to the target of the interrupt.

An endpoint execution context must use the `FFA_MSG_SEND_DIRECT_RESP` interface (see [10.3 FFA_MSG_SEND_DIRECT_RESP](#)) to hand control to the SPM or Hypervisor and complete a managed exit.

Use of this mechanism leaves the endpoint execution context in an *idle* state. This means that when it is next run on a PE, it can either start processing a new request or resume processing the preempted request.

A managed exit could be used for the following reasons.

1. It enables other application threads in an endpoint to make progress while one or more application threads have been preempted.
2. It ensures that the CPU cycles allocated to an endpoint execution context are used to process the request that the calling endpoint has issued instead of a request from another endpoint.
3. It ensures that critical events can be conveyed to the endpoint in time.

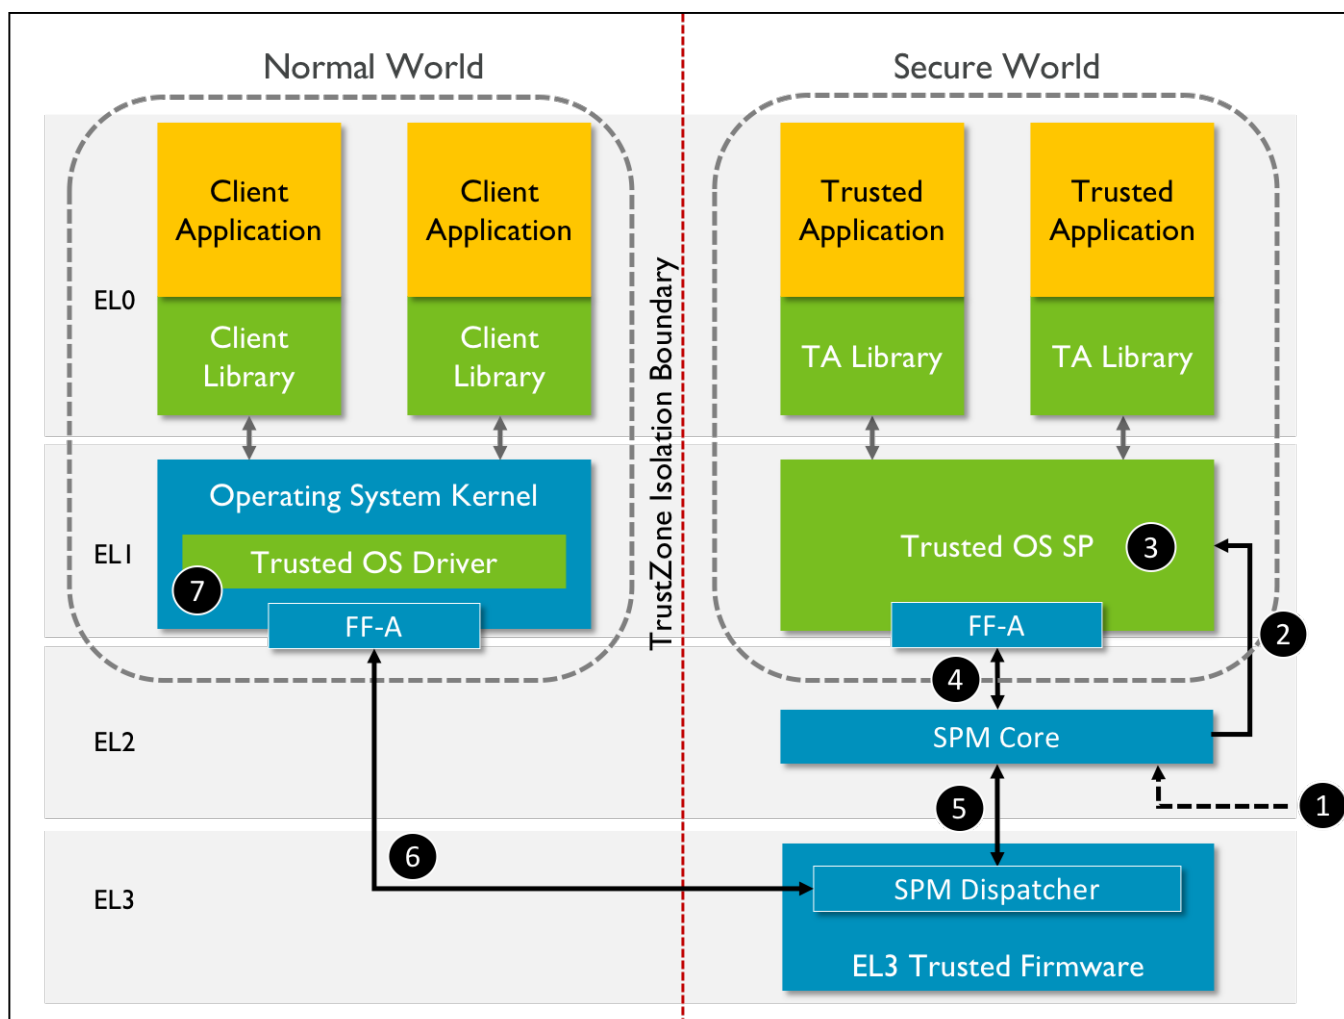
For example, the OS could issue a power state transition event on a PE. Endpoint execution contexts pinned to that PE might need to be notified about this event. An endpoint execution context could be in a state where it can only resume the request that was previously preempted. This could lead to an unacceptable delay before the endpoint processes the event.

4. It enables application threads in an MP endpoint with pinned execution contexts to be migrated to a different PE. They could be then resumed under the execution context pinned on that PE. This is in contrast to the endpoint execution context, and therefore all its application threads, remaining in a preempted state, on the PE where they were preempted until later resumed.

An endpoint could either depend on the SPM or the Hypervisor for performing a managed exit or implement it on its own. This depends on the architectural environment the endpoint is executing in as follows.

Accesses made by an Endpoint to the GIC (see [\[8\]](#)) could be virtualized for example, while running in a virtual machine under the control of the Hypervisor or SPM. In this case, the endpoint can only manage virtual interrupts.

Physical interrupts are targeted to the Exception level that the Hypervisor or SPM are running in. If a physical interrupt must be handled in another FF-A component, the SPM or the Hypervisor must notify the preempted endpoint that a managed exit is required through an IMPLEMENTATION DEFINED mechanism. This flow is illustrated in [Figure 4.11](#).

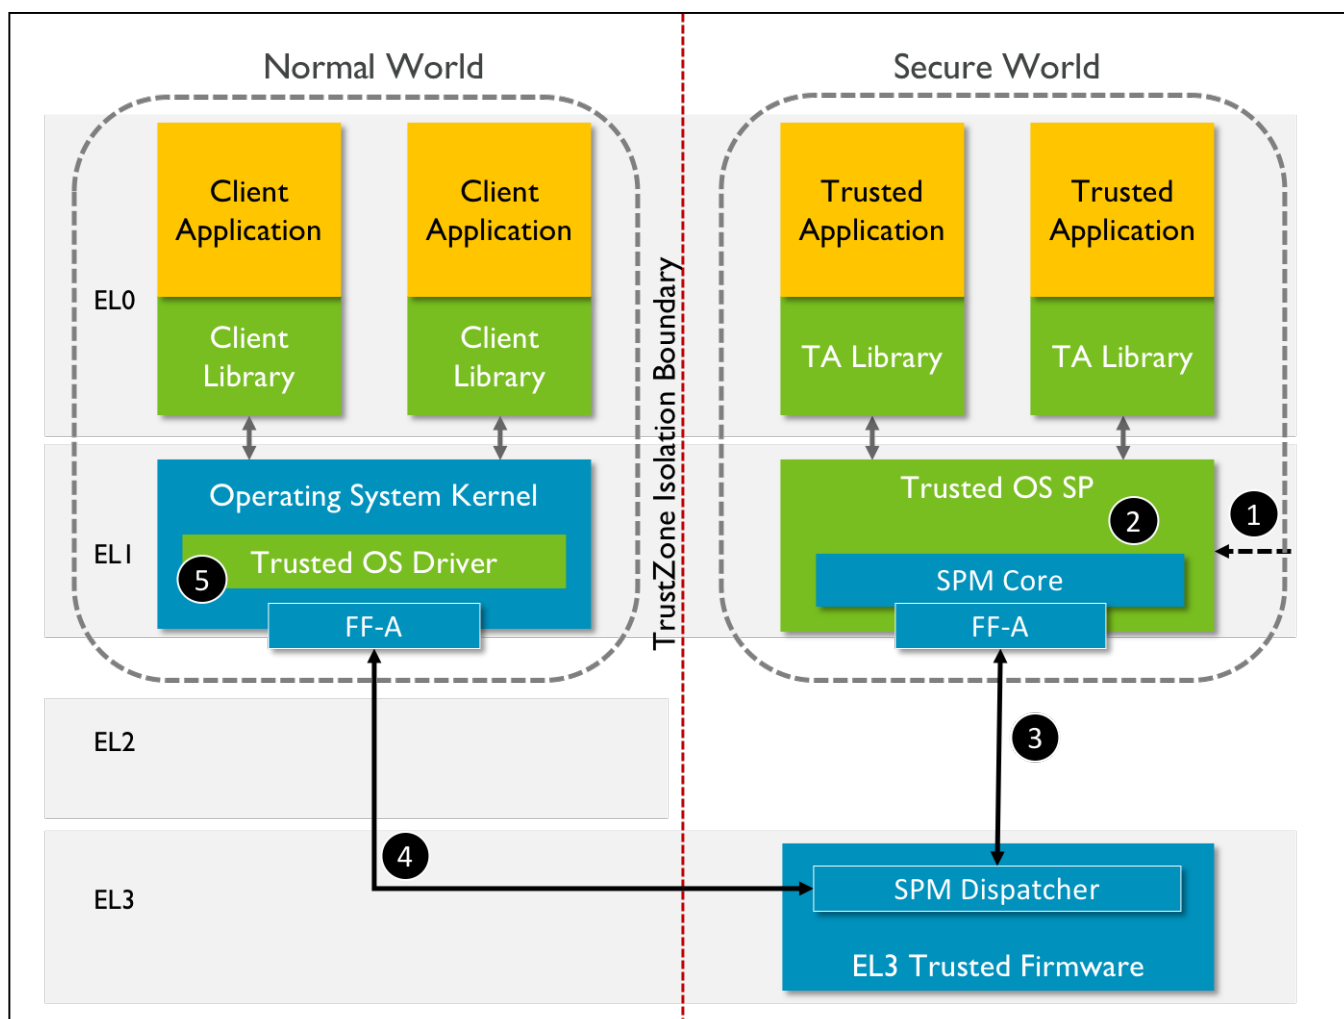


1	Physical non-secure interrupt arrives at SPM Core in S-EL2
2	SPM Core masks non-secure interrupt, pends "managed exit" virtual interrupt and resumes SP
3	Trusted OS SP saves application context
4	Trusted OS SP returns control to SPM Core by invoking the FFA_MSG_SEND_DIRECT_RESP ABI
5	SPM Core unmarks non-secure interrupt and returns control to SPM Dispatcher by forwarding the ABI invocation
6	SPM Dispatcher returns control to Normal world by forwarding the ABI invocation
7	Non-secure interrupt is handled in Normal world OS

Figure 4.11: Example managed exit flow with GIC virtualization

An endpoint could access the physical GIC for example, a Trusted OS running in S-EL1 under the control of the SPM. The endpoint can manage physical interrupts.

An interrupt targeted to another FF-A component can be used by the endpoint as a signal to perform a managed exit. This flow is illustrated in [Figure 4.12](#).



1	Physical non-secure interrupt arrives at SPM Core in S-EL1
2	Trusted OS SP saves application context
3	SPM Core in Trusted OS SP returns control to SPM Dispatcher by invoking the FFA_MSG_SEND_DIRECT_RESP ABI
4	SPM Dispatcher returns control to Normal world by forwarding the ABI invocation
5	Non-secure interrupt is handled in Normal world OS

Figure 4.12: Example managed exit flow without GIC virtualization

In both cases, the pending state of the interrupt that is targeted to another FF-A component, and triggers a managed exit in the currently running endpoint, must not change in the GIC through any software action until the managed exit has completed.

The use of a managed exit by an endpoint is optional. This mechanism can be used by an endpoint only if it fulfills the following constraints:

- The endpoint is capable of receiving direct requests.
- The endpoint runs in a privileged Exception level. This is one of the following.
 - EL1.
 - Secure EL1.

- Supervisor mode.
- Secure Supervisor mode.

An endpoint that fulfills these constraints must specify whether it uses the managed exit mechanism in its partition manifest (see [Table 3.1](#) in [3.2 Partition manifest at virtual FF-A instance](#)).

[Figure 4.13](#) illustrates an example managed exit flow where *Client 0* in a NS-Endpoint sends a direct message to MP capable SP. The SP has access to the virtual GIC and two execution contexts *EC0* and *EC1* which are pinned to *CPU0* and *CPU1* respectively. SP *EC0* stops message processing and performs a managed exit in response to a Non-secure physical interrupt. Message processing is later resumed on *CPU1* by the NS-Endpoint.

Chapter 4. Message passing
4.5. Partition message processing

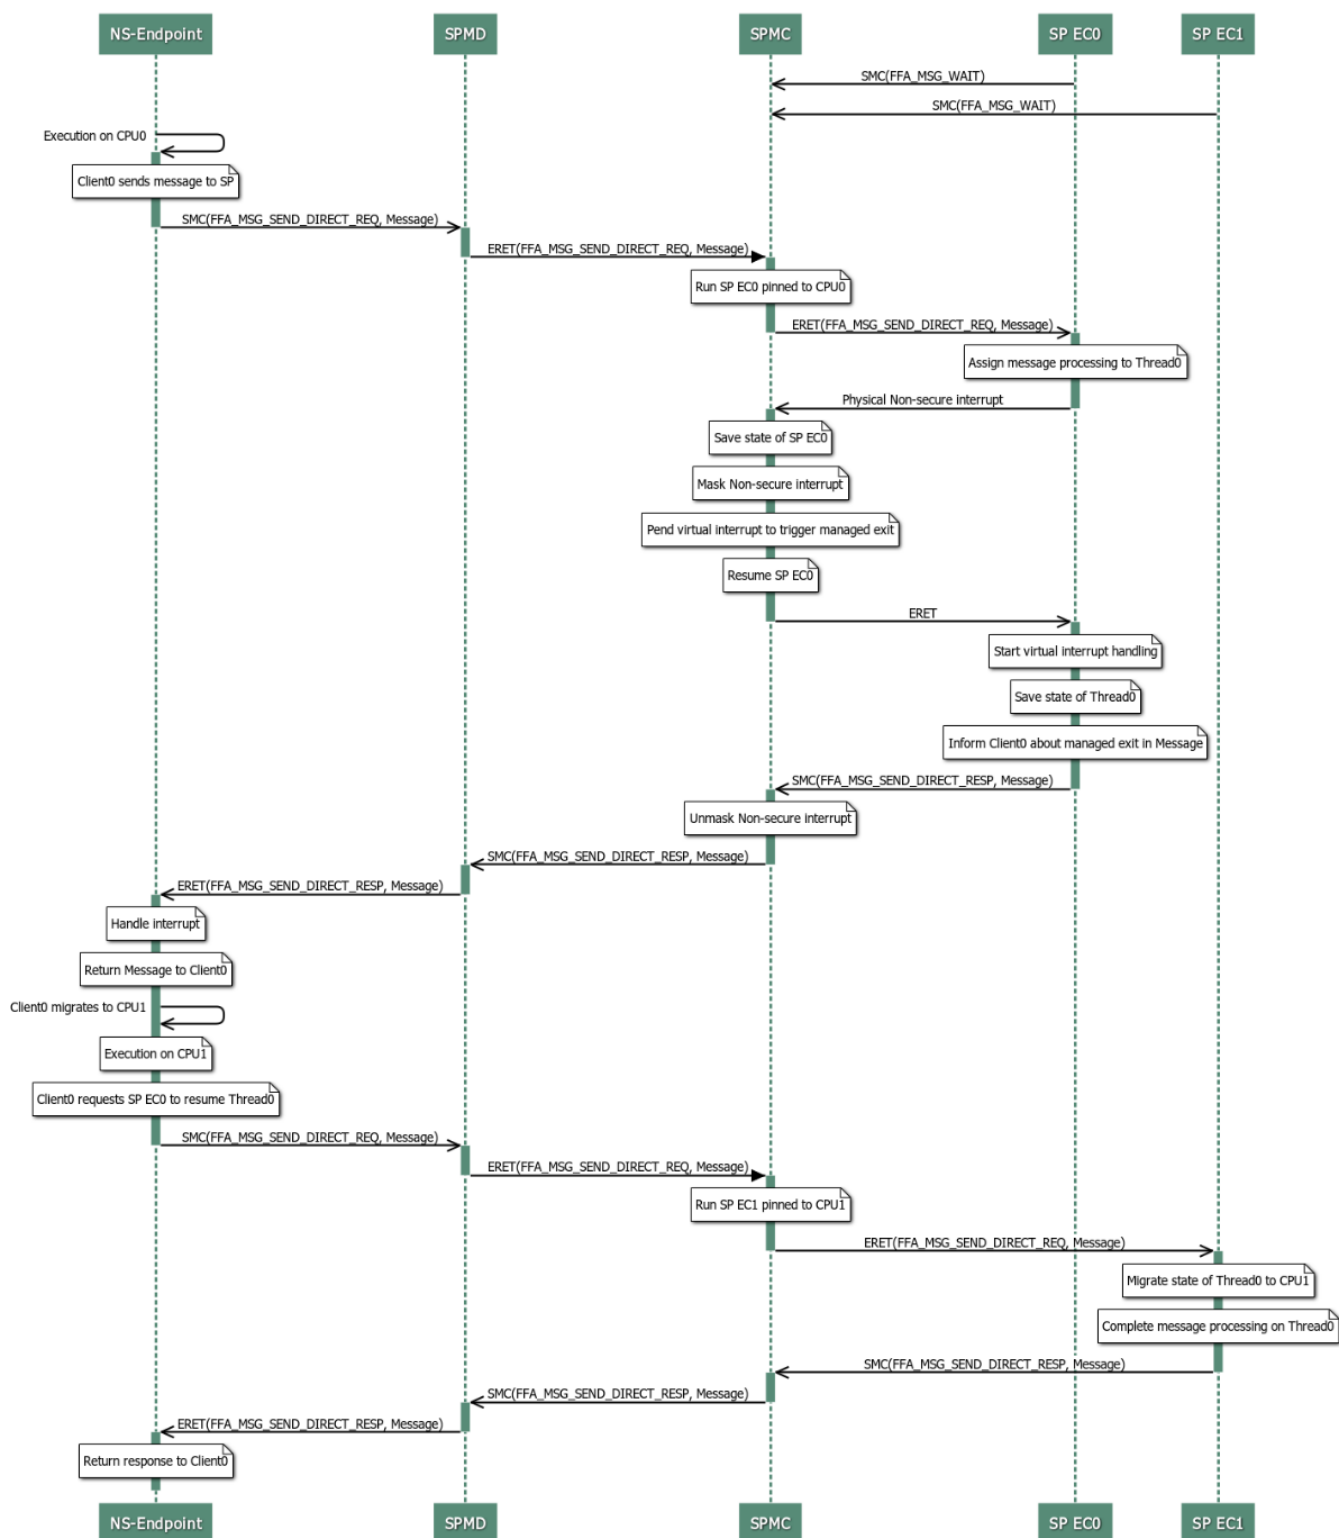


Figure 4.13: Example managed exit flow

Chapter 5

Memory Management

5.1 Overview

The Firmware Framework describes mechanisms and interfaces that enable FF-A components to manage access and ownership of memory regions in the physical address space to fulfill use cases such as:

- DRM protected video path.
- Communication with a VM with pre-configured machine learning frameworks,
- Biometric authentication and Secure payments.

FF-A components can use a combination of Framework and Partition messages to manage memory regions in the following ways.

1. The Owner of a memory region can transfer its ownership to another FF-A component.
2. The Owner of a memory region can relinquish access to it and grant access to one or more FF-A components.
3. The Owner of a memory region can share access to it with one or more FF-A components.
4. The Owner of a memory region can reclaim access to it by requesting FF-A components to relinquish access to the memory region.

5.2 Direct memory access

The Framework enables FF-A components to manage access to the physical address space from a device that is upstream of an SMMU using the memory management transactions described in [5.5 Memory management transactions](#).

As per the Arm® SMMU architecture, each transaction generated by a device is associated with a *Stream ID*. This Stream ID could be one of many that a the device is configured to use. A Stream ID is used to determine the stage 1 and stage 2 address translations that must be used for the transaction. It is also possible that one or both stages of translation could be bypassed for a Stream ID in the SMMU.

If enabled, the stage 2 translations corresponding to a Stream ID control access to the physical address space that the device has. A set of stage 2 translation tables could map to one or more Stream IDs. The Framework manages stage 2 translations in the SMMU as described in [5.3 Address translation regimes](#).

The Hypervisor programs the SMMU to create and manage the association between a Non-secure Stream ID and the stage 2 translations its transactions must use.

The SPM programs the SMMU to create and manage the association between a Secure Stream ID and the stage 2 translations its transactions must use.

The Framework does not manage the stage 1 translations and their association with Stream IDs in the SMMU on behalf of the device. This should be done by an endpoint through an IMPLEMENTATION DEFINED mechanism.

5.2.1 Stream endpoint

A set of SMMU stage 2 translations maintained by a partition manager is called a *Stream endpoint*. Each Stream endpoint is assigned a 16-bit ID called the *Stream endpoint ID* or *SEPID*.

Stream endpoints associated with a Secure Stream ID are called *Secure SEPIDs*

Stream endpoints associated with a Non-secure Stream ID are called *Non-secure SEPIDs*

Endpoints that run on a PE are referred to as *PE endpoints* to differentiate them from Stream endpoints. The term *endpoint* is used when it is not required to distinguish between these types of endpoints.

There is a 1:N ($N \geq 1$) mapping between a SEPID and Stream IDs assigned to different devices that is, the stage 2 translations corresponding to the SEPID could be *shared* by one or more Stream IDs.

SEPIDs are used in memory management transactions to:

- Grant and revoke access to a physical memory region to a device.
- Transfer ownership of a physical memory region from or to a device.

SEPID values must be distinct from those assigned to PE endpoints. A SEPID is not discoverable through the **FFA_ID_GET** interface (also see [8.7 FFA_ID_GET](#)).

This version of the Framework considers two types of devices.

1. Devices that can act as initiators and recipients of memory management transactions. These devices are called *independent peripheral devices*. Each device must specify the following information in its partition manifest (see [3.4 Independent peripheral device manifest](#)).
 - A SEPID assigned to the device at boot time.
 - The SMMU ID that the device is upstream of.
 - Each Stream ID the device can generate.
 - Regions in the physical address space that must be mapped in the translation tables corresponding to the SEPID at boot time.

This information enables the partition manager to create an association between a device and a SEPID at boot time.

A partition manager or PE endpoint and an *independent* device must use an IMPLEMENTATION DEFINED mechanism to notify each other about a memory management transaction targeted to a SEPID used by the device (see [5.5.2 Transaction life cycle](#)).

2. Devices that cannot act as initiators and recipients of memory management transactions. These devices are called *dependent* peripheral devices. They rely on a PE endpoint to initiate and receive memory management transactions on their behalf. The PE endpoint is called a *proxy endpoint*.

A dependent device could be *assigned* to a PE endpoint. This implies,

- Access to its MMIO regions is assigned to the endpoint during boot (see [2.10 System resource management](#) & [Table 3.3](#)).
- The endpoint manages the association between Stream IDs generated by the device and stage 1 translations in the SMMU that the device is upstream of (see [Table 3.3](#)).

The device could be either assigned to its *proxy endpoint* or a different PE endpoint.

When assigned to its proxy endpoint, this version of the Framework assumes that all the Stream IDs generated by the device have the same visibility of the physical address space as the endpoint. The stage 2 translations in the SMMU for these Stream IDs are the same as those maintained by the partition manager on behalf of the endpoint. They are not assigned a SEPID. The partition ID of the *proxy endpoint* is used instead. All memory management transactions with this partition ID effect both sets of translations.

When assigned to a different endpoint, the partition manifest of the *proxy endpoint* (see [3.2 Partition manifest at virtual FF-A instance](#)) must specify the following information to enable the partition manager to create an association between a device and a SEPID at boot time.

- The SMMU ID that the device is upstream of.
- Each Stream ID the device can generate.
- The SEPID corresponding to each Stream ID.

The partition ID of the proxy endpoint must be distinct from the SEPID allocated to manage the preceding association. The SEPID must be specified in the partition manifest of the proxy endpoint (see [Table 3.1](#)).

The stage 2 translations corresponding to the SEPID are configured at boot time with no access to the physical address space.

A memory management transaction targeted to the SEPID must be allowed to complete only if it is either initiated or authorized by the *proxy endpoint* for the device (see [5.5.2 Transaction life cycle](#)).

The SEPIDs used by an *independent* device must be distinct from the SEPIDs used by a *dependent* device. This constraint avoids the scenario where a memory management transaction is allowed to change the stage 2 translations before the *proxy endpoint* has authorized it.

5.3 Address translation regimes

Memory management relies on the two fundamental operations of mapping and un-mapping a memory region from the stage of a translation regime managed by a partition manager on behalf of a partition. The translation regime and the stage depends on the type of partition as follows.

1. The Hypervisor creates and manages stage 2 translations on behalf of a EL1 PE endpoint, in the Non-secure EL1&0 translation regime, when EL2 is enabled.
2. The Hypervisor creates and manages stage 2 translations for a Non-secure Stream ID assigned to an independent or dependent peripheral device, in the Non-secure EL1&0 translation regime in the SMMU. A SEPID is used to identify the stage 2 translation tables (see [5.2.1 Stream endpoint](#)).
3. The SPMC creates and manages stage 2 translations on behalf of a S-EL1 PE endpoint in the Secure EL1&0 translation regime, when S-EL2 is enabled.
4. The SPMC creates and manages stage 1 translations on behalf of a S-EL0 PE endpoint in the Secure EL1&0 translation regime, when S-EL2 is disabled.
5. The SPMC creates and manages stage 2 translations for a Secure Stream ID assigned to an independent or dependent peripheral device in the Secure EL1&0 translation regime in the SMMU. A SEPID is used to identify the stage 2 translation tables (see [5.2.1 Stream endpoint](#)).

5.4 Ownership and access attributes

The Hypervisor, SPM, and all endpoints have *access* and *ownership* attributes associated with every memory region in the physical address space.

Access determines the data and instruction access permissions to the memory region. A component can have the following access permissions to a memory region.

- No access.
- Read-only, Execute-never.
- Read-only, Executable.
- Read/write, Execute-never.

Access control must be enforced through an IMPLEMENTATION DEFINED mechanism and/or by encoding these permissions in the translation regime of an endpoint managed by the partition manager (see [5.3 Address translation regimes](#)).

Ownership is a software attribute that determines if a component can grant access to a memory region to another component. A component that has access to a memory region without ownership is called the *Borrower*. A component that lends access to a memory region it owns is called the *Lender*.

Ownership of a memory region is initially assigned to the component that it is allocated to. At boot time all memory regions are owned by Secure firmware. A memory region could be configured as Secure or normal memory either statically at reset, or by Secure firmware during boot. Secure firmware transfers ownership of normal memory to Normal world software. It sub-divides Secure memory such that:

- It owns and has exclusive access to some memory regions.
- It owns but grants access to some memory regions to SPs.
- It transfers ownership of some memory regions to SPs.

If virtualization is enabled in the Normal world, the Hypervisor divides a subset of normal memory among VMs and transfers ownership to them. In the absence of virtualization, all normal memory donated by the Secure world is owned by the OS kernel.

An endpoint requests access to and/or ownership of a memory region through its partition manifest (also see [Table 3.2](#)).

5.4.1 Ownership and access rules

The SPM and Hypervisor must enforce the following general ownership and access rules to memory regions.

1. The size of a memory region to which ownership and access rules apply must be a multiple of the smallest translation granule size supported on the system.
 - It is 4K in the AArch32 Execution state.
 - A EL1 or S-EL1 partition must discover this by reading the ID_AA64MMFR0_EL1 System register in the AArch64 Execution state.
 - A S-EL0 SP must determine this through an IMPLEMENTATION DEFINED discovery mechanism for example, DT or ACPI tables.
2. A normal memory region must be mapped with the Non-secure security attribute in any component that is granted access to it.
3. A Secure memory region must be mapped with the Secure security attribute in any component that is granted access to it.
4. Each memory region in the physical address space must have a single Owner.
5. A FF-A component must have access to a memory region it owns unless it has granted exclusive access to the region to another FF-A component.

6. Only the Owner of a memory region can grant access to it to one or more Borrowers in the system.
7. Only the Owner of a memory region can transfer its ownership to another endpoint in the system.
8. If an SP is terminated 'cause of a fatal error condition, access to the memory regions of the SP is transferred to the SPM.
9. If a VM is terminated 'cause of a fatal error condition, access to the memory regions of the VM and their ownership are transferred to the Hypervisor.
10. If the Hypervisor or OS kernel are terminated 'cause of a fatal error condition, access to the their memory regions and ownership are transferred to the SPM.
11. The number of distinct components to whom an Owner can grant access to a memory region is IMPLEMENTATION DEFINED.
12. The Owner of a memory region must not be able to change its ownership or access attributes until all Borrowers have relinquished access to it.

5.4.2 Ownership and access states

Table 5.1 describes the ownership states applicable to an FF-A component for a memory region.

Table 5.1: Ownership states

No.	Ownership state	Acronym	Description
1	Owner	Owner	Component owns the memory region.
2	Not Owner	!Owner	Component does not own the memory region.

Table 5.2 describes the access states applicable to an FF-A component for a memory region.

Table 5.2: Access states

No.	Access state	Acronym	Description
1	No access	NA	A component has <i>no</i> access to a memory region. It is not mapped in its translation regime.
2	Exclusive access	EA	A component has exclusive access to a memory region. It is mapped only in its translation regime.
3	Shared access	SA	A component has shared access to a memory. It is mapped in its translation regime and the translation regime of at least one other component

Table 5.3 describes the valid combination of access and ownership states applicable to an FF-A component for a memory region.

Table 5.3: Valid combinations of ownership and access states

No.	Ownership state	Access state	Acronym	Description
1	Not Owner	No access	!Owner-NA	Component has neither ownership nor access to the memory region.
2	Not Owner	Exclusive access	!Owner-EA	Component has exclusive access without ownership of the memory region.
3	Not Owner	Shared access	!Owner-SA	Component has shared access with one or more components without ownership of the memory region.
4	Owner	No access	Owner-NA	Component owns the memory region and has granted: <ul style="list-style-type: none"> • Either exclusive access to the memory region to another component. • Or shared access to the memory region among other components.
5	Owner	Exclusive access	Owner-EA	Component owns the memory region and has exclusive access to it.
6	Owner	Shared access	Owner-SA	Component owns the memory region and shares access to it with one or more components.

For two FF-A components *A* and *B* and a memory region, valid combinations of states defined in [Table 5.3](#) are described in [Table 5.4](#). Other combinations of states are considered invalid.

Table 5.4: Valid combinations of ownership and access states between two components

No.	Component A state	Component B state	Description
1	Owner-EA	!Owner-NA	Component A has exclusive access and ownership of a memory region that is inaccessible from component B.
2	Owner-NA	!Owner-NA	Component A has granted exclusive access to a memory region it owns to another component. It is inaccessible from component B.
3	Owner-NA	!Owner-EA	Component A has granted exclusive access to a memory region it owns to component B.
4	Owner-NA	!Owner-SA	Component A has relinquished access to a memory region it owns. Access to the memory region is shared between component B and at least one other component
5	Owner-SA	!Owner-NA	Component A shares access to a region of memory it owns with another component. Component B cannot access the memory region.

No.	Component A state	Component B state	Description
6	Owner-SA	!Owner-SA	Component A shares access to a region of memory it owns with component B and possibly other components.

U

Implementation Note

To fulfill the use cases and enforce the rules listed earlier, FF-A components should track the state of a memory region. This could be done as follows,

- An Owner tracks the level of access it has to a memory region.
- An Owner tracks the level of access that Borrowers have to a memory region along with the identity of the Borrowers.
- A Borrower tracks the level of access the Owner has to a memory region along with the identity of the Owner.
- A Borrower tracks the level of access it has to a memory region.
- A Borrower tracks the level of access that other Borrowers have to a memory region along with the identity of the Borrowers.
- For each memory region, the SPM and Hypervisor track the following.
 - The identity of each Borrower.
 - The identity of the Owner.
 - The level of access of each Borrower.
 - The level of access of the Owner.

5.5 Memory management transactions

This version of the Framework describes transactions that enable endpoints to manage access and ownership of physical memory regions.

- Transitions between states described in [5.4.2 Ownership and access states](#) happen in response to transactions described in [Table 5.5](#). Each transaction involves exchange of one or more Framework and partition messages.
- Each transition is described as a transaction involving two endpoints (*A* and *B*) and a memory region. Endpoint *A* is the Owner of the memory region.

Table 5.5: Memory region transactions

No.	Transaction	Description
1.	Donate	<ul style="list-style-type: none">• Endpoint <i>A</i> transfers ownership of a memory region it owns to endpoint <i>B</i>. See 5.6 Donate memory transaction.
2.	Lend	<ul style="list-style-type: none">• Endpoint <i>A</i> relinquishes access to a memory region and grants it to only endpoint <i>B</i>. Endpoint <i>B</i> gains exclusive access to the memory region.• Endpoint <i>A</i> relinquishes access to a memory region and grants it to endpoint <i>B</i> and at least one other endpoint simultaneously. Endpoint <i>B</i> gains shared access to the memory region.• See 5.7 Lend memory transaction.
3.	Share	<ul style="list-style-type: none">• Endpoint <i>A</i> grants access to a memory region to endpoint <i>B</i> and optionally to other endpoints simultaneously. See 5.8 Share memory transaction.
4.	Relinquish	<ul style="list-style-type: none">• Endpoint <i>B</i> relinquishes access to a memory region granted to it by Endpoint <i>A</i>. Endpoint <i>A</i> reclaims exclusive access to the memory region. See 5.9 Relinquish memory transaction.

5.5.1 Component roles

In this version of the Framework, endpoints can fulfill the role of an Owner, Lender or Borrower (see [5.4 Ownership and access attributes](#)).

The Hypervisor and SPM participate in memory management transactions to validate and transmit them from a *Sender* endpoint to a *Receiver* endpoint. They are also responsible for managing the translation regime of an endpoint and tracking the ownership and access attributes of a memory region. This collective role is termed as a *Relayer*.

[Table 5.6](#) specifies the roles each FF-A component can play in a memory management transaction.

In the absence of the Hypervisor, the OS Kernel subsumes the role of the Relayer. Its roles as the Relayer, Owner, Lender and Borrower are considered to be logically separate from each other. The interface used by internal components that implement these roles to exchange memory management transactions is IMPLEMENTATION DEFINED.

The roles of the SPMD and SPMC components of the SPM (see [2.2 SPM architecture](#)) as Relayers are as follows.

- In SPM configurations where the SPMD and SPMC reside in separate Exception levels (see [Table 2.1 & Table 2.2](#)):

- The SPMD component must forward memory management transactions between the Secure and Non-secure physical FF-A instances.
- The SPMC component must handle outbound and inbound transactions on behalf of the Sender and Receiver.
- In the SPM configuration where the SPMC coexists with an SP at S-EL1 or Secure Supervisor mode (see [Table 2.3](#)), the roles of the SPMC as the Relayer and the SP as the Owner, Borrower or Lender are considered to be logically separate. The interface used by internal components that implement these roles to exchange memory management transactions is IMPLEMENTATION DEFINED. The SP and SPM must still appear as separate FF-A components to software in the Normal world and SPs at the Secure virtual FF-A instance. Also see [5.5.2 Transaction life cycle](#).

Table 5.6: FF-A component roles in a memory management transaction

Config No.	FF-A component	Owner	Lender	Borrower	Relayer
1.	NS-Endpoint	Yes	Yes	Yes	No
2.	S-Endpoint	Yes	Yes	Yes	No
3.	SEPID	Yes	Yes	Yes	No
4.	Hypervisor	No	No	No	Yes
5.	SPM	No	No	No	Yes

In all transactions, an endpoint must be a Sender or Receiver. This depends on the type of transaction as follows.

- In a transaction to donate ownership of a memory region, the Sender is the current Owner and the Receiver is the new Owner.
- In a transaction to lend or share access to a memory region, the Sender is the Lender and the Receiver is the Borrower.
- In a transaction to relinquish access to a memory region, the Sender is the Borrower and the Receiver is the Lender.

Valid combinations of component roles in a transaction to donate, lend or share memory are listed in [Table 5.7](#). A FF-A component can use one or more combinations in a memory management transaction as the Sender.

Valid combinations of component roles in a transaction to relinquish memory are listed in [Table 5.8](#).

Table 5.7: Valid role combinations in donate, lend or share memory transactions

Config No.	Sender	Receiver	Relayer
1.	VM	VM	Hypervisor
2.	VM	NS SEPID	Hypervisor

Config No.	Sender	Receiver	Relayer
3.	NS-Endpoint	Secure SEPID	Hypervisor (if present) and SPM
4.	NS-Endpoint	SP	Hypervisor (if present) and SPM
5.	SP	Secure SEPID	SPM
6.	SP	SP	SPM

Table 5.8: Valid role combinations in relinquish memory transactions

Config No.	Sender	Receiver	Relayer
1.	VM	VM	Hypervisor
2.	NS-SEPID	VM	Hypervisor
3.	Secure SEPID	NS-Endpoint	Hypervisor (if present) and SPM
4.	SP	NS-Endpoint	Hypervisor (if present) and SPM
5.	Secure SEPID	SP	SPM
6.	SP	SP	SPM

5.5.2 Transaction life cycle

Each transaction described in [Table 5.5](#) takes place in three steps as follows and illustrated in [Figure 5.1](#).

1. The Sender sends a Framework message to the Relayer to start a transaction involving one or more Receivers.
2. The Sender sends a Partition message requesting each Receiver to complete the transaction.
3. Each Receiver sends a Framework message to the Relayer to complete the transaction.

A transaction could be targeted to a *dependent* peripheral device identified by a SEPID (see [5.2.1 Stream endpoint](#)). In this case, the partition message in *step 2* is sent to the *proxy* endpoint of the device. The proxy endpoint sends a Framework message in *step 3* to validate, authorize and complete the transaction of behalf of the device.

A transaction could be targeted to an *independent* peripheral device identified by a SEPID (see [5.2.1 Stream endpoint](#)). In this case, an IMPLEMENTATION DEFINED message is sent to this device in *step 2*. The device uses an IMPLEMENTATION DEFINED mechanism to communicate with the Relayer to complete the transaction in *step 3*.

In the SPM configuration where the SPMC coexists with an SP at S-EL1 or Secure Supervisor mode (see [Table 2.3](#)), the Relayer, Sender and Receiver components are implemented in the same Exception level and software image. The transaction life-cycle for this configuration is as follows.

- When the SP is the Sender, it must use an IMPLEMENTATION DEFINED interface to deliver the Framework message to the SPMC in step 1.
- When the SP is the Receiver, the Framework message sent in step 1 is received by the SPMC. It must be delivered to the SP at the Secure physical FF-A instance through an IMPLEMENTATION DEFINED interface between them.
 - A successful completion of the interface used in step 1 by the SPMC must indicate completion of the entire transaction to the Sender.
 - Steps 2 & 3 are not required as the SP is made aware of the ongoing transaction in Step 1.
 - The following aspects of the memory management transaction in this scenario are IMPLEMENTATION DEFINED.
 - * How the Sender discovers the presence of this SPM configuration.
 - * How the SPMC delivers the Framework message to the Receiver.
 - * How the Receiver interacts with the SPMC to complete the transaction.

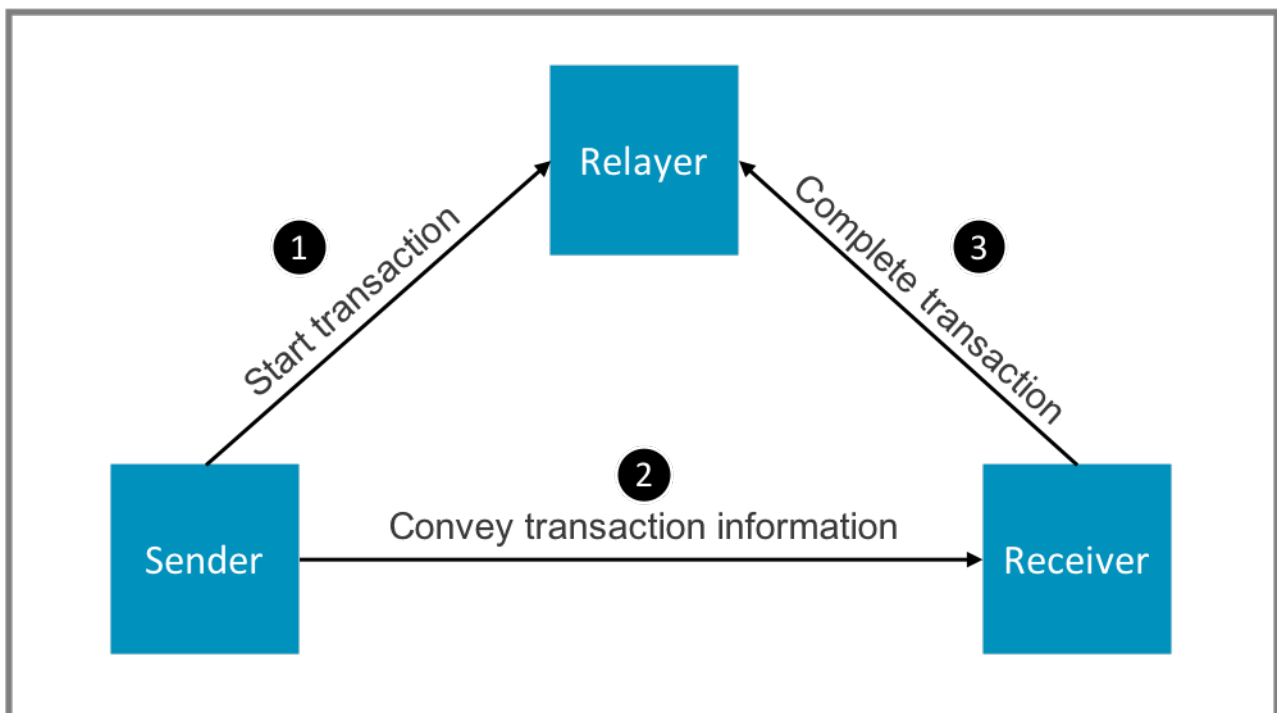


Figure 5.1: Memory management transaction lifecycle

5.6 Donate memory transaction

This transaction is used to transfer the ownership of a memory region from the endpoint that owns it to another endpoint. A list of valid combinations of roles played by various FF-A components in this transaction is specified in [Table 5.7](#).

5.6.1 Donate memory state machine

[Table 5.9](#) describes the state machine for donating a memory region from the perspective of two endpoints *A* & *B*. *A* owns the memory region. It attempts to donate the memory region to *B*. Valid and invalid state transitions in response to this transaction have been listed.

In each valid transition,

- *A* loses both ownership and access to the memory region and enters the *!Owner-NA* state.
- *B* gains ownership and exclusive access to the memory region and enters the *Owner-EA* state.

Table 5.9: Donate memory transaction state machine

No.	Current Endpoint A state	Current Endpoint B state	Next Endpoint A state	Next Endpoint B state	Description
1	Owner-EA	!Owner-NA	!Owner-NA	Owner-EA	• Owner has exclusive access to the memory region and transfers ownership to endpoint B.
2	Owner-NA	!Owner-NA	Error	–	• Owner does not have exclusive access to the memory region. It cannot transfer its ownership.
3	Owner-NA	!Owner-SA	Error	–	• Owner has lent access to the memory region to endpoint B and possibly other endpoints. It cannot transfer its ownership.
4	Owner-SA	!Owner-NA	Error	–	• Owner has shared access to the memory region with one or more endpoints. It cannot transfer its ownership.
5	Owner-SA	!Owner-SA	Error	–	• Owner has shared access to the memory region with endpoint B and possibly other endpoints. It cannot transfer its ownership.

5.6.2 Donate memory transaction lifecycle

This transaction takes place as follows (also see [5.5.2 Transaction life cycle](#)).

1. The Owner uses the *FFA_MEM_DONATE* interface to describe the memory region and convey the identity of the Receiver to the Relayer as specified in [Table 5.19](#). This interface is described in [11.1 FFA_MEM_DONATE](#).

2. If the Receiver is a *PE endpoint* or a *SEPID* associated with a dependent peripheral device,
 1. The Owner uses a Partition message to request the Receiver to retrieve the donated memory region. This message contains a description of the memory region relevant to the Receiver.
 2. The Receiver uses the *FFA_MEM_RETRIEVE_REQ* and *FFA_MEM_RETRIEVE_RESP* interfaces to map the memory region in its translation regime and complete the transaction. These interfaces are described in [11.4 FFA_MEM_RETRIEVE_REQ](#) & [11.5 FFA_MEM_RETRIEVE_RESP](#) respectively.

In case of an error, the Sender can abort the transaction before the Receiver retrieves the memory region by calling the *FFA_MEM_RECLAIM* ABI (see [11.7 FFA_MEM_RECLAIM](#)).
3. If the Receiver is a *SEPID* associated with an independent peripheral device, an IMPLEMENTATION DEFINED mechanism is used by the Sender and Relayer to map and describe the memory region to the Receiver (see [11.1.1 Component responsibilities for FFA_MEM_DONATE](#)).

5.7 Lend memory transaction

This transaction is used by an Owner to relinquish its access to a memory region and grant access to it to one or more Borrowers.

- If the region is lent to a single Borrower, it is granted exclusive access to it.
- If the region is lent to more than one Borrower, they are granted shared access to it.

A list of valid combinations of roles played by various FF-A components in this transaction is specified in [Table 5.7](#).

5.7.1 Lend memory transaction state machine

[Table 5.10](#) describes the state machine for lending a memory region from the perspective of two components *A* & *B*. *A* owns the memory region. It attempts to relinquish its access to the memory region and grant shared or exclusive access to it to *B*. Valid and invalid state transitions in response to this transaction have been listed.

Table 5.10: Lend memory transaction state machine

No.	Current Endpoint A state	Current Endpoint B state	Next Endpoint A state	Next Endpoint B state	Description
1	Owner-EA	!Owner-NA	Owner-NA	!Owner-EA or !Owner-SA	• Owner has exclusive access to the memory region and relinquishes access to it to one or more Borrowers including endpoint B.
2	Owner-NA	!Owner-NA	Error	–	• Owner has already lent the memory region to one or more endpoints. It cannot lend it to endpoint B.
3	Owner-NA	!Owner-EA	Error	–	• Owner has already lent the memory region to endpoint B with exclusive access.
4	Owner-NA	!Owner-SA	Error	–	• Owner has already lent the memory region to endpoint B and other endpoints.
5	Owner-SA	!Owner-NA	Error	–	• Owner has already shared the memory region with one or more endpoints. It cannot lend it to endpoint B.
6	Owner-SA	!Owner-SA	Error	–	• Owner has already shared the memory region with endpoint B and possibly other endpoints.

5.7.2 Lend memory transaction lifecycle

This transaction takes place as follows (also see [5.5.2 Transaction life cycle](#)).

1. The Lender uses the *FFA_MEM_LEND* interface to describe the memory region and convey the identities of the Borrowers to the Relayer as specified in [Table 5.19](#). This interface is described in [11.2 FFA_MEM_LEND](#).
2. If a Borrower is a *PE endpoint* or a *SEPID* associated with a dependent peripheral device,
 1. The Lender uses a Partition message to request each Borrower to retrieve the lent memory region. This message contains a description of the memory region relevant to the Borrower.
 2. Each Borrower uses the *FFA_MEM_RETRIEVE_REQ* and *FFA_MEM_RETRIEVE_RESP* interfaces to map the memory region in its translation regime and complete the transaction. These interfaces are described in [11.4 FFA_MEM_RETRIEVE_REQ](#) & [11.5 FFA_MEM_RETRIEVE_RESP](#) respectively.
3. If the Borrower is a *SEPID* associated with an independent peripheral device, an IMPLEMENTATION DEFINED mechanism is used by the Lender and Relayer to map and describe the memory region to the Borrower (see [11.2.1 Component responsibilities for FFA_MEM_LEND](#)).
4. In case of an error, the Lender can abort the transaction before the Borrower retrieves the memory region by calling the *FFA_MEM_RECLAIM* ABI (see [11.7 FFA_MEM_RECLAIM](#)).

5.8 Share memory transaction

This transaction is used by a Owner of a memory region to share access to it with one or more Borrowers.

A list of valid combinations of roles played by various FF-A components in this transaction is specified in [Table 5.7](#).

5.8.1 Share memory transaction state machine

[Table 5.11](#) describes the state machine for sharing a memory region from the perspective of two components *A* & *B*. *A* owns the memory region. It attempts to share the memory region with *B*. Valid and invalid state transitions in response to this transaction have been listed.

Table 5.11: Share memory transaction state machine

No.	Current Endpoint A state	Current Endpoint B state	Next Endpoint A state	Next Endpoint B state	Description
1	Owner-EA	!Owner-NA	Owner-SA	!Owner-SA	• Owner has exclusive access to the memory region and grants access to it to one or more Borrowers including endpoint B.
2	Owner-NA	!Owner-NA	Error	–	• Owner has already lent the memory region to one or more endpoints. It cannot share it with endpoint B.
3	Owner-NA	!Owner-EA	Error	–	• Owner has already lent the memory region to endpoint B with exclusive access.
4	Owner-NA	!Owner-SA	Error	–	• Owner has already lent the memory region to endpoint B and other endpoints.
5	Owner-SA	!Owner-NA	Error	–	• Owner has already shared the memory region with one or more endpoints. It cannot share it with endpoint B.
6	Owner-SA	!Owner-SA	Error	–	• Owner has already shared the memory region with endpoint B and possibly other endpoints.

5.8.2 Share memory transaction lifecycle

This transaction takes place as follows (also see [5.5.2 Transaction life cycle](#)).

1. The Lender uses the *FFA_MEM_SHARE* interface to describe the memory region and convey the identities of the Borrowers to the Relayer as specified in [Table 5.19](#). This interface is described in [11.2 FFA_MEM_LEND](#).
2. If a Borrower is a *PE endpoint* or a *SEPID* associated with a dependent peripheral device,

1. The Lender uses a Partition message to request each Borrower to retrieve the shared memory region. This message contains a description of the memory region relevant to the Borrower.
2. Each Borrower uses the *FFA_MEM_RETRIEVE_REQ* and *FFA_MEM_RETRIEVE_RESP* interfaces to map the memory region in its translation regime and complete the transaction. These interfaces are described in [11.4 FFA_MEM_RETRIEVE_REQ](#) & [11.5 FFA_MEM_RETRIEVE_RESP](#) respectively.
3. If the Borrower is a *SEPID* associated with an independent peripheral device, an IMPLEMENTATION DEFINED mechanism is used by the Lender and Relayer to map and describe the memory region to the Borrower (see [11.3.1 Component responsibilities for FFA_MEM_SHARE](#)).
4. In case of an error, the Lender can abort the transaction before the Borrower retrieves the memory region by calling the *FFA_MEM_RECLAIM* ABI (see [11.7 FFA_MEM_RECLAIM](#)).

5.9 Relinquish memory transaction

This transaction is used by one or more Borrowers to relinquish their access to a memory region so that the Lender can reclaim exclusive access to it. The Lender starts this transaction by requesting each Borrower through a partition message to relinquish access. It reclaims access once all Borrowers have done so.

A list of valid combinations of roles played by various FF-A components in this transaction is specified in [Table 5.8](#).

5.9.1 Relinquish memory access state machine

[Table 5.12](#) describes the state machine for relinquishing a memory region from the perspective of two components *A* & *B*. *A* owns the memory region and could have lent or shared it with *B*. Alternatively, *B* might not have access to the memory region. *B* attempts to relinquish access to this memory region. Valid and invalid state transitions in response to this transaction have been listed.

Table 5.12: Relinquish and reclaim memory state machine

No.	Current Endpoint A state	Current Endpoint B state	Next Endpoint A state	Next Endpoint B state	Description
1	Owner-EA	!Owner-NA	Error	–	• Endpoint B tries to relinquish access to a memory region that the Owner has exclusive access to.
2	Owner-NA	!Owner-NA	Error	–	• Endpoint B tries to relinquish access to a memory region that the Owner has granted shared or exclusive access to one or more other Borrowers.
3	Owner-NA	!Owner-EA	Owner-EA	!Owner-NA	• Endpoint B relinquishes exclusive access to the memory region and transfers it back to the Owner.
4	Owner-NA	!Owner-SA	Owner-EA	!Owner-NA	• Endpoint B relinquishes access to the memory region that it shares with other Borrowers. Owner reclaims exclusive access once all Borrowers have relinquished access.
5	Owner-SA	!Owner-NA	Error	–	• Endpoint B tries to give up access to a memory region that the Owner shares with one or more other Borrowers.

No.	Current Endpoint A state	Current Endpoint B state	Next Endpoint A state	Next Endpoint B state	Description
6	Owner-SA	!Owner-SA	Owner-EA	!Owner-NA	<ul style="list-style-type: none"> Endpoint B relinquishes access to the memory region that it shares with the Owner and possibly other Borrowers. Owner reclaims exclusive access once all Borrowers have relinquished access.

5.9.2 Relinquish memory transaction lifecycle

This transaction takes place as follows (also see [5.5.2 Transaction life cycle](#)). It is assumed that the memory region was originally lent or shared by the Lender to the Borrowers. This transaction must not be used on a memory region owned by an endpoint.

1. If a Borrower is a *PE endpoint* or a *SEPID* associated with a dependent peripheral device,
 1. The Lender could use a Partition message to request each Borrower to relinquish access to the memory region. This message contains a description of the memory region relevant to the Borrower.
 2. Each Borrower uses the *FFA_MEM_RELINQUISH* interface (see [11.6 FFA_MEM_RELINQUISH](#)) to unmap the memory region from its translation regime. This could be done in response to the message from the Lender or independently.
 3. Each Borrower uses a Partition message to inform the Lender that it has relinquished access to the memory region.

In case of an error, the Borrower can abort the transaction before the Lender reclaims the memory region by calling the *FFA_MEM_RETRIEVE_REQ* ABI (see [11.4 FFA_MEM_RETRIEVE_REQ](#)).

2. If the Borrower is a *SEPID* associated with an independent peripheral device,
 1. The Lender could use an IMPLEMENTATION DEFINED mechanism to request each Borrower to relinquish access to the memory region.
 2. Each Borrower uses an IMPLEMENTATION DEFINED mechanism to request the Relayer to unmap the memory region from its translation regime (see [11.7.1 Component responsibilities for FFA_MEM_RECLAIM](#)). This could be done in response to the message from the Lender or independently.
 3. Each Borrower uses an IMPLEMENTATION DEFINED mechanism to inform the Lender that it has relinquished access to the memory region.
3. Once all Borrowers have relinquished access to the memory region, the Lender uses the *FFA_MEM_RECLAIM* interface to reclaim exclusive access to the memory region. This interface is described in [11.7 FFA_MEM_RECLAIM](#).

5.10 Memory region description

A memory region is described in a memory management transaction either through a *Composite memory region descriptor* (see [5.10.1 Composite memory region descriptor](#)) or a globally unique *Handle* (see [5.10.2 Memory region handle](#)).

The former is used to describe a memory region when a transaction to share, lend or donate memory is initiated by the Owner and when the memory region is retrieved by the Receiver.

The latter is used to describe a memory region when it is retrieved by a Receiver, relinquished by a Borrower and reclaimed by the Owner.

5.10.1 Composite memory region descriptor

A memory region is described in a memory management transaction by specifying the list and count of 4K sized pages that constitute it (see [Table 5.13](#)).

Table 5.13: Composite memory region descriptor

Field	Byte length	Byte offset	Description
Total page count	4	0	<ul style="list-style-type: none">• Size of the memory region described as the count of 4K pages.• Must be equal to the sum of page counts specified in each constituent memory region descriptor. See Table 5.14.
Address range count	4	4	<ul style="list-style-type: none">• Count of address ranges specified using constituent memory region descriptors.
Reserved (MBZ)	8	8	
Address range array	–	16	<ul style="list-style-type: none">• Array of address ranges specified using constituent memory region descriptors.

The list is specified by using one or more constituent memory region descriptors (see [Table 5.14](#)). Each descriptor specifies the base address and size of a virtually or physically contiguous memory region.

Table 5.14: Constituent memory region descriptor

Field	Byte length	Byte offset	Description
Address	8	–	<ul style="list-style-type: none">• Base VA, PA or IPA of constituent memory region aligned to the page size (4K) granularity.
Page count	4	8	<ul style="list-style-type: none">• Number of 4K pages in constituent memory region.
Reserved (MBZ)	4	12	

The pages are addressed using VAs, IPAs or PAs depending on the FF-A instance at which the transaction is taking place. This is as follows.

- VAs are used at a Secure virtual FF-A instance if the partition runs in Secure EL0 or Secure User mode.
- IPAs are used at a virtual FF-A instance if the partition runs in one of the following Exception levels.
 - Secure EL1.
 - Secure Supervisor mode.
 - EL1.
 - Supervisor mode.
- PAs are used at all physical FF-A instances.

Figure 5.2 describes a virtually contiguous memory region range *VA_0* of size 16K through its composite memory region descriptors at the virtual and physical FF-A instances. *VA_0* was allocated through a dynamic memory management mechanism inside an endpoint for example, malloc. It is composed of:

- Two constituent IPA regions *IPA_0* and *IPA_1* of size 8K each at the virtual FF-A instance.
- *IPA_0* is comprised of two PA regions *PA_0* and *PA_1* at the physical FF-A instance. Each PA region is of size 4K.
- *IPA_1* is comprised of two PA regions *PA_2* and *PA_3* at the physical FF-A instance. Each PA region is of size 4K.




Figure 5.2: Example memory region description

5.10.2 Memory region handle

- A 64-bit *Handle* is used to identify a composite memory region description for example, *VA_0* described in [Figure 5.2](#)
- The *Handle* is allocated by the *Relayer* as follows.
 - The SPM must allocate the *Handle* if any Receiver participating in the memory management transaction is an SP or SEPID associated with a Secure Stream ID in the SMMU.
 - The Hypervisor must allocate the *Handle* if no Receiver participating in the memory management transaction is an SP or SEPID associated with a Secure Stream ID in the SMMU.
- A *Handle* is allocated once a transaction to lend, share or donate memory is successfully initiated by the *Owner*.
- Each *Handle* identifies a single unique composite memory region description that is, there is a 1:1 mapping between the two.
- A *Handle* is freed by the *Relayer* after it has been reclaimed by its *Owner* at the end of a successful transaction to relinquish the corresponding memory region description.
- Encoding of a *Handle* is as follows.
 - Bit[63]: *Handle* allocator.
 - * b'0: Allocated by SPM.
 - * b'1: Allocated by Hypervisor.
 - Bit[62:0]: IMPLEMENTATION DEFINED.
- A *Handle* must be encoded as a register parameter in any ABI that requires it as follows.
 - Two 32-bit general-purpose registers must be used such that if *Rx* and *Ry* are used, such that $x < y$,
 - * $Rx = Handle[31:0]$.
 - * $Ry = Handle[63:32]$.

5.11 Memory region properties

The properties of a memory region are as follows.

- *Instruction and data access permissions* describe the type of access permitted on the memory region.
- *One or more endpoint IDs* that have access to the memory region specified by a combination of access permissions and memory region attributes.
- *Memory region attributes* control the memory type, accesses to the caches, and whether the memory region is Shareable and therefore is coherent.

There is a 1:1 association between an endpoint and the permissions with which it can access a memory region. This is specified in [Table 5.15](#).

Table 5.15: Memory access permissions descriptor

Field	Byte length	Byte offset	Description
Endpoint ID	2	–	<ul style="list-style-type: none"> • 16-bit ID of endpoint to which the memory access permissions apply.
Memory access permissions	1	2	<ul style="list-style-type: none"> • Permissions used to access a memory region. <ul style="list-style-type: none"> – bits[7:4]: Reserved (MBZ). – bits[3:2]: Instruction access permission. <ul style="list-style-type: none"> * b'00: Not specified and must be ignored. * b'01: Not executable. * b'10: Executable. * b'11: Reserved. Must not be used. – bits[1:0]: Data access permission. <ul style="list-style-type: none"> * b'00: Not specified and must be ignored. * b'01: Read-only. * b'10: Read-write. * b'11: Reserved. Must not be used.
Flags	1	3	<ul style="list-style-type: none"> • ABI specific flags as described in 5.11.1 ABI-specific flags usage.

[Table 5.16](#) specifies the data structure that is used in memory management transactions to create an association between an endpoint, memory access permissions and a composite memory region description.

This data structure must be included in other data structures that are used in memory management transactions instead of being used as a stand alone data structure (see [5.12 Lend, donate, and share transaction descriptor](#)). A composite memory region description is referenced by specifying an offset to it as described in [Table 5.16](#). This enables one or more endpoints to be associated with the same memory region but with different memory access permissions for example, SP0 could have *RO* data access permission and SP1 could have *RW* data access permission to the same memory region.

Table 5.16: Endpoint memory access descriptor

Field	Byte length	Byte offset	Description
Memory access permissions descriptor	4	–	<ul style="list-style-type: none"> Memory access permissions descriptor as specified in Table 5.15.
Composite memory region descriptor offset	4	4	<ul style="list-style-type: none"> Offset to the composite memory region descriptor to which the endpoint access permissions apply (see Table 5.13). Offset must be calculated from the base address of the data structure this descriptor is included in. An offset value of 0 indicates that the endpoint access permissions apply to a memory region description identified by the <i>Handle</i> parameter specified in the data structure that includes this one.
Reserved (MBZ)	8	8	

5.11.1 ABI-specific flags usage

An endpoint can specify properties specific to the memory management ABI being invoked through this field.

In this version of the Framework, the *Flags* field MBZ and is reserved in an invocation of the following ABIs.

- FFA_MEM_DONATE.
- FFA_MEM_LEND.
- FFA_MEM_SHARE.

The *Flags* field must be encoded by the Receiver and the Relayer as specified in [Table 5.17](#) in an invocation of the following ABIs.

- FFA_MEM_RETRIEVE_REQ.
- FFA_MEM_RETRIEVE_RESP.

The Relayer must return *INVALID_PARAMETERS* if the *Flags* field has been incorrectly encoded.

Table 5.17: Flags usage in FFA_MEM_RETRIEVE_REQ and FFA_MEM_RETRIEVE_RESP ABIs

Field	Description
Bit[0]	<ul style="list-style-type: none"> Non-retrieval Borrower flag. <ul style="list-style-type: none"> In a memory management transaction with multiple Borrowers, during the retrieval of the memory region, this flag specifies if the memory region must be or was retrieved on behalf of this endpoint or if the endpoint is another Borrower. <ul style="list-style-type: none"> b'0: Memory region must be or was retrieved on behalf of this endpoint. b'1: Memory region must not be or was not retrieved on behalf of this endpoint. It is another Borrower of the memory region. This field MBZ if this endpoint: <ul style="list-style-type: none"> Is the only PE endpoint Borrower/Receiver in the transaction. Is a <i>Stream endpoint</i> and the caller of the <i>FFA_MEM_RETRIEVE_REQ</i> ABI is its <i>proxy endpoint</i>.
Bit[7:1]	<ul style="list-style-type: none"> Reserved (MBZ).

5.11.2 Data access permissions usage

An endpoint could have either *Read-only* or *Read-write* data access permission to a memory region from the highest Exception level it runs in.

- Read-write* permission is more permissive than *Read-only* permission.
- Data access permission is specified by setting *Bits[1:0]* in [Table 5.15](#) to the appropriate value.

This access control is used in memory management transactions as follows.

- In a transaction to lend or share memory,
 - The Lender must specify the level of access that the Borrower is permitted to have on the memory region. This is done while invoking the *FFA_MEM_SHARE* or *FFA_MEM_LEND* ABIs.
 - The Relayer must validate the permission specified by the Lender as follows. This is done in response to an invocation of the *FFA_MEM_SHARE* or *FFA_MEM_LEND* ABIs. The Relayer must return the *DENIED* error code if the validation fails.
 - At the Non-secure physical FF-A instance, an IMPLEMENTATION DEFINED mechanism is used to perform validation.
 - At any virtual FF-A instance, if the endpoint is running in EL1 or S-EL1 in either Execution state, the permission specified by the Lender is considered valid only if it is the same or less permissive than the permission used by the Relayer in the *S2AP* field in the stage 2 translation table descriptors for the memory region in one of the following translation regimes:
 - Secure EL1&0 translation regime, when S-EL2 is enabled.
 - Non-secure EL1&0 translation regime, when EL2 is enabled.
 - At the Secure virtual FF-A instance, if the endpoint is running in S-EL0 in either Execution state, the permission specified by the Lender is considered valid only if it is the same or less permissive than the permission used by the Relayer in the *AP[1]* field in the stage 1 translation table descriptors for the memory region in one of the following translation regimes:
 - Secure EL1&0 translation regime, when EL2 is disabled.
 - Secure PL1&0 translation regime, when EL2 is disabled.
 - Secure EL2&0 translation regime, when Armv8.1-VHE is enabled.

If the Borrower is an independent peripheral device, then the validated permission is used to map the memory region into the address space of the device.

- The Borrower (if a PE or Proxy endpoint) should specify the level of access that it would like to have on the memory region.

In a transaction to share or lend memory with more than one Borrower, each Borrower (if a PE or Proxy endpoint) could also specify the level of access that other Borrowers have on the memory region.

This is done while invoking the *FFA_MEM_RETRIEVE_REQ* ABI.

- The Relayer must validate the permissions, if specified by the Borrower (if a PE or Proxy endpoint) in response to an invocation of the *FFA_MEM_RETRIEVE_REQ* ABI.

It must ensure that the permission of the Borrower is the same or less permissive than the permission that was specified by the Lender and validated by the Relayer.

It must ensure that the permissions for other Borrowers are the same as those specified by the Lender and validated by the Relayer.

The Relayer must return the *DENIED* error code if the validation fails.

2. In a transaction to donate memory,

- Whether the Owner is allowed to specify the level of access that the Receiver is permitted to have on the memory region depends on the type of Receiver.
 - If the Receiver is a PE or Proxy endpoint, the Owner must not specify the level of access.
 - If the Receiver is an independent peripheral device, the Owner could specify the level of access.

The Owner must specify its choice in an invocation of the *FFA_MEM_DONATE* ABI.

- The value of data access permission field specified by the Owner must be interpreted by the Relayer as follows. This is done in response to an invocation of the *FFA_MEM_DONATE* ABI.
 - If the Receiver is a PE or Proxy endpoint, the Relayer must return *INVALID_PARAMETERS* if the value is not *b'00*.
 - If the Receiver is an independent peripheral device and the value is not *b'00*, the Relayer must take one of the following actions.
 - * Return *DENIED* if the permission is determined to be invalid through an IMPLEMENTATION DEFINED mechanism.
 - * Use the permission specified by the Owner to map the memory region into the address space of the device.
 - If the Receiver is an independent peripheral device and the value is *b'00*, the Relayer must determine the permission value through an IMPLEMENTATION DEFINED mechanism.
- The Receiver (if a PE or Proxy endpoint) should specify the level of access that it would like to have on the memory region. This is done while invoking the *FFA_MEM_RETRIEVE_REQ* ABI.
- The Relayer must validate the permission specified by the Receiver to ensure that it is the same or less permissive than the permission determined by the Relayer through an IMPLEMENTATION DEFINED mechanism. This is done in response to an invocation of the *FFA_MEM_RETRIEVE_REQ* ABI. The Relayer must return the *DENIED* error code if the validation fails.

3. The Relayer must specify the permission that was used to map the memory region in the translation regime of the Receiver or Borrower. This must be done in an invocation of the *FFA_MEM_RETRIEVE_RESP* ABI.

4. In a transaction to relinquish memory that was lent to one or more Borrowers, the memory region must be mapped back into the translation regime of the Lender with the same data access permission that was used at the start of the transaction to lend the memory region. This is done in response to an invocation of the *FFA_MEM_RECLAIM* ABI.

5.11.3 Instruction access permissions usage

An endpoint could have either *Execute (X)* or *Execute-never (XN)* instruction access permission to a memory region from the highest Exception level it runs in.

- *Execute* permission is more permissive than *Execute-never* permission.
- Instruction access permission is specified by setting *Bits[3:2]* in [Table 5.15](#) to the appropriate value.

This access control is used in memory management transactions as follows.

1. Only XN permission must be used in the following transactions.

- In a transaction to share memory with one or more Borrowers.
- In a transaction to lend memory to more than one Borrower.

Bits[3:2] in [Table 5.15](#) must be set to *b'00* as follows.

- By the Lender in an invocation of *FFA_MEM_SHARE* or *FFA_MEM_LEND* ABIs.
- By the Borrower in an invocation of the *FFA_MEM_RETRIEVE_REQ* ABI.

The Relayer must set *Bits[3:2]* in [Table 5.15](#) to *b'01* while invoking the *FFA_MEM_RETRIEVE_RESP* ABI.

2. In a transaction to donate memory or lend memory to a single Borrower,

- Whether the Owner/Lender is allowed to specify the level of access that the Receiver is permitted to have on the memory region depends on the type of Receiver.
 - If the Receiver is a PE or Proxy endpoint, the Owner must not specify the level of access.
 - If the Receiver is an independent peripheral device, the Owner could specify the level of access.

The Owner must specify its choice in an invocation of the *FFA_MEM_DONATE* or *FFA_MEM_LEND* ABIs.

- The value of instruction access permission field specified by the Owner/Lender must be interpreted by the Relayer as follows. This is done in response to an invocation of the *FFA_MEM_DONATE* or *FFA_MEM_LEND* ABIs.
 - If the Receiver is a PE or Proxy endpoint, the Relayer must return *INVALID_PARAMETERS* if the value is not *b'00*.
 - If the Receiver is an independent peripheral device and the value is not *b'00*, the Relayer must take one of the following actions.
 - * Return *DENIED* if the permission is determined to be invalid through an IMPLEMENTATION DEFINED mechanism.
 - * Use the permission specified by the Owner to map the memory region into the address space of the device.
 - If the Receiver is an independent peripheral device and the value is *b'00*, the Relayer must determine the permission value through an IMPLEMENTATION DEFINED mechanism.
- The Receiver (if a PE or Proxy endpoint) should specify the level of access that it would like to have on the memory region. This must be done in an invocation of the *FFA_MEM_RETRIEVE_REQ* ABI.
- The Relayer must validate the permission specified by the Receiver (if a PE or Proxy endpoint) to ensure that it is the same or less permissive than the permission determined by the Relayer through an IMPLEMENTATION DEFINED mechanism.
 - For example, the Relayer could deny executable access to a Borrower on a memory region of Device memory type.

This is done in response to an invocation of the *FFA_MEM_RETRIEVE_REQ* ABI. The Relayer must return the *DENIED* error code if the validation fails.

If the invocation of `FFA_MEM_RETRIEVE_REQ` succeeds, the Relayer must set `Bits[3:2]` in [Table 5.15](#) to either `b'01` or `b'10` while invoking the `FFA_MEM_RETRIEVE_RESP` ABI.

3. In a transaction to relinquish memory that was lent to one or more Borrowers, the memory region must be mapped back into the translation regime of the Lender with the same instruction access permission that was used at the start of the transaction to lend the memory region. This is done in response to an invocation of the `FFA_MEM_RECLAIM` ABI.

5.11.4 Memory region attributes usage

An endpoint can access a memory region by specifying attributes as follows.

- *Memory type*. This could be Device or Normal. Device memory type could be one of the following types.
 - Device-nGnRnE.
 - Device-nGnRE.
 - Device-nGRE.
 - Device-GRE.

The precedence rules for memory types are as follows. < should be read as *is less permissive than*.

- Device-nGnRnE < Device-nGnRE < Device-nGRE < Device-GRE < Normal.

- *Cacheability attribute*. This could be one of the following types.
 - Non-cacheable.
 - Write-Back Cacheable.

These attributes are used to specify both inner and outer cacheability. The precedence rules are as follows.

- Non-cacheable < Write-Back Cacheable.

- *Shareability attribute*. This could be one of the following types.
 - Non-shareable.
 - Outer Shareable.
 - Inner Shareable.

The precedence rules are as follows.

- Non-Shareable < Inner Shareable < Outer shareable.

Memory region attributes are specified by an endpoint by setting `Bits[5:0]` in [Table 5.18](#) to appropriate values.

The data structure to encode memory region attributes is specified in [Table 5.18](#).

Table 5.18: Memory region attributes descriptor

Field	Byte length	Byte offset	Description
Memory region attributes	1	–	<ul style="list-style-type: none"> Attributes used to access a memory region. <ul style="list-style-type: none"> bits[7:6]: Reserved (MBZ). bits[5:4]: Memory type. <ul style="list-style-type: none"> b'00: Not specified and must be ignored. b'01: Device memory. b'10: Normal memory. b'11: Reserved. Must not be used. bits[3:2]: <ul style="list-style-type: none"> Cacheability attribute if bit[5:4] = b'10. <ul style="list-style-type: none"> b'00: Reserved. Must not be used. b'01: Non-cacheable. b'10: Reserved. Must not be used. b'11: Write-Back. Device memory attributes if bit[5:4] = b'01. <ul style="list-style-type: none"> b'00: Device-nGnRnE. b'01: Device-nGnRE. b'10: Device-nGRE. b'11: Device-GRE. bits[1:0]: <ul style="list-style-type: none"> Shareability attribute if bit[5:4] = b'10. <ul style="list-style-type: none"> b'00: Non-shareable. b'01: Reserved. Must not be used. b'10: Outer Shareable. b'11: Inner Shareable. Reserved & MBZ if bit[5:4] = b'01.

Memory region attributes are used in memory management transactions as follows.

- In a transaction to share memory with one or more Borrowers and to lend memory to more than one Borrower,
 - The Lender specifies the attributes that each Borrower must access the memory region with. This is done by invoking the *FFA_MEM_SHARE* or *FFA_MEM_LEND* ABIs. The same attributes are used for all Borrowers.
 - The Relayer validates the attributes specified by the Lender as follows. This is done in response to an invocation of the *FFA_MEM_SHARE* or *FFA_MEM_LEND* ABIs. The Relayer must return the *DENIED* error code if the validation fails.
 - At the Non-secure physical FF-A instance, an IMPLEMENTATION DEFINED mechanism is used.
 - At any virtual FF-A instance, if the endpoint is running in EL1 or S-EL1 in either Execution state, the attributes specified by the Lender are considered valid only if they are the same or less permissive than the attributes used by the Relayer in the stage 2 translation table descriptors for the memory region in one of the following translation regimes:
 - Secure EL1&0 translation regime, when S-EL2 is enabled.
 - Non-secure EL1&0 translation regime, when EL2 is enabled.
 - At the Secure virtual FF-A instance, if the endpoint is running in S-EL0 in either Execution state, the attributes specified by the Lender are considered valid only if they are either the same or less

permissive than the attributes used by the Relayer in the stage 1 translation table descriptors for the memory region in one of the following translation regimes:

- * Secure EL1&0 translation regime, when EL2 is disabled.
- * Secure PL1&0 translation regime, when EL2 is disabled.
- * Secure EL2&0 translation regime, when Armv8.1-VHE is enabled.

If the Borrower is an independent peripheral device, then the validated attributes are used to map the memory region into the address space of the device.

- The Borrower (if a PE or Proxy endpoint) should specify the attributes that it would like to access the memory region with. This is done by invoking the *FFA_MEM_RETRIEVE_REQ* ABI.
- The Relayer must validate the attributes specified by the Borrower (if a PE or Proxy endpoint) to ensure that they are the same or less permissive than the attributes that were specified by the Lender and validated by the Relayer. This is done in response to an invocation of the *FFA_MEM_RETRIEVE_REQ* ABI. The Relayer must return the *DENIED* error code if the validation fails.

2. In a transaction to donate memory or lend memory to a single Borrower,

- Whether the Owner/Lender is allowed to specify the memory region attributes that the Receiver must use to access the memory region depends on the type of Receiver.
 - If the Receiver is a PE or Proxy endpoint, the Owner must not specify the attributes.
 - If the Receiver is an independent peripheral device, the Owner could specify the attributes.

The Owner must specify its choice in an invocation of the *FFA_MEM_DONATE* or *FFA_MEM_LEND* ABIs.

- The values in the memory region attributes field specified by the Owner/Lender must be interpreted by the Relayer as follows. This is done in response to an invocation of the *FFA_MEM_DONATE* or *FFA_MEM_LEND* ABIs.
 - If the Receiver is a PE or Proxy endpoint, the Relayer must return *INVALID_PARAMETERS* if the value in *bits[5:4] != b'00*.
 - If the Receiver is an independent peripheral device and the value is not *b'00*, the Relayer must take one of the following actions.
 - * Return *DENIED* if the attributes are determined to be invalid through an IMPLEMENTATION DEFINED mechanism.
 - * Use the attributes specified by the Owner to map the memory region into the address space of the device.
 - If the Receiver is an independent peripheral device and the value is *b'00*, the Relayer must determine the attributes through an IMPLEMENTATION DEFINED mechanism.
- The Receiver (if a PE or Proxy endpoint) should specify the memory region attributes it would like to use to access the memory region. This is done while invoking the *FFA_MEM_RETRIEVE_REQ* ABI.
- The Relayer must validate the attributes specified by the Receiver (if a PE or Proxy endpoint) to ensure that they are the same or less permissive than the attributes determined by the Relayer through an IMPLEMENTATION DEFINED mechanism.

This is done in response to an invocation of the *FFA_MEM_RETRIEVE_REQ* ABI. The Relayer must return the *DENIED* error code if the validation fails.

3. The Relayer must specify the memory region attributes that were used to map the memory region in the translation regime of the Receiver or Borrower. This must be done while invoking the *FFA_MEM_RETRIEVE_RESP* ABI.

4. In a transaction to relinquish memory that was lent to one or more Borrowers, the memory region must be mapped back into the translation regime of the Lender with the same attributes that were used at the start of the transaction to lend the memory region. This is done in response to an invocation of the *FFA_MEM_RECLAIM* ABI.

5.12 Lend, donate, and share transaction descriptor

[Table 5.19](#) specifies the data structure that must be used by the Owner/Lender and a Borrower/Receiver in a transaction to donate, lend or share a memory region. It specifies the memory region description (see [5.10 Memory region description](#)), properties (see [5.11 Memory region properties](#)) and other transaction attributes in an invocation of the following ABIs.

- FFA_MEM_DONATE.
- FFA_MEM_LEND.
- FFA_MEM_SHARE.
- FFA_MEM_RETRIEVE_REQ.
- FFA_MEM_RETRIEVE_RESP.

The interpretation of some fields in [Table 5.19](#) depends on the ABI this table is used with. This variance in behavior is also specified in [Table 5.19](#).

Table 5.19: Lend, donate or share memory transaction descriptor

Field	Byte length	Byte offset	Description
Sender endpoint ID	2	0	<ul style="list-style-type: none"> • ID of the Owner endpoint.
Memory region attributes	1	2	<ul style="list-style-type: none"> • Attributes must be encoded as specified in 5.11.4 Memory region attributes usage. • Attribute usage is subject to validation at the virtual and physical FF-A instances as specified in 5.11.4 Memory region attributes usage.
Reserved (MBZ)	1	3	
Flags	4	4	<ul style="list-style-type: none"> • Flags must be encoded as specified in 5.12.4 Flags usage.
Handle	8	8	<ul style="list-style-type: none"> • Memory region handle in ABI invocations specified in 5.12.1 Handle usage.
Tag	8	16	<ul style="list-style-type: none"> • This field must be encoded as specified in 5.12.2 Tag usage.
Reserved (MBZ)	4	24	
Endpoint memory access descriptor count	4	28	<ul style="list-style-type: none"> • Count of endpoint memory access descriptors.
Endpoint memory access descriptor array	–	32	<ul style="list-style-type: none"> • Each entry in the array must be encoded as specified in 5.12.3 Endpoint memory access descriptor array usage. See Table 5.16 for the encoding of the endpoint memory access descriptor.

5.12.1 Handle usage

- This field must be zero (MBZ) in an invocation of the following ABIs.
 - FFA_MEM_DONATE.
 - FFA_MEM_LEND.
 - FFA_MEM_SHARE.
- A successful invocation of each of the preceding ABIs returns a *Handle* (see [5.10.2 Memory region handle](#)) to identify the memory region in the transaction.
- The Sender must convey the *Handle* to the Receiver through a Partition message.
- This field must be used by the Receiver to encode this *Handle* in an invocation of the FFA_MEM_RETRIEVE_REQ ABI.
- A Relayer must validate this field in an invocation of the FFA_MEM_RETRIEVE_REQ ABI as follows.
 - Ensure that it holds a *Handle* value that was previously allocated and has not been reclaimed by the Owner.
 - Ensure that the *Handle* identifies a memory region that was shared, lent or donated to the Receiver.
 - Ensure that the *Handle* was allocated to the Owner specified in the *Sender endpoint ID* field of the transaction descriptor.

It must return *INVALID_PARAMETERS* if the validation fails.

- This field must be used by the Relayer to encode the *Handle* in an invocation of the FFA_MEM_RETRIEVE_RESP ABI.

5.12.2 Tag usage

- This 64-bit field must be used to specify an IMPLEMENTATION DEFINED value associated with the transaction and known to participating endpoints.
- The Sender must specify this field to the Relayer in an invocation of the following ABIs.
 - FFA_MEM_DONATE.
 - FFA_MEM_LEND.
 - FFA_MEM_SHARE.
- The Sender must convey the *Tag* to the Receiver through a Partition message.
- This field must be used by the Receiver to encode the *Tag* in an invocation of the FFA_MEM_RETRIEVE_REQ ABI.
- The Relayer must ensure the *Tag* value specified by the Receiver is equal to the value that was specified by the Sender. It must return *INVALID_PARAMETERS* if the validation fails.
- This field must be used by the Relayer to encode the *Tag* value in an invocation of the FFA_MEM_RETRIEVE_RESP ABI.

5.12.3 Endpoint memory access descriptor array usage

5.12.3.1 Sender usage

A Sender must use this field to specify one or more Receivers and the access permissions each should have on the memory region it is donating, lending or sharing through an invocation of one of the following ABIs.

- FFA_MEM_DONATE.
- FFA_MEM_LEND.
- FFA_MEM_SHARE.

The access permissions and flags are subject to validation at the virtual and physical FF-A instances as specified in [5.11.3 Instruction access permissions usage](#), [5.11.2 Data access permissions usage](#) and [5.11.1 ABI-specific flags usage](#).

In a FFA_MEM_SHARE ABI invocation, the Sender could request the memory region to be mapped with different data access permission in its own translation regime. It must specify these permissions and its endpoint ID in a separate Endpoint memory access descriptor.

A Sender must describe the memory region in a composite memory region descriptor (see [Table 5.13](#)) with the following non-exhaustive list of checks.

- Ensure that the address ranges specified in the composite memory region descriptor do not overlap each other.
- *Total page count* is equal to the sum of the *Page count* fields in each *Constituent memory region descriptor*.

The offset to this descriptor from the base of [Table 5.19](#) must be specified in the *Offset* field of the Endpoint memory access descriptor as follows.

- In a FFA_MEM_DONATE ABI invocation,
 - The *Endpoint memory access descriptor count* field in the transaction descriptor must be set to 1. This implies that the Owner must specify a single Receiver endpoint in a transaction to donate memory.
 - The *Offset* field of the Endpoint memory access descriptor must be set to the offset of the composite memory region descriptor
- In a FFA_MEM_LEND and FFA_MEM_SHARE ABI invocation,
 - The *Endpoint memory access descriptor count* field in the transaction descriptor must be set to a non-zero value. This implies that the Owner must specify at least a single Borrower endpoint in a transaction to lend or share memory.
 - The *Offset* field in the Endpoint memory access descriptor of each Borrower must be set to the offset of the composite memory region descriptor. This implies that all values of the *Offset* field must be equal.

5.12.3.2 Receiver usage

A Receiver must use this field to specify the access permissions it should have on the memory region being donated, lent or shared in an invocation of the FFA_MEM_RETRIEVE_REQ ABI. This is specified in [5.11.3 Instruction access permissions usage](#) and [5.11.2 Data access permissions usage](#).

- A Receiver could do this on its behalf if it is a PE endpoint.
- A Receiver could do this on the behalf of its dependent peripheral devices if it is a proxy endpoint.

A Receiver could specify the address ranges that must be used to map the memory region in its translation regime by describing them in a composite memory region descriptor. The Receiver must perform the same checks as a Sender. These checks are described in [5.12.3.1 Sender usage](#).

The offset to this descriptor from the base of [Table 5.19](#) must be specified in the *Offset* field of the corresponding endpoint memory access descriptor in the array. This implies that all values of the *Offset* field could be different from each other.

A Receiver could let the Relay allocate the address ranges that must be used to map the memory region in its translation regime and optionally provide an alignment hint (see *Address range alignment hint* in [Table 5.21](#)). The value 0 must be specified in the *Offset* field of the corresponding endpoint memory access descriptor in the array. This implies that the *Handle* specified in [Table 5.19](#) must be used to identify the memory region (see [5.12.1 Handle usage](#)).

A memory management transaction could be to lend or share memory with multiple Borrowers. The Receiver must use this field to specify:

- The SEPIDs and data access permissions of any dependent peripheral devices (if any) that the Receiver is a *proxy endpoint* for.

If the Relayer must allocate the address ranges to map the memory region in the *Stream endpoints*, the value 0 must be specified in the *Offset* field of the corresponding endpoint memory access descriptor in the array.

If the Receiver specifies the address ranges to map the memory region in the *Stream endpoints*, then it must follow the preceding guidance to specify the address ranges that must be used to map the memory region in its translation regime.

- The identity of any other Borrowers and their data access permissions on the memory region (see [5.11.2 Data access permissions usage](#) and [5.11.1 ABI-specific flags usage](#)).

The value 0 must be specified in the *Offset* field of the corresponding endpoint memory access descriptor in the array.

5.12.3.3 Relayer usage

A Relayer must validate the *Endpoint memory access descriptor count* and each entry in the *Endpoint memory access descriptor array* as follows.

- The Relayer could support memory management transactions targeted to only a single Receiver endpoint.
It must return *INVALID_PARAMETERS* if the Sender or Receiver specifies an *Endpoint memory access descriptor count* $\neq 1$.
- It must ensure that these fields have been populated by the Sender as specified in [5.12.3.1 Sender usage](#) in an invocation of any of the following ABIs.
 - FFA_MEM_DONATE.
 - FFA_MEM_LEND.
 - FFA_MEM_SHARE.

The Relayer must return *INVALID_PARAMETERS* in case of an error.

- It must ensure that the *Endpoint ID* field in each Memory access permissions descriptor specifies a valid endpoint. The Relayer must return *INVALID_PARAMETERS* in case of an error.
- In an invocation of the FFA_MEM_RETRIEVE_REQ ABI,
 - It must ensure that these fields have been populated by the Receiver as specified in [5.12.3.2 Receiver usage](#).
 - If the memory region has been lent or shared with multiple Borrowers, the Relayer must ensure that the identity of each Borrower specified by the Receiver is the same as that specified by the Sender.
 - If one or more Borrowers are dependent peripheral devices, the Relayer must ensure that the Receiver is their proxy endpoint.
 - If the Receiver specifies the address ranges that must be used to map the memory region in its translation regime, the Relayer must ensure that the size of the memory region is equal to that specified by the Sender.

The Relayer must return *INVALID_PARAMETERS* in case of an error.

- It must validate the access permissions in the *Memory access permissions descriptor* in each *Endpoint memory access descriptor* as specified in [5.11.2 Data access permissions usage](#) and [5.11.3 Instruction access permissions usage](#).

A Relayer must use this field in an invocation of the FFA_MEM_RETRIEVE_RESP ABI in response to successful validation of a FFA_MEM_RETRIEVE_REQ ABI invocation as follows.

- To specify the access permissions with which the memory region has been mapped in the translation regime of the Receiver.
- A Receiver could let the Relayer allocate the address ranges to map the memory region. In this case, the Relayer must describe the address ranges in a composite memory region descriptor. The Relayer must perform the same checks as a Sender. These checks are described in [5.12.3.1 Sender usage](#).

The offset to this descriptor from the base of [Table 5.19](#) must be specified in the *Offset* field of the corresponding endpoint memory access descriptor in the array. This implies that all values of the *Offset* field could be different from each other.

- A Receiver could specify the address ranges that must be used to map the memory region in its translation regime. The Relayer must specify the value 0 in the *Offset* field of the corresponding endpoint memory access descriptor in the array.

5.12.4 Flags usage

- Flags are used to govern the behavior of a memory management transaction.
- Usage of the Flags field in an invocation of the following ABIs is specified in [Table 5.20](#).
 - FFA_MEM_DONATE.
 - FFA_MEM_LEND.
 - FFA_MEM_SHARE.
- Usage of the Flags field in an invocation of the FFA_MEM_RETRIEVE_REQ ABI is specified in [Table 5.21](#).
- Usage of the Flags field in an invocation of the FFA_MEM_RETRIEVE_RESP ABI is specified in [Table 5.22](#).

5.12.4.1 Zero memory flag

In some ABI invocations, the caller could set a flag to request the Relayer to zero a memory region. To do this, the Relayer must:

- Map the memory region in its translation regime once it is not mapped in the translation regime of any other FF-A component.

The caller must ensure that the memory region fulfills the size and alignment requirements listed in [2.7 Memory granularity and alignment](#) to allow the Relayer to map this memory region. It must discover these requirements by invoking the *FFA_FEATURES* interface with the function ID of the *FFA_RXTX_MAP* interface (see [8.2 FFA_FEATURES](#)).

The Relayer must return *INVALID_PARAMETERS* if the memory region does not meet these requirements.

- Zero the memory region and perform cache maintenance such that the memory regions contents are coherent between any PE caches, system caches and system memory.
- Unmap the memory region from its translation regime before it is mapped in the translation regime of any other FF-A component.

Table 5.20: Flags usage in FFA_MEM_DONATE, FFA_MEM_LEND and FFA_MEM_SHARE ABIs

Field	Description
Bit[0]	<ul style="list-style-type: none"> • Zero memory flag. <ul style="list-style-type: none"> – In an invocation of FFA_MEM_DONATE or FFA_MEM_LEND, this flag specifies if the memory region contents must be zeroed by the Relayer after the memory region has been unmapped from the translation regime of the Owner. <ul style="list-style-type: none"> * b'0: Relayer must not zero the memory region contents. * b'1: Relayer must zero the memory region contents. – MBZ in an invocation of FFA_MEM_SHARE, else the Relayer must return <i>INVALID_PARAMETERS</i>. – MBZ if the Owner has Read-only access to the memory region , else the Relayer must return <i>DENIED</i>.

Field	Description
Bit[1]	<ul style="list-style-type: none"> Operation time slicing flag. <ul style="list-style-type: none"> In an invocation of FFA_MEM_DONATE, FFA_MEM_LEND or FFA_MEM_SHARE, this flag specifies if the Relayer can time slice this operation. <ul style="list-style-type: none"> b'0: Relayer must not time slice this operation. b'1: Relayer can time slice this operation. MBZ if the Relayer does not support time slicing of memory management operations (see 12.2.3 Time slicing of memory management operations), else the Relayer must return <i>INVALID_PARAMETERS</i>.
Bit[31:2]	<ul style="list-style-type: none"> Reserved (MBZ).

Table 5.21: Flags usage in FFA_MEM_RETRIEVE_REQ ABI

Field	Description
Bit[0]	<ul style="list-style-type: none"> Zero memory before retrieval flag. <ul style="list-style-type: none"> In an invocation of FFA_MEM_RETRIEVE_REQ, during a transaction to lend or donate memory, this flag is used by the Receiver to specify whether the memory region must be retrieved with or without zeroing its contents first. <ul style="list-style-type: none"> b'0: Retrieve the memory region irrespective of whether the Sender requested the Relayer to zero its contents prior to retrieval. b'1: Retrieve the memory region only if the Sender requested the Relayer to zero its contents prior to retrieval by setting the <i>Bit[0]</i> in Table 5.20. MBZ in a transaction to share a memory region, else the Relayer must return <i>INVALID_PARAMETER</i>. If the Sender has Read-only access to the memory region and the Receiver sets <i>Bit[0]</i>, the Relayer must return <i>DENIED</i>. MBZ if the Receiver has previously retrieved this memory region, else the Relayer must return <i>INVALID_PARAMETERS</i> (see 11.4.2 Support for multiple retrievals by a Borrower).
Bit[1]	<ul style="list-style-type: none"> Operation time slicing flag. <ul style="list-style-type: none"> In an invocation of FFA_MEM_RETRIEVE_REQ, this flag specifies if the Relayer can time slice this operation. <ul style="list-style-type: none"> b'0: Relayer must not time slice this operation. b'1: Relayer can time slice this operation. MBZ if the Relayer does not support time slicing of memory management operations (see 12.2.3 Time slicing of memory management operations), else the Relayer must return <i>INVALID_PARAMETERS</i>.

Field	Description
Bit[2]	<ul style="list-style-type: none"> Zero memory after relinquish flag. <ul style="list-style-type: none"> In an invocation of FFA_MEM_RETRIEVE_REQ, this flag specifies whether the Relayer must zero the memory region contents after unmapping it from the translation regime of the Borrower or if the Borrower crashes. <ul style="list-style-type: none"> b'0: Relayer must not zero the memory region contents. b'1: Relayer must zero the memory region contents. If the memory region is lent to multiple Borrowers, the Relayer must clear memory region contents after unmapping it from the translation regime of each Borrower, if any Borrower including the caller sets this flag. This flag could be overridden by the Receiver in an invocation of FFA_MEM_RELINQUISH (see <i>Flags</i> field in Table 11.25). MBZ if the Receiver has Read-only access to the memory region, else the Relayer must return <i>DENIED</i>. The Receiver could be a PE endpoint or a dependent peripheral device. MBZ in a transaction to share a memory region, else the Relayer must return <i>INVALID_PARAMETER</i>.
Bit[4:3]	<ul style="list-style-type: none"> Memory management transaction type flag. <ul style="list-style-type: none"> In an invocation of FFA_MEM_RETRIEVE_REQ, this flag is used by the Receiver to either specify the memory management transaction it is participating in or indicate that it will discover this information in the invocation of FFA_MEM_RETRIEVE_RESP corresponding to this request. <ul style="list-style-type: none"> b'00: Relayer must specify the transaction type in FFA_MEM_RETRIEVE_RESP. b'01: Share memory transaction. b'10: Lend memory transaction. b'11: Donate memory transaction. Relayer must return <i>INVALID_PARAMETERS</i> if the transaction type specified by the Receiver is not the same as that specified by the Sender for the memory region identified by the <i>Handle</i> value specified in the transaction descriptor.
Bit[9:5]	<ul style="list-style-type: none"> Address range alignment hint. <ul style="list-style-type: none"> In an invocation of FFA_MEM_RETRIEVE_REQ, this flag is used by the Receiver to specify the boundary, expressed as multiples of 4KB, to which the address ranges allocated by the Relayer to map the memory region must be aligned. Bit[9]: Hint valid flag. <ul style="list-style-type: none"> b'0: Relayer must choose the alignment boundary. Bits[8:5] are reserved and MBZ. b'1: Relayer must use the alignment boundary specified in Bits[8:5]. Bit[8:5]: Alignment hint. <ul style="list-style-type: none"> If the value in this field is n, then the address ranges must be aligned to the $2^n \times 4KB$ boundary. MBZ if the Receiver specifies the IPA or VA address ranges that must be used by the Relayer to map the memory region, else the Relayer must return <i>INVALID_PARAMETERS</i>. Relayer must return <i>DENIED</i> if it is not possible to allocate the address ranges at the alignment boundary specified by the Receiver. Relayer must return <i>INVALID_PARAMETERS</i> if a reserved value is specified by the Receiver.

Field	Description
Bit[31:10]	<ul style="list-style-type: none"> Reserved (MBZ).

Table 5.22: Flags usage in FFA_MEM_RETRIEVE_RESP ABI

Field	Description
Bit[0]	<ul style="list-style-type: none"> Zero memory before retrieval flag. <ul style="list-style-type: none"> In an invocation of FFA_MEM_RETRIEVE_RESP during a transaction to lend or donate memory, this flag is used by the Relayer to specify whether the memory region was retrieved with or without zeroing its contents first. <ul style="list-style-type: none"> b'0: Memory region was retrieved without zeroing its contents. b'1: Memory region was retrieved after zeroing its contents. MBZ in a transaction to share a memory region. MBZ if the Sender has Read-only access to the memory region.
Bit[2:1]	<ul style="list-style-type: none"> Reserved (MBZ).
Bit[4:3]	<ul style="list-style-type: none"> Memory management transaction type flag. <ul style="list-style-type: none"> In an invocation of FFA_MEM_RETRIEVE_RESP, this flag is used by the Relayer to specify the memory management transaction the Receiver is participating in. <ul style="list-style-type: none"> b'00: Reserved. b'01: Share memory transaction. b'10: Lend memory transaction. b'11: Donate memory transaction.
Bit[31:5]	<ul style="list-style-type: none"> Reserved (MBZ).

Chapter 6

Interface overview

The interfaces used by FF-A components for communication at an FF-A instance are described in the following sections.

- Interfaces for reporting status of execution of other interfaces are described in [Chapter 7 Status reporting interfaces](#).
- Interfaces for partition setup and discovery using Framework messages are described in [Chapter 8 Setup and discovery interfaces](#).
- Interfaces to implement memory management transactions using Framework messages are described in [Chapter 11 Memory management interfaces](#).
- Interfaces to manage CPU cycles allocated to an endpoint are described in [Chapter 9 CPU cycle management interfaces](#).
- Interfaces to implement exchange of direct and indirect Partition messages between endpoints are described in [Chapter 10 Messaging interfaces](#).

Each interface can be invoked using one more conduits described in [2.4 Conduits](#). Each interface is based on the AArch64 and AArch32 SMC calling convention described in [\[4\]](#). Usage of only those architectural registers that are relevant to an interface is specified. The values of all other architectural registers must be ignored.

The following standard Secure service call identifier ranges have been reserved for FF-A interfaces in the SMCCC [\[4\]](#).

1. **0x84000060-0x8400007F**: FF-A 32-bit calls.
 2. **0xC4000060-0xC400007F**: FF-A 64-bit calls.
- If the caller is in the AArch32 Execution state, it must use the function identifiers for 32-bit calls.
 - If the callee is in the AArch64 Execution state, it could use a function identifier for 32-bit or 64-bit calls.

6.1 Compliance requirements

This section lists the ABIs and other functionality that a minimally compliant implementation of this version of the Framework must implement.

- The following FF-A components are in the scope of the implementation.
 - NS-Endpoints.
 - Hypervisor.
 - SPM.
 - S-Endpoints.
- Each component could implement an ABI as its *caller*, *callee*, or both.

6.1.1 Common compliance requirements

All implementations of the Framework must comply with the following rules,

- The SPM must implement support for partition setup through its manifest as specified in [Chapter 3 Partition setup](#).
- The Hypervisor (if present) must implement support for partition setup through its manifest as specified in [Chapter 3 Partition setup](#) or an IMPLEMENTATION DEFINED mechanism.
- It must be possible for the FF-A components at an FF-A instance to discover the presence and version number of their Framework implementations through the *FFA_VERSION* interface (see [8.1 FFA_VERSION](#)).
- All FF-A components must implement support for *Status reporting interfaces* described in [Chapter 7 Status reporting interfaces](#). These interfaces are as follows.
 - FFA_SUCCESS.
 - FFA_ERROR.
 - FFA_INTERRUPT.
- All FF-A components must implement support for *Setup and discovery interfaces* described in [Chapter 8 Setup and discovery interfaces](#). These interfaces are as follows.
 - FFA_VERSION.
 - FFA_FEATURES.
 - FFA_RX_RELEASE.
 - FFA_RXTX_MAP.
 - FFA_RXTX_UNMAP.
 - FFA_PARTITION_INFO_GET.
 - FFA_ID_GET.

Each interface must be implemented at an FF-A instance where it can be invoked through a valid conduit as specified in the interface description.

An invocation of any of these interfaces must be completed by invoking the *FFA_ERROR* interface with the *NOT_SUPPORTED* error code in the following scenarios.

- The interface was invoked at an FF-A instance where it cannot be invoked through any conduit.
- The interface was invoked through an invalid conduit at an FF-A instance where it can be invoked.
- It must be possible for an FF-A component, at the lower EL at an FF-A instance to use the *FFA_FEATURES* interface (see [8.2 FFA_FEATURES](#)) to discover if an FF-A interface is implemented by the FF-A component at the higher EL.

6.1.2 Compliance requirements for deploying an SP

If one or more SPs are deployed in the system, then the implementation of the Framework must comply with the following rules,

1. An NS-Endpoint must implement the following ABIs.
 - FFA_RUN.
 - As caller only if endpoint implements the primary scheduler.
 - FFA_MSG_SEND_DIRECT_REQ.
 - As caller.
 - FFA_MSG_SEND_DIRECT_RESP.
 - As callee.
2. A Hypervisor (if present) must implement the following ABIs.
 - FFA_RUN.
 - As caller.
 - As callee only if a VM implements the primary scheduler.
 - FFA_MSG_SEND_DIRECT_REQ.
 - As both caller and callee.
 - FFA_MSG_SEND_DIRECT_RESP.
 - As both caller and callee.
3. A S-Endpoint must implement the following ABIs.
 - FFA_MSG_SEND_DIRECT_REQ.
 - As callee.
 - FFA_MSG_SEND_DIRECT_RESP.
 - As caller.
4. A SPM must implement the following ABIs.
 - FFA_NORMAL_WORLD_RESUME.
 - As caller in SPMC.
 - As callee in SPMD.
 - FFA_RUN.
 - As both caller and callee in SPMD.
 - As callee in SPMC.
 - FFA_MSG_SEND_DIRECT_REQ.
 - As both caller and callee.
 - FFA_MSG_SEND_DIRECT_RESP.
 - As both caller and callee.

6.1.3 Compliance requirements for deploying a VM

If one or more VMs are deployed in the system, then the implementation of the Framework must comply with the following rules,

1. A VM must implement one or both messaging methods (see [Chapter 4 Message passing](#)).
2. If a VM implements *Indirect messaging*, then the implementation of the Framework must comply with the following rules,
 1. A VM must implement the following ABIs.
 - FFA_MSG_SEND.
 - As caller.
 - As callee only if VM implements the primary scheduler.
 - FFA_MSG_POLL.
 - As caller.

- FFA_MSG_WAIT.
 - As caller.
 - As callee only if VM implements the primary scheduler.
- FFA_YIELD.
 - As caller only if VM does not implement the primary scheduler.
 - As callee only if VM implements the primary scheduler.
- FFA_RUN.
 - As caller only if VM implements the primary scheduler.

2. Hypervisor must implement the following ABIs.

- FFA_RUN.
 - As callee only if a VM implements the primary scheduler.
- FFA_MSG_SEND.
 - As callee.
 - As caller only if VM implements the primary scheduler.
- FFA_MSG_POLL.
 - As callee.
- FFA_MSG_WAIT.
 - As callee.
 - As caller only if VM implements the primary scheduler.
- FFA_YIELD.
 - As callee.
 - As caller only if VM implements the primary scheduler.

3. If a VM implements *Direct messaging*, then it and the Hypervisor must comply with the same rules listed for these components in [6.1.2 Compliance requirements for deploying an SP](#).

6.1.4 Compliance requirements for memory management

If memory management ABIs are implemented, then all components in the Framework implementation must comply with the following rules.

1. Memory management must be supported between PE endpoints as per the deployment choice in [6.1.3 Compliance requirements for deploying a VM](#) or [6.1.2 Compliance requirements for deploying an SP](#).
2. At least one of the following ABIs must be implemented:
 - FFA_MEM_DONATE.
 - FFA_MEM_LEND.
 - FFA_MEM_SHARE.
3. The following ABIs must be implemented if any of the preceding ABIs are implemented:
 - FFA_MEM_RETRIEVE_REQ.
 - FFA_MEM_RETRIEVE_RESP.
 - FFA_MEM_RECLAIM.
 - FFA_MEM_RELINQUISH.

Chapter 7

Status reporting interfaces

7.1 Overview

Interfaces described in this section are used to report the status of a previous FF-A ABI invocation. The status indicates successful or unsuccessful completion or preemption of the ABI invocation. This ABI must be one that is listed in the following sections.

- Interfaces for partition setup and discovery¹ in [Chapter 8 Setup and discovery interfaces](#).
- Interfaces to implement memory management transactions in [Chapter 11 Memory management interfaces](#).
- Interfaces to manage CPU cycles in [Chapter 9 CPU cycle management interfaces](#).
- Interfaces to implement messaging between endpoints in [Chapter 10 Messaging interfaces](#).

¹The *FFA_VERSION* interface (see [8.1 FFA_VERSION](#)) is used for discovering the presence of a Framework implementation. It does not use the status reporting interfaces.

7.2 FFA_ERROR

Description

- Returns error code in response to a previous invocation of an FF-A function.
- [Table 7.2](#) defines the values for status codes used with FF-A functions. All values are considered to be 32-bit signed integers.
- Valid FF-A instances and conduits are listed in [Table 7.3](#).
- Syntax of this function is described in [Table 7.4](#).
- [Figure 7.1](#) illustrates example usage of this function with the following assumptions.
 - Component A makes an invalid request to Component B through an FF-A function described in this specification.
 - Component B uses the FFA_ERROR function to return the error code to Component A.
 - The FF-A function used by component A can be invoked through the SMC and ERET conduits.
 - Both components could be interacting at any FF-A instance support by the FF-A function. The two possible scenarios have been considered.
 - * Component A is at a lower EL than component B at the FF-A instance.
 - * Component A is at a higher EL than component B at the FF-A instance.

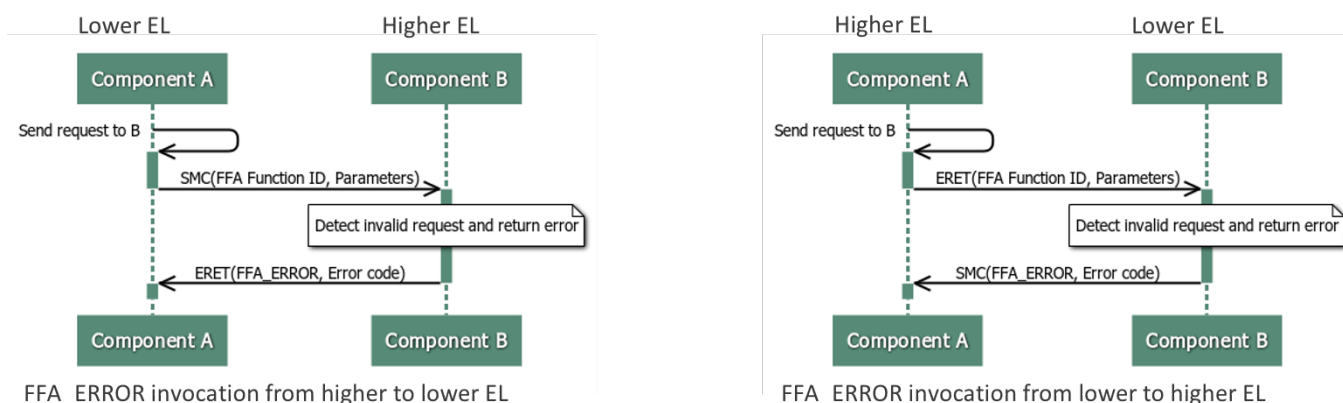


Figure 7.1: Example usage of FFA_ERROR

Table 7.2: Error status codes

Status code	Description
-1	NOT_SUPPORTED
-2	INVALID_PARAMETERS
-3	NO_MEMORY
-4	BUSY
-5	INTERRUPTED
-6	DENIED
-7	RETRY

Status code	Description
-8	ABORTED

Table 7.3: FFA_ERROR instances and conduits

Config No.	FF-A instance	Valid conduits
1	Secure and Non-secure physical	SMC, ERET
2	Non-secure virtual	SMC, HVC, ERET
3	Secure virtual	SMC, ERET

Table 7.4: FFA_ERROR function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none"> 0x84000060.
uint32 Target information	w1	<ul style="list-style-type: none"> Information to identify target SP/VM. <ul style="list-style-type: none"> Valid only when SMC conduit is used at the Non-secure virtual FF-A instance. MBZ otherwise. Bits[31:16]: ID of SP/VM. Bits[15:0]: ID of vCPU of SP/VM to deliver error to.
int32 Error code	w2	<ul style="list-style-type: none"> FF-A function specific error code. See function definition for applicable error codes .
Other Parameter registers	w3-w7 x3-x7	<ul style="list-style-type: none"> Reserved (MBZ).

7.3 FFA_SUCCESS

Description

- Returns results on successful completion of a previous invocation of an FF-A function.
- Valid FF-A instances and conduits are listed in [Table 7.6](#).
- Syntax of this function is described in [Table 7.7](#).
- [Figure 7.2](#) illustrates example usage of this function with the following assumptions.
 - Component A makes an valid request to Component B through an FF-A function described in this specification.
 - Component B uses the FFA_SUCCESS function to return the results to Component A.
 - The FF-A function used by component A can be invoked through the SMC and ERET conduits.
 - Both components could be interacting at any FF-A instance support by the FF-A function. The two possible scenarios have been considered.
 - * Component A is at a lower EL than component B at the FF-A instance.
 - * Component A is at a higher EL than component B at the FF-A instance.

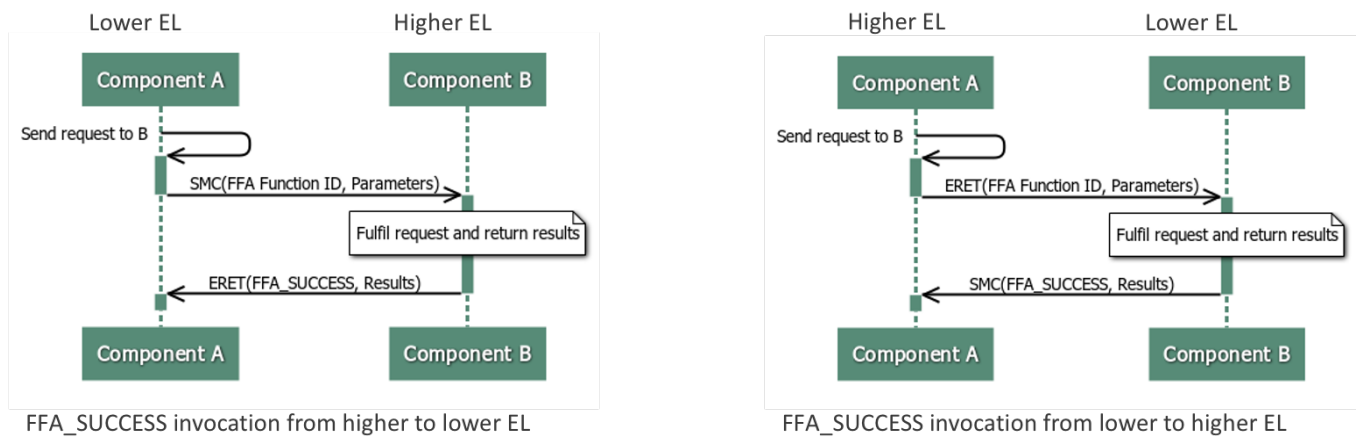


Figure 7.2: Example usage of FFA_SUCCESS

Table 7.6: FFA_SUCCESS instances and conduits

Config No.	FF-A instance	Valid conduits
1	Secure and Non-secure physical	SMC, ERET
2	Non-secure virtual FF-A	SMC, HVC, ERET
3	Secure virtual FF-A	SMC, ERET

Table 7.7: FFA_SUCCESS function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none"> • 0x84000061. • 0xC4000061.
uint32 Target information	w1	<ul style="list-style-type: none"> • Information to identify target SP/VM. <ul style="list-style-type: none"> – Valid only when SMC conduit is used at the Non-secure virtual FF-A instance. MBZ otherwise. – Bits[31:16]: ID of SP/VM. – Bits[15:0]: ID of vCPU of SP/VM to deliver results to.
uint32/uint64 Result registers	w2-w7 x2-x7	<ul style="list-style-type: none"> • FF-A function specific return results. See function definition for result encoding. MBZ if not explicitly specified.

7.4 FFA_INTERRUPT

Description

- Returns control from the caller to the callee in response to an interrupt that must be:
 - Either handled by the callee.
 - Or handled by another FF-A component reachable only through the callee.
- Valid FF-A instances and conduits are listed in [Table 7.9](#).
- Syntax of this function is described in [Table 7.10](#).
- Example usage of this interface is illustrated in [Figure 4.10](#).

Table 7.9: FFA_INTERRUPT instances and conduits

Config No.	FF-A instance	Valid conduits
1	Non-secure physical	ERET
2	Secure and Non-secure virtual	ERET
3	Secure physical	SMC

Table 7.10: FFA_INTERRUPT function syntax

Parameter	Register	Value
uint32 Function ID	w0	• 0x84000062.
uint32 Endpoint/vCPU IDs	w1	• Endpoint and vCPU IDs of the caller. Only valid at a physical FF-A instance. Else MBZ <ul style="list-style-type: none"> – Bits[31:16]: Endpoint ID. – Bits[15:0]: vCPU ID.
uint32 Interrupt ID	w2	• Interrupt ID. Only valid at Secure virtual FF-A instance with a S-EL0 SP as callee.
Other parameter registers	w3-w7 x3-x7	• Reserved (MBZ).

Chapter 8

Setup and discovery interfaces

8.1 FFA_VERSION

Description

- Returns version of the Firmware Framework implementation at an FF-A instance as described in [8.1.1 Overview](#).
- Valid FF-A instances and conduits are listed in [Table 8.2](#).
- Syntax of this function is described in [Table 8.3](#).
- Encoding of a version number in return parameters is described in [Table 8.4](#).
- Encoding of error codes in return parameters is described in [Table 8.5](#).

Table 8.2: FFA_VERSION instances and conduits

Config No.	FF-A instance	Valid conduits
1	Secure and Non-secure physical	SMC
2	Secure and Non-secure virtual	SMC, HVC, SVC

Table 8.3: FFA_VERSION function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none"> • 0x84000063.
uint32 Input version number	w1	<ul style="list-style-type: none"> • Version number specified by the caller as follows. <ul style="list-style-type: none"> – Bit[31]: Must be 0. – Bit[30:16] Major Version number. – Bit[15:0] Minor Version number.
Other Parameter registers	w2-w7 x2-x7	<ul style="list-style-type: none"> • Reserved (MBZ).

Table 8.4: Encoding of a version number

Parameter	Register	Value
int32 Output version number	w0	<ul style="list-style-type: none"> • On a successful return, the format of the value is as follows. <ul style="list-style-type: none"> – Bit[31]: Must be 0. – Bit[30:16] Major Version: Must be 1 for this revision of FF-A. – Bit[15:0] Minor Version: Must be 0 for this revision of FF-A.

Parameter	Register	Value
Other Result registers	w1-w7 x1-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 8.5: Encoding of error codes

Parameter	Register	Value
int32 Error code	w0	<ul style="list-style-type: none"> NOT_SUPPORTED: A Firmware Framework implementation does not exist at this FF-A instance.

8.1.1 Overview

The version number of a Firmware Framework implementation is a 31-bit unsigned integer, with the upper 15 bits denoting the major revision, and the lower 16 bits denoting the minor revision.

If this function returns a valid version number:

- All the functions that are described in this specification must be implemented, unless it is explicitly stated that a function is optional.
- A partition manager could implement an optional interface and make it available to a subset of endpoints it manages.

The following rules apply to the version numbering.

- Different major revision values indicate possibly incompatible functions.
- For two revisions, A and B, for which the major revision values are identical, if the minor revision value of revision B is greater than the minor revision value of revision A, then every function in revision A must work in a compatible way with revision B. However, it is possible for revision B to have a higher function count than revision A.

In an invocation of this function, the compatibility of the version number ($x.y$) of the caller with the version number ($a.b$) of the callee can also be as follows.

1. If $x \neq a$, then the versions are incompatible.
 - The caller cannot inter-operate with the callee.
2. If $x == a$ and $y > b$, then the versions are incompatible.
 - The caller can inter-operate with the callee only if it downgrades its minor revision such that $y \leq b$.
3. If $x == a$ and $y \leq b$, then the versions are compatible.

A version number ($x.y$) is less than a version number ($a.b$) if one of the following conditions is true.

- $x < a$.
- $y < b$ if $x == a$.

8.1.2 Usage

This function enables the caller to determine if the callee implements the Firmware Framework and the version number of the implementation. The function must be invoked as follows.

- The caller must specify a version number in the *Input version number* parameter.
- The callee must take one of the following actions.
 - If it supports a Firmware Framework implementation that is compatible with the version number specified by the caller, it must return the version number of the implementation.

- If it only supports a Firmware Framework implementation that is incompatible with and at a greater version number than specified by the caller, it must either return the version number of this implementation or the NOT_SUPPORTED error code.
- If it supports a Firmware Framework implementation that is incompatible with and at a lesser version number than specified by the caller, it must return the highest version number of this implementation.
- If it does not support any version of the Firmware Framework, it must return the NOT_SUPPORTED error code.
- The caller must use the preceding compatibility rules to determine if it can inter-operate with the version number returned by the callee.

Each FF-A instance must support this call and return its version number. For this revision of FF-A, the major version is 1 and the minor version is 0.

This interface returns a version number of the Framework at the FF-A instance where it is invoked. It is possible that version numbers of the Framework at different FF-A instances differ. These versions must be supported in accordance to the preceding major and minor version number compatibility rules.

8.1.3 SPM usage

In SPM configurations where the SPMD and SPMC reside in separate Exception levels (see [Table 2.1](#) & [Table 2.2](#)), the versions of these two components could differ. The following constraints must be met to avoid a version mismatch.

- The SPMC must specify the version that it implements to the SPMD through an IMPLEMENTATION DEFINED mechanism.
- The SPMD must compare the version specified by the SPMC with the version it implements.
 - If the versions are not compatible as per the preceding compatibility rules, the SPMD must not initialize the SPMC.
 - If the versions are compatible, the SPMD must report the lowest compatible version in response to an invocation of FFA_VERSION at either physical FF-A instance.

8.2 FFA_FEATURES

Description

- This interface is used by an FF-A component at the lower EL at an FF-A instance to query:
 - If an FF-A interface is implemented by the FF-A component at the higher EL.
 - If an implemented FF-A interface also implements any optional features described in its interface definition.
 - Any implementation details exported by an implemented FF-A interface as described in its interface definition.
 - This interface can be invoked at the FF-A instances through the conduits listed in [Table 8.7](#).
 - Syntax of this function is described in [Table 8.8](#).
 - If the FF-A interface that was queried is implemented, the callee completes this call with an invocation of the *FFA_SUCCESS* interface as described in [Table 8.9](#).
 - If the FF-A interface that was queried is not implemented or invalid, the callee completes this call with an invocation of the *FFA_ERROR* interface with the *NOT_SUPPORTED* error code.
-

Table 8.7: FFA_FEATURES instances and conduits

Config No.	FF-A instance	Valid conduits
1	Non-secure physical	SMC
2	Secure physical	ERET, SMC
3	Secure and Non-secure virtual	SMC, HVC, SVC

Table 8.8: FFA_FEATURES function syntax

Parameter	Register	Value
uint32 Function ID	w0	• 0x84000064.
uint32 FF-A function ID	w1	• Function ID of the FF-A interface whose implementation must be queried. • If an interface defines both SMC32 and SMC64 FIDs, then either FID could be used.
Other Parameter registers	w2-w7 x2-x7	• Reserved (MBZ).

Table 8.9: FFA_SUCCESS encoding

Parameter	Register	Value
uint32 Interface properties	w2-w3	<ul style="list-style-type: none"> Used to encode any optional features implemented or any implementation details exported by the queried interface. <ul style="list-style-type: none"> FF-A interfaces that use these parameters and the encodings of their properties are listed in Table 8.10. MBZ if no optional features are implemented or no implementation details are exported by the queried interface.
Other Result registers	w4-w7 x4-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 8.10: Encoding of interface properties parameters

FF-A Function ID	Return parameters
FFA_RXTX_MAP	<ul style="list-style-type: none"> w2 : Bits[31:2] are reserved (MBZ) . <ul style="list-style-type: none"> Bit[1:0]: Minimum buffer size and alignment boundary (see 4.2.2.3 Buffer attributes). <ul style="list-style-type: none"> b'00: 4K. b'01: 64K. b'10: 16K. b'11: Reserved. w3/x3 : Reserved (MBZ).
FFA_MEM_DONATE	<ul style="list-style-type: none"> w2 : Bits[31:1] are reserved (MBZ) . <ul style="list-style-type: none"> Bit[0]: Dynamically allocated buffer support. See 12.2.1 Transmission of transaction descriptor in dynamically allocated buffers. <ul style="list-style-type: none"> b'0: Partition manager does not support transmission of a memory transaction descriptor in a buffer dynamically allocated by the endpoint. b'1: Partition manager supports transmission of a memory transaction descriptor in a buffer dynamically allocated by the endpoint. w3/x3 : Reserved (MBZ).
FFA_MEM_LEND	<ul style="list-style-type: none"> Same as FFA_MEM_DONATE.
FFA_MEM_SHARE	<ul style="list-style-type: none"> Same as FFA_MEM_DONATE.

FF-A Function ID	Return parameters
FFA_MEM_RETRIEVE_REQ	<ul style="list-style-type: none"> • Same as FFA_MEM_DONATE. Also adds: • $w3$: Outstanding retrievals field. <ul style="list-style-type: none"> – Bit[31:8]: Reserved MBZ. – Bit[7:0]: Number of times a Receiver is allowed to retrieve a memory region before relinquishing it. The value specified is interpreted as $((1U << (value + 1)) - 1$.

8.3 FFA_RX_RELEASE

Description

- Relinquish ownership of a RX buffer after reading a message from it (see [4.2.2.4 Buffer synchronization](#)).
 - Valid FF-A instances and conduits are listed in [Table 8.12](#).
 - Syntax of this function is described in [Table 8.13](#).
 - Returns FFA_SUCCESS without any further parameters on successful completion.
 - Encoding of error code in the FFA_ERROR function is described in [Table 8.14](#).
-

Table 8.12: FFA_RX_RELEASE instances and conduits

Config No.	FF-A instance	Valid conduits
1	Non-secure Physical	SMC
2	Secure Physical	ERET
3	Secure virtual	SMC, HVC, SVC
4	Non-secure virtual	SMC, HVC, SVC, ERET

Table 8.13: FFA_RX_RELEASE function syntax

Parameter	Register	Value
uint32 Function ID	w0	• 0x84000065.
Other Parameter registers	w1-w7 x1-x7	• Reserved (MBZ).

Table 8.14: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none">• DENIED: Caller did not have ownership of the RX buffer.• NOT_SUPPORTED: This function is not implemented at this FF-A instance.

8.4 FFA_RXTX_MAP

Description

- Maps the RX/TX buffer pair in the translation regime of the callee on behalf of an endpoint or Hypervisor.
 - A SP describes the VA or IPA contiguous pages allocated for each buffer in the pair to the SPM.
 - A VM describes the VA or IPA contiguous pages allocated for each buffer in the pair to the Hypervisor.
 - Hypervisor or OS Kernel describe the physically contiguous pages allocated for each buffer in the pair to the SPM.
 - Hypervisor forwards the description of pages allocated for each buffer in the pair by a VM to the SPM.
 - * Description of buffer pair is populated in the TX buffer of the Hypervisor as described in [Table 8.19](#).
 - Both Hypervisor and SPM must ensure the caller has exclusive access and ownership of the RX/TX buffer memory regions.
- Valid FF-A instances and conduits are listed in [Table 8.16](#).
- Syntax of this function is described in [Table 8.17](#).
- Returns FFA_SUCCESS without any further parameters on successful completion.
- Encoding of error code in the FFA_ERROR function is described in [Table 8.18](#).

Table 8.16: FFA_RXTX_MAP instances and conduits

Config No.	FF-A instance	Valid conduits
1	Non-secure physical	SMC
2	Secure physical	ERET
3	Virtual	SMC, HVC, SVC

Table 8.17: FFA_RXTX_MAP function syntax

Parameter	Register	Value
uint32 Function ID	w0/x0	<ul style="list-style-type: none"> • 0x84000066. • 0xC4000066.
uint32/uint64 TX address	w1/x1	<ul style="list-style-type: none"> • Base address of the TX buffer if invoked by an endpoint or Hypervisor to register its buffer pair. <ul style="list-style-type: none"> – Address is a IPA or VA at the virtual FF-A instance. – Address is a PA at the physical FF-A instance. • MBZ if Hypervisor is forwarding this call on behalf of an endpoint. <ul style="list-style-type: none"> – Description of RX/TX buffer and identity of endpoint is specified in the TX buffer of the Hypervisor.

Parameter	Register	Value
uint32/uint64 RX address	w2/x2	<ul style="list-style-type: none"> Base address of the RX buffer. <ul style="list-style-type: none"> Address is a IPA or VA at the virtual FF-A instance. Address is a PA at the physical FF-A instance. MBZ if Hypervisor is forwarding this call on behalf of an endpoint. <ul style="list-style-type: none"> Description of RX/TX buffer and identity of endpoint is specified in the TX buffer of the Hypervisor.
uint32 RX/TX page count	w3/x3	<ul style="list-style-type: none"> Bit[31:6]: Reserved (MBZ). Bit[5:0]: Number of contiguous 4K pages allocated for each buffer.
Other Parameter registers	w4-w7 x4-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 8.18: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS: One or more fields in input parameters is incorrectly encoded. NO_MEMORY: <ul style="list-style-type: none"> Not enough memory to map the buffers in the translation regime of the callee. Not enough memory in TX buffer of Hypervisor to describe caller buffer pair to SPM. DENIED: Buffer pair already registered for the FF-A component with specified ID. NOT_SUPPORTED: This function is not implemented at this FF-A instance.

Table 8.19: Endpoint RX/TX descriptor

Field	Byte length	Byte offset	Description
Endpoint ID	2	0	<ul style="list-style-type: none"> ID of endpoint that allocated the RX/TX buffer.
Reserved	2	2	<ul style="list-style-type: none"> MBZ.
RX address range count	4	4	<ul style="list-style-type: none"> Count of address ranges specified using constituent memory descriptors for the RX buffer.

Field	Byte length	Byte offset	Description
TX address range count	4	8	<ul style="list-style-type: none"> Count of address ranges specified using constituent memory descriptors for the TX buffer.
RX address range array	–	12	<ul style="list-style-type: none"> Array of address ranges allocated for the RX buffer that the callee must map in its translation regime. See Table 5.14 for how the address ranges are encoded.
TX address range array	–	–	<ul style="list-style-type: none"> Array of address ranges allocated for the TX buffer that the callee must map in its translation regime. See Table 5.14 for how the address ranges are encoded.

8.5 FFA_RXTX_UNMAP

Description

- Unmaps the RX/TX buffer pair of an endpoint or Hypervisor from the translation regime of the callee.
 - A SP invokes this interface to unmap its buffer pair from the translation regime of the SPM.
 - A VM invokes this interface to unmap its buffer pair from the translation regime of the Hypervisor.
 - Hypervisor or OS Kernel invoke this interface to unmap their buffer pair from the translation regime of the SPM.
 - Hypervisor forwards an invocation of this interface by a VM to the SPM.
 - * Identity of VM is specified in *w1*.
- Valid FF-A instances and conduits are listed in [Table 8.21](#).
- Syntax of this function is described in [Table 8.22](#).
- Returns FFA_SUCCESS without any further parameters on successful completion.
- Encoding of error code in the FFA_ERROR function is described in [Table 8.23](#).

Table 8.21: FFA_RXTX_UNMAP instances and conduits

Config No.	FF-A instance	Valid conduits
1	Non-secure physical	SMC
2	Secure physical	ERET
3	Virtual	SMC, HVC, SVC

Table 8.22: FFA_RXTX_UNMAP function syntax

Parameter	Register	Value
uint32 Function ID	w0/x0	• 0x84000067.
uint32 ID	w1	• ID of FF-A component that allocated the RX/TX buffer. <ul style="list-style-type: none"> – Bit[31:16]: ID. – Bit[15:0]: Reserved MBZ.
Other Parameter registers	w2-w7 x2-x7	• Reserved (MBZ).

Table 8.23: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none">INVALID_PARAMETERS: There is no buffer pair registered on behalf of the caller.NOT_SUPPORTED: This function is not implemented at this FF-A instance.

8.6 FFA_PARTITION_INFO_GET

Description

- Request Hypervisor and SPM to return information about partitions instantiated in the system as follows:
 - A NS-Endpoint can request information for all partitions in the system including the caller by specifying the Nil UUID.
 - * If the Nil UUID is specified at the Non-secure virtual FF-A instance, the Hypervisor must provide information for partitions resident in both Security states.
 - A NS-Endpoint can request information for a subset of partitions in the system by specifying the non-Nil UUID.
 - A S-Endpoint can request information for all SPs in the system including the caller by specifying the Nil UUID.
 - * If the Nil UUID is specified at the Secure virtual FF-A instance, the SPM must only provide information for all SPs.
 - A S-Endpoint can request information for a subset of SPs in the system by specifying the non-Nil UUID.
 - Information returned for each partition is described in [Table 8.25](#).
 - Partition information is returned in the RX buffer of the caller as an array of partition information descriptors.
 - Count of partition information descriptors is returned in *w2*.
- Valid FF-A instances and conduits are listed in [Table 8.26](#).
- Syntax of this function is described in [Table 8.27](#).
- Encoding of result parameters in the FFA_SUCCESS function is described in [Table 8.28](#).
- Encoding of error code in the FFA_ERROR function is described in [Table 8.29](#).

Table 8.25: Partition information descriptor

Field	Byte length	Byte offset	Description
Partition ID	2	0	<ul style="list-style-type: none"> • 16-bit ID of the partition.
Execution context count	2	2	<ul style="list-style-type: none"> • Number of execution contexts implemented by this partition (also see 2.9 Execution context).

Field	Byte length	Byte offset	Description
Partition properties	4	4	<ul style="list-style-type: none"> Flags to determine partition properties. <ul style="list-style-type: none"> Bit[0] has the following encoding: <ul style="list-style-type: none"> b'0: Does not support receipt of direct requests b'1: Supports receipt of direct requests. Count of execution contexts must be either 1 or equal to the number of PEs in the system (also see 4.4 Direct messaging usage). bit[1] has the following encoding: <ul style="list-style-type: none"> b'0: Cannot send direct requests. b'1: Can send direct requests. bit[2] has the following encoding: <ul style="list-style-type: none"> b'0: Cannot send and receive indirect messages. MBZ for an SP. b'1: Can send and receive indirect messages. bit[31:3]: Reserved (MBZ).

Table 8.26: FFA_PARTITION_INFO_GET instances and conduits

Config No.	FF-A instance	Valid conduits
1	Non-secure physical	SMC
2	Secure physical	ERET
3	Non-secure virtual	SMC, HVC
4	Secure virtual	SMC, HVC, SVC

Table 8.27: FFA_PARTITION_INFO_GET function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none"> 0x84000068.
uint128 UUID	w1-w4	<ul style="list-style-type: none"> Specified as described in Section 5.3 of [4].
Other Parameter registers	w5-w7 x5-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 8.28: FFA_SUCCESS encoding

Parameter	Register	Value
uint32 Count	w2	<ul style="list-style-type: none"> Count of partition information descriptors populated in RX buffer of caller.

Parameter	Register	Value
Other Result registers	w3-w7 x3-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 8.29: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> BUSY: RX buffer of the caller is not free. INVALID_PARAMETERS: Unrecognized UUID. NO_MEMORY: Results cannot fit in RX buffer of the caller. DENIED: Callee is not in a state to handle this request. NOT_SUPPORTED: This function is not implemented at this FF-A instance.

8.7 FFA_ID_GET

Description

- Returns 16-bit ID of calling FF-A component.
 - ID value 0 must be returned at the Non-secure physical FF-A instance (see [2.8 Partition identification and discovery](#)).
- Valid FF-A instances and conduits are listed in [Table 8.31](#).
- Syntax of this function is described in [Table 8.32](#).
- Encoding of result parameters in the FFA_SUCCESS function is described in [Table 8.33](#).
- Encoding of error code in the FFA_ERROR function is described in [Table 8.34](#).

Table 8.31: FFA_ID_GET instances and conduits

Config No.	FF-A instance	Valid conduits
1	Physical FF-A instance	SMC
2	Virtual FF-A instance	SMC, HVC, SVC

Table 8.32: FFA_ID_GET function syntax

Parameter	Register	Value
uint32 Function ID	w0	• 0x84000069.
Other Parameter registers	w1-w7 x1-x7	• Reserved (MBZ).

Table 8.33: FFA_SUCCESS encoding

Parameter	Register	Value
uint32 ID	w2	<ul style="list-style-type: none"> • ID of the caller. <ul style="list-style-type: none"> – Bit[31:16]: Reserved (MBZ). – Bit[15:0]: ID.
Other Result registers	w3-w7 x3-x7	• Reserved (MBZ).

Table 8.34: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none">• NOT_SUPPORTED: This function is not implemented at this FF-A instance.

Chapter 9

CPU cycle management interfaces

9.1 FFA_MSG_WAIT

Description

- Blocks the caller until one of the following conditions is true.
 - A message is available in the RX buffer of the caller.
 - A message is available in the parameter registers of the caller.
 - The call is interrupted by an interrupt targeted to the caller.
 - Execution is returned to the primary scheduler if none of the preceding conditions is true.
 - An optional 64-bit timeout could be provided to the primary scheduler in EL1 by the Hypervisor in EL2.
 - The scheduler must run the caller endpoint after the timeout expires.
 - The execution context of the caller enters the *idle* state (also see [2.12 Run-time states](#)).
 - Valid FF-A instances and conduits are listed in [Table 9.2](#).
 - ERET conduit is used only to return execution to the primary scheduler.
 - Syntax of this function is described in [Table 9.3](#).
 - Successful completion of this function is indicated through the invocation of the following functions by the callee. Each function ID encodes enough information for the caller to retrieve the message.
 - *FFA_INTERRUPT* at any FF-A instance.
 - Any Framework message function at the Secure physical FF-A instance.
 - Message transmission functions at the Secure physical FF-A instance.
 - *FFA_MSG_SEND* at any virtual FF-A instance if the partition supports receipt of indirect messages.
 - *FFA_MSG_SEND_DIRECT_REQ* at any virtual FF-A instance if the partition supports receipt of direct messages.
 - Encoding of error code in the *FFA_ERROR* function is described in [Table 9.4](#).
-

Table 9.2: FFA_MSG_WAIT instances and conduits

Config No.	FF-A instance	Valid conduits
1	Non-secure physical	ERET
2	Secure physical	SMC
3	Non-secure virtual	SMC, HVC, ERET
4	Secure virtual	SMC, HVC, SVC

Table 9.3: FFA_MSG_WAIT function syntax

Parameter	Register	Value
uint32 Function ID	w0	• 0x8400006B.
uint32 Endpoint/vCPU IDs	w1	• Endpoint and vCPU IDs of the caller. Only valid with the ERET conduit at the Non-secure virtual FF-A instance else MBZ. <ul style="list-style-type: none">– Bit[31:16]: Endpoint ID.– Bit[15:0]: vCPU ID.

Parameter	Register	Value
uint32 TimeoutLo	w2	<ul style="list-style-type: none"> Bits[31:0] of an interval measured in nanoseconds after which vCPU of the endpoint specified in <i>w1</i> must be run. Only valid with the ERET conduit at the Non-secure virtual FF-A instance else MBZ. This parameter MBZ if the caller does not specify a timeout.
uint32 TimeoutHi	w3	<ul style="list-style-type: none"> Bits[63:32] of an interval measured in nanoseconds after which vCPU of the endpoint specified in <i>w1</i> must be run. Only valid with the ERET conduit at the Non-secure virtual FF-A instance else MBZ. This parameter MBZ if the caller does not specify a timeout.
Other Parameter registers	w4-w7 x4-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 9.4: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS: Unrecognized endpoint or vCPU ID specified at Non-secure physical or virtual FF-A instance. DENIED: Callee is not in a state to handle this request. NOT_SUPPORTED: This function is not implemented at this FF-A instance.

9.1.1 Component responsibilities for FFA_MSG_WAIT

This section describes the common responsibilities that the participating FF-A components must fulfill during an invocation of the *FFA_MSG_WAIT* interface. These components are:

1. SP or VM.
2. Relayers.
3. Primary scheduler.

This interface is used by an endpoint during message processing to block their execution and enter the *idle* state by passing control back to the primary scheduler. Execution of the endpoint is resumed when a new direct or indirect message becomes available.

The Relayers are responsible for returning control to the primary scheduler in response to an invocation of this interface. Their responsibilities in this regard are influenced by the location of the primary scheduler relative to theirs. These are described in [9.1.1.1 Relayer responsibilities](#).

9.1.1.1 Relayer responsibilities

9.1.1.1.1 Invocation from VM

1. The Hypervisor and primary scheduler are co-resident. It must use an IMPLEMENTATION DEFINED mechanism to hand control to the scheduler.
2. The Hypervisor and primary scheduler are not co-resident. It must forward the *FFA_MSG_WAIT* call to the primary scheduler through the ERET conduit on the PE where the call is made. The ID of the endpoint and its caller execution context must be passed in the *w1* parameter register.

9.1.1.1.2 Invocation from SP

1. The SPM must forward the *FFA_MSG_WAIT* call to the primary scheduler through the ERET conduit on the PE where the call is made. The ID of the endpoint and its caller execution context must be passed in the *w1* parameter register.
2. If the Hypervisor is present, then its responsibilities are the same as those described in [9.1.1.1.1 Invocation from VM](#).

9.2 FFA_YIELD

Description

- Relinquish execution back to the scheduler on current physical CPU from the calling VM.
 - Used by a VM to avoid a busy wait for a shared resource for example, an internal lock that is not currently available.
 - Allows other software components to make progress on the PE until the shared resource is not available for the caller.
 - Must not be invoked when the caller is processing a direct request.
 - An optional 64-bit timeout could be provided to the primary scheduler in EL1 by the Hypervisor in EL2.
 - The scheduler must run the caller endpoint after the timeout expires.
- Valid FF-A instances and conduits are listed in [Table 9.6](#).
 - ERET conduit is used only to return execution to the primary scheduler.
- Syntax of this function is described in [Table 9.7](#).
- Successful completion of this function is indicated through an invocation of the following functions by the callee:
 - *FFA_RUN* to indicate that the execution context of the VM has been scheduled by the primary scheduler.
 - *FFA_INTERRUPT* to indicate that the call was interrupted.
- Encoding of error code in the FFA_ERROR function is described in [Table 9.8](#).

Table 9.6: FFA_YIELD instances and conduits

Config No.	FF-A instance	Valid conduits
1	Non-secure virtual	SMC, HVC, ERET

Table 9.7: FFA_YIELD function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none"> • 0x8400006C.
uint32 Endpoint/vCPU IDs	w1	<ul style="list-style-type: none"> • Endpoint and vCPU IDs of the caller. Only valid with the ERET conduit at the Non-secure virtual FF-A instance else MBZ. <ul style="list-style-type: none"> – Bit[31:16]: Endpoint ID. – Bit[15:0]: vCPU ID.
uint32 TimeoutLo	w2	<ul style="list-style-type: none"> • Bits[31:0] of an interval measured in nanoseconds after which vCPU of the endpoint specified in <i>w1</i> must be run. Only valid with the ERET conduit at the Non-secure virtual FF-A instance else MBZ. • This parameter MBZ if the caller does not specify a timeout.

Parameter	Register	Value
uint32 TimeoutHi	w3	<ul style="list-style-type: none"> Bits[63:32] of an interval measured in nanoseconds after which vCPU of the endpoint specified in <i>w1</i> must be run. Only valid with the ERET conduit at the Non-secure virtual FF-A instance else MBZ. This parameter MBZ if the caller does not specify a timeout.
Other Parameter registers	w4-w7 x4-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 9.8: FFA_ERROR encoding

Parameter	Register	Value
int32 Status	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS: Unrecognized endpoint or vCPU ID. Only valid with the ERET conduit. DENIED: Callee is not in a state to handle this request. NOT_SUPPORTED: This function is not implemented at this FF-A instance.

9.2.1 Component responsibilities for FFA_YIELD

This section describes the common responsibilities that the participating FF-A components must fulfill during an invocation of the *FFA_YIELD* interface. These components are:

1. VM.
2. Hypervisor.
3. Primary scheduler.

This interface is used by a VM during indirect message processing to block their execution and pass control back to the primary scheduler when an internal dependency of caller cannot be fulfilled. Execution of the caller is resumed by the primary scheduler in response to a notification that the dependency has been fulfilled. The caller remains in the *busy* state during an invocation of this interface.

The Hypervisor is responsible for returning control to the primary scheduler in response to an invocation of this interface. Its responsibilities in this regard are influenced by the relative location of the primary scheduler. These are as follows.

1. The Hypervisor and primary scheduler are co-resident. It must use an IMPLEMENTATION DEFINED mechanism to hand control to the scheduler.
2. The Hypervisor and primary scheduler are not co-resident. It must forward the *FFA_YIELD* call to the primary scheduler through the ERET conduit on the PE where the call is made. The ID of the endpoint and its caller execution context must be passed in the *w1* parameter register.

9.3 FFA_RUN

Description

- Run an endpoint execution context on the current PE.
 - Initial invocation must be from the FF-A component that implements the primary scheduler or is allowed to send direct messages.
 - An invocation must be forwarded to an endpoint by the Relayer.
 - Valid FF-A instances and conduits are listed in [Table 9.10](#).
 - Syntax of this function is described in [Table 9.11](#).
 - Successful completion of this function is indicated through the invocation of any FF-A function. FFA_RUN is used to allocate CPU cycles to an endpoint for message processing. The endpoint could invoke any FF-A function while processing a message.
 - Encoding of error code in the FFA_ERROR function is described in [Table 9.12](#).
-

Table 9.10: FFA_RUN instances and conduits

Config No.	FF-A instance	Valid conduits
1	Non-secure physical	SMC
2	Secure physical	ERET
3	Non-secure virtual	SMC, HVC, ERET
4	Secure virtual	ERET

Table 9.11: FFA_RUN function syntax

Parameter	Register	Value
uint32 Function ID	w0	• 0x8400006D.
uint32 Target information	w1	• Information to identify target SP/VM. <ul style="list-style-type: none">– Bits[31:16]: ID of SP/VM.– Bits[15:0]: ID of vCPU of SP/VM to run.
Other Parameter registers	w2-w7 x2-x7	• Reserved (MBZ).

Table 9.12: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS: Unrecognized endpoint or vCPU ID. NOT_SUPPORTED: This function is not implemented at this FF-A instance. DENIED: Callee is not in a state to handle this request. BUSY: vCPU is busy and caller must retry later. ABORTED: vCPU or VM ran into an unexpected error and has aborted.

9.3.1 Component responsibilities for FFA_RUN

This section describes the common responsibilities that the participating FF-A components must fulfill during an invocation of the *FFA_RUN* function (also see [4.3.3 Scheduling the Receiver](#)) by the primary scheduler to run the Receiver of a message. The location of the Receiver relative to the primary scheduler is one of the following.

1. The primary scheduler resides in a separate VM from the Receiver.
2. The primary scheduler resides in the same EL as the Hypervisor and the Receiver is an SP.
3. The Hypervisor is not present. The primary scheduler resides in the OS kernel and the Receiver is an SP.

[Table 9.13](#) lists the valid combinations of the Receiver and primary scheduler location.

Table 9.13: Valid configurations for FFA_RUN

Config no.	Primary Scheduler location	Receiver location
1.	VM	VM
2.	VM	SP
3.	Hypervisor	SP
4.	OS Kernel	SP

The use of the *FFA_RUN* interface with the configurations listed in [Table 9.13](#) is as follows.

1. In configs 3 & 4, the *FFA_RUN* invocation from the Hypervisor or OS kernel must be intercepted by the SPM. It must assume the responsibility of running the Receiver.
2. In configs 1 & 2, the *FFA_RUN* invocation from the primary scheduler must be intercepted by the Hypervisor. It must assume the responsibility of running the Receiver.

In config 1, the Hypervisor must program a return to the VM as described in the list entry 2 in [4.3.3 Scheduling the Receiver](#).

In config 2, the Hypervisor must invoke the *FFA_RUN* interface to request the SPM to run the SP.

In both 1 & 2, execution of *FFA_RUN* will result in the execution of the SPM. The SPM must run the SP in the same manner as the Hypervisor runs a VM as described in the list entry 2 in [4.3.3 Scheduling the Receiver](#).

The SPM and Hypervisor must return *NOT_SUPPORTED* if an endpoint invokes this interface when it only supports indirect messaging and does not implement the primary scheduler.

The SPM and Hypervisor must return *BUSY* if an endpoint invokes this interface when the specified vCPU of the endpoint is busy processing a request on another PE.

9.4 FFA_NORMAL_WORLD_RESUME

Description

- Request SPMD to resume execution of Normal world on current PE. See [9.4.1 Overview](#) for details.
- Valid FF-A instances and conduits are listed in [Table 9.15](#).
- Syntax of this function is described in [Table 9.16](#).
- Successful completion of this function is indicated through the invocation of any FF-A function.
- Encoding of error code in the FFA_ERROR function is described in [Table 9.17](#).

Table 9.15: FFA_NORMAL_WORLD_RESUME instances and conduits

Config No.	FF-A instance	Valid conduits
1	Secure physical	SMC

Table 9.16: FFA_WORLD_RESUME function syntax

Parameter	Register	Value
uint32 Function ID	• w0	• 0x8400007C.
Other Parameter registers	• w1-w7 • x1-x7	• Reserved (MBZ).

Table 9.17: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> • DENIED: Callee is not in a state to handle this request. • NOT_SUPPORTED: This function is not implemented at this FF-A instance.

9.4.1 Overview

Execution in Normal world could be preempted in response to an exception for example, a Secure physical interrupt. As per the Armv8-A architecture, the exception will be delivered to EL3 in the AArch64 Execution state or Monitor mode in the AArch32 Execution state. The exception could be handled in the Secure state at a lower Exception level than EL3 or Monitor mode.

This function must be used by the SPMC in S-EL2 (see [2.2.1 SPM architecture with Secure EL2](#)), S-EL1 (see

[2.2.2.1 S-EL1 SPM core component](#)) or Secure Supervisor mode (see [2.2.2.2 Secure Supervisor mode SPM core component](#)) to request the SPMD to resume Normal world execution once the exception has been handled.

The SPMD must ensure that the Normal world execution is resumed with exactly the same PE state that was saved when it was preempted.

The SPMD must return *DENIED* if this function is invoked at the Secure physical FF-A instance and the Normal world execution was not preempted.

The partition manager must return *NOT_SUPPORTED* if this function is invoked at any other FF-A instance.

An invocation of this function at the Secure physical FF-A instance could be completed through a valid invocation of any FF-A function through the ERET conduit.

Chapter 10

Messaging interfaces

10.1 FFA_MSG_SEND

Overview

- Send a Partition message to a VM through the RX/TX buffers by using indirect messaging.
 - Message is copied by Hypervisor from the TX buffer of Sender NS-Endpoint to the RX buffer of Receiver NS-endpoint.
 - The scheduler is informed about the pending message in the RX buffer of the Receiver.
 - Message will be read when the Receiver endpoint is scheduled to run.
 - See [10.1.2 Component responsibilities for FFA_MSG_SEND](#) for caller and callee roles and responsibilities.
 - Must not be invoked when the caller is processing a direct request.
 - Valid FF-A instances and conduits are listed in [Table 10.2](#).
 - Is used with the ERET conduit in the following scenarios.
 - * Inform an endpoint that a message is available in its RX buffer.
 - * Inform the primary scheduler that the Receiver has a pending message in its RX buffer.
 - Syntax of this function is described in [Table 10.3](#).
 - Successful completion of this function call is indicated as follows.
 - *w0* contains *FFA_SUCCESS* function ID.
 - *w1/x1-w7/x7* are reserved and MBZ.
 - Successful completion of this function does not imply that the message has been read by the Receiver endpoint.
 - Encoding of error code in the FFA_ERROR function is described in [Table 10.4](#).
 - See [10.1.1 Target availability notification](#) for behavior when BUSY is returned and caller must be notified about availability of TX buffer.
-

Table 10.2: FFA_MSG_SEND instances and conduits

Config No.	FF-A instance	Valid conduits
1	Non-secure virtual	SMC, HVC, ERET

Table 10.3: FFA_MSG_SEND function syntax

Parameter	Register	Value
uint32 Function ID	w0	• 0x8400006E.
uint32 Sender/Receiver IDs	w1	• Sender and Receiver endpoint IDs. <ul style="list-style-type: none">– Bit[31:16]: Sender endpoint ID.– Bit[15:0]: Receiver endpoint ID.
uint32/uint64 Reserved	w2/x2	• Reserved for future use (MBZ).

Parameter	Register	Value
uint32 Message size	w3	<ul style="list-style-type: none"> Length of message payload in the RX buffer. This is an optional field when used with the <i>ERET</i> conduit at the Non-secure virtual FF-A instance and the callee is not the Receiver of the message. It MBZ in this case.
uint32 Flags	w4	<ul style="list-style-type: none"> Message flags. <ul style="list-style-type: none"> Must be ignored by callee when SVC conduit is used. Bit[0]: Blocking behavior. <ul style="list-style-type: none"> b'0: Return BUSY if message cannot be delivered to Receiver. b'1: Return BUSY if message cannot be delivered to Receiver and notify when delivery is possible. Bit[31:1]: Reserved (MBZ).
uint32 Sender vCPU ID	w5	<ul style="list-style-type: none"> Information to identify execution context or vCPU of Sender endpoint. <ul style="list-style-type: none"> Only valid when ERET conduit is used. MBZ and ignored by callee otherwise. Bits[31:16]: Reserved (MBZ). Bits[15:0]: vCPU ID of Sender endpoint.
Other Parameter registers	w6-w7 x6-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 10.4: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS: A field in input parameters is incorrectly encoded. BUSY: Receiver RX buffer is not free. DENIED: Callee is not in a state to handle this request. NO_MEMORY: Insufficient memory to handle this request. NOT_SUPPORTED: This function is not implemented at this FF-A instance.

10.1.1 Target availability notification

When this interface is invoked, it is possible that the callee determines that the RX buffer of the Receiver VM cannot be written to. This can happen if either another instance of a Producer is writing to the RX buffer or the Receiver VM is reading from it as a Consumer (see [4.2.2.4 Buffer synchronization](#)). The callee must complete the interface invocation with a *BUSY* error code in this case.

A VM running in EL1 in either Security state can request to be notified when the RX buffer becomes available again by setting *bit[0] = 1* in the *Flags* parameter. In this case, the Hypervisor must:

1. Determine when the RX buffer is available as per the ownership rules described in [4.2.2.4 Buffer synchronization](#).
2. Notify each caller about the RX buffer availability.

The Hypervisor must describe the interrupt to indicate availability of the Receiver VM RX buffer to each VM respectively through an IMPLEMENTATION DEFINED mechanism. This could be done through a platform discovery mechanism like ACPI or Device tree.

A Consumer that is, OS kernel or VM must indicate the availability of its RX buffer by using a mechanism listed in [4.2.2.4 Buffer synchronization](#) for example, through the **FFA_RX_RELEASE** interface.

10.1.2 Component responsibilities for FFA_MSG_SEND

This section describes the common responsibilities that the participating FF-A components must fulfill during transmission of Partition messages between VMs through the *FFA_MSG_SEND* interface. This interface is used in the scenarios listed in [4.1.1 Indirect messaging](#).

10.1.2.1 Sender VM responsibilities

1. Must acquire ownership of empty TX buffer (see [4.2.2.4 Buffer synchronization](#)).
2. Must write Partition message payload to TX buffer.
3. Must specify length of Partition message payload.
4. Must specify blocking behavior in *Flags* parameter.
5. Must specify Sender and Receiver VM IDs.
6. Must implement support for handling all error status codes that can be returned on completion of these interfaces.
7. See [10.1.2.2 Hypervisor responsibilities](#) for Hypervisor responsibilities in this message transmission.

10.1.2.2 Hypervisor responsibilities

1. Must validate Sender and Receiver VM IDs and return *INVALID PARAMETER* if either is invalid.
2. Must check that reserved bits are 0 in *Flags* parameter. Return *INVALID PARAMETER* if this check fails.
3. Must check that reserved and unused parameter registers are 0. Return *INVALID PARAMETER* if this check fails.
4. Must check that the size of the *Receiver* RX buffer is large enough to accommodate the message. Must return *NO_MEMORY* if this is not true.
5. Must lock TX buffer of *Sender* from concurrent accesses before copying the message.
6. Must determine availability of RX buffer of *Receiver*.
 1. Return *BUSY* if RX buffer is not available.
 1. Save *Sender* ID if it wants the target availability interrupt when the RX buffer becomes free.
 2. Arrange for target availability interrupt to be delivered to Sender.
 2. Mark RX buffer as unavailable if it is available.
7. Must protect RX buffer of *Receiver* from concurrent accesses.
8. Must copy message from *Sender* TX buffer to *Receiver* RX buffer.
9. Must unlock TX buffer of *Sender* after copying the message.
10. Must unlock RX buffer of *Receiver* after copying the message.
11. Must inform primary scheduler that *Receiver* has a pending message as described in [10.1.3 Mechanism for scheduler notification](#).
12. Must return *SUCCESS* to *Sender* if message is successfully transmitted.
13. Must mark the RX buffer as available when the Receiver releases it.

10.1.2.3 Receiver VM responsibilities

1. Copy message from RX buffer.
2. Transfer ownership of the RX buffer by invoking the *FFA_RX_RELEASE* interface.

10.1.3 Mechanism for scheduler notification

This section describes how the primary scheduler must be notified depending on its location relative to the message Sender.

1. A VM is the Sender. The primary scheduler and Hypervisor are co-resident. The Hypervisor must use an IMPLEMENTATION DEFINED mechanism to notify the scheduler in response to the *FFA_MSG_SEND* call.
2. A VM is the Sender.
 1. The primary scheduler is resident in another VM.
 1. The Hypervisor must forward the *FFA_MSG_SEND* call to the primary scheduler using the *ERET* conduit on the PE where the call is made.
 2. Primary scheduler must respond to the forwarded *FFA_MSG_SEND* call with either a *FFA_SUCCESS* or *FFA_ERROR* invocation through the SMC conduit.
 2. The primary scheduler and Sender are co-resident. The Sender must use an IMPLEMENTATION DEFINED mechanism to notify the scheduler.

10.2 FFA_MSG_SEND_DIRECT_REQ

Description

- Send a Partition message in parameter registers as a request to a target endpoint, run the endpoint and block until a response is available.
 - Valid FF-A instances and conduits are listed in [Table 10.6](#).
 - Syntax of this function is described in [Table 10.7](#).
 - Successful completion of this function is indicated through an invocation of the following interfaces by the callee:
 - *FFA_MSG_SEND_DIRECT_RESP* to provide a response to the direct request.
 - *FFA_INTERRUPT* to indicate that the direct request was interrupted and must be resumed through the *FFA_RUN* interface.
 - *FFA_SUCCESS* to indicate completion of the direct request without a corresponding direct response.All other parameter registers MBZ.
 - Encoding of error code in the *FFA_ERROR* function is described in [Table 10.8](#).
-

Table 10.6: FFA_MSG_SEND_DIRECT_REQ instances and conduits

Config No.	FF-A instance	Valid conduits
1	Physical	SMC, ERET
2	Non-secure virtual	SMC, HVC, ERET
3	Secure virtual	SMC, HVC, SVC, ERET

Table 10.7: FFA_MSG_SEND_DIRECT_REQ function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none">• 0x8400006F.• 0xC400006F.
uint32 Source/Destination IDs	w1	<ul style="list-style-type: none">• Source and destination endpoint IDs.<ul style="list-style-type: none">– Bit[31:16]: Source endpoint ID.– Bit[15:0]: Destination endpoint ID.
uint32/uint64 Reserved	w2/x2	<ul style="list-style-type: none">• Reserved for future use (MBZ).
Other Parameter registers	w3-w7 x3-x7	<ul style="list-style-type: none">• IMPLEMENTATION DEFINED values. Is

Table 10.8: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS: Invalid endpoint ID or non-zero reserved register. DENIED: Callee is not in a state to handle this request. NOT_SUPPORTED: This function is not implemented at this FF-A instance. BUSY: Message target is busy. ABORTED: Message target ran into an unexpected error and has aborted.

10.2.1 Component responsibilities for FFA_MSG_SEND_DIRECT_REQ

This section describes the common responsibilities that the participating FF-A components must fulfill during transmission of Partition messages between endpoints through the *FFA_MSG_SEND_DIRECT_REQ* interface. This interface is used in the scenarios listed in [Table 4.9](#).

10.2.1.1 Sender responsibilities

10.2.1.1.1 Send from NS-Endpoint to S-Endpoint

1. Must write Partition message payload to parameter registers.
2. Must specify Sender and Receiver endpoint IDs.
3. Must implement support for handling all error status codes that can be returned on completion of this interface.
4. See [10.2.1.2.2 Relay from VM to S-Endpoint](#) & [10.2.1.3.3 Relay from NS-Endpoint to S-Endpoint](#) for Relayer responsibilities in this message transmission.

10.2.1.1.2 Send from VM to VM

1. Same as Sender responsibilities while sending message from NS-Endpoint to S-Endpoint as listed in [10.2.1.1.1 Send from NS-Endpoint to S-Endpoint](#).
2. See [10.2.1.2.1 Relay from VM to VM](#) for Relayer responsibilities in this message transmission.

10.2.1.1.3 Send from SP to SP

1. Same as Sender responsibilities while sending message from NS-Endpoint to S-Endpoint as listed in [10.2.1.1.1 Send from NS-Endpoint to S-Endpoint](#).
2. See [10.2.1.3.1 Relay from SP to SP](#) for Relayer responsibilities in this message transmission.

10.2.1.1.4 Send from S-Endpoint to NS-Endpoint

1. Same as Sender responsibilities while sending message from NS-Endpoint to S-Endpoint as listed in [10.2.1.1.1 Send from NS-Endpoint to S-Endpoint](#).
2. See [10.2.1.3.2 Relay from S-Endpoint to NS-Endpoint](#) for Relayer responsibilities in this message transmission.

10.2.1.2 Hypervisor responsibilities

10.2.1.2.1 Relay from VM to VM

1. Must validate that the Sender is allowed to send direct messages. Invoke *FFA_ERROR* with *NOT_SUPPORTED* as status if this is not the case.
2. Must validate Sender and Receiver endpoint IDs. Invoke *FFA_ERROR* with *INVALID PARAMETER* as status if either is invalid.

3. Must check that reserved parameter registers are 0. Invoke *FFA_ERROR* with *INVALID PARAMETER* as status if either is invalid.
4. Must ensure that target endpoint supports receipt of direct messages. Invoke *FFA_ERROR* with *DENIED* as status if this is not the case.
5. Must determine availability of an idle target endpoint execution context on this PE. Invoke *FFA_ERROR* with *BUSY* as status if not available.
6. Must ensure invocation of this interface by the Sender is completed only in response to an invocation of the *FFA_MSG_SEND_DIRECT_RESP* interface.
7. Must copy parameter registers from *Sender* execution context to *Receiver* execution context.
8. Must complete the invocation of the interface the Receiver had used to enter the idle state with an invocation of *FFA_MSG_SEND_DIRECT_REQ* through the *ERET* conduit.

10.2.1.2.2 Relay from VM to S-Endpoint

1. Same as 1-3 in [10.2.1.2.1 Relay from VM to VM](#).
2. Invoke *FFA_MSG_SEND_DIRECT_REQ* at physical FF-A instance through the *SMC* conduit with the same parameters as specified by the *Sender*. See [10.2.1.3.3 Relay from NS-Endpoint to S-Endpoint](#) for responsibilities of the SPM as the Relayer.

10.2.1.2.3 Relay from S-Endpoint to VM

1. Same as 1-8 in [10.2.1.2.1 Relay from VM to VM](#).

10.2.1.3 SPM responsibilities

10.2.1.3.1 Relay from SP to SP

1. Same as 1-8 in [10.2.1.2.1 Relay from VM to VM](#).

10.2.1.3.2 Relay from S-Endpoint to NS-Endpoint

1. Same as 1-3 in [10.2.1.2.1 Relay from VM to VM](#).
2. Invoke *FFA_MSG_SEND_DIRECT_REQ* at physical FF-A instance through the *ERET* conduit with the same parameters as specified by the *Sender*. See [10.2.1.2.3 Relay from S-Endpoint to VM](#) for responsibilities as the Relayer.

10.2.1.3.3 Relay from NS-Endpoint to S-Endpoint

1. Same as 1-8 in [10.2.1.2.1 Relay from VM to VM](#).

10.2.1.4 Receiver responsibilities

All Receivers have the same responsibilities irrespective of the origin of the message and the role of the Relayers in transmitting the message. These are as follows.

1. Copy message from parameter registers and process it.
2. Use the *FFA_MSG_SEND_DIRECT_RESP* interface to return the results of message processing..

10.3 FFA_MSG_SEND_DIRECT_RESP

Description

- Send a Partition message in parameter registers as a response to a target endpoint, run the endpoint and block until a new message is available.
- Valid FF-A instances and conduits are listed in [Table 10.10](#).
- Syntax of this function is described in [Table 10.11](#).
- Successful completion of this function is indicated in the same manner as that of the *FFA_MSG_WAIT* function (also see [9.1 FFA_MSG_WAIT](#)).
- Encoding of error code in the FFA_ERROR function is described in [Table 10.12](#).

Table 10.10: FFA_MSG_SEND_DIRECT_RESP instances and conduits

Config No.	FF-A instance	Valid conduits
1	Physical	SMC, ERET
2	Non-secure virtual	SMC, HVC, ERET
3	Secure virtual	SMC, HVC, SVC, ERET

Table 10.11: FFA_MSG_SEND_DIRECT_RESP function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none"> • 0x84000070. • 0xC4000070.
uint32 Source/Destination IDs	w1	<ul style="list-style-type: none"> • Source and destination endpoint IDs. <ul style="list-style-type: none"> – Bit[31:16]: Source endpoint ID. – Bit[15:0]: Destination endpoint ID.
uint32/uint64 Reserved	w2/x2	<ul style="list-style-type: none"> • Reserved for future use (MBZ).
Other Parameter registers	w3-w7 x3-x7	<ul style="list-style-type: none"> • IMPLEMENTATION DEFINED values.

Table 10.12: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS: Unrecognized endpoint or non-zero reserved register. DENIED: Callee is not in a state to handle this request. NOT_SUPPORTED: This function is not implemented at this FF-A instance. BUSY: Message target is busy.

10.3.1 Component responsibilities for FFA_MSG_SEND_DIRECT_RESP

This section describes the common responsibilities that the participating FF-A components must fulfill during transmission of Partition messages between endpoints through the *FFA_MSG_SEND_DIRECT_RESP* interface. This interface is used in the scenarios listed in [Table 4.9](#).

10.3.1.1 Sender responsibilities

10.3.1.1.1 Send from NS-Endpoint to S-Endpoint

1. Must write Partition message payload to parameter registers.
2. Must specify Sender and Receiver endpoint IDs.
3. Must implement support for handling all error status codes that can be returned on completion of this interface.
4. See [10.3.1.2.2 Relay from VM to S-Endpoint](#) & [10.3.1.3.3 Relay from NS-Endpoint to S-Endpoint](#) for Relayer responsibilities in this message transmission.

10.3.1.1.2 Send from VM to VM

1. Same as Sender responsibilities while sending message from NS-Endpoint to S-Endpoint as listed in [10.3.1.1.1 Send from NS-Endpoint to S-Endpoint](#).
2. See [10.3.1.2.1 Relay from VM to VM](#) for Relayer responsibilities in this message transmission.

10.3.1.1.3 Send from SP to SP

1. Same as Sender responsibilities while sending message from NS-Endpoint to S-Endpoint as listed in [10.3.1.1.1 Send from NS-Endpoint to S-Endpoint](#).
2. See [10.3.1.3.1 Relay from SP to SP](#) for Relayer responsibilities in this message transmission.

10.3.1.1.4 Send from S-Endpoint to NS-Endpoint

1. Same as Sender responsibilities while sending message from NS-Endpoint to S-Endpoint as listed in [10.3.1.1.1 Send from NS-Endpoint to S-Endpoint](#).
2. See [10.3.1.3.2 Relay from S-Endpoint to NS-Endpoint](#) for Relayer responsibilities in this message transmission.

10.3.1.2 Hypervisor responsibilities

10.3.1.2.1 Relay from VM to VM

1. Must validate that the Sender is allowed to send direct messages. Invoke *FFA_ERROR* with *NOT_SUPPORTED* as status if this is not the case.
2. Must validate Sender and Receiver endpoint IDs. Invoke *FFA_ERROR* with *INVALID PARAMETER* as status if either is invalid.
3. Must check that reserved parameter registers are 0. Invoke *FFA_ERROR* with *INVALID PARAMETER* as status if either is invalid.

4. Must ensure that target endpoint supports receipt of direct messages. Invoke *FFA_ERROR* with *DENIED* as status if this is not the case..
5. Must determine availability of an idle target endpoint execution context on this PE. Invoke *FFA_ERROR* with *BUSY* as status if not available.
6. Must ensure invocation of this interface by the Sender is completed only in response to an invocation of the *FFA_MSG_SEND_DIRECT_RESP* interface.
7. Must copy parameter registers from *Sender* execution context to *Receiver* execution context.
8. Must complete the invocation of the *FFA_MSG_SEND_DIRECT_REQ* interface that the Receiver had used to send the request to which the response is being provided.

10.3.1.2.2 Relay from VM to S-Endpoint

1. Same as 1-3 in [10.3.1.2.1 Relay from VM to VM](#).
2. Invoke *FFA_MSG_SEND_DIRECT_RESP* at physical FF-A instance through the *SMC* conduit with the same parameters as specified by the *Sender*. See [10.3.1.3.3 Relay from NS-Endpoint to S-Endpoint](#) for responsibilities of the SPM as the Relayer.

10.3.1.2.3 Relay from S-Endpoint to VM

1. Same as 1-8 in [10.3.1.2.1 Relay from VM to VM](#).

10.3.1.3 SPM responsibilities

10.3.1.3.1 Relay from SP to SP

1. Same as 1-7 in [10.3.1.2.1 Relay from VM to VM](#).

10.3.1.3.2 Relay from S-Endpoint to NS-Endpoint

1. Same as 1-3 in [10.3.1.2.1 Relay from VM to VM](#).
2. Invoke *FFA_MSG_SEND_DIRECT_RESP* at physical FF-A instance through the *ERET* conduit with the same parameters as specified by the *Sender*. See [10.3.1.2.3 Relay from S-Endpoint to VM](#) for responsibilities as the Relayer.

10.3.1.3.3 Relay from NS-Endpoint to S-Endpoint

1. Same as 1-8 in [10.3.1.2.1 Relay from VM to VM](#).

10.3.1.4 Receiver responsibilities

All Receivers have the same responsibilities irrespective of the origin of the message and the role of the Relayers in transmitting the message. These are as follows.

1. Copy response from parameter registers and process it.

10.4 FFA_MSG_POLL

Description

- Poll if a message is available in the RX buffer of the caller. Execution is returned to the caller if no message is available.
 - Must not be invoked when the caller is processing a direct request.
 - Valid FF-A instances and conduits are listed in [Table 9.2](#).
 - Syntax of this function is described in [Table 10.15](#).
 - Successful completion of this function is indicated through the invocation of the FFA_MSG_SEND interface (see [10.1 FFA_MSG_SEND](#)).
 - Encoding of error code in the FFA_ERROR function is described in [Table 10.16](#).
-

Table 10.14: FFA_MSG_POLL instances and conduits

Config No.	FF-A instance	Valid conduits
1	Non-secure virtual	SMC, HVC

Table 10.15: FFA_MSG_POLL function syntax

Parameter	Register	Value
uint32 Function ID	w0	• 0x8400006A.
Other Parameter registers	w1-w7 x1-x7	• Reserved (MBZ).

Table 10.16: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none">• RETRY: Message is not available in the caller's RX buffer.• DENIED: Callee is not in a state to handle this request.• NOT_SUPPORTED: This function is not implemented at this FF-A instance.

Chapter 11

Memory management interfaces

11.1 FFA_MEM_DONATE

Description

- Starts a transaction to transfer of ownership of a memory region from a Sender endpoint to a Receiver endpoint.
- Transaction details are described in a memory transaction descriptor (see [Table 5.19](#)).
- Descriptor is populated in the TX buffer of the Owner by default.
- Valid FF-A instances and conduits are listed in [Table 11.2](#).
- Syntax of this function is described in [Table 11.3](#).
- Encoding of result parameters in the FFA_SUCCESS function is described in [Table 11.4](#).
- Encoding of error code in the FFA_ERROR function is described in [Table 11.5](#).

Table 11.2: FFA_MEM_DONATE instances and conduits

Config No.	FF-A instance	Valid conduits
1	Secure and Non-secure physical	SMC, ERET
2	Secure and Non-secure virtual	SMC, HVC, SVC

Table 11.3: FFA_MEM_DONATE function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none"> • 0x84000071. • 0xC4000071.
uint32 Total length	w1	<ul style="list-style-type: none"> • Total length of the memory transaction descriptor in bytes.
uint32 Fragment length	w2	<ul style="list-style-type: none"> • Length in bytes of the memory transaction descriptor passed in this ABI invocation. • <i>Fragment length</i> must be \leq <i>Total length</i>. • If <i>Fragment length</i> < <i>Total length</i> then see 12.2.2 Transmission of transaction descriptor in fragments about how the remainder of the descriptor will be transmitted.
uint32/uint64 Address	w3/x3	<ul style="list-style-type: none"> • Base address of a buffer allocated by the Owner and distinct from the TX buffer. See 12.2.1 Transmission of transaction descriptor in dynamically allocated buffers. • MBZ if the TX buffer is used.

Parameter	Register	Value
uint32 Page count	w4	<ul style="list-style-type: none"> Number of 4K pages in the buffer allocated by the Owner and distinct from the TX buffer. See 12.2.1 Transmission of transaction descriptor in dynamically allocated buffers. MBZ if the TX buffer is used.
Other Parameter registers	w5-w7 x5-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 11.4: FFA_SUCCESS encoding

Parameter	Register	Value
uint64 Handle	w2/w3	<ul style="list-style-type: none"> Globally unique Handle to identify the memory region on successful transmission of the transaction descriptor.
Other Result registers	w4-w7 x4-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 11.5: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS. DENIED. NO_MEMORY. BUSY. ABORTED.

11.1.1 Component responsibilities for FFA_MEM_DONATE

This interface is used to initiate a transaction to donate a memory region to a single Receiver endpoint (also see [5.6.2 Donate memory transaction lifecycle](#)). Only the Owner and Relayer participate in this stage of the transaction. Responsibilities of the:

- Owner are listed in [11.1.1.1 Owner responsibilities](#).
- Relayer are listed in [11.1.1.2 Relayer responsibilities](#).

The transaction descriptor could be populated in a buffer dynamically allocated by the Owner as specified in [12.2.1 Transmission of transaction descriptor in dynamically allocated buffers](#).

Transmission of the transaction descriptor in fragments must be implemented by the Owner and Relayer as specified in [12.2.2 Transmission of transaction descriptor in fragments](#).

Time slicing of this ABI invocation must be implemented by the Owner and Relayer as specified in [12.2.3 Time slicing of memory management operations](#).

11.1.1.1 Owner responsibilities

1. Must ensure it is a *PE endpoint* and Owner of the memory region.
2. Must ensure the memory region is in an access state suitable for donation (see [Table 5.9](#)).
3. Must ensure the memory region fulfills the applicable rules stated in [5.4.1 Ownership and access rules](#).
4. Must describe memory region in the descriptor specified in [Table 5.19](#) with a single endpoint memory access descriptor (also see [5.12.3.1 Sender usage](#)).
5. Must implement support for handling all error status codes that can be returned on completion of this interface.
6. If the invocation of this interface completes successfully, then must send at least the following information to the Receiver in a Partition message:
 1. Globally unique Handle returned by the Relayer.
 2. Owner endpoint ID.

If the Receiver specified in the memory transaction descriptor is a SEPID, then the message must be sent to:

- Either the *proxy endpoint* for the SEPID (see [5.2 Direct memory access](#)) through a Partition message.
- Or the *independent* peripheral device associated with the SEPID through an IMPLEMENTATION DEFINED mechanism.

Provision of any other information from the transaction descriptor is IMPLEMENTATION DEFINED.

7. If the Receiver rejects the request in step 6, the Sender should use the *FFA_MEM_RECLAIM* interface with the Handle returned by the Relayer to reclaim ownership of the memory region. It must treat the memory region as being inaccessible until the *FFA_MEM_RECLAIM* invocation completes.

11.1.1.2 Relayer responsibilities

1. Must validate the *Total length* input parameter to ensure that the length of the transaction descriptor does not exceed the size of the buffer it has been populated in. Must return *INVALID_PARAMETERS* in case of an error.
2. Must validate the *Sender endpoint ID* field in the transaction descriptor to ensure that the Sender is the Owner of the memory region and a *PE endpoint*. Must return *DENIED* in case of an error.
3. Must ensure that a request by an SP to donate Secure memory to a NS-Endpoint is rejected by returning the *DENIED* error code.
4. Must ensure that the memory region is in the *Owner-EA* state for the Owner (see [Table 5.9](#)). It must return *DENIED* in case of an error.
5. Must validate that the *Endpoint memory access descriptor count & Endpoint memory access descriptor array* fields in the transaction descriptor as specified in [5.12.3.3 Relayer usage](#).
6. Must validate the *Memory region attributes* field in the transaction descriptor as specified in [5.11.4 Memory region attributes usage](#).
7. Must validate the *Flags* field specified in the transaction descriptor as specified in [5.12.4 Flags usage](#).
8. Must validate the *Handle* field specified in the transaction descriptor as specified in [5.12.1 Handle usage](#).
9. Unmap the memory region from the translation regime of the Owner, if managed by the Relayer as specified in [5.3 Address translation regimes](#).
10. If the Receiver is a *PE endpoint* or a *Stream endpoint* with a *proxy endpoint* managed by the Relayer, then the Relayer must:
 1. Save the transaction descriptor information so that it can be validated when retrieved through invocations of the *FFA_MEM_RETRIEVE_REQ* & *FFA_MEM_RETRIEVE_RESP* interfaces.
 2. Return *NO_MEMORY* if there is not enough memory to complete this operation.

11. If the Receiver is a *Stream endpoint* associated with an *independent* device managed by the Relayer, then the Relayer must:
 1. Allocate an IPA range and map the memory region in the translation regime of the Receiver managed by the Relayer as specified in [5.3 Address translation regimes](#).

The mapping must be created with the memory region attributes and permissions specified in the transaction descriptor.
 2. Describe the memory region to the device using the SEPID through an IMPLEMENTATION DEFINED mechanism.
12. If the call executes successfully, the Relayer must:
 1. Ensure that the state of the memory region in the participating FF-A components is observed as follows:
 1. If the Receiver is a *PE endpoint* or a *SEPID* associated with a dependent peripheral device, then:
 - *Owner-NA* for the Owner.
 - *!Owner-NA* for the Receiver.
 2. If the Receiver is a *SEPID* associated with an independent peripheral device, then:
 - *!Owner-NA* for the Owner.
 - *Owner-EA* for the Receiver.
 2. Allocate and return a *Handle* as described in [5.10.2 Memory region handle](#).
13. If the Owner is a VM and the Receiver is an SP or SEPID associated with a Secure Stream ID, the Hypervisor must forward the memory transaction descriptor to the SPM. This must be done by invoking this interface at the Non-secure physical FF-A instance as follows.
 1. The fields of the transaction descriptor must be unchanged apart from the following exception.
 1. The memory region must be described as composed of physically addressed constituent 4K pages in one or more *Constituent memory region descriptors*.

This must be done by performing the VA or IPA to PA translation of the memory region described by the Owner at the non-secure virtual FF-A instance.

The order in which the address ranges are specified by the Owner must be preserved by the Hypervisor.
 2. The *Constituent memory region descriptors* must be described in a *Composite memory region descriptor* which must be referenced by the *Endpoint memory access descriptor* included in the transaction descriptor.
 2. The updated transaction descriptor must be copied into the TX buffer shared between the Hypervisor and SPM.

If the TX buffer is busy, the Hypervisor must return *BUSY*.

If the TX buffer is too small and it is not possible to use the optional features to transmit the descriptor listed in [12.2.2 Transmission of transaction descriptor in fragments](#) and [12.2.1 Transmission of transaction descriptor in dynamically allocated buffers](#), the Hypervisor must return *NO_MEMORY*

The SPM must fulfill the Relayer responsibilities listed in this section.

11.2 FFA_MEM_LEND

Description

- Starts a transaction to transfer access to a memory region from its Owner to one or more Borrowers.
 - Transaction details are described in a memory transaction descriptor (see [Table 5.19](#)).
 - Descriptor is populated in the TX buffer of the Owner by default.
 - Valid FF-A instances and conduits are listed in [Table 11.7](#).
 - Syntax of this function is described in [Table 11.8](#).
 - Encoding of result parameters in the FFA_SUCCESS function is described in [Table 11.9](#).
 - Encoding of error code in the FFA_ERROR function is described in [Table 11.10](#).
-

Table 11.7: FFA_MEM_LEND instances and conduits

Config No.	FF-A instance	Valid conduits
1	Secure and Non-secure physical	SMC, ERET
2	Secure and Non-secure virtual	SMC, HVC, SVC

Table 11.8: FFA_MEM_LEND function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none">• 0x84000072.• 0xC4000072.
uint32 Total length	w1	<ul style="list-style-type: none">• Total length of the memory transaction descriptor in bytes.
uint32 Fragment length	w2	<ul style="list-style-type: none">• Length in bytes of the memory transaction descriptor passed in this ABI invocation.• <i>Fragment length</i> must be \leq <i>Total length</i>.• If <i>Fragment length</i> < <i>Total length</i> then see 12.2.2 Transmission of transaction descriptor in fragments about how the remainder of the descriptor will be transmitted.
uint32/uint64 Address	w3/x3	<ul style="list-style-type: none">• Base address of a buffer allocated by the Owner and distinct from the TX buffer. See 12.2.1 Transmission of transaction descriptor in dynamically allocated buffers.• MBZ if the TX buffer is used.

Parameter	Register	Value
uint32 Page count	w4	<ul style="list-style-type: none"> Number of 4K pages in the buffer allocated by the Owner and distinct from the TX buffer. See 12.2.1 Transmission of transaction descriptor in dynamically allocated buffers. MBZ if the TX buffer is used.
Other Parameter registers	w5-w7 x5-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 11.9: FFA_SUCCESS encoding

Parameter	Register	Value
uint64 Handle	w2/w3	<ul style="list-style-type: none"> Globally unique Handle to identify the memory region on successful transmission of the transaction descriptor. MBZ otherwise (see 5.10.2 Memory region handle).
Other Result registers	w4-w7 x4-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 11.10: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS. DENIED. NO_MEMORY. BUSY. ABORTED.

11.2.1 Component responsibilities for FFA_MEM_LEND

This interface is used to initiate a transaction to lend a memory region to one or more Borrower endpoints (also see [5.7.2 Lend memory transaction lifecycle](#)). Only the Lender and Relayer participate in this stage of the transaction. Responsibilities of the:

- Lender are listed in [11.2.1.1 Lender responsibilities](#).
- Relayer are listed in [11.2.1.2 Relayer responsibilities](#).

The transaction descriptor could be populated in a buffer dynamically allocated by the Lender as specified in [12.2.1 Transmission of transaction descriptor in dynamically allocated buffers](#).

Transmission of the transaction descriptor in fragments must be implemented by the Lender and Relayer as specified in [12.2.2 Transmission of transaction descriptor in fragments](#).

Time slicing of this ABI invocation must be implemented by the Lender and Relayer as specified in [12.2.3 Time slicing of memory management operations](#).

11.2.1.1 Lender responsibilities

1. Must ensure it is a *PE endpoint* and Owner of the memory region.
2. Must ensure the memory region is in an access state suitable for lending (see [Table 5.10](#)).
3. Must ensure the memory region fulfills the applicable rules stated in [5.4.1 Ownership and access rules](#).
4. Must describe memory region in the descriptor specified in [Table 5.19](#) with an endpoint memory access descriptor for each Borrower (also see [5.12.3.1 Sender usage](#)).
5. Must implement support for handling all error status codes that can be returned on completion of this interface.
6. If the invocation of this interface completes successfully, then must send at least the following information to each Borrower in a Partition message:
 1. Globally unique Handle returned by the Relayer.
 2. Lender endpoint ID.

If the Borrower specified in the memory transaction descriptor is a SEPID, then the message must be sent to:

- Either the *proxy endpoint* for the SEPID (see [5.2 Direct memory access](#)) through a Partition message.
- Or the *independent* peripheral device associated with the SEPID through an IMPLEMENTATION DEFINED mechanism.

Provision of any other information from the transaction descriptor is IMPLEMENTATION DEFINED.

7. If the Borrower rejects the request in step 6, the Lender should use the *FFA_MEM_RECLAIM* interface with the Handle returned by the Relayer to reclaim ownership of the memory region. It must treat the memory region as being inaccessible until the *FFA_MEM_RECLAIM* invocation completes.

11.2.1.2 Relayer responsibilities

1. Must validate the *Total length* input parameter to ensure that the length of the transaction descriptor does not exceed the size of the buffer it has been populated in. Must return *INVALID_PARAMETERS* in case of an error.
2. Must validate the *Sender endpoint ID* field in the transaction descriptor to ensure that the Lender is the Owner of the memory region and a *PE endpoint*. Must return *DENIED* in case of an error.
3. Must ensure that a request by an SP to lend Secure memory to a NS-Endpoint is rejected by returning the *DENIED* error code.
4. Must validate that the memory region is in the *Owner-EA* state for the Lender (see [Table 5.10](#)). It must return *DENIED* in case of an error.
5. Must validate that the *Endpoint memory access descriptor count & Endpoint memory access descriptor array* fields in the transaction descriptor as specified in [5.12.3.3 Relayer usage](#).
6. Must validate the *Memory region attributes* field in the transaction descriptor as specified in [5.11.4 Memory region attributes usage](#).
7. Must validate the *Flags* field specified in the transaction descriptor as specified in [5.12.4 Flags usage](#).
8. Must validate the *Handle* field specified in the transaction descriptor as specified in [5.12.1 Handle usage](#).
9. Unmap the memory region from the translation regime of the Lender, if managed by the Relayer as specified in [5.3 Address translation regimes](#).
10. If the Borrower is a *PE endpoint* or a *Stream endpoint* with a *proxy endpoint* managed by the Relayer, then the Relayer must:
 1. Save the transaction descriptor information so that it can be validated when retrieved through invocations of the *FFA_MEM_RETRIEVE_REQ* & *FFA_MEM_RETRIEVE_RESP* interfaces.
 2. Return *NO_MEMORY* if there is not enough memory to complete this operation.

11. If the Borrower is a *Stream endpoint* associated with an *independent* device managed by the Relayer, then the Relayer must:
 1. Allocate an IPA range and map the memory region in the translation regime of the Borrower managed by the Relayer as specified in [5.3 Address translation regimes](#).

The mapping must be done with the memory region attributes and permissions specified in the transaction descriptor.
 2. Describe the memory region to the device using the SEPID through an IMPLEMENTATION DEFINED mechanism.
12. If the call executes successfully, the Relayer must:
 1. Ensure that the state of the memory region in the participating FF-A components is observed as follows:
 1. If a Borrower is a *PE endpoint* or a *SEPID* associated with a dependent peripheral device, then:
 - *Owner-NA* for the Lender.
 - *!Owner-NA* for the Borrower.
 2. If a Borrower is a *SEPID* associated with an independent peripheral device, then:
 - *Owner-NA* for the Lender.
 - *!Owner-EA* for the Borrower, if the count of Borrowers in the transaction is = 1.
 - *Owner-SA* for the Borrower, if the count of Borrowers in the transaction is > 1.
 2. Allocate and return a *Handle* as described in [5.10.2 Memory region handle](#).
13. If the Lender is a VM and the Borrower is an SP or SEPID associated with a Secure Stream ID, the Hypervisor must forward the memory transaction descriptor to the SPM. This must be done by invoking this interface at the Non-secure physical FF-A instance as follows.
 1. The fields of the transaction descriptor must be unchanged apart from the following exception.
 1. The memory region must be described as composed of physically addressed constituent 4K pages in one or more *Constituent memory region descriptors*.

This must be done by performing the VA or IPA to PA translation of the memory region described by the Owner at the non-secure virtual FF-A instance.

The order in which the address ranges are specified by the Lender must be preserved by the Hypervisor.
 2. The *Constituent memory region descriptors* must be described in a *Composite memory region descriptor* which must be referenced by the *Endpoint memory access descriptor* included in the transaction descriptor.
 2. The updated transaction descriptor must be copied into the TX buffer shared between the Hypervisor and SPM.

If the TX buffer is busy, the Hypervisor must return *BUSY*.

If the TX buffer is too small and it is not possible to use the optional features to transmit the descriptor listed in [12.2.2 Transmission of transaction descriptor in fragments](#) and [12.2.1 Transmission of transaction descriptor in dynamically allocated buffers](#), the Hypervisor must return *NO_MEMORY*

The SPM must fulfill the Relayer responsibilities listed in this section.

11.3 FFA_MEM_SHARE

Description

- Starts a transaction to grant access to a memory region to one or more Borrowers.
- Transaction details are described in a memory transaction descriptor (see [Table 5.19](#)).
- Descriptor is populated in the TX buffer of the Owner by default.
- Valid FF-A instances and conduits are listed in [Table 11.12](#).
- Syntax of this function is described in [Table 11.13](#).
- Encoding of result parameters in the FFA_SUCCESS function is described in [Table 11.14](#).
- Encoding of error code in the FFA_ERROR function is described in [Table 11.15](#).

Table 11.12: FFA_MEM_SHARE instances and conduits

Config No.	FF-A instance	Valid conduits
1	Secure and Non-secure physical	SMC, ERET
2	Secure and Non-secure virtual	SMC, HVC, SVC

Table 11.13: FFA_MEM_SHARE function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none"> • 0x84000073. • 0xC4000073.
uint32 Total length	w1	<ul style="list-style-type: none"> • Total length of the memory transaction descriptor in bytes.
uint32 Fragment length	w2	<ul style="list-style-type: none"> • Length in bytes of the memory transaction descriptor passed in this ABI invocation. • <i>Fragment length</i> must be \leq <i>Total length</i>. • If <i>Fragment length</i> < <i>Total length</i> then see 12.2.2 Transmission of transaction descriptor in fragments about how the remainder of the descriptor will be transmitted.
uint32/uint64 Address	w3/x3	<ul style="list-style-type: none"> • Base address of a buffer allocated by the Owner and distinct from the TX buffer. See 12.2.1 Transmission of transaction descriptor in dynamically allocated buffers. • MBZ if the TX buffer is used.

Parameter	Register	Value
uint32 Page count	w4	<ul style="list-style-type: none"> Number of 4K pages in the buffer allocated by the Owner and distinct from the TX buffer. See 12.2.1 Transmission of transaction descriptor in dynamically allocated buffers. MBZ if the TX buffer is used.
Other Parameter registers	w5-w7 x5-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 11.14: FFA_SUCCESS encoding

Parameter	Register	Value
uint64 Handle	w2/w3	<ul style="list-style-type: none"> Globally unique Handle to identify the memory region on successful transmission of the transaction descriptor. MBZ otherwise (see 5.10.2 Memory region handle).
Other Result registers	w4-w7 x4-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 11.15: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS. DENIED. NO_MEMORY. BUSY. ABORTED.

11.3.1 Component responsibilities for FFA_MEM_SHARE

This interface is used to initiate a transaction to share a memory region with one or more Receiver endpoints (also see [5.8.2 Share memory transaction lifecycle](#)). Only the Owner and Relayer participate in this stage of the transaction. Responsibilities of the:

- Owner are listed in [11.3.1.1 Owner responsibilities](#).
- Relayer are listed in [11.3.1.2 Relayer responsibilities](#).

The transaction descriptor could be populated in a buffer dynamically allocated by the Lender as specified in [12.2.1 Transmission of transaction descriptor in dynamically allocated buffers](#).

Transmission of the transaction descriptor in fragments must be implemented by the Lender and Relayer as specified in [12.2.2 Transmission of transaction descriptor in fragments](#).

Time slicing of this ABI invocation must be implemented by the Lender and Relayer as specified in [12.2.3 Time slicing of memory management operations](#).

11.3.1.1 Owner responsibilities

1. Must ensure it is a *PE endpoint* and Owner of the memory region.
2. Must ensure the memory region is in an access state suitable for sharing (see [Table 5.11](#)).
3. Must ensure the memory region fulfills the applicable rules stated in [5.4.1 Ownership and access rules](#).
4. Must describe memory region in the descriptor specified in [Table 5.19](#) with an endpoint memory access descriptor for each Borrower (also see [5.12.3.1 Sender usage](#)).
5. Must implement support for handling all error status codes that can be returned on completion of this interface.
6. If the invocation of this interface completes successfully, then must send at least the following information to each Borrower in a Partition message:
 1. Globally unique Handle returned by the Relayer.
 2. Owner endpoint ID.

If the Borrower specified in the memory transaction descriptor is a SEPID, then the request must be sent to:

- Either the *proxy endpoint* for the SEPID (see [5.2 Direct memory access](#)) through a Partition message.
- Or the *independent* peripheral device associated with the SEPID through an IMPLEMENTATION DEFINED mechanism.

Provision of any other information from the transaction descriptor is IMPLEMENTATION DEFINED.

7. If the Receiver rejects the request in step 6, the Sender should use the *FFA_MEM_RECLAIM* interface with the Handle returned by the Relayer to reclaim ownership of the memory region. It must treat the memory region as being inaccessible until the *FFA_MEM_RECLAIM* invocation completes.

11.3.1.2 Relayer responsibilities

1. Must validate the *Total length* input parameter to ensure that the length of the transaction descriptor does not exceed the size of the buffer it has been populated in. Must return *INVALID_PARAMETERS* in case of an error.
2. Must validate the *Sender endpoint ID* field in the transaction descriptor to ensure that the Lender is the Owner of the memory region and a *PE endpoint*. Must return *DENIED* in case of an error.
3. Must ensure that a request by an SP to share Secure memory to a NS-Endpoint is rejected by returning the *DENIED* error code.
4. Must validate that the memory region is in an access state suitable for sharing (see [Table 5.11](#)) and return *DENIED* in case of an error.
5. Must validate that the *Endpoint memory access descriptor count & Endpoint memory access descriptor array* fields in the transaction descriptor as specified in [5.12.3.3 Relayer usage](#).
6. Must validate the *Memory region attributes* field in the transaction descriptor as specified in [5.11.4 Memory region attributes usage](#).
7. Must validate the *Flags* field specified in the transaction descriptor as specified in [5.12.4 Flags usage](#).
8. Must validate the *Handle* field specified in the transaction descriptor as specified in [5.12.1 Handle usage](#).
9. If the Lender has specified a different data access permission to access the memory region in its translation regime, then the Relayer must validate the permission as specified in [5.11.2 Data access permissions usage](#), save the current permission and update the translation tables to reflect the new permission.
10. If the Borrower is a *PE endpoint* or a *Stream endpoint* with a *proxy endpoint* managed by the Relayer, then the Relayer must:
 1. Save the transaction descriptor information so that it can be validated when retrieved through invocations of the *FFA_MEM_RETRIEVE_REQ* & *FFA_MEM_RETRIEVE_RESP* interfaces.

2. Return *NO_MEMORY* if there is not enough memory to complete this operation.
11. If the Borrower is a *Stream endpoint* associated with an *independent* device managed by the Relayer, then the Relayer must:
 1. Allocate an IPA range and map the memory region in the translation regime of the Borrower managed by the Relayer as specified in [5.3 Address translation regimes](#).

The mapping must be done with the memory region attributes and permissions specified in the transaction descriptor.
 2. Describe the memory region to the device using the SEPID through an IMPLEMENTATION DEFINED mechanism.
12. If the call executes successfully, the Relayer must:
 1. Ensure that the state of the memory region in the participating FF-A components is observed as follows:
 1. If a Borrower is a *PE endpoint* or a *SEPID* associated with a dependent peripheral device, then:
 - *Owner-SA* for the Lender.
 - *!Owner-NA* for the Borrower.
 2. If a Borrower is a *SEPID* associated with an independent peripheral device, then:
 - *Owner-SA* for the Lender.
 - *!Owner-SA* for the Borrower.
 2. Allocate and return a *Handle* as described in [5.10.2 Memory region handle](#).
13. If the Lender is a VM and the Borrower is an SP or SEPID associated with a Secure Stream ID, the Hypervisor must forward the memory transaction descriptor to the SPM. This must be done by invoking this interface at the Non-secure physical FF-A instance as follows.
 1. The fields of the transaction descriptor must be unchanged apart from the following exception.
 1. The memory region must be described as composed of physically addressed constituent 4K pages in one or more *Constituent memory region descriptors*.

This must be done by performing the VA or IPA to PA translation of the memory region described by the Owner at the non-secure virtual FF-A instance.

The order in which the address ranges are specified by the Lender must be preserved by the Hypervisor.
 2. The *Constituent memory region descriptors* must be described in a *Composite memory region descriptor* which must be referenced by the *Endpoint memory access descriptor* included in the transaction descriptor.
 2. The updated transaction descriptor must be copied into the TX buffer shared between the Hypervisor and SPM.

If the TX buffer is busy, the Hypervisor must return *BUSY*.

If the TX buffer is too small and it is not possible to use the optional features to transmit the descriptor listed in [12.2.2 Transmission of transaction descriptor in fragments](#) and [12.2.1 Transmission of transaction descriptor in dynamically allocated buffers](#), the Hypervisor must return *NO_MEMORY*.

The SPM must fulfill the Relayer responsibilities listed in this section.

11.4 FFA_MEM_RETRIEVE_REQ

Description

- Requests completion of a donate, lend or share memory management transaction.
 - Transaction details are described in a memory transaction descriptor (see [Table 5.19](#)).
 - Descriptor is populated in the TX buffer of the Receiver by default.
 - Valid FF-A instances and conduits are listed in [Table 11.17](#).
 - Syntax of this function is described in [Table 11.18](#).
 - Encoding of error code in the FFA_ERROR function is described in [Table 11.19](#).
 - Successful transmission of the transaction descriptor is indicated by an invocation of the *FFA_MEM_RETRIEVE_RESP* function (see [11.5 FFA_MEM_RETRIEVE_RESP](#)).
-

Table 11.17: FFA_MEM_RETRIEVE_REQ instances and conduits

Config No.	FF-A instance	Valid conduits
1	Secure and Non-secure physical	SMC, ERET
2	Secure and Non-secure virtual	SMC, HVC, SVC

Table 11.18: FFA_MEM_RETRIEVE_REQ function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none">• 0x84000074.• 0xC4000074.
uint32 Total length	w1	<ul style="list-style-type: none">• Total length of the memory transaction descriptor in bytes.
uint32 Fragment length	w2	<ul style="list-style-type: none">• Length in bytes of the memory transaction descriptor passed in this ABI invocation.• <i>Fragment length</i> must be \leq <i>Total length</i>.• If <i>Fragment length</i> < <i>Total length</i> then see 12.2.2 Transmission of transaction descriptor in fragments about how the remainder of the descriptor will be transmitted.
uint32/uint64 Address	w3/x3	<ul style="list-style-type: none">• Base address of a buffer allocated by the Owner and distinct from the TX buffer. See 12.2.1 Transmission of transaction descriptor in dynamically allocated buffers.• MBZ if the TX buffer is used.

Parameter	Register	Value
uint32 Page count	w4	<ul style="list-style-type: none"> Number of 4K pages in the buffer allocated by the Owner and distinct from the TX buffer. See 12.2.1 Transmission of transaction descriptor in dynamically allocated buffers. MBZ if the TX buffer is used.
Other Parameter registers	w5-w7 x5-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 11.19: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS. DENIED. NO_MEMORY. BUSY. ABORTED.

11.4.1 Component responsibilities for FFA_MEM_RETRIEVE_REQ

This ABI is used by a Receiver to retrieve a memory region. Retrieval implies a request to the Relayer to map the memory region in the translation regime of the Receiver. The Receiver must use the transaction descriptor (see [Table 5.19](#)) to identify the memory region and specify its properties. The Receiver could:

1. Retrieve a memory region that was shared, lent or donated by an Owner. For this scenario, responsibilities of the:
 - Receiver are listed in [11.4.1.1 Receiver responsibilities](#).
 - Relayer are listed in [11.4.1.2 Relayer responsibilities](#).
2. Retrieve a memory region that it had relinquished but has not been reclaimed by the Owner yet (see [11.4.2 Support for multiple retrievals by a Borrower](#)).

It is also possible for a Hypervisor to use this interface to retrieve a memory region description on its behalf. This scenario is described in [11.4.3 Support for retrieval by the Hypervisor](#).

In all cases, a successful retrieval is indicated by an invocation of the *FFA_MEM_RETRIEVE_RESP* ABI by the Relayer.

The transaction descriptor could be populated in a buffer dynamically allocated by the Receiver as specified in [12.2.1 Transmission of transaction descriptor in dynamically allocated buffers](#).

Transmission of the transaction descriptor in fragments must be implemented by the caller and Relayer as specified in [12.2.2 Transmission of transaction descriptor in fragments](#).

Time slicing of this ABI invocation must be implemented by the Receiver and Relayer as specified in [12.2.3 Time slicing of memory management operations](#).

11.4.1.1 Receiver responsibilities

1. Must populate a transaction descriptor with the *Handle* (see [5.12.1 Handle usage](#)) that identifies the memory region, the *Endpoint ID* that identifies the Owner and the *Tag* (see [5.12.2 Tag usage](#)) associated with the

transaction.

Could populate other fields of the transaction descriptor as follows. This depends on how much information the Sender shares with the Receiver about the memory management transaction.

- See [5.12.3.2 Receiver usage](#) for usage of the *Endpoint memory access descriptor array* field.
- See [5.12.4 Flags usage](#) for usage of the *Flags* field.
- See [5.11.4 Memory region attributes usage](#) for usage of the *Memory region attributes* field.

The Relayer must validate the information provided by the Sender. It must reject the transaction by sending a Partition message to the Sender if the validation fails.

The protocol between the Sender and Receiver to convey transaction information and to reject the transaction is IMPLEMENTATION DEFINED.

2. Must implement support for handling all error status codes that can be returned on completion of this interface.

11.4.1.2 Relayer responsibilities

1. Must validate the *Total length* input parameter to ensure that the length of the transaction descriptor does not exceed the size of the buffer it has been populated in. Must return *INVALID_PARAMETERS* in case of an error.
2. Must validate the Sender endpoint ID field in the transaction descriptor to ensure that the Sender is the Owner of the memory region or the proxy endpoint acting on behalf of a Stream endpoint. Must return *DENIED* in case of an error.
3. Must validate the *Memory region attributes* field in the transaction descriptor as specified in [5.11.4 Memory region attributes usage](#).
4. Must validate the *Flags* field specified in the transaction descriptor as specified in [5.12.4 Flags usage](#).
5. Must validate the *Handle* field specified in the transaction descriptor as specified in [5.12.1 Handle usage](#).
6. Must validate the *Tag* field specified in the transaction descriptor as specified in [5.12.2 Tag usage](#).
7. Must validate that the *Endpoint memory access descriptor count* & *Endpoint memory access descriptor array* fields in the transaction descriptor as specified in [5.12.3.3 Relayer usage](#).
8. Must map the memory region in the translation regime of the Receiver managed by the Relayer as specified in [5.3 Address translation regimes](#).

If the Receiver is a proxy endpoint for one or more Stream endpoints then the memory region must be mapped in the stage 2 translation tables corresponding to each SEPID. The memory region must not be mapped in the translation regime of the proxy endpoint.

The order in which the address ranges are specified by the Lender must be preserved by the Hypervisor.

The Relayer must return *NO_MEMORY* if there is not enough memory to map the memory region.

9. Must return *BUSY*, if *FFA_MEM_RETRIEVE_RESP* cannot be invoked because the Receiver RX buffer is busy.
10. Must return *NO_MEMORY* if *FFA_MEM_RETRIEVE_RESP* cannot be invoked because there is not enough memory to allocate a transaction descriptor to describe the memory region.
11. If the call executes successfully, the Relayer must ensure that the state of the memory region in the participating FF-A components is observed as follows:
 - If the transaction type is *FFA_MEM_DONATE*,
 - *!Owner-NA* for the Owner.
 - *Owner-EA* for the Receiver.

- If the transaction type is FFA_MEM_LEND, and the count of Borrowers in the transaction is = I ,
 - *Owner-NA* for the Lender.
 - *!Owner-EA* for the Borrower.
- If the transaction type is FFA_MEM_LEND, and the count of Borrowers in the transaction is $> I$,
 - *Owner-SA* for the Lender.
 - *!Owner-SA* for the Borrower.
- If the transaction type is FFA_MEM_SHARE,
 - *Owner-SA* for the Lender.
 - *!Owner-SA* for the Borrower.

11.4.2 Support for multiple retrievals by a Borrower

After a Receiver relinquishes access to a memory region (see [11.6 FFA_MEM_RELINQUISH](#)) that was lent or shared, a Relayer must allow the Receiver to retrieve the memory region again as long as it has not been reclaimed by its Owner. To support this mechanism, it must:

1. Allow the Owner to reclaim the memory region only if all Borrowers have relinquished it as many times as they have retrieved it.
2. Unmap the memory region from the translation regime of the Borrower only after it has been relinquished as many times as it was retrieved.
3. Ensure that the address ranges used to describe the memory region on each retrieval are the same if the memory region is already mapped in the translation regime of the Receiver.

The number of times a Receiver is allowed to retrieve a memory region without relinquishing it first is I by default. A Receiver must use the FFA_FEATURES ABI (see [8.2 FFA_FEATURES](#)) to determine the number of outstanding retrievals supported by the Relayer. The Relayer must return *DENIED* if a Receiver exceeds the retrieval count.

11.4.3 Support for retrieval by the Hypervisor

In a transaction to donate, share or lend a memory region between a Owner VM and a Receiver SP, the SPM is responsible for allocating the *Handle* to identify the memory region (see [5.10.2 Memory region handle](#)).

The Hypervisor implementation could maintain an association between the Handle and the memory region. For example, to map the memory region back into the translation regime of the Owner in response to a FFA_MEM_RECLAIM ABI.

A Hypervisor implementation could choose to rely on the SPM to manage the association between the Handle and the memory region. For example, to avoid memory costs associated with tracking this state over a period of time.

In this case, the Hypervisor could use the FFA_MEM_RETRIEVE_REQ ABI to obtain the memory region description by specifying its Handle. It would use this description to map or unmap the memory region depending on the operation requested by a VM. For example, an operation to reclaim a memory region would follow these steps.

1. Lender VM calls FFA_MEM_RECLAIM.
2. Hypervisor uses the Handle to call FFA_MEM_RETRIEVE_REQ and obtain the memory region description.
3. Hypervisor forwards FFA_MEM_RECLAIM to the SPM to ensure all Borrowers have stopped using the memory region.
4. On a successful return from the SPM, the Hypervisor uses the memory region description to map the region in the translation regime of the Lender VM.
5. Hypervisor completes the invocation of FFA_MEM_RECLAIM from the Lender VM successfully.

If it chooses to use this mechanism, the Hypervisor must populate the transaction descriptor as follows.

1. It must specify the Handle specified by the VM in the *Handle* field.
2. It must ensure that all other fields in the transaction descriptor are zeroed.

From the perspective of an SPM, an invocation of the FFA_MEM_RETRIEVE_REQ ABI at the Non-secure physical FF-A instance could,

- Either originate from the Hypervisor as described above.
- Or originate from a Borrower VM. It was forwarded by the Hypervisor.

In the former case, the SPM must not update the ownership and access state associated with the memory region as it would do in the latter case (see [11.4.1.2 Relayer responsibilities](#)). To do this, the SPM must distinguish between the two types of invocation as follows.

- In the former case, the *Endpoint memory access descriptor count* in the transaction descriptor must be 0.
- In the latter case, the *Endpoint memory access descriptor count* in the transaction descriptor must be ≥ 1 .

In the former case, the SPM must also validate the *Handle* field specified in the transaction descriptor as follows.

- Ensure that it identifies a memory region that was either shared or lent to at least a single VM or is owned by a VM.
- Ensure that it was previously allocated and has not been reclaimed by the Owner.

The SPM must provide the memory region description to the Hypervisor through an invocation of the FFA_MEM_RETRIEVE_RESP ABI as follows.

- The memory region must be described as composed of physically addressed constituent 4K pages in one or more *Constituent memory region descriptors*.
- The *Constituent memory region descriptors* must be described in a *Composite memory region descriptor*.
- The *Composite memory region descriptor* must be referenced by a single *Endpoint memory access descriptor* included in the transaction descriptor.
- The *Sender endpoint ID* field must be set to the Lender or Owner VM ID in the transaction descriptor.
- The *Handle* field must be set to the input *Handle*.

U

Implementation Note

This feature allows the Hypervisor to retrieve the physical address ranges of a memory region that must be either mapped or unmapped from the stage 2 translation descriptors of a VM.

It is possible that the Hypervisor implementation maintains mappings in the stage 2 translation descriptors for a VM such that a $IPA \neq PA$. In this case, it must track the original IPA ranges through an IMPLEMENTATION DEFINED mechanism to be able to correctly map or unmap the retrieved memory region.

Furthermore, in the case where the Hypervisor must map the memory region in the stage 2 translation descriptors for a VM, it must track the original memory access permissions and attributes of the memory region through an IMPLEMENTATION DEFINED mechanism.

11.5 FFA_MEM_RETRIEVE_RESP

Description

- A Relay uses this interface to describe a memory region and its properties in response to the latest successful invocation of the *FFA_MEM_RETRIEVE_REQ* interface by an endpoint or Hypervisor.
- Transaction details are described in a transaction descriptor specified in [Table 5.19](#).
- Descriptor is populated in the RX buffer of the Receiver by default.
- Valid FF-A instances and conduits are listed in [Table 11.21](#).
- Syntax of this function is described in [Table 11.22](#).
- Encoding of error code in the FFA_ERROR function is described in [Table 11.23](#).
- Successful transmission of the transaction descriptor is indicated by an invocation of any FF-A function by the Receiver.

Table 11.21: FFA_MEM_RETRIEVE_RESP instances and conduits

Config No.	FF-A instance	Valid conduits
1	Secure and Non-secure physical	SMC, ERET
2	Secure and Non-secure virtual	ERET

Table 11.22: FFA_MEM_RETRIEVE_RESP function syntax

Parameter	Register	Value
uint32 Function ID	w0	• 0x84000075.
uint32 Total length	w1	• Total length of the memory transaction descriptor in bytes.
uint32 Fragment length	w2	• Length in bytes of the memory transaction descriptor passed in this ABI invocation. • <i>Fragment length</i> must be \leq <i>Total length</i> . • If <i>Fragment length</i> < <i>Total length</i> then see 12.2.2 Transmission of transaction descriptor in fragments about how the remainder of the descriptor will be transmitted.
uint32/uint64 Parameter	w3/x3	• Reserved (MBZ).
uint32/uint64 Parameter	w4/x4	• Reserved (MBZ).
Other Parameter registers	w5-w7 x5-x7	• Reserved (MBZ).

Table 11.23: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS. NO_MEMORY.

11.5.1 Component responsibilities for FFA_MEM_RETRIEVE_RESP

A Relayer invokes this interface as the caller with an endpoint as the callee in the following scenarios. It must fulfill the responsibilities listed in [11.5.1.1 Relayer responsibilities](#).

- The endpoint calls *FFA_MEM_RETRIEVE_REQ* to retrieve a memory region that was donated, lent or shared with it by the Owner. The Relayer completes the transaction by invoking the *FFA_MEM_RETRIEVE_RESP* interface.
- The endpoint calls *FFA_MEM_RETRIEVE_REQ* to retrieve a memory region it had relinquished earlier. The Relayer fulfills the request by invoking the *FFA_MEM_RETRIEVE_RESP* interface (see [11.4.2 Support for multiple retrievals by a Borrower](#)).

The SPM invokes this interface as the caller with the Hypervisor as the callee in the scenario described in [11.4.3 Support for retrieval by the Hypervisor](#).

In all scenarios, the Relayer must populate a transaction descriptor specified in [Table 5.19](#) to describe the memory region and its properties. This must be done in one of the following buffers.

- The RX buffer of the callee must be used if the callee used its TX buffer in the counterpart invocation of the *FFA_MEM_RETRIEVE_REQ* ABI earlier.
- If the callee used a dynamically allocated buffer in the counterpart invocation of the *FFA_MEM_RETRIEVE_REQ* ABI earlier, then the same buffer must be used (see [12.2.1 Transmission of transaction descriptor in dynamically allocated buffers](#)).

Transmission of the transaction descriptor in fragments in all scenarios must be implemented by the Relayer and callee as specified in [12.2.2 Transmission of transaction descriptor in fragments](#).

In all scenarios, the callee (endpoint or Hypervisor) must fulfill the responsibilities listed in [11.5.1.2 Callee responsibilities](#). It must use the error codes listed in [Table 11.23](#) to report an error back to the Relayer.

11.5.1.1 Relayer responsibilities

- Must populate the Sender endpoint ID field in the transaction descriptor with the endpoint ID of the Owner.
- Must populate the *Memory region attributes* field in the transaction descriptor as specified in [5.11.4 Memory region attributes usage](#).
- Must populate the *Flags* field specified in the transaction descriptor as specified in [5.12.4 Flags usage](#).
- Must populate the *Handle* field specified in the transaction descriptor as specified in [5.12.1 Handle usage](#).
- Must populate the *Tag* field specified in the transaction descriptor as specified in [5.12.2 Tag usage](#).
- Must populate the *Endpoint memory access descriptor count & Endpoint memory access descriptor array* fields in the transaction descriptor as specified in [5.12.3.3 Relayer usage](#).

11.5.1.2 Callee responsibilities

- Must return *INVALID_PARAMETERS* if any field in the transaction descriptor has been incorrectly encoded.
- Must return *NO_MEMORY* if there is not enough memory to use the memory region description provided by the Relayer.

11.6 FFA_MEM_RELINQUISH

Description

- Starts a transaction to transfer access to a shared or lent memory region from a Borrower back to its Owner.
 - Valid FF-A instances and conduits are listed in [Table 11.26](#).
 - Syntax of this function is described in [Table 11.27](#).
 - Successful completion of this function is indicated through the invocation of the FFA_SUCCESS function by the callee without any further parameters.
 - Encoding of error code in the FFA_ERROR function is described in [Table 11.28](#).
-

Table 11.25: Descriptor to relinquish a memory region

Field	Byte length	Byte offset	Description
Handle	8	0	<ul style="list-style-type: none">• Globally unique Handle to identify a memory region.
Flags	4	8	<ul style="list-style-type: none">• Bit[0]: Zero memory after relinquish flag.<ul style="list-style-type: none">– This flag specifies if the Relayer must clear memory region contents after unmapping it from from the translation regime of the Borrower.<ul style="list-style-type: none">* b'0: Relayer must not zero the memory region contents.* b'1: Relayer must zero the memory region contents.– If the memory region was lent to multiple Borrowers, the Relayer must clear memory region contents after unmapping it from the translation regime of each Borrower, if any Borrower including the caller sets this flag.– MBZ if the memory region was shared, else the Relayer must return <i>INVALID_PARAMETERS</i>.– MBZ if the Borrower has Read-only access to the memory region, else the Relayer must return <i>DENIED</i>.– Relayer must fulfill memory zeroing requirements listed in 5.12.4 Flags usage.

Field	Byte length	Byte offset	Description
			<ul style="list-style-type: none"> Bit[1]: Operation time slicing flag. <ul style="list-style-type: none"> This flag specifies if the Relayer can time slice this operation. <ul style="list-style-type: none"> b'0: Relayer must not time slice this operation. b'1: Relayer can time slice this operation. MBZ if the Relayer does not support time slicing of memory management operations (see 12.2.3 Time slicing of memory management operations). Bit[31:2]: Reserved (MBZ).
Endpoint count	4	12	<ul style="list-style-type: none"> Count of endpoints.
Endpoint array	–	16	<ul style="list-style-type: none"> Array of endpoint IDs. Each entry contains a 16-bit ID.

Table 11.26: FFA_MEM_RELINQUISH instances and conduits

Config No.	FF-A instance	Valid conduits
1	Secure and Non-secure physical	SMC, ERET
2	Secure and Non-secure virtual	SMC, HVC, SVC

Table 11.27: FFA_MEM_RELINQUISH function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none"> 0x84000076.
Other Parameter registers	w1-w7 x1-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 11.28: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS. DENIED. NO_MEMORY. ABORTED.

11.6.1 Component responsibilities for FFA_MEM_RELINQUISH

This interface is used by a Borrower endpoint to inform the Relayer that it is relinquishing access to a memory region that was lent or shared with it earlier. The memory region is identified by its *Handle*.

Transaction details are populated in the descriptor specified in [Table 11.25](#) as follows.

- The Handle and list of Borrower endpoints is populated in the descriptor described in [Table 11.25](#) in the TX buffer of the caller.

If the caller is a *proxy endpoint*, then the identity and count of the *Stream* endpoints on whose behalf it is relinquishing the memory region must be specified in the *Endpoint count* and *Endpoint array* fields in the descriptor.

If the caller is a *PE endpoint* Borrower, then it must specify its ID in the *Endpoint array* field in the descriptor.

- The caller could use the *Flags* field to request the Relayer to zero the memory region after it has been unmapped from its translation regime or time slice the unmapping operation.

Responsibilities of the:

- Borrower are listed in [11.6.1.1 Borrower responsibilities](#).
- Relayer are listed in [11.6.1.2 Relayer responsibilities](#).

Time slicing of this ABI invocation must be implemented by the Borrower and Relayer as specified in [12.2.3 Time slicing of memory management operations](#).

11.6.1.1 Borrower responsibilities

- Must ensure it has access to the memory region identified by the *Handle* parameter.
- Must ensure it is either the Borrower of the memory region or the proxy endpoint acting on behalf of one or more Stream endpoints who are the Borrowers instead.
- Must implement support for handling all error status codes that can be returned on completion of this interface.

11.6.1.2 Relayer responsibilities

- Must ensure that the *Handle* provided by the Borrower is valid and associated with a memory region it can access. Must return *INVALID_PARAMETERS* in case of an error.
- Must ensure that the *Flags* parameter is correctly encoded in the descriptor and the identities of Borrower endpoints are valid. Must return *INVALID_PARAMETERS* in case of an error.
- Must ensure that the *Endpoint count* field has a value > 0 . Must return *INVALID_PARAMETERS* in case of an error.
- Must ensure that the memory region is in the *!Owner-SA* or *!Owner-EA* state (see [Table 5.4](#)) for all Borrower endpoints specified by the caller. Must return *DENIED* in case of an error.

5. Must ensure that the memory region is unmapped from the translation regime of the Borrower (that is, it enters the *!Owner-NA* state (see [Table 5.4](#))) only if it has been relinquished as many times as it has been retrieved by the Borrower.

The memory region must be unmapped from the translation regime of the Borrower managed by the Relayer as specified in [5.3 Address translation regimes](#).

If the caller is a proxy endpoint for a Stream endpoint then the memory region must be unmapped from the stage 2 translation tables corresponding to the SEPID.

The Relayer must update internal state of the Borrower associated with the memory region to *!Owner-NA*.

The Relayer must return *NO_MEMORY* if there is not enough memory to unmap the memory region.

6. Must clear the contents of the memory region after unmapping it if *bit[0]* is set in the *Flags* parameter.

11.7 FFA_MEM_RECLAIM

Description

- Restores exclusive access to a memory region back to its Owner.
 - Valid FF-A instances and conduits are listed in [Table 11.30](#).
 - Syntax of this function is described in [Table 11.31](#).
 - Successful completion of this function is indicated through the invocation of the FFA_SUCCESS function by the callee.
 - Encoding of error code in the FFA_ERROR function is described in [Table 11.32](#).
-

Table 11.30: FFA_MEM_RECLAIM instances and conduits

Config No.	FF-A instance	Valid conduits
1	Secure and Non-secure physical	SMC, ERET
2	Secure and Non-secure virtual	SMC, HVC, SVC

Table 11.31: FFA_MEM_RECLAIM function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none">• 0x84000077.
uint64 Handle	w1/w2	<ul style="list-style-type: none">• Globally unique Handle to identify the memory region (see 5.10.2 Memory region handle).
uint32 Flags	w3	<ul style="list-style-type: none">• Bit[0]: Zero memory before reclaim flag.<ul style="list-style-type: none">– This flag specifies if the Relayer must clear memory region contents before mapping it in the Owner translation regime.<ul style="list-style-type: none">* b'0: Relayer must not zero the memory region contents.* b'1: Relayer must zero the memory region contents.– Relayer must fulfill memory zeroing requirements listed in 5.12.4 Flags usage.– MBZ if the Owner has Read-only access to the memory region, else the Relayer must return <i>DENIED</i>.

Parameter	Register	Value
		<ul style="list-style-type: none"> Bit[1]: Operation time slicing flag. <ul style="list-style-type: none"> This flag specifies if the Relayer can time slice this operation. <ul style="list-style-type: none"> b'0: Relayer must not time slice this operation. b'1: Relayer can time slice this operation. MBZ if the Relayer does not support time slicing of memory management operations (see 12.2.3 Time slicing of memory management operations). Bit[31:2]: Reserved (MBZ).
Other Parameter registers	w4-w7 x4-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 11.32: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS. DENIED. NO_MEMORY. ABORTED.

11.7.1 Component responsibilities for FFA_MEM_RECLAIM

This interface is used in the following ways.

- To complete a transaction to relinquish a memory region owned by the caller endpoint. Borrowers use the *FFA_MEM_RELINQUISH* interface to relinquish access to the memory region. The Owner uses this interface to reclaim exclusive access to the memory region.
- To abort an in-progress transaction to donate, lend or share a memory region owned by the caller endpoint. If any Receiver endpoint is unable to accept the transaction and the memory region is not mapped into the translation regime of any other Receiver endpoint, the Owner can use this transaction to reclaim exclusive access to the memory region.

Responsibilities of the:

- Owner are listed in [11.7.1.1 Owner responsibilities](#).
- Relayer are listed in [11.7.1.2 Relayer responsibilities](#).

Time slicing of this ABI invocation must be implemented by the Owner and Relayer as specified in [12.2.3 Time slicing of memory management operations](#).

11.7.1.1 Owner responsibilities

- Must ensure it is the Owner of the memory region identified by the *Handle* parameter.
- Must ensure that access to the memory region has been relinquished by all Borrowers.

3. Must implement support for handling all error status codes that can be returned on completion of this interface.

11.7.1.2 Relay responsibilities

1. Must ensure that the *Handle* provided by the Owner is valid and associated with a memory region it owns. Must return *INVALID_PARAMETERS* in case of an error.
2. Must ensure that the *Flags* parameter is correctly encoded. Must return *INVALID_PARAMETERS* in case of an error.
3. Must ensure that the memory region is in the *!Owner-NA* state (see [Table 5.4](#)) for all the Receiver endpoints associated with the memory region.

If one or more Borrowers are Stream endpoints associated with an *independent* peripheral device then in this case:

1. Each Borrower must relinquish access to the memory region through an IMPLEMENTATION DEFINED mechanism.
2. The Relay must unmap the memory region from the stage 2 translation tables identified by the SEPID.

Must return *DENIED* in case of an error.

4. Must clear the contents of the memory region if *bit[0]* is set in the *Flags* parameter.
5. If the state of the memory region for the Owner is *Owner-NA*, this implies that the region was lent. The Relay must map the memory region in the translation regime of the Owner as specified in [5.3 Address translation regimes](#).

The mapping must be created at the same address range and with the same memory region properties as those when the *FFA_MEM_LEND* interface was invoked.

Must return *NO_MEMORY* in case there is not enough memory to create the mapping in the Owner translation regime.

6. If the state of the memory region for the Owner is *Owner-SA* this implies that the region was shared. The Relay must map the memory region in the translation regime of the Owner as specified in [5.3 Address translation regimes](#).

The mapping must be created at the same address range and with the same memory region properties as those when the *FFA_MEM_SHARE* interface was invoked.

Must return *NO_MEMORY* in case there is not enough memory to change the mapping in the Owner translation regime.

7. If a VM is the Owner and the Borrower is an SP or SEPID associated with a Secure Stream ID, the Hypervisor must forward an invocation of this interface to the SPM.

This must be done by invoking this interface at the Non-secure physical FF-A instance with the same parameter values specified by the Owner at the Non-secure virtual FF-A instance.

8. If the call executes successfully, the state of the memory region for the Owner must transition to *Owner-EA*.

Chapter 12

Appendix

12.1 S-EL0 & User mode partitions

S-EL0 & Secure User mode partitions are used to achieve isolation among Secure services on Armv8.3 and earlier architecture versions. They could host one or more device drivers to control hardware that is only accessible from the Secure world. Normal world accesses these drivers through the message passing interfaces described in this specification. An example use case of S-EL0 partitions is described in [12.1.1 UEFI PI Standalone Management Mode partitions](#).

12.1.1 UEFI PI Standalone Management Mode partitions

Standalone management mode (STMM) is described in [9] as a processor architecture agnostic, sandboxed secure execution environment. It is meant to be used for device drivers that cannot be implemented in the OS kernel but are required during run-time.

On Armv8-A systems, STMM is implemented in a S-EL0 partition to constraint its visibility of the system address map and physical interrupts. This isolation enables a more robust Secure firmware implementation. This design is better from a security perspective than a design where STMM drivers are implemented in EL3.

Furthermore, execution in EL3 always runs to completion. Isolation of STMM drivers in an SP enables Secure firmware to transparently preempt them in response to OS Kernel interrupts and resume them once the interrupt has been handled. For some use cases, this prevents an adverse impact on OS responsiveness that could happen with a run to completion model.

12.1.1.1 FF-A usage to access STMM services

This section provides guidance around how services that would be typically implemented in EL3, can be implemented in multiple STMM S-EL0 partitions and accessed through FF-A interfaces. This guidance is based on certain assumptions about the Standalone management mode as follows.

- A STMM driver is neither reentrant nor thread safe but its single execution context can run on any PE in the system. Hence, a STMM S-EL0 partition is considered to be a *UP migrate capable* partition.
- STMM services are accessed from the UEFI runtime environment in the Normal world through direct Partition messages (see [4.4 Direct messaging usage](#)). A component called the MM communication driver is used for this purpose.
- STMM services can be accessed in response to an interrupt targeted to EL3 apart from the UEFI runtime environment.
- There are no dependencies between STMM partitions. One partition does not access services of another partition.
- A STMM partition processes one request at a time and is incapable of having multiple outstanding requests at any point of time.

The MM interface specification [10] specifies the **MM_COMMUNICATE** interface that enables the Normal world to access driver services implemented in a single STMM S-EL0 partition. FF-A interfaces can be used to access such services implemented in more than one STMM S-EL0 partitions as follows.

- **FFA_MSG_SEND_DIRECT_REQ** & **FFA_MSG_SEND_DIRECT_RESP** interfaces must be used to send and receive Partition messages directly.
- **FFA_RUN** interface must be used by the MM communication driver to resume a preempted STMM SP.
- A STMM SP can use any interface defined in [Chapter 6 Interface overview](#) apart from the following:
 - **FFA_MSG_POLL**.
 - **FFA_YIELD**.
 - **FFA_MSG_SEND**.

This implies that STMM SPs must neither use indirect messaging for communication nor relinquish execution to the primary scheduler.

Based on the assumptions about Standalone management mode, the following runtime models are applicable to a STMM partition.

- A STMM SP implements services that are invoked in response to interrupts targeted to the SPM. The execution of these services must run to completion.
- A STMM SP implements services that are invoked in response to FF-A interface invocations. The execution of these services could be required to:
 - Either run to completion.
 - Or allowed to be transparently preempted and resumed.

For simplicity and due to architectural limitations of S-EL0, a STMM SP that implements services that can be preempted must not implement services that are invoked through interrupts. Preemptible services must be implemented in separate STMM SPs. Each SP manifest must specify the expected runtime behavior when one of its services is invoked (see [3.2 Partition manifest at virtual FF-A instance](#)).

The SPM must enable run to completion semantics through an IMPLEMENTATION DEFINED mechanism. These semantics imply that once a service has been invoked, the STMM SP must run at a priority level such that:

- It cannot be preempted by a Non-secure interrupt.
- It cannot be preempted by a Secure interrupt of the same or lower priority.
- It can be transparently preempted and resumed by the SPM due to an interrupt of higher priority.

The SPM must also ensure that a STMM SP process only a single request at a time to preserve its non reentrant and non-thread safe run-time model. This is implicitly guaranteed if the request processing runs to completion.

Some example flows to illustrate common aspects of interaction with a STMM SP based on the preceding concepts are as follows.

- [Figure 12.1](#) describes how the MM communication driver can discover presence of STMM SPs and their properties. It is assumed that:
 - All STMM SPs share a MM service UUID. This UUID is used by MM communication driver to discover all the STMM SPs.
 - Each STMM SP specifies this UUID, its run-time model, memory regions, devices etc in its partition manifest.
 - The MM communication buffer for each STMM SP is allocated by the EFI MM communication driver.
- [Figure 12.2](#) describes how the MM communication driver and a STMM SP can communicate using direct Partition messages and the communication buffer shared between them.
- [Figure 12.3](#) describes how the STMM SP can be invoked in response to an interrupt.

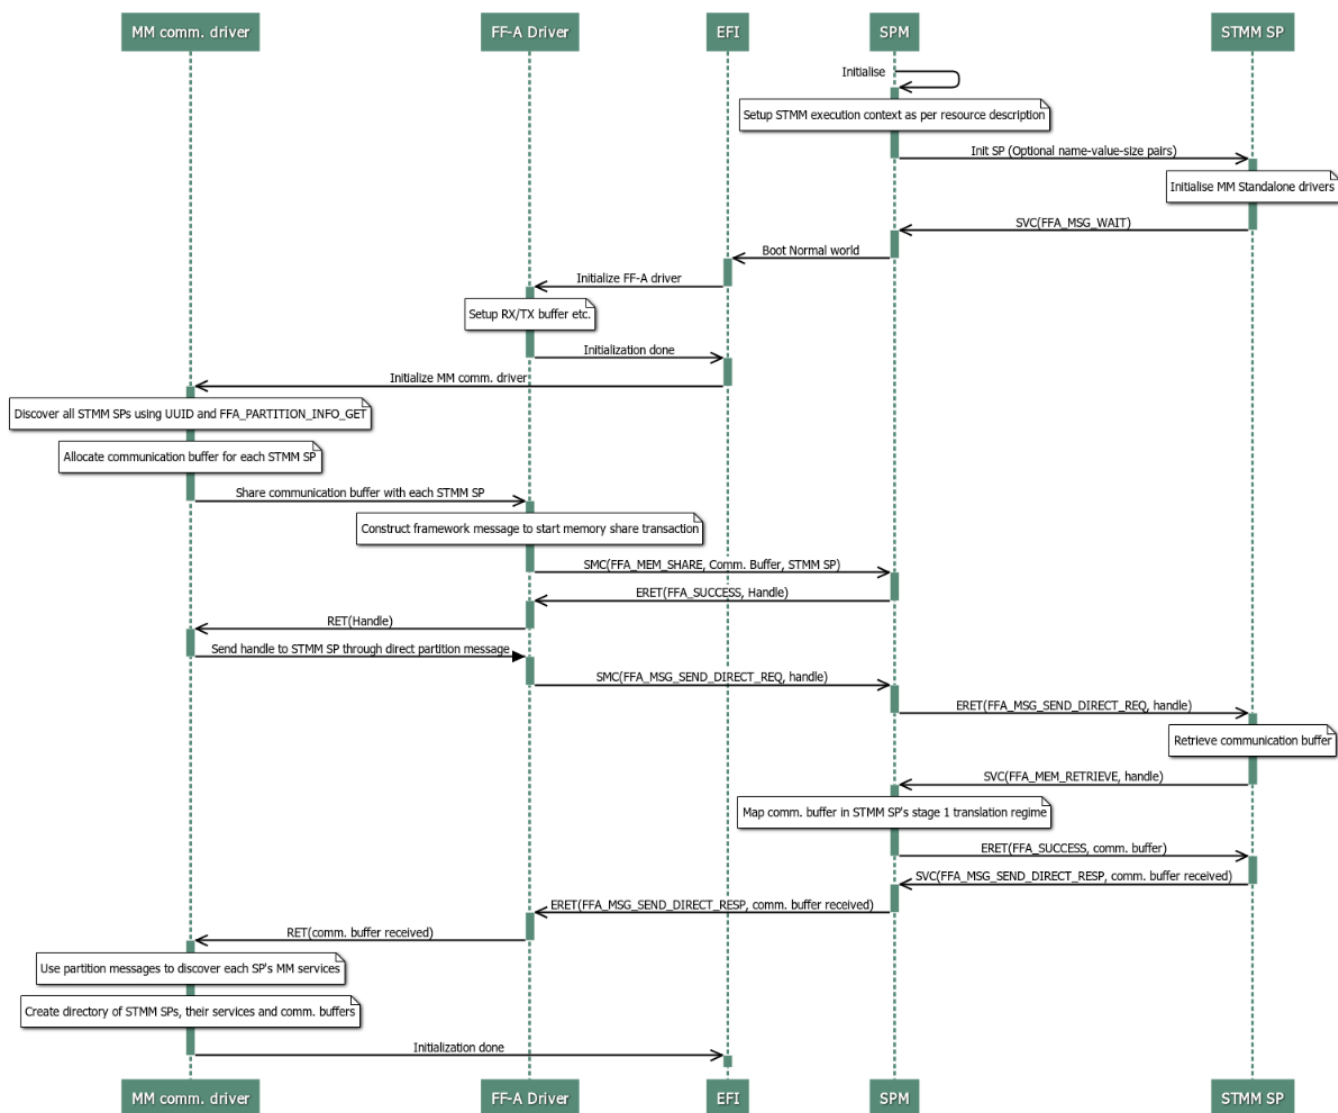


Figure 12.1: MM communication driver initialization

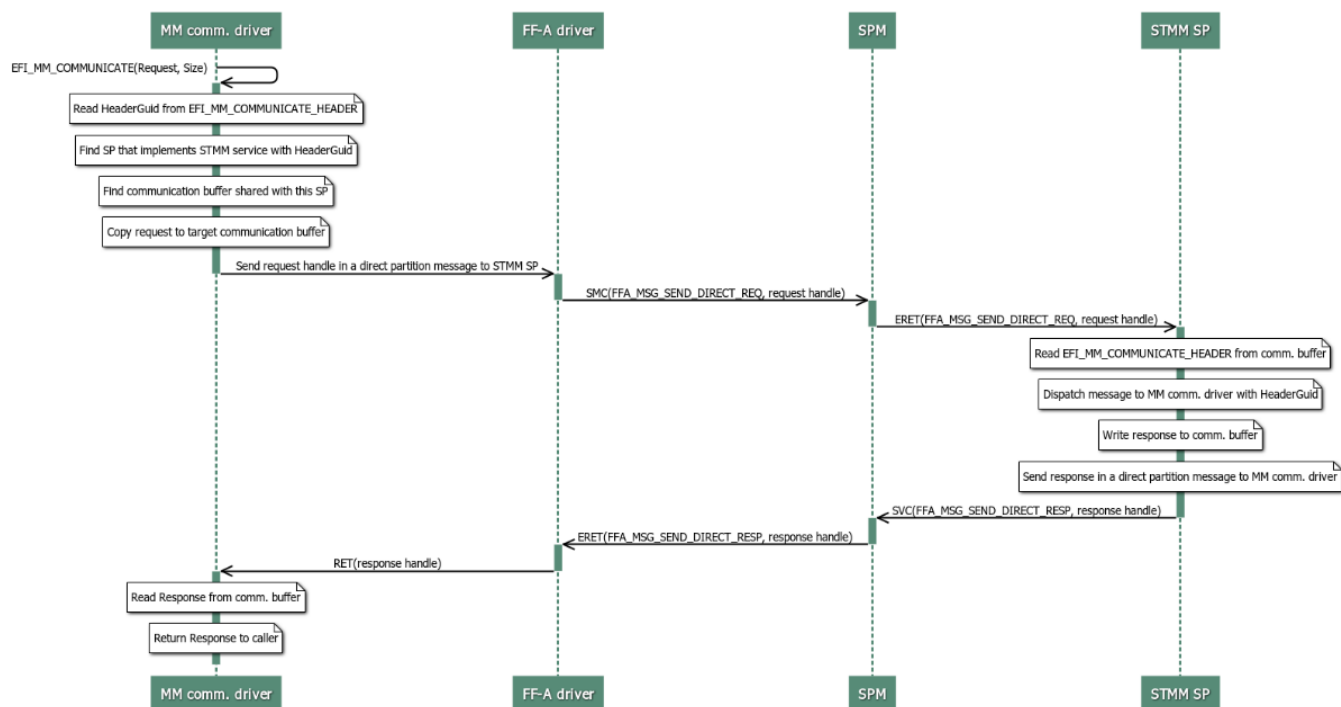


Figure 12.2: Message exchange between a STMM SP and MM communication driver




Figure 12.3: Invocation of a STMM SP in response to an interrupt

12.2 Additional memory management features

12.2.1 Transmission of transaction descriptor in dynamically allocated buffers

12.2.1.1 Rationale

The transaction descriptor (see [Table 5.19](#)) is transmitted from the caller to the callee in an invocation of the following ABIs.

- *FFA_MEM_DONATE*. See [11.1 FFA_MEM_DONATE](#).
- *FFA_MEM_LEND*. See [11.2 FFA_MEM_LEND](#).
- *FFA_MEM_SHARE*. See [11.3 FFA_MEM_SHARE](#).
- *FFA_MEM_RETRIEVE_REQ*. See [11.4 FFA_MEM_RETRIEVE_REQ](#).
- *FFA_MEM_RETRIEVE_RESP*. See [11.5 FFA_MEM_RETRIEVE_RESP](#).

This version of the Framework assumes that by default, the transaction descriptor is populated in the RX/TX buffers of an endpoint, Hypervisor or SPM as follows.

- The TX buffer is used to transmit the descriptor from an endpoint to the Hypervisor or SPM.
- The TX buffer is used to transmit the descriptor from the Hypervisor to the SPM.
- The RX buffer is used to transmit the descriptor from the Hypervisor or SPM to an endpoint.
- The RX buffer is used to transmit the descriptor from the SPM to the Hypervisor.

It is possible that the size of the descriptor is larger than the RX or TX buffer. For example, an endpoint memory access descriptor entry (see [Table 5.16](#)) in the transaction descriptor could reference one or more composite memory region descriptors (see [Table 5.13](#)). The total size of the composite memory region descriptors could be larger than the RX or TX buffer.

Each FF-A component is allowed to share only a single RX/TX pair with another FF-A component (see [4.2.2 RX/TX buffers](#)). It is possible that an endpoint or a partition manager cannot tolerate the latency in acquiring access to these buffers for a memory management operation on a busy system. It is also possible that other users cannot tolerate the latency to acquire access to these buffers due to an ongoing memory management operation.

12.2.1.2 Overview

This version of the Framework supports an optional feature that allows an endpoint to:

1. Dynamically allocate a separate buffer, instead of using the TX buffer, to transmit the transaction descriptor in an invocation of the following ABIs.
 - *FFA_MEM_DONATE*.
 - *FFA_MEM_LEND*.
 - *FFA_MEM_SHARE*.
 - *FFA_MEM_RETRIEVE_REQ*.
2. Use this buffer instead of the RX buffer to transmit the transaction descriptor in an invocation of the *FFA_MEM_RETRIEVE_RESP* ABI.

12.2.1.3 Description

The ability of an endpoint to use this feature depends on whether its partition manager implements support to map the dynamically allocated buffer into its translation regime. An endpoint can discover the availability of this support through the *FFA_FEATURES* interface (see [8.2 FFA_FEATURES](#)).

An endpoint must follow these rules while allocating a buffer dynamically.

- The dynamically allocated buffer must use the same attributes as RX/TX buffers that are specified in [4.2.2.3 Buffer attributes](#).
- The dynamically allocated buffer must be contiguous in the address space where it is allocated.

- The dynamically allocated buffer must fulfill the size and alignment requirements listed in [2.7 Memory granularity and alignment](#) to allow the partition manager to map it. The endpoint must discover these requirements by invoking the *FFA_FEATURES* interface with the function ID of the *FFA_RXTX_MAP* interface (see [8.2 FFA_FEATURES](#)).

The address and size of a dynamically allocated buffer must be specified in an invocation of the following ABIs.

- *FFA_MEM_DONATE*.
- *FFA_MEM_LEND*.
- *FFA_MEM_SHARE*.
- *FFA_MEM_RETRIEVE_REQ*.

The syntax for specifying the address and size is as follows.

- The *w3/x3* register must be used to specify the VA, IPA or PA of the dynamically allocated buffer.
A value of 0 must be specified to indicate that the TX buffer is being used.
- The *w4* register must be used to specify the size of the dynamically allocated buffer as a count of the contiguous 4K pages that constitute it.
A value of 0 must be specified if the TX buffer is being used.

In an invocation of the *FFA_MEM_RETRIEVE_RESP* ABI:

- The partition manager must use the same buffer that was used in the counterpart *FFA_MEM_RETRIEVE_REQ* ABI invocation.
- A value of 0 must be specified in the *w3/x3* register since there is no need to specify which buffer is being used.
- A value of 0 must be specified in the *w4* register since there is no need to specify the size of the buffer is being used.

If dynamically allocated buffers are supported, a partition manager must map the dynamically allocated buffer in its translation regime on invocation, and unmap it on completion of the following ABIs.

- *FFA_MEM_DONATE*.
- *FFA_MEM_LEND*.
- *FFA_MEM_SHARE*.

The buffer must be mapped by the partition manager on invocation of the *FFA_MEM_RETRIEVE_REQ* ABI. It must be unmapped after the complete transaction descriptor has been transmitted through the invocation of the counterpart *FFA_MEM_RETRIEVE_RESP* ABI.

A partition manager must return:

- *INVALID_PARAMETERS* in the following scenarios:
 - It does not support this feature and an endpoint attempts to use it as described above.
 - The address or size of the dynamically allocated buffer is invalid.
- *NO_MEMORY* if it does not have enough memory to map the dynamically allocated buffer in its translation regime.

12.2.2 Transmission of transaction descriptor in fragments

12.2.2.1 Rationale

The size of a memory transaction descriptor (see [Table 5.19](#)) could exceed the size of the buffer used by an endpoint to transmit it. This is possible in the following scenarios.

1. An endpoint or partition manager does not implement support for dynamically allocated buffers (see [12.2.1 Transmission of transaction descriptor in dynamically allocated buffers](#)). The RX/TX buffers must be used instead and cannot always accommodate the memory transaction descriptor.
2. An endpoint or partition manager do implement support for dynamically allocated buffers. In some memory management operations, the size of the memory transaction descriptor exceeds the size of the dynamically allocated buffer.

12.2.2.2 Overview

This version of the Framework supports an optional feature that:

- Allows the Sender of the transaction descriptor, to break the descriptor into equal or variable sized *fragments*, such that each fragment fits into the RX, TX or a dynamically allocated buffer.
- Adds support to transmit the first fragment of a transaction descriptor instead of the entire descriptor to the following ABIs.
 - `FFA_MEM_DONATE`. See [11.1 FFA_MEM_DONATE](#).
 - `FFA_MEM_LEND`. See [11.2 FFA_MEM_LEND](#).
 - `FFA_MEM_SHARE`. See [11.3 FFA_MEM_SHARE](#).
 - `FFA_MEM_RETRIEVE_REQ`. See [11.4 FFA_MEM_RETRIEVE_REQ](#).
 - `FFA_MEM_RETRIEVE_RESP`. See [11.5 FFA_MEM_RETRIEVE_RESP](#).
- Defines the following ABIs to transmit the remaining fragments from the Sender to the Receiver.
 - `FFA_MEM_FRAG_RX`. See [12.2.2.4 FFA_MEM_FRAG_RX](#).
 - `FFA_MEM_FRAG_TX`. See [12.2.2.5 FFA_MEM_FRAG_TX](#).

A Sender can invoke these interfaces as many times as there are fragments to transmit the complete descriptor to the Receiver.

12.2.2.3 Description

The ability of an endpoint to use this feature depends on whether its partition manager implements support for receipt and transmission of fragments of the memory transaction descriptor. An endpoint can discover the availability of this support through the `FFA_FEATURES` interface (see [8.2 FFA_FEATURES](#)). An endpoint must support this feature if its partition manager supports it.

It is strongly recommended that endpoint and partition manager implementations include support for this feature.

An endpoint and partition manager must implement the following protocol to use this feature.

1. An endpoint is the *Sender* and the partition manager is the *Receiver* of fragments in an invocation of the following ABIs.
 - `FFA_MEM_DONATE`.
 - `FFA_MEM_LEND`.
 - `FFA_MEM_SHARE`.
 - `FFA_MEM_RETRIEVE_REQ`.
2. The partition manager is the *Sender* and an endpoint is the *Receiver* of fragments in an invocation of the `FFA_MEM_RETRIEVE_RESP` ABI.
3. A *Sender* must use these ABIs to transmit the first fragment of the memory transaction descriptor as follows.
 - The `w2` register must be used to specify the length the first fragment.

- The buffer used to transmit the first fragment depends on the ABI being invoked as follows.
 - The *Sender* must either use its TX buffer or a dynamically allocated buffer (if supported by the *Receiver*) in an invocation of the following ABIs.
 - * *FFA_MEM_DONATE*.
 - * *FFA_MEM_LEND*.
 - * *FFA_MEM_SHARE*.
 - * *FFA_MEM_RETRIEVE_REQ*.
 - The buffer used by the *Sender* in an invocation of the *FFA_MEM_RETRIEVE_RESP* ABI must be one of the following.
 - * The RX buffer of the *Receiver* if it used its TX buffer in the earlier counterpart invocation of the *FFA_MEM_RETRIEVE_REQ* ABI.
 - * The dynamically allocated buffer that was used by the *Receiver* in the earlier counterpart invocation of the *FFA_MEM_RETRIEVE_REQ* ABI.
- A partition manager as the *Receiver* must return *INVALID_PARAMETERS* if it does not support this feature or the length of the fragment is invalid.
- 4. After receiving the first fragment, a *Receiver* must allocate a *Handle* (see [5.10.2 Memory region handle](#)) and use it to associate the remaining fragments with the current instance of the ABI invocation.

The same *Handle* must be used to identify the memory region description once all the fragments have been received.

- 5. A *Receiver* must request the *Sender* to transmit the next fragment through an invocation of the *FFA_MEM_FRAG_RX* ABI. See [12.2.2.4 FFA_MEM_FRAG_RX](#) for a description of this ABI and its parameters.

The *Receiver* must use this interface to request retransmission of a fragment as well. This could happen if it was unable to receive the previous fragment due to an IMPLEMENTATION DEFINED reason.

The *Receiver* must populate the *w4* parameter register at a physical FF-A instance as follows.

1. With the endpoint ID of the *Owner* of the memory region, if the fragment is being transmitted in response to the following ABI invocations.
 - *FFA_MEM_DONATE*.
 - *FFA_MEM_LEND*.
 - *FFA_MEM_SHARE*.
 - *FFA_MEM_RETRIEVE_REQ* on behalf of the Hypervisor see [11.4.3 Support for retrieval by the Hypervisor](#).
 - *FFA_MEM_RETRIEVE_RESP* on behalf of the Hypervisor see [11.4.3 Support for retrieval by the Hypervisor](#).
2. With the endpoint ID of the *Borrower* of the memory region, if the fragment is being transmitted in response to the following ABI invocations.
 - *FFA_MEM_RETRIEVE_REQ* by the Hypervisor on behalf of a Borrower VM.
 - *FFA_MEM_RETRIEVE_RESP* by the Hypervisor on behalf of a Borrower VM.
6. A *Sender* must transmit the next fragment to the *Receiver* through an invocation of the *FFA_MEM_FRAG_TX* ABI. See [12.2.2.5 FFA_MEM_FRAG_TX](#) for a description of this ABI and its parameters.

The buffer used to transmit the fragment must be the same as the one used to transmit the first fragment.

The *Sender* must populate the *w4* parameter register at a physical FF-A instance with the endpoint ID that was populated in the same register in the counterpart invocation of *FFA_MEM_FRAG_RX* by the *Receiver*.

7. A *Receiver* must acknowledge receipt of the final fragment. It must do this by completing the invocation of the ABI that was invoked to transmit the first fragment. For example, *FFA_MEM_SHARE* must be completed with the *FFA_SUCCESS* function as described in [11.3 FFA_MEM_SHARE](#).
8. A *Receiver* could abort the memory management operation while transmission of fragments is in-progress due to IMPLEMENTATION DEFINED reasons. The operation is identified by the ABI used to transmit the first fragment. The invocation of this ABI must be completed to signal to the *Sender* that the operation has been aborted.

The mechanism to do this depends on the type of *Receiver* and the FF-A instance it resides at as follows.

- The *Receiver* is a partition manager at a virtual FF-A instance. It must invoke the *FFA_ERROR* function with the *ABORTED* error code.
- The *Receiver* is a partition manager at a physical FF-A instance. It must invoke the *FFA_ERROR* function with the *ABORTED* error code.
- The *Receiver* is an endpoint at a virtual FF-A instance. In this version of the Framework, this scenario is possible only during the invocation of the *FFA_MEM_RETRIEVE_RESP* ABI. The *Receiver* must invoke the *FFA_MEM_RELINQUISH* ABI (see [11.6 FFA_MEM_RELINQUISH](#)) to abort the operation.

In all cases, the *Receiver* must restore any globally observable state associated with the memory region described by the transaction descriptor to what it was prior to receipt of the first fragment.

In all cases, the *Sender* must restore any globally observable state associated with the memory region described by the transaction descriptor to what it was prior to transmission of the first fragment.

9. A *Sender* at a virtual FF-A instance must not abort the memory management operation while transmission of fragments is in-progress.

The Hypervisor could abort a operation as the *Sender* at the Non-secure physical FF-A instance. It must invoke the *FFA_ERROR* function with the *ABORTED* error code to do this.

[Figure 12.4](#) illustrates an example where the *FFA_MEM_RETRIEVE_REQ* and *FFA_MEM_RETRIEVE_RESP* interfaces are used to retrieve a transaction descriptor in fragments at a virtual FF-A instance. The following assumptions have been made.

- The memory region is shared with only a single Borrower.
- The RX/TX buffers of the Borrower are used by these interfaces.
- In invocations of both *FFA_MEM_RETRIEVE_REQ* and *FFA_MEM_RETRIEVE_RESP*, the descriptor in [Table 5.19](#) is split into two fragments to be delivered to the Relayer and Borrower respectively.
- In invocations of both *FFA_MEM_RETRIEVE_REQ* and *FFA_MEM_RETRIEVE_RESP*, only parameters relevant to fragment transmission have been illustrated.

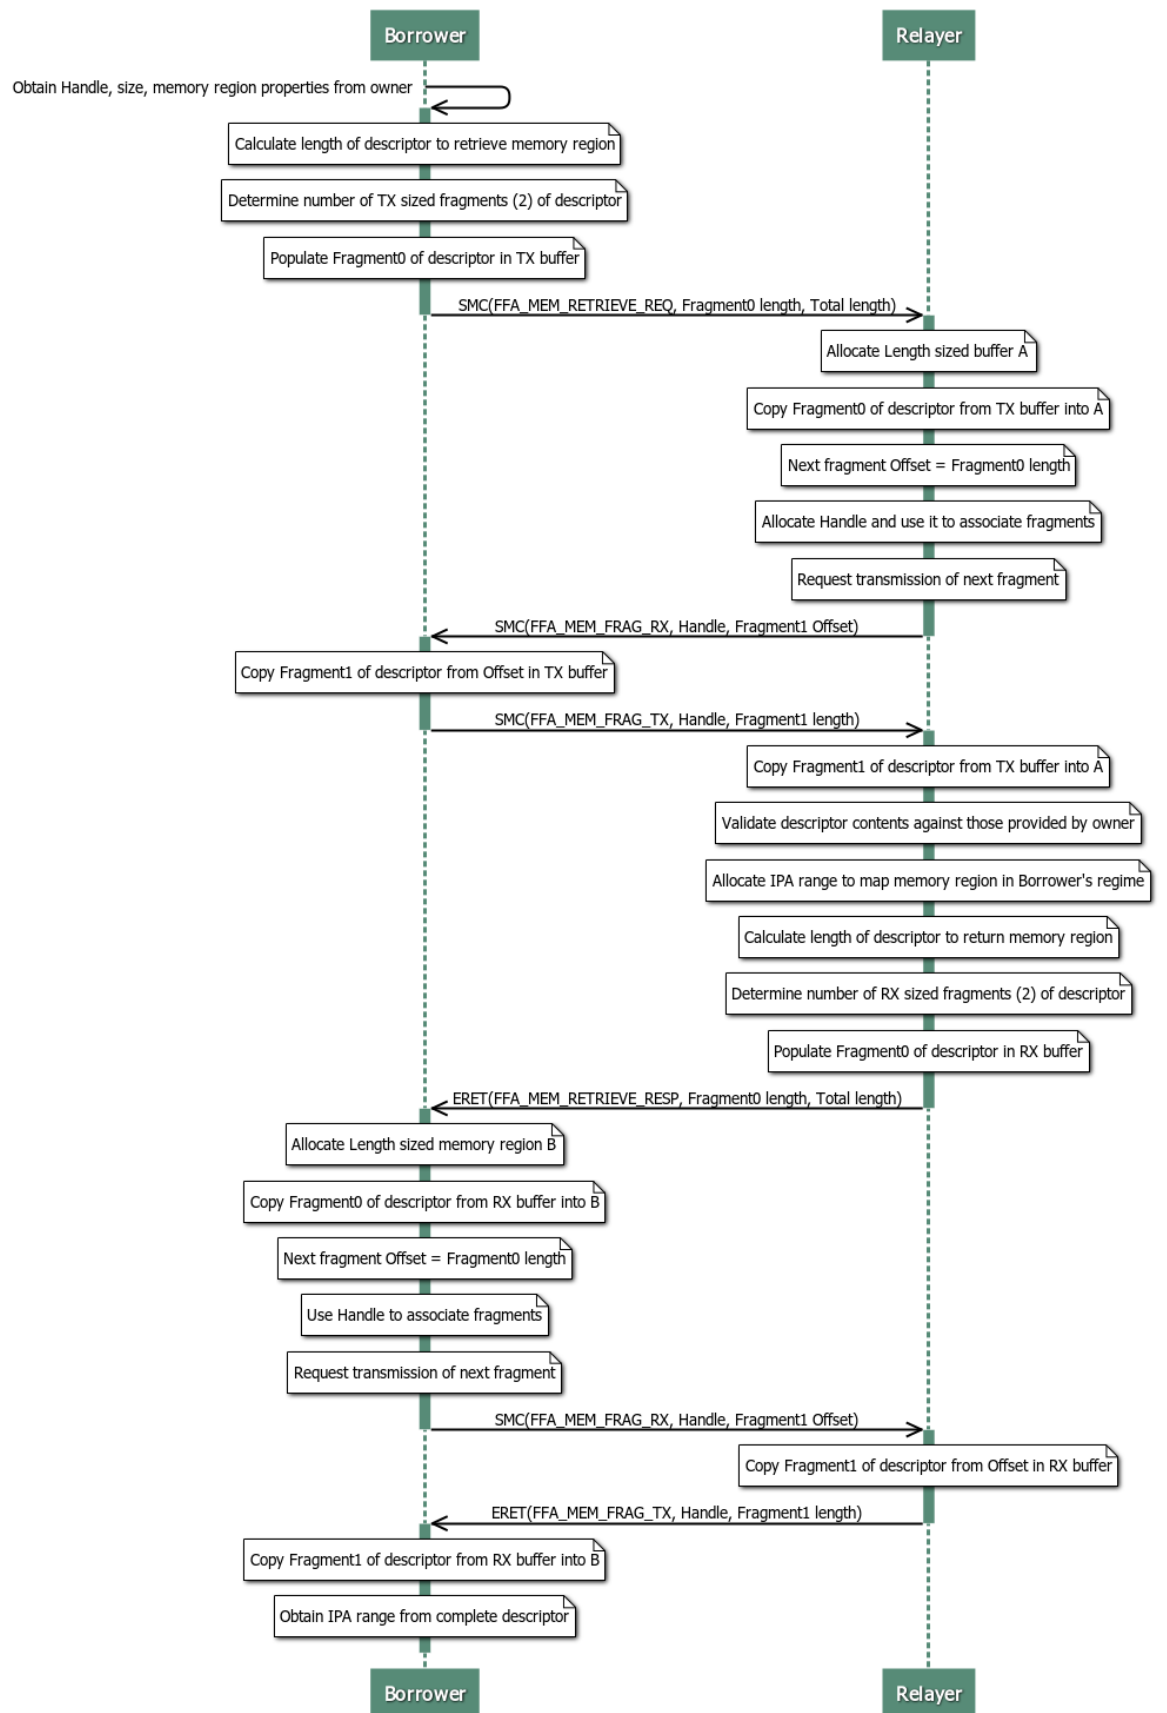


Figure 12.4: Example of fragment transmission while retrieving memory
Copyright © 2020 Arm Limited or its affiliates. All rights reserved.
Non-confidential

12.2.2.4 FFA_MEM_FRAG_RX

Description

- A caller uses this interface to request the callee to transmit the next fragment of the memory transaction descriptor.
- Valid FF-A instances and conduits are listed in [Table 12.2](#).
- Syntax of this function is described in [Table 12.3](#).
- Successful completion of this function is indicated by an invocation of the *FFA_MEM_FRAG_TX* function (see [12.2.2.5 FFA_MEM_FRAG_TX](#)).
- Encoding of error code in the *FFA_ERROR* function is described in [Table 12.4](#).

Table 12.2: FFA_MEM_FRAG_RX instances and conduits

Config No.	FF-A instance	Valid conduits
1	Secure and Non-secure physical	SMC, ERET
2	Secure and Non-secure virtual	SMC, HVC, SVC, ERET

Table 12.3: FFA_MEM_FRAG_RX function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none"> • 0x8400007A.
uint64 Handle	w1/w2	<ul style="list-style-type: none"> • <i>Handle</i> value to associate the fragment with the transaction descriptor of the ongoing memory management transaction with the callee.
uint32 Fragment offset	w3	<ul style="list-style-type: none"> • Byte offset from where the next fragment to be transmitted must start. • Offset must be calculated from the base of the transaction descriptor being transmitted . • Offset must be equal to one of the following: <ul style="list-style-type: none"> – The number of bytes of the transaction descriptor transmitted prior to the invocation of this interface. – The offset used in the previous invocation of this interface. This allows the Sender to retransmit the previous fragment if the Receiver could not receive it due to an IMPLEMENTATION DEFINED reason.
uint32 Endpoint ID	w4	<ul style="list-style-type: none"> • ID of the Owner or Borrower endpoint. <ul style="list-style-type: none"> – Bit[31:16]: Endpoint ID. – Bit[15:0]: Reserved (MBZ). • Reserved (MBZ) at any virtual FF-A instance.

Parameter	Register	Value
Other parameter registers	w5-w7 w5-x7	<ul style="list-style-type: none"> Reserved (MBZ).

Table 12.4: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> INVALID_PARAMETERS: Invalid Handle, fragment offset or endpoint ID value. NOT_SUPPORTED: This function is not implemented at this FF-A instance. ABORTED. Sender aborted transmission of fragments.

12.2.2.5 FFA_MEM_FRAG_TX

Description

- A caller uses this interface to transmit the next fragment of the transaction descriptor to the callee.
- Valid FF-A instances and conduits are listed in [Table 12.6](#).
- Syntax of this function is described in [Table 12.7](#).
- Successful completion of this function is indicated by an invocation of the *FFA_MEM_FRAG_RX* function (see [12.2.2.4 FFA_MEM_FRAG_RX](#)).
- Encoding of error code in the *FFA_ERROR* function is described in [Table 12.8](#).

Table 12.6: FFA_MEM_FRAG_TX instances and conduits

Config No.	FF-A instance	Valid conduits
1	Secure and Non-secure physical	SMC, ERET
2	Secure and Non-secure virtual	SMC, HVC, SVC, ERET

Table 12.7: FFA_MEM_FRAG_TX function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none"> • 0x8400007B.
uint64 Handle	w1/w2	<ul style="list-style-type: none"> • <i>Handle</i> value to associate the fragment with the transaction descriptor of the ongoing memory management transaction with the callee.
uint32 Fragment length	w3	<ul style="list-style-type: none"> • Length of the fragment being transmitted.
uint32 Endpoint ID	w4	<ul style="list-style-type: none"> • ID of the Owner or Borrower endpoint. <ul style="list-style-type: none"> – Bit[31:16]: Endpoint ID. – Bit[15:0]: Reserved (MBZ). • Reserved (MBZ) at any virtual FF-A instance.
Other parameter registers	w5-w7 w5-x7	<ul style="list-style-type: none"> • Reserved (MBZ).

Table 12.8: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none">• INVALID_PARAMETERS: Invalid Handle, fragment length or endpoint ID value.• NOT_SUPPORTED: This function is not implemented at this FF-A instance.• ABORTED. Receiver aborted transmission of fragments.

12.2.3 Time slicing of memory management operations

12.2.3.1 Rationale

In each FF-A memory management ABI as follows, the partition manager is responsible for mapping or unmapping a memory region from the translation regime of an endpoint that invokes the ABI.

- *FFA_MEM_DONATE*. This interface is described in [11.1 FFA_MEM_DONATE](#).
- *FFA_MEM_LEND*. This interface is described in [11.2 FFA_MEM_LEND](#).
- *FFA_MEM_SHARE*. This interface is described in [11.3 FFA_MEM_SHARE](#).
- *FFA_MEM_RETRIEVE_REQ*. This interface is described in [11.4 FFA_MEM_RETRIEVE_REQ](#).
- *FFA_MEM_RELINQUISH*. This interface is described in [11.6 FFA_MEM_RELINQUISH](#).
- *FFA_MEM_RECLAIM*. This interface is described in [11.7 FFA_MEM_RECLAIM](#).

The duration of a mapping or unmapping operation on a set of translation tables depends on factors such as the size of the memory region, number of translation table entries it requires, number of cache and TLB maintenance operations etc. The operation runs to completion in the partition manager. This could prevent progress of the endpoint that requested the operation. In some scenarios, an endpoint might not be able to tolerate this delay for example, if it is prevented from processing its pending interrupts.

12.2.3.2 Overview

This version of the Framework supports an optional feature that allows the partition manager to divide the translation table operations into *time slices*.

An operation runs for the duration of a time slice. Once the time slice is over, the partition manager relinquishes control back to the endpoint. The endpoint resumes the operation later. On resumption the partition manager runs the operation for another time slice. The process repeats itself until the operation completes. The duration of a time slice and its discovery by a partition manager is IMPLEMENTATION DEFINED.

This optional feature enables both the endpoint and the partition manager to make progress during a long running memory management ABI invocation.

12.2.3.3 Description

The ability of an endpoint to use this feature depends on whether its partition manager implements support for time-slicing memory management ABI invocations. An endpoint can discover the availability of this support through the *FFA_FEATURES* interface (see [8.2 FFA_FEATURES](#)). This feature is only available to EL1 and S-EL1 endpoints.

An endpoint and its partition manager must implement the following protocol to use this feature.

1. An endpoint must request its partition manager to use time-slicing by setting the *Operation time slicing* flag in the *Flags* field as follows:
 - In the transaction descriptor (see [5.12.4 Flags usage](#)) in an invocation of the following ABIs.
 - *FFA_MEM_DONATE*.
 - *FFA_MEM_LEND*.
 - *FFA_MEM_SHARE*.
 - *FFA_MEM_RETRIEVE_REQ*.
 - In [Table 11.25](#) in an invocation of the *FFA_MEM_RELINQUISH* ABI.
 - In *w3* in an invocation of the *FFA_MEM_RECLAIM* ABI.
 - A partition manager must return *INVALID_PARAMETERS* if an endpoint sets the *Operation time slicing* flag and it does not support this feature.
2. A partition manager must divide the translation table operations required by the invoked ABI, if their duration is expected to exceed the time slice duration.

This must be done only after the partition manager has received the entire transaction descriptor in an invocation of the following ABIs.

- FFA_MEM_DONATE.
- FFA_MEM_LEND.
- FFA_MEM_SHARE.
- FFA_MEM_RETRIEVE_REQ.

Once the time slice duration expires, the partition manager must:

- Save enough state to resume the ABI invocation at a later point of time and not prevent progress of other FF-A functions before the paused ABI invocation is resumed.
 - Cater for the scenario where the ABI invocation is paused on one PE and resumed by the endpoint on another PE.
 - Use the *Handle* (see 5.10.2 *Memory region handle*) to identify the current instance of the ABI invocation when it is resumed later. A new *Handle* must be allocated if this was not done previously.
 - Invoke the *FFA_MEM_OP_PAUSE* interface (see 12.2.3.4 *FFA_MEM_OP_PAUSE*) to inform the endpoint that the current ABI invocation has been paused and must be resumed later.
3. An endpoint must use the *FFA_MEM_OP_RESUME* interface (see 12.2.3.5 *FFA_MEM_OP_RESUME*) to resume the paused ABI invocation identified by the *Handle*.

The endpoint could invoke other FF-A functions before resuming the paused ABI invocation.

4. A partition manager could abort the ABI invocation when it is resumed later due to IMPLEMENTATION DEFINED reasons. It must signal to the endpoint that the ABI invocation has been aborted by invoking the *FFA_ERROR* function with the *ABORTED* error code.

Figure 12.5 illustrates an example where the *FFA_MEM_RETRIEVE_REQ* and *FFA_MEM_RETRIEVE_RESP* interfaces are used by an endpoint to retrieve a transaction descriptor at a virtual FF-A instance. The following assumptions have been made.

- The operation to map the memory region in the translation regime of the endpoint is expected to take longer than the time slice value known to the partition manager.
- The endpoint has requested time slicing by setting the *Operation time slicing* flag.
- The partition manager divides the *FFA_MEM_RETRIEVE_REQ* function invocation into 3 time slices through the *FFA_MEM_OP_PAUSE* and *FFA_MEM_OP_RESUME* interfaces.

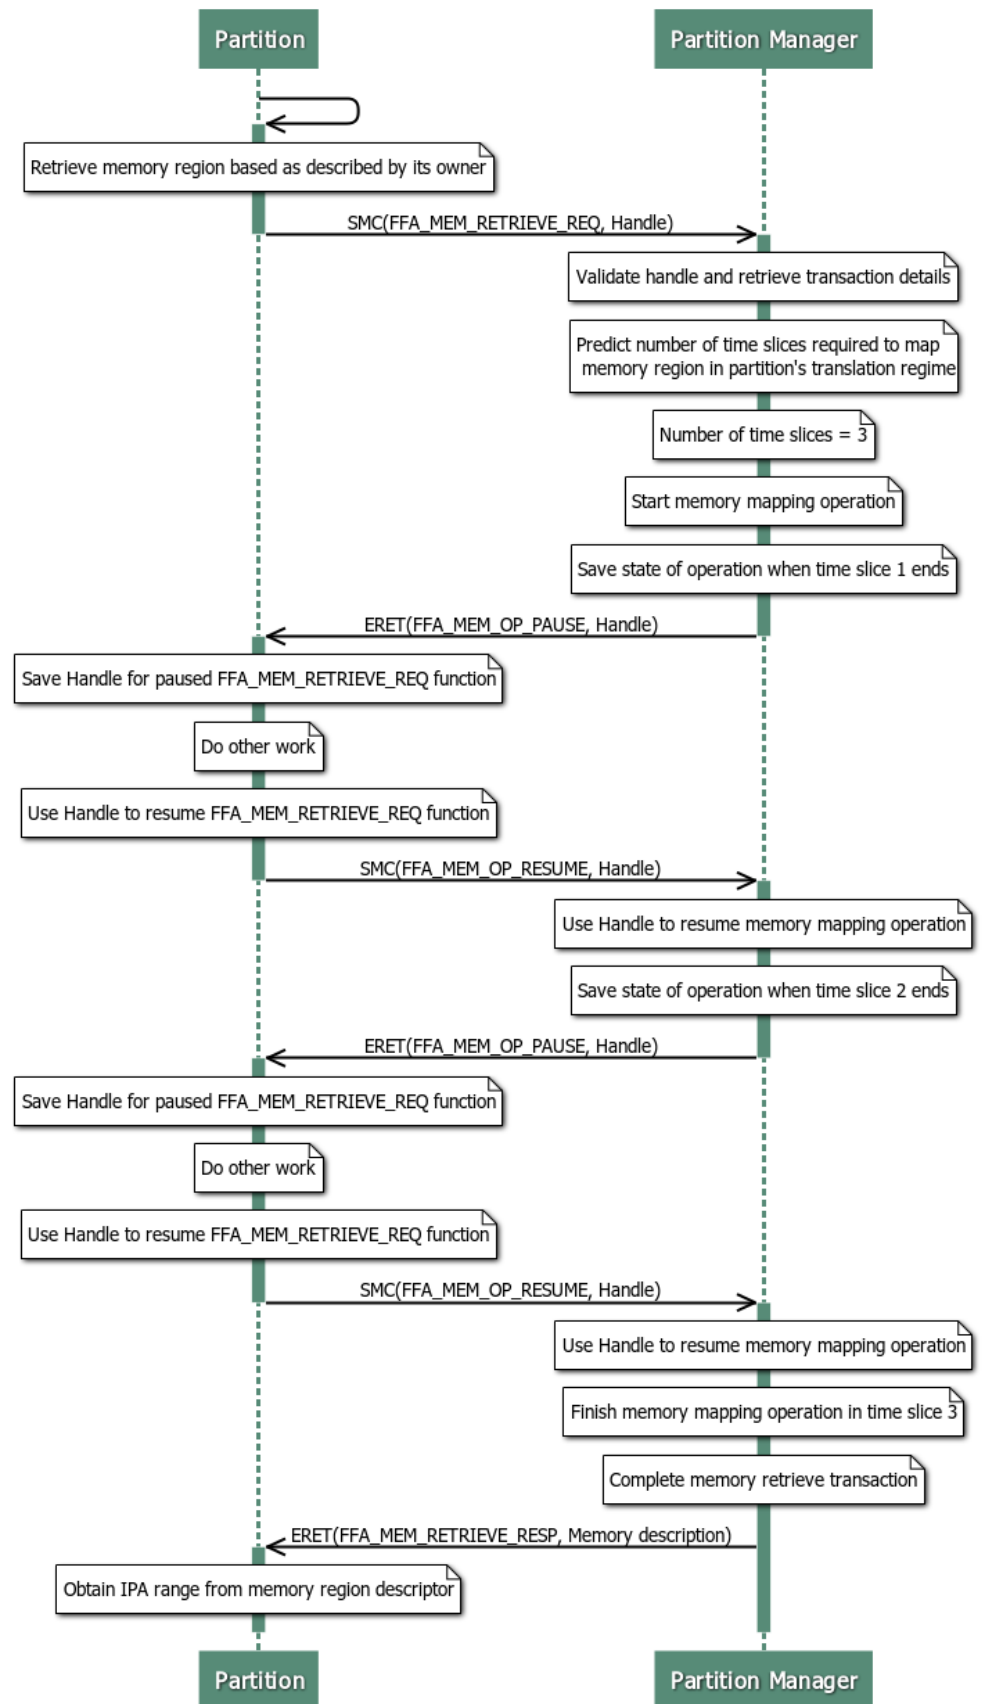


Figure 12.5: Example of time slicing during FFA_MEM_RETRIEVE_REQ
Copyright © 2020 Arm Limited or its affiliates. All rights reserved.
Non-confidential

12.2.3.4 FFA_MEM_OP_PAUSE

Description

- A partition manager uses this interface to pause the execution of a memory management ABI invoked by an endpoint. Execution is returned to the endpoint.
- Valid FF-A instances and conduits are listed in [Table 12.10](#).
- Syntax of this function is described in [Table 12.11](#).
- Encoding of error code in the FFA_ERROR function is described in [Table 12.12](#).

Table 12.10: FFA_MEM_OP_PAUSE instances and conduits

Config No.	FF-A instance	Valid conduits
1	Non-secure physical	ERET
2	Secure physical	SMC
3	Secure and Non-secure virtual	ERET

Table 12.11: FFA_MEM_OP_PAUSE function syntax

Parameter	Register	Value
uint32 Function ID	w0	<ul style="list-style-type: none"> • 0x84000078.
uint64 Handle	w1/w2	<ul style="list-style-type: none"> • <i>Handle</i> value to identify the paused memory management operation.
Other parameter registers	w3-w7 x3-x7	<ul style="list-style-type: none"> • Reserved (MBZ).

Table 12.12: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> • INVALID_PARAMETERS: Invalid handle value. • NOT_SUPPORTED: This function is not implemented at this FF-A instance.

12.2.3.5 FFA_MEM_OP_RESUME

Description

- An endpoint uses this interface to request the partition manager to resume execution of a paused memory management ABI. The paused operation is identified by the supplied *Handle*.
- Valid FF-A instances and conduits are listed in [Table 12.14](#).
- Syntax of this function is described in [Table 12.15](#).
- Encoding of error code in the FFA_ERROR function is described in [Table 12.16](#).

Table 12.14: FFA_MEM_OP_RESUME instances and conduits

Config No.	FF-A instance	Valid conduits
1	Non-secure physical	SMC
2	Secure physical	ERET
3	Secure and Non-secure virtual	SMC, HVC, SVC

Table 12.15: FFA_MEM_OP_RESUME function syntax

Parameter	Register	Value
uint32 Function ID	w0	• 0x84000079.
uint64 Handle	w1/w2	• <i>Handle</i> value to identify the paused memory management operation.
Other parameter registers	w3-w7 x3-x7	• Reserved (MBZ).

Table 12.16: FFA_ERROR encoding

Parameter	Register	Value
int32 Error code	w2	<ul style="list-style-type: none"> • INVALID_PARAMETERS: Invalid Handle value. • NOT_SUPPORTED: This function is not implemented at this FF-A instance.

Terms and abbreviations

ABI	Application Binary Interface
DMA	Direct Memory Access
DSP	Digital Signal Processor
FF-A	Firmware Framework for A-profile
GIC	Generic Interrupt Controller
HVC	Hypervisor Call
MBP	Must be preserved
MBZ	Must be zero
MM	Management Mode
MMIO	Memory Mapped Input Output
MP	Multi-processing
OS	Operating System
PE	Processing Element
PPI	Private Peripheral Interrupt
PSA	Platform Security Architecture
SGI	Software Generated Interrupt
SMC	

	Secure Monitor Call
SMCCC	
	SMC Calling Convention
SMMU	
	System Memory Management Unit
SP	
	Secure Partition
SPCI	
	Secure Partition Client Interface
SPI	
	Shared Peripheral Interrupt
SPM	
	Secure Partition Manager
SPRT	
	Secure Partition Run Time
STMM	
	Standalone Management Mode
SVC	
	Supervisor Call
TCB	
	Trusted Computing Base
TEE	
	Trusted Execution Environment
UUID	
	Unique Universal Identifier
VCPU	
	Virtual CPU
VHE	
	Virtualization Host Extensions
VM	
	Virtual Machine
VMSA	
	Virtual Memory System Architecture