Arm® Corstone™ SSE-710 Subsystem

Revision: r0p0

Technical Reference Manual

Issue 02
102342_0000_02_en

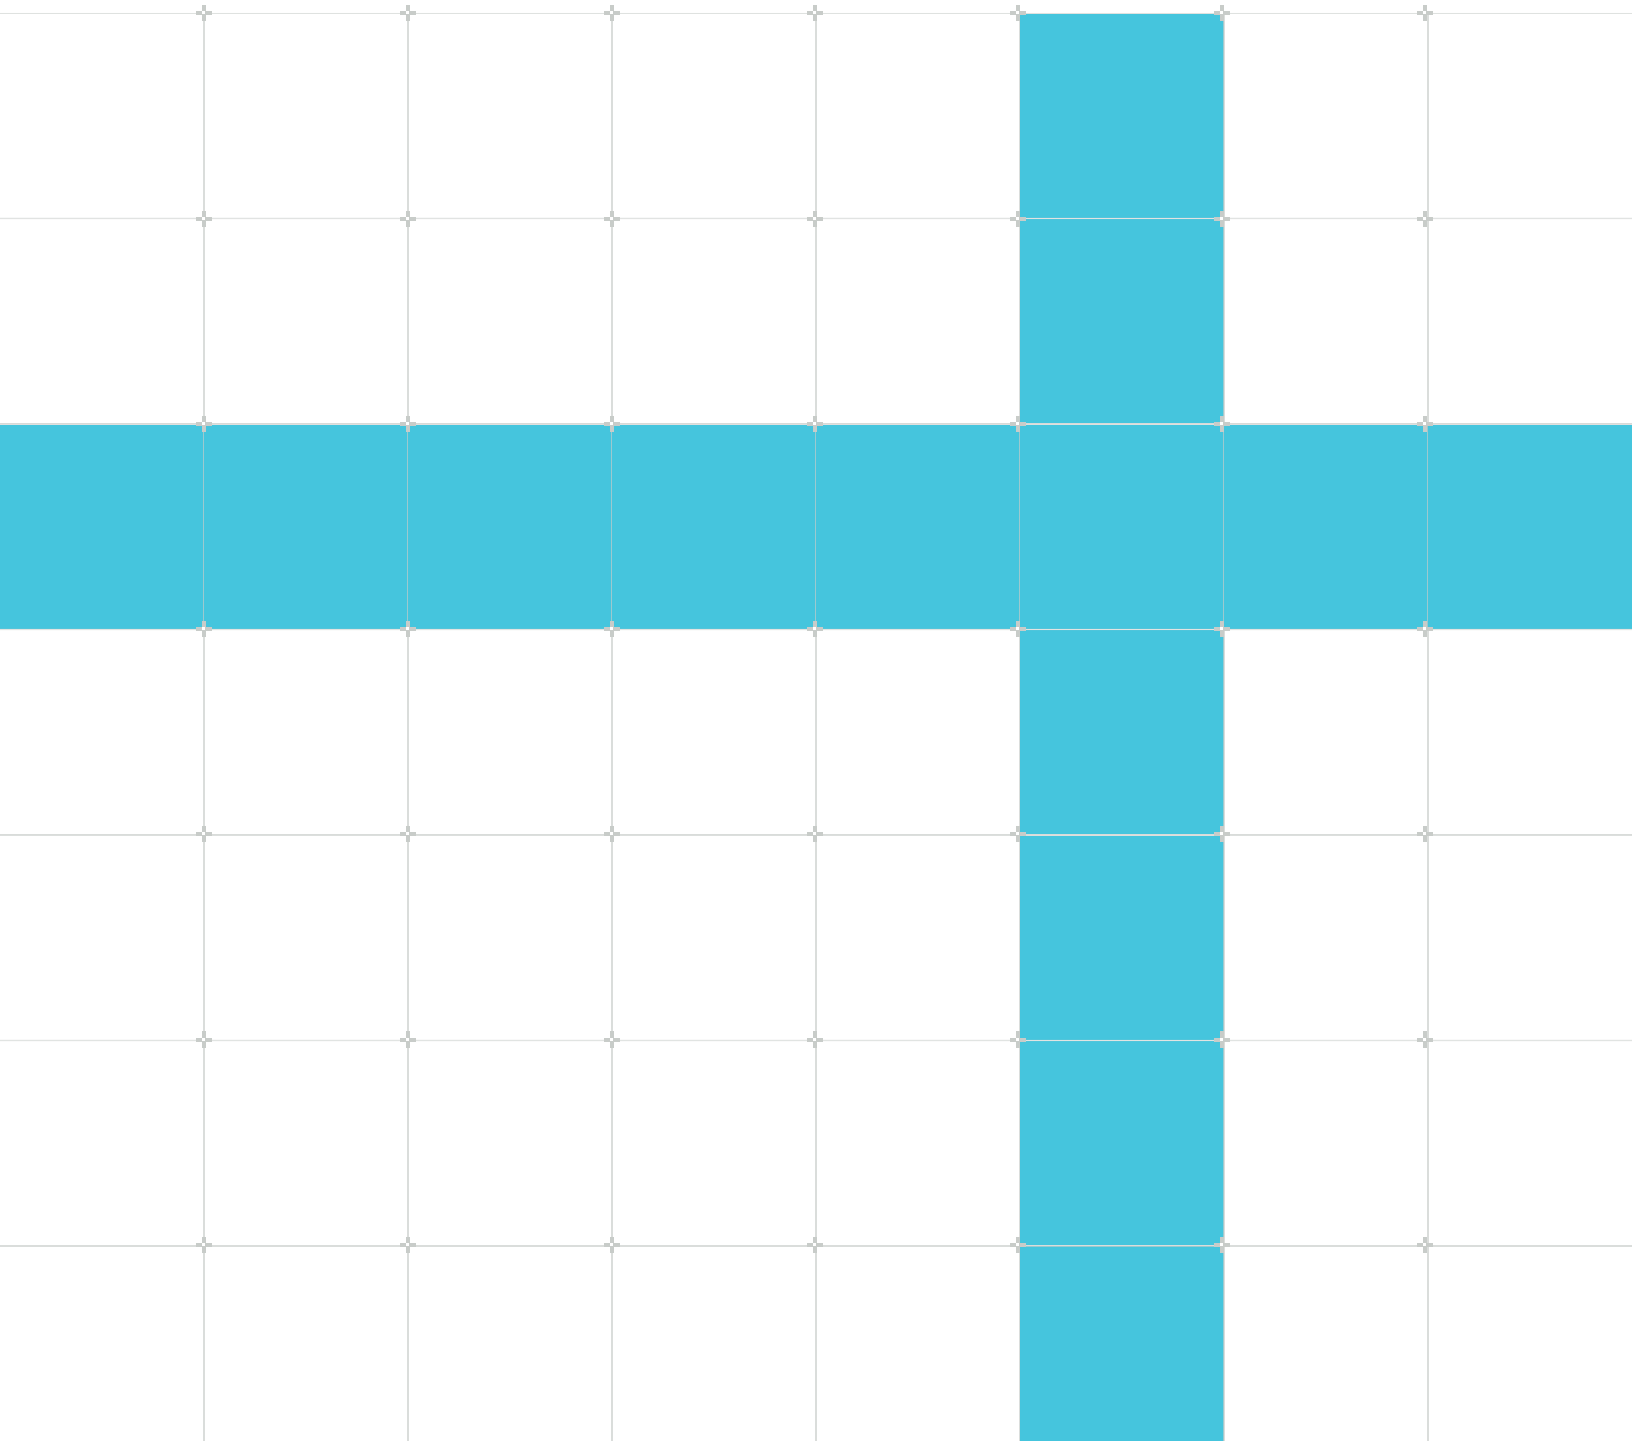

Arm® Corstone™ SSE-710 Subsystem

Technical Reference Manual

Release Information

Document history

Issue	Date	Confidentiality	Change
0000-01	28 April 2021	Non-Confidential	First release for r0p0 EAC
0000-02	15 June 2022	Non-Confidential	First release for r0p0 REL

Proprietary Notice

Arm Limited. Company 02557590 registered in England.

110 Fulbourn Road, Cambridge, England CB1 9NJ.

(LES-PRE-20349)

Confidentiality Status

This document is Non-Confidential. The right to use, copy and disclose this document may be subject to license restrictions in accordance with the terms of the agreement entered into by Arm and the party that Arm delivered this document to.

Unrestricted Access is an Arm internal classification.

Product Status

The information in this document is Final, that is for a developed product.

Feedback

Arm welcomes feedback on this product and its documentation. To provide feedback on the product, create a ticket on https://support.developer.arm.com.

To provide feedback on the document, fill the following survey: https://developer.arm.com/documentation-feedback-survey.

Inclusive language commitment

Arm values inclusive communities. Arm recognizes that we and our industry have used language that can be offensive. Arm strives to lead the industry and create change.

This document includes language that can be offensive. We will replace this language in a future issue of this document.

To report offensive language in this document, email terms@arm.com.

Contents

1. Introduction

1.1 Product revision status

The r*x*p*y* identifier indicates the revision status of the product described in this manual, for example, r*1*p*2*, where:

r*x* Identifies the major revision of the product, for example, r1.

p*y* Identifies the minor revision or modification status of the product, for example, p2.

1.2 Intended audience

This book is written for system designers, system integrators, and programmers who are designing or programming a *System-on-Chip* (SoC) that uses the Corstone™ SSE-710 Subsystem.

1.3 Conventions

The following subsections describe conventions used in Arm documents.

Glossary

The Arm® Glossary is a list of terms used in Arm documentation, together with definitions for those terms. The Arm Glossary does not contain terms that are industry standard unless the Arm meaning differs from the generally accepted meaning.

See the Arm Glossary for more information: developer.arm.com/glossary.

Typographic conventions

Convention	Use
italic	Citations.
bold	Interface elements, such as menu names. Signal names. Terms in descriptive lists, where appropriate.
`monospace`	Text that you can enter at the keyboard, such as commands, file and program names, and source code.
`monospace bold`	Language keywords when used outside example code.

Convention	Use
`monospace underline`	A permitted abbreviation for a command or option. You can enter the underlined text instead of the full command or option name.
`<and>`	Encloses replaceable terms for assembler syntax where they appear in code or code fragments. For example: `MRC p15, 0, <Rd>, <CRn>, <CRm>, <Opcode_2>`
SMALL CAPITALS	Terms that have specific technical meanings as defined in the *Arm® Glossary*. For example, **IMPLEMENTATION DEFINED**, **IMPLEMENTATION SPECIFIC**, **UNKNOWN**, and **UNPREDICTABLE**.

Recommendations. Not following these recommendations might lead to system failure or damage.

Requirements for the system. Not following these requirements might result in system failure or damage.

Requirements for the system. Not following these requirements will result in system failure or damage.

An important piece of information that needs your attention.

A useful tip that might make it easier, better or faster to perform a task.

A reminder of something important that relates to the information you are reading.

Timing diagrams

The following figure explains the components used in timing diagrams. Variations, when they occur, have clear labels. You must not assume any timing information that is not explicit in the diagrams.

Shaded bus and signal areas are undefined, so the bus or signal can assume any value within the shaded area at that time. The actual level is unimportant and does not affect normal operation.

Figure 1-1: Key to timing diagram conventions

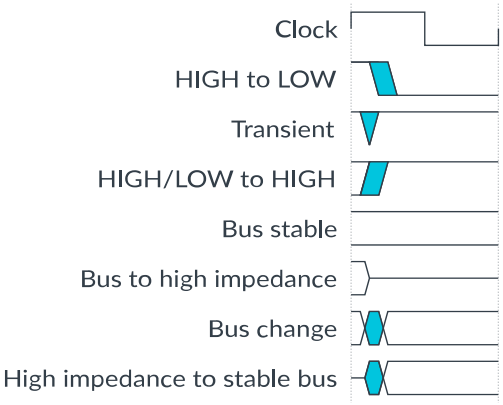

Signals

The signal conventions are:

Signal level

The level of an asserted signal depends on whether the signal is active-HIGH or active-LOW. Asserted means:

- HIGH for active-HIGH signals.
- LOW for active-LOW signals.

Lowercase n

At the start or end of a signal name, n denotes an active-LOW signal.

1.4 Additional reading

This document contains information that is specific to this product. See the following documents for other relevant information:

Table 1-2: Arm Publications

Document name	Document ID	Licensee only Y/N
AMBA® APB Protocol Specification Version 2.0	IHI 0024	N
AMBA® 4 ATB Protocol Specification ATBv1.0 and ATBv1.1	IHI 0032	N
AMBA® 5 AHB Protocol Specification	IHI 0033	N
AMBA® AXI and ACE Protocol Specification	IHI 0022	N

Document name	Document ID	Licensee only Y/N
AMBA® 4 AXI4-Stream Protocol Specification	IHI 0051	N
AMBA® Low Power Interface Specification, Arm® Q-Channel and P-Channel Interfaces	IHI 0068	N
CoreLink™ GIC-400 Generic Interrupt Controller Technical Reference Manual	DDI 0471	N
Arm® CoreLink™ NIC-450 Network Interconnect Technical Overview	100459	N
Arm® CoreLink™ PCK-600 Power Control Kit Technical Reference Manual	101150	N
Arm® CoreSight™ SDC-600 Secure Debug Channel Technical Reference Manual	101130	N
Arm® CoreSight™ System-on-Chip SoC-600 Technical Reference Manual	100806	N
Arm® CoreSight™ STM-500 System Trace Macrocell Technical Reference Manual	DDI 0528	N
Arm® CoreSight™ Architecture Specification v3.0	IHI 0029	N
Cortex®-M0+ Technical Reference Manual	DDI 0484	N
Arm® Cortex®-M System Design Kit Technical Reference Manual	DDI 0479	N
Arm® Cortex®-A32 Processor Technical Reference Manual	100241	N
Arm® Cortex®-A35 Processor Technical Reference Manual	100236	N
Arm® Cortex®-A53 MPCore Processor Technical Reference Manual	DDI 0500	N
Arm® Debug Interface Architecture Specification ADIv6.0	IHI 0074	N
Arm® Generic Interrupt Controller Architecture Specification, architecture version 2.0	IHI 0048	N
Arm® Power Policy Unit Architecture Specification, version 1.1	DEN 0051	N
Arm® Server Base System Architecture 5.0	DEN 0029	N
Arm®v6-M Architecture Reference Manual	DDI 0419	N
Arm®v7-M Architecture Reference Manual	DDI 0403	N
Arm® Architecture Reference Manual Armv8, for Armv8-A architecture profile	DDI 0487	N
PrimeCell UART (PL011) Technical Reference Manual	DDI 0183	N
Arm® Corstone™ SSE-710 Subsystem Configuration and Integration Manual	102343	Y
Arm® CoreSight™ System-on-Chip SoC-600 Configuration and Integration Manual	100807	Y
Arm® Cortex®-A32 Processor Configuration and Sign-off Guide	100244	Y
Arm® Cortex®-A32 Processor Integration Manual	100245	Y
Arm® Cortex®-A35 Processor Integration Manual	100240	Y
Arm® Cortex®-A35 Processor Configuration and Sign-off Guide	100239	Y
Arm® Cortex®-A53 MPCore Processor Integration Manual	DIT 0036	Y
Arm® Cortex®-A53 MPCore Processor Configuration and Sign-off Guide	DII 0281	Y
Cortex®-M0+ Integration and Implementation Manual	DIT 0032	Y
Arm® Debug Interface Architecture Specification ADIv6.0	IHI 0074	Y
Arm® Power Control System Architecture Specification	DEN 0050	Y

Other Publications

None

Arm tests its PDFs only in Adobe Acrobat and Acrobat Reader. Arm cannot guarantee the quality of its documents when used with any other PDF reader.

Adobe PDF reader products can be downloaded at http://www.adobe.com

2. Overview

This chapter introduces the Corstone™ SSE-710 Subsystem (SSE-710).

2.1 About SSE-710 subsystem

The SSE-710 subsystem (SSE-710) is a flexible subsystem designed to provide a secure solution for rich *Internet of Things* (IoT) applications based on Arm® supported Cortex®-A processors, Cortex®-M, or other managers that are present in External Systems.

SSE-710 is designed to cover a range of *Power, Performance, and Area* (PPA) applications, and enable extension for use-case specific applications, for example, sensors, cloud connectivity, and edge compute.

Connected embedded devices are exposed to many different security threats. SSE-710 implements the reference subsystem for an SoC that targets secure, rich IoT applications. Built-in security is one of its key features.

SSE-710 consists of:

- A reference subsystem
- An example integration layer
- An example Cortex®-A32, Cortex®-A35, Cortex®-A53 (the supported Cortex®-A processors for SSE-710), and Arm® GIC-400 integration.
- A set of documentation including PSA L2 certification Guidance document for Secure Enclave.

To create a SoC, the SSE-710 subsystem must be extended. A complete system typically contains the following components:

Compute subsystem

In SSE-710 the compute subsystem consists of one to four supported Cortex®-A processors, processor cores and associated bus, debug, controller, peripherals, and interface logic.

Memory and peripherals

An extended SoC requires extra memory and peripheral components beyond the minimum subsystem components. Flash memory, for example, is not provided with the SSE-710, but can be added through implemented interfaces.

Sensors and actuators

The reference design can be extended by adding sensors or actuator logic, such as temperature input or motor control output.

Software development environment

Arm® provides a complete software development environment, which includes the Arm® Mbed™ *Operating System* (OS) and Linux, Mbed™ or GNU (GCC) compilers and debuggers, and firmware.

Custom peripherals typically require corresponding third-party firmware that can be integrated into the software stack.

2.2 Features of SSE-710

SSE-710 provides the minimum set of features for an SoC.

You can configure some aspects of the SSE-710 design so that you can meet the specific requirements of your SoC and intended application. For example, you can configure the number of shared interrupt inputs or Host System firewalls.

SSE-710 standardizes the following product features:

- Memory and interrupt maps of the Host System and Secure Enclave system.
- Hardware address compartmentalization.
- Boot and security, including a hardware isolated *Root of Trust* (RoT).
- Communication between different parts of the SoC.
- Timer and watchdog infrastructure.
- Debug requirements.
- Power, clock, and reset control.
- Requirements for adding new or existing External Systems for use-case specific applications.
- Requirements for adding use-case specific managers and peripherals to the Host System.

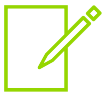
Note

For more information on the topology of SSE-710 and the different types of systems, such as Host System, Secure Enclave system, and External System, see 2.3 SSE-710 topology on page 22.

In this document, the term *integrator* refers to the person who integrates the SSE-710 design into an SoC. The integrator implements the rest of the SoC following the requirements made by SSE-710.

The SSE-710 includes the following key components:

- A Cortex®-A processor cluster and *Generic Interrupt Controller* (GIC) socket to support up to four Cortex®-A cores, and a GIC-400.
- A Secure Enclave socket into which you can integrate a Crypto Accelerator. The socket is based on a Cortex®-M0 processor with dedicated SRAM, ROM, and peripherals, such as timers and watchdogs.
- Secure Enclave contains a Crypto Accelerator socket with Arm® CryptoCell™-312 pre-integrated, but other cryptographic engines can also be integrated.
- Two External System Harnesses support integration of External Systems into SSE-710.

- *Message Handling Units* (MHUs) for communication between the different systems in the SSE-710. An MHU provides a unidirectional communication channel, therefore MHUs are provided in pairs to allow for bidirectional communication:
  - Two pairs of MHUs for communication between supported Cortex®-A processors and Secure Enclave.
  - Two pairs of MHUs for communication between supported Cortex®-A processors and each External System.
  - Two pairs of MHUs for communication between Secure Enclave and each External System.
- Interrupt Router to handle routing interrupts from shared peripherals to the interrupt controller of a specific system
- Firewall to provide:
  - Hardware compartmentalization of the Host System address space
  - Translation between the address space of the Secure Enclave system and the External System into the address space of the Host System
- Common debug infrastructure supporting single and multi-system debug, by self-hosted and external debug agents.
- Certificate-based debug authentication is supported by CoreSight™ SDC-600.
- Advanced hardware autonomous power control design.
- Shared peripherals, such as counters, timers, and watchdog, that any system in the SSE-710 can use.

The integrator can tailor the final implementation to the target use-cases by using:

- The supported configuration options, see 3. Configuration options on page 30.
- The sockets for the Host processor, Host GIC, and External Systems.
- The expansion interfaces of SSE-710.

2.3 SSE-710 topology

SSE-710 design consists of several components:

- A Host System.
- A Secure Enclave system.
- Two External System harnesses.

Figure 2-1: SSE-710 topology on page 23 shows the topology of the SSE-710. An SSE-710-based SoC is formed of three different types of system:

- Secure Enclave: green box.
- External Systems: red box.
- Host System: the rest of the SSE-710 subsystem.

The Secure Enclave socket and External System harnesses are connected, with all other master and peripherals added by the integrator, to the interconnect inside the Host System. Each system provides different functionality to the overall SoC. They operate as self-contained entities, independent of the other systems in the SoC.

Figure 2-1: SSE-710 topology

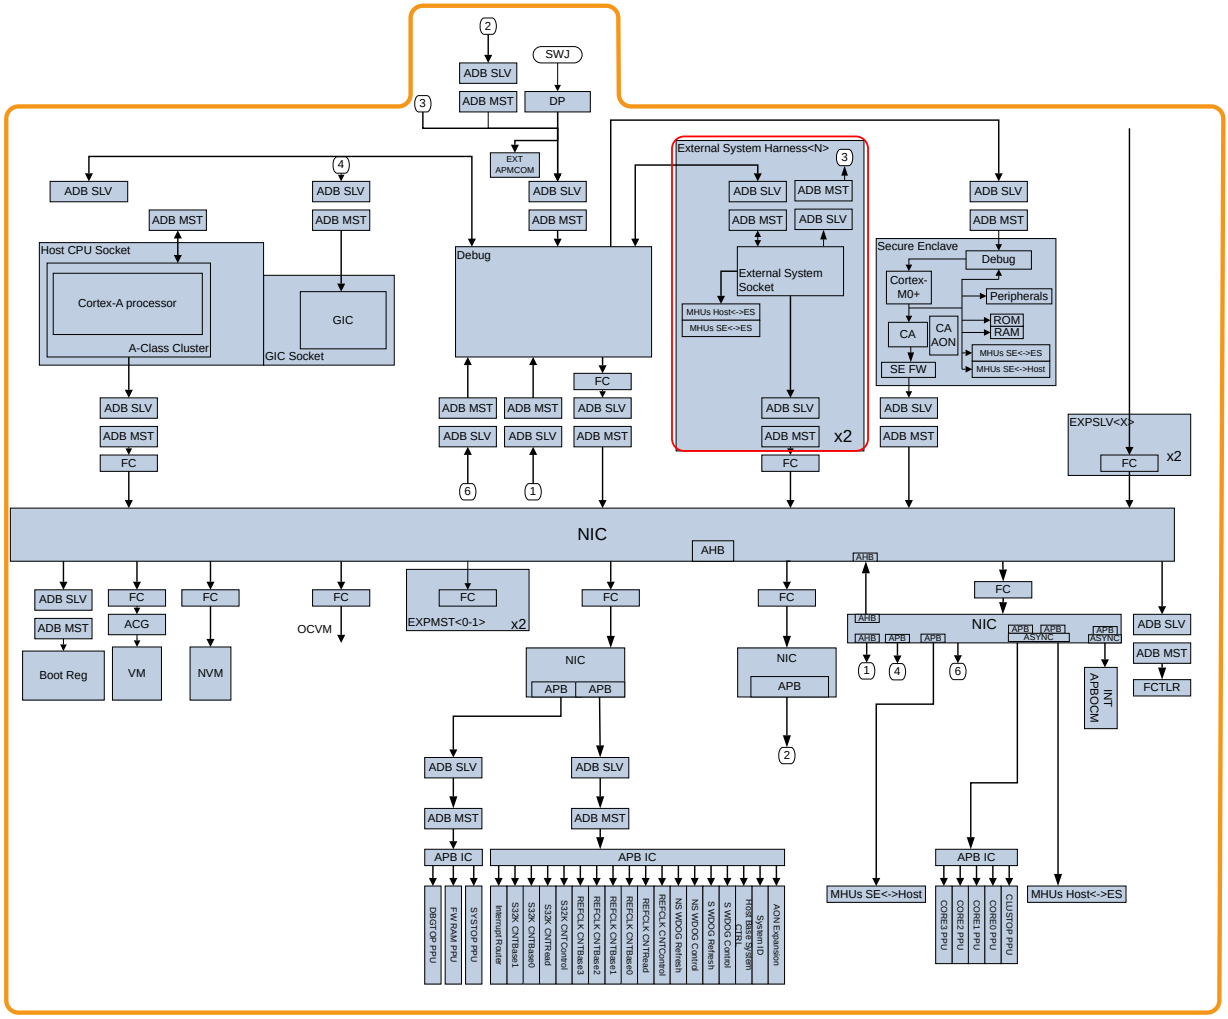

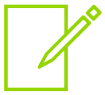
Note

In Figure 2-1: SSE-710 topology on page 23, the Secure Enclave and External Systems are not part of the Host System.

2.3.1 Host System

The Host System is based on an Arm® A-profile processor with standard peripherals that enable booting a Rich OS.

The Host System provides access to resources that can be shared with the Secure Enclave or External System, such as volatile and non-volatile memory. The Host System includes a firewall that provides hardware compartmentalization of the Host System address space. This compartmentalization enables software to create sandboxes between the different systems, and to allow only specific regions to be used for communication between the systems.

The Host System provides:

- Common timestamps for use with Generic Timer and Watchdog as defined in the *Arm® Architecture Reference Manual Armv8, for Armv8-A architecture profile*.
- UARTs.
- MHUs for communication with the Secure Enclave and External Systems.
- Routing of interrupts, from shared peripherals, to the interrupt controller associated with the software entity controlling the peripheral.
- Shared secured debug infrastructure with support for:
  - Trace
  - Self-hosted and External debug of all systems

There is no self-hosted debug on the Secure Enclave because Cortex®-M0+ does not support it.

  - Functional I/O debug
  - Multi-system debug

The Host System defines a Host processor and *Generic Interrupt Controller* (GIC) Socket. These are the interfaces that allow integration of the following supported A-profile processors or GICs:

- Cortex®-A32, Cortex®-A35, or Cortex®-A53 (the supported Cortex®-A processors for SSE-710)
- CoreLink™ GIC-400

2.3.2 Secure Enclave

The Secure Enclave within the SSE-710, is a Cortex®-M0+ based security subsystem that acts as the root-of-trust for the system.

The Secure Enclave holds and generates keys, and provides cryptographic services and security controls to the Host System. For example:

- Authenticating the firmware of the Secure Enclave itself, the Host System and the External System
- Enabling/disabling debug capabilities based on the secure state of the device

Hardware isolates the software running on the Secure Enclave system. This reduces the complexity in security reviews. Communication with the Secure Enclave system is achieved using MHUs.

At power-on reset, it is the first system to boot and performs initial configuration of its own system and other components of the SSE-710, such as the Host System firewall.

The Secure Enclave consists of:

- A Cortex®-M0+ Processor, with an in-built *Nested Vectored Interrupt Controller* (NVIC)
- Dedicated ROM and SRAM
- PL011 UART
- *Message Handling Units* (MHUs) for communication with other systems.
- Two CMSDK Timers
- Two CMSDK watchdogs:
  - Secure Enclave watchdog
  - SoC watchdog
- Secure Enclave Base System Control registers controlling the SSE-710 subsystem
- Secure Enclave System Control registers controlling the Secure Enclave
- Independent Clock and Power Control infrastructure
- *Security Control Bits* (SCB) controlling security access across the SSE-710 subsystem
- A firewall that permits the Secure Enclave to access any location in the Host System address space
- A Crypto Accelerator socket

For more information about the firewall used in Secure Enclave, see section 9.1.6.10 Secure Enclave Firewall on page 156.

The following figure is a high-level block diagram of the components within the Secure Enclave.

Figure 2-2: Secure Enclave block diagram

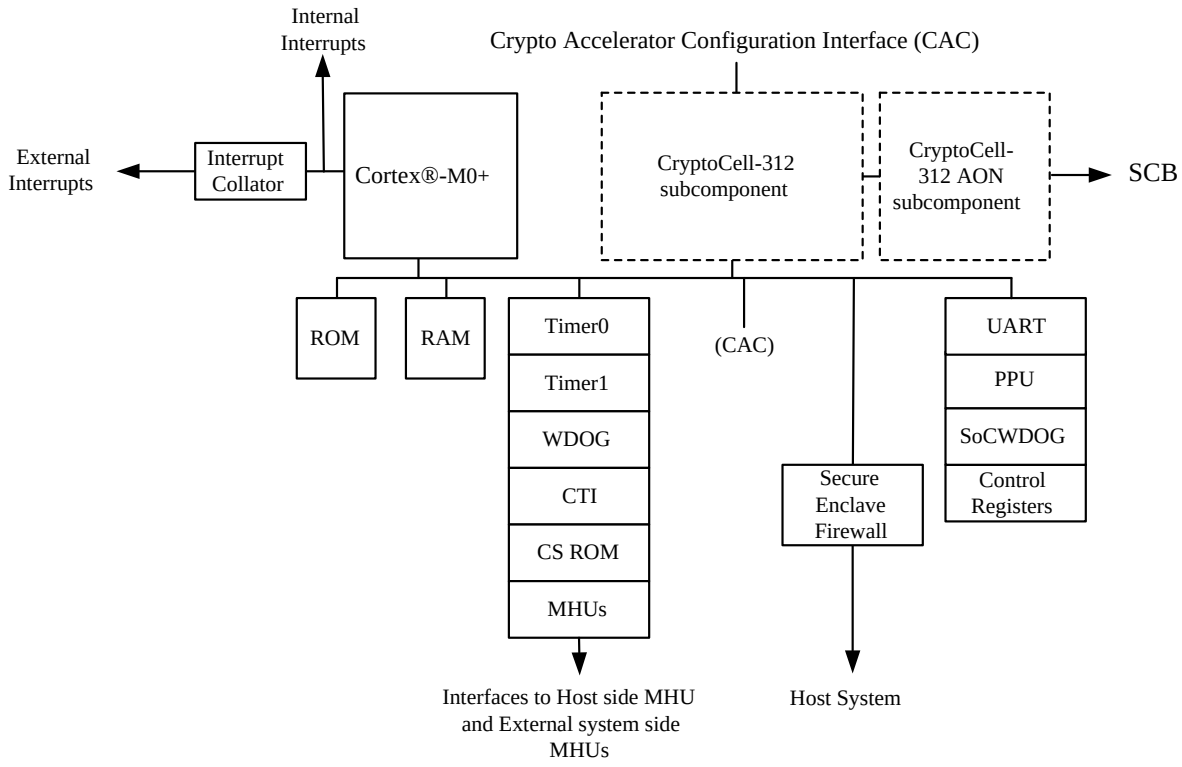

The integrator must integrate a Crypto Accelerator into the Crypto Accelerator socket and Crypto Accelerator AON domain. Therefore, the algorithms supported by the Crypto Accelerator are **IMPLEMENTATION DEFINED**.

To reduce complexity in security reviews, software running on the Secure Enclave is isolated by hardware. Communication between the Host System, External Systems and the Secure Enclave is done by MHUs. For more information about MHUs in the Secure Enclave, see 9.1.6.5 Secure Enclave MHUs on page 155.

The Secure Enclave consists of two sections:

- An always-on section in the AONTOP power domain. This contains the following components:
  - SoC Watchdog
  - Secure Enclave UART
  - SECENCTOP PPU and PCSM
  - Crypto Accelerator Socket Always-on (AON)
  - Secure Enclave System and Base System Control

- A switchable section in the SECENCTOP power domain. For more details, see 6.4.6 SECENCTOP on page 96.

12.2.2 Secure Enclave interrupt map on page 202 and 12.1.3 Secure Enclave memory map on page 195 define the interrupt and address maps of the Secure Enclave.

The 12.3.2.1 Secure Enclave Base System Control register summary on page 236 and 12.3.2.2 Secure Enclave System Control register summary on page 249 define the registers in the Base System Control and Secure Enclave System Control blocks.

2.3.3 External Systems

An External System lets the SoC provide functionality for a specific use-case.

One example of an External System is a pre-existing MCU, which provides connectivity or sensor monitoring functionality.

The SSE-710 design does not specify what can be included as part of an External System. However, the External System must meet several requirements to integrate into an SoC. To support integration, SSE-710 has two External System harnesses to integrate External Systems. Each External System is considered an independent system and can contain any processors, peripherals, and masters required to meet the use cases of the External System.

The External System harness interfaces enable the External System to:

- Access the address space of the Host System
- Communicate with the Host and Secure Enclave systems
- Connect to the shared SoC debug infrastructure

For more information on the interfaces of the External System harness, see 4.3 External System Harness interface on page 43.

2.4 Compliance

The SSE-710 complies with, or includes components that comply with multiple specifications.


Note

This *Technical Reference Manual* (TRM) complements the TRMs for included components, Architecture Reference Manuals, architecture specifications, protocol specifications, and relevant external standards. It does not duplicate information from these sources.

Arm® Architecture

The supported Cortex®-A processors supported by SSE-710, implement the Armv8-A architecture, which executes the A32 and T32 instruction sets.The Cortex®-M0+ processor in the SSE-710 implements the Armv6-M architecture profile.

For more information, see the *Arm®v6-M Architecture Reference Manual* and *Arm® Architecture Reference Manual Armv8, for Armv8-A architecture profile*.

Debug

The SSE-710 implements Arm® CoreSight™ SoC-600 and complies with the following specifications:

- *Arm® CoreSight™ System-on-Chip SoC-600 Technical Reference Manual*
- *Arm® Debug Interface Architecture Specification ADIv6.0*
- *Arm® CoreSight™ Architecture Specification v3.0*

Interrupt controller architecture

The GIC-400 interrupt controller supported by the SSE-710 complies with the following specification:

Arm® Generic Interrupt Controller Architecture Specification, GIC architecture version 2.0

Power Policy Unit (PPU) architecture

The CoreLink™ PCK-600 power management infrastructure in the SSE-710 complies with the following specifications:

- *Arm® Power Policy Unit Architecture Specification, version 1.1*
- *Arm® Power Control System Architecture Specification*
- *AMBA® Low Power Interface Specification, Arm® Q-Channel and P-Channel Interfaces*

Advanced Microcontroller Bus Architecture (AMBA)

The SSE-710 complies with the:

- *AMBA® AXI and ACE Protocol Specification*
- *AMBA® APB Protocol Specification Version 2.0*
- *AMBA® 5 AHB Protocol Specification*
- *AMBA® 4 ATB Protocol Specification ATBv1.0 and ATBv1.1*
- *AMBA® Low Power Interface Specification, Arm® Q-Channel and P-Channel Interfaces*

2.5 Product documentation and design flow

Each Corstone™ SSE-710 Subsystem document has an intended audience and is associated with specific tasks in the design flow. These documents do not reproduce SSE-710 architecture and protocol information.

For relevant protocol and architectural information that relates to this product, see Additional reading.

The SSE-710 documentation is as follows:

Technical Reference Manual

The TRM describes the functionality and the effects of functional options on the behavior of the SSE-710. It is required at all stages of the design flow. The choices that are made in the design flow can mean that some behaviors that the TRM describes are not relevant.

If you are programming the SSE-710 contact:

- The implementer to determine:
  - The build configuration of the implementation.
  - What integration, if any, was performed before implementing the SSE-710.
- The integrator to determine the signal configuration of the device that you use.

The TRM complements architecture and protocol specifications and relevant external standards. It does not duplicate information from these sources.

Configuration and Integration Manual

The CIM describes:

- The available build configuration options.
- How to configure the *Register Transfer Level* (RTL) with the build configuration options.
- How to integrate the SSE-710 into an SoC.
- How to implement the SSE-710 into your design.
- The processes to validate the configured design.

The Arm product deliverables include reference scripts and information about using them to implement your design.

The CIM is a confidential book that is only available to licensees.

2.6 Product revisions

There can be differences in functionality between different product revisions. Arm records these differences in this section.

r0p0 First release

3. Configuration options

This chapter describes the SSE-710 subsystem configurable options that are visible by the Software.

For complete list of supported configuration options refer to *Arm® Corstone™ SSE-710 Subsystem Configuration and Integration Manual* .

3.1 SSE-710 configuration

You can select certain configuration options of the SSE-710 product.

Table 3-1: SSE-710 configuration

Configuration option	Legal values	Default	Description
`NUM_EXP_SHD_INT`	`0-395`	64	Number of supported shared interrupt inputs
`SI{x}_ICI_DST`	`0x00 – 0x3F`	-	Allowed destinations, which shared interrupt `x` can be routed to
`SI{x}_DEF_ICI`	`0x00 – 0x3F`	-	Default routing, for shared interrupt `x`
`NUM_ACLK_QCH`	`0-8`	8	Number of expansion **AOCLK** Q-Channels, in addition to those for the expansion slave and master, and OCVM interfaces
`NUM_DBGCLK_QCH`	`1-8`	8	Number of expansion **DBGCLK** Q-Channels

3.2 Host system configuration

You must define certain configuration elements of the Host system.

Table 3-2: Host System configuration

Configuration option	Legal values	Default	Description
`HOST_CPU_TYPE`	1, 2, 3	1	Selects the type of Host CPU: • 1 – Cortex®-A32 • 2 – Cortex®-A35 • 3 – Cortex®-A53
`HOST_CPU_NUM_CORES`	1-4	4	Selects the number of Host processor cores.
`HOST_EXP_ROM_ENTRY`	`32'00000000`, `32h'01000007` - `32h'01FF_F007`	`32'h00000000`	CoreSight™ ROM table entry in the Host ROM table to point to expansion debug logic. This can be an entry to another CoreSight™ ROM table or a single debug component: • `32'h00000000`: No additional debug logic is added. • `32'h01000007` - `32'h01FFF007`: Base address of the additional debug component or ROM table added on HOSTDBGEXP interface.

Configuration option	Legal values	Default	Description
OCVM_EN	0, 1	1	Selects whether the OCVM interface is included or not: • 0 – OCVM is not included • 1 – OCVM is included

SSE-710 supports up to four supported Cortex®-A processors cores.


Note

The only *Generic Interrupt Controller of* (GIC) supported by the SSE-710 is CoreLink™ GIC-400. No other GICs are supported.

3.3 Secure Enclave configuration

You can define certain configuration elements of the Secure Enclave.

Table 3-3: Secure Enclave configuration

Configuration option	Legal values	Default value	Description
SEC_ENC_ROM_SIZE	16, 32, 64, 128	32	Size of Secure Enclave ROM in KB.
SEC_ENC_RAM_SIZE	8, 16, 32, 64, 128, 256, 512, 1024	128	Size of Secure Enclave RAM in KB.

3.4 External system configuration

You must define certain configuration elements of the external system.

Table 3-4: External System configuration

Configuration option	Legal values	Default	Description
NUM_EXT_SYS	2	2	Configures the number of External Systems
EXT_SYS0_TZ_SPT	0, 1	1	Configures support for Arm® TrustZone® for External System0: • 0 – No TrustZone® support. • 1 – TrustZone® support. This parameter is ignored for External Systems that are not implemented.
EXT_SYS1_TZ_SPT	0, 1	1	Configures support for TrustZone® for External System1: • 0 – No TrustZone® support. • 1 – TrustZone® support. This parameter is ignored for External Systems that are not implemented.

Configuration option	Legal values	Default	Description
EXT_SYS0_ROM_ENTRY	0x0000_0000, 0x001D_0017 -0x002C_F017	0x001D_0017	CoreSight™ ROM Table entry in the EXTDBG ROM to point to the External System's debug logic. This can be an entry to a CoreSight™ ROM table or a single debug component: • 0x0000_0000: No debug logic is provided by the External System 0 • 0x001D_0017 -0x002C_F017: Base address of the additional debug component or ROM table added by the External System using the EXTSYS0DBG interface.
EXT_SYS1_ROM_ENTRY	0x0000_0000, 0x002D_0027 - 0x003C_F027	0x002D_0027	CoreSight™ ROM Table entry in the EXTDBG ROM to point to the External System's debug logic. This can be an entry to a CoreSight™ ROM table or a single debug component: • 0x0000_0000: No debug logic is provided by the External System 1 • 0x002D_0027 - 0x003C_F027: Base address of the additional debug component or ROM table added by the External System using the EXTSYS1DBG interface.

3.5 Host system firewall configuration

You must define certain configuration elements of the host system firewall.

Table 3-5: Host system firewall configuration

Configuration option	Legal values	Default value	Description
FIREWALL_F0_CFG_SSE710 _ AONPERIPH_FC_RGN{x}_BASE_ADDR	0x1A60_0000 - 0x1A6F_FFFF	0x1A60_0000 - 0x1A6F_FFFF	AONPERIPH Firewall Component region x base address, where x is 24-39.
FIREWALL_F0_CFG_SSE710 _ AONPERIPH_FC_RGN{x}_SIZE	0x00, 0x0C-0x14	0x10	AONPERIPH Firewall Component region x size, where x is 24-39
XNVM_NUM_RGN	16, 32, 48, 64	32	Number of regions for the XNVM Firewall Component.
XNVM_RSE_LVL	0, 1	1	Level of RSE implemented by the XNVM Firewall Component.
CVM_NUM_RGN	16, 32, 48, 64	32	Number of regions for the CVM Firewall Component.
CVM_RSE_LVL	0, 1	1	Level of RSE implemented by the CVM Firewall Component.
DBG_NUM_RGN	4, 8	8	Number of regions for the DBG Firewall Component.
EXTSYS0_NUM_RGN	8, 16	8	Number of regions for the EXTSYS0 Firewall Component.
EXTSYS1_NUM_RGN	8, 16	8	Number of regions for the EXTSYS1 Firewall Component
EXPSLV0_NUM_RGN	8, 16, 32	8	Number of regions for the EXPSLV0 Firewall Component.
EXPSLV1_NUM_RGN	8, 16, 32	8	Number of regions for the EXPSLV1 Firewall Component.
EXPMST0_PE_LVL	1, 2	2	Level of PE implemented by the EXPMST0 Firewall Component.
EXPMST1_PE_LVL	1, 2	2	Level of PE implemented by the EXPMST1 Firewall Component.

Configuration option	Legal values	Default value	Description
EXPMST0_RSE_LVL	0, 1	1	Level of RSE implemented by the EXPMST0 Firewall Component. This value must be 0 when EXPMST0_PE_LVL is 1.
EXPMST1_RSE_LVL	0, 1	1	Level of RSE implemented by the EXPMST1 Firewall Component. This value must be 0 when EXPMST0_PE_LVL is 1.
EXPMST0_NUM_RGN	When PE.1 implemented: 1-64 When PE.2 implemented: 8, 16, 32	32	Number of regions for EXPMST0 Firewall Component.
EXPMST1_NUM_RGN	When PE.1 implemented: 1-64 When PE.2 implemented: 8, 16, 32	32	Number of regions for EXPMST1 Firewall Component.
FIREWALL_F0_CFG_SSE710 _EXPMST0_ FC_RGN{x}_BASE_ADDR	0x4000_0000 - 0x5FFF_FFFF	-	EXPMST0 region x base address, where x is between 0 and EXPMST0_NUM_RGN-1. This value is ignored when EXPMST0_PE_LVL is 2.
FIREWALL_F0_CFG_SSE710 _EXPMST0_ FC_RGN{x}_SIZE	0, 0x0C-0x1D	-	EXPMST0 region x is defined as a power of 2 or multiple of the MNRS, where x is between 0 and EXPMST0_NUM_RGN-1. This value is ignored when EXPMST0_PE_LVL is 2.
EXPMST0_MXRS	29	29	MXRS value for EXPMST0 Firewall Component.
EXPMST1_MXRS	29	29	MXRS value for EXPMST1 Firewall Component.
OCVM_NUM_RGN	16, 32, 64	32	Number of regions for the OCVM Firewall Component.
OCVM_RSE_LVL	0, 1	1	Level of RSE implemented by the OCVM Firewall Component.
FC_NUM_MST_ID	7-231	9	Number of StreamIDs which have a unique value set by FIREWALL_F0_CFG_GLOBAL_ SSE710_FC_MST_ID{x}_VAL and FIREWALL_ F0_CFG_GLOBAL_SSE710_ ERR_RESP_ PER_MST_ID{x}_VAL configuration options. This value is a globally defined value for the Host system firewall network.
FC_ERR_RESP_DEF **Note:** The different firewall instances implemented in SSE-710 support different memory path widths, which could be 32-bit, 64-bit, or 128-bit. For 64-bit or 128-bit firewalls, the defined 32-bit default data for error responses is duplicated to all 32-bit word lanes	0x0000_0000 - 0xFFFF_FFFF	0xDEADDEAD	Default data value for error response when StreamID value does not match the pre-configured values, and provides data value for error response when SINGLE_MST is set to 1 for a firewall. For more information, see 10.4 StreamID and CPUID on page 178.

Configuration option	Legal values	Default value	Description
FIREWALL_F0_CFG_GLOBAL_SSE710_FC_MST_ID7_VAL	0x20 - 0xFF	8'd32	Array of parameters containing the list of supported StreamID values. The size of array is defined by FC_NUM_MST_ID. FC_MST_ID<0-6>_VAL is used by masters of SSE-710. FC_MST_ID<7 - (FC_NUM_MST_ID-1)> should be used to give unique StreamID value to each master added by SoC integrator The configured value is StreamID of the issuing master.
FIREWALL_F0_CFG_GLOBAL_SSE710_FC_MST_ID8_VAL		8'd33	
FIREWALL_F0_CFG_GLOBAL_SSE710_ERR_RESP_PER_MST_ID7_VAL	0x0000_0000 - 0xFFFF_FFFF	32'ha5a5a5a5	Unique error response data value to be returned as part of read response corresponding to the StreamID value defined in array. ERR_RESP_PER_MST_ID<0-6> is defined internally in SSE-710,ERR_RESP_PER_MST_ID<7 - FC_NUM_MST_ID-1> should be defined by SoC integrator to provide unique error response data to master *x* that have StreamID defined by FC_NUM_MST_ID<x>_VAL.
FIREWALL_F0_CFG_GLOBAL_SSE710_ERR_RESP_PER_MST_ID8_VAL		32'ha5a5a5a5	

3.6 MHU configuration

You must define the number of channels supported on an MHU.

Table 3-6: MHU configuration

Configuration option	Legal values	Default value	Description
MHU_{x}_NUM_CH	1-32	2	Number of channels supported on an MHU, where x is the MHU short name as listed in 11.5 MHU on page 184.

3.7 IP configuration

Some of the IPs that are used in the SSE-710 subsystem provide configuration options.

The following IP configuration options are supported:

CoreLink™ ADB-400

The size of the FIFOs of ADBs can be set to any value to support the required bandwidth.

CoreLink™ GIC-400

All configuration options are supported. The number of SPIs must be between 73-480.

Cortex®-A32, Cortex®-A35, Cortex®-A53 (the supported Cortex®-A processors for SSE-710)

The SSE-710 supports all allowed supported Cortex®-A processors configurations, with the following conditions:

- *Accelerator Coherency Ports* (ACPs) cannot be included.

- Master interfaces must be in AXI4 mode.
- ETM must be included.
- GIC CPU interface must not be present.

PCK-600

The following configuration options of the *Power Policy Unit* (PPU) are supported:

- `DEV_PREQ_DLY`
- `PCSM_PREQ_DLY`
- `ISO_CLKEN_DLY_CFG`
- `CLKEN_RST_DLY_CFG`
- `RST_HWSTAT_DLY_CFG`
- `CLKEN_ISO_DLY_CFG`
- `ISO_RST_DLY_CFG`

The CoreLink™ GIC-400 and supported Cortex®-A processors detailed configurations are not described in the SSE-710 product documentation.

For configuration information about the supported Cortex®-A processors and CoreLink™ GIC-400, see the respective IP documents listed in 1.4 Additional reading on page 17.

4. Interfaces

This chapter describes the SSE-710 subsystem interfaces.

4.1 Interfaces overview

The SSE-710 interfaces are defined with associated properties, such as the address and data width. All interfaces are defined with a clock, power, and reset domain, except for the clock and reset interfaces.

For more information on the clock, power, and reset domains of an SSE-710, see sections 5. Clocks on page 72, 6.1 Power overview on page 82, and 7.1 Reset overview on page 108.

The following conventions are used:

Input and output definitions

Inputs and Outputs are defined with respect to the SSE-710. For example: An interface that is defined as an input, is an input of the SSE-710 and an output is an output of SSE-710.

An AMBA® interface that is described as a master interface

An interface where SSE-710 is the master. It is connected to a slave interface of a component that is outside the SSE-710.

For an AXI master interface, the **ARVALID** signal is an output and **ARREADY** is an input.

An AMBA® interface that is described as a slave interface

An interface where SSE-710 is the slave. It must be connected to a master interface of a component external to SSE-710.

For an AXI slave interface, the **ARVALID** signal is an input and **ARREADY** is an output.

A Q-Channel interface that is described as a control interface

An interface where SSE-710 is the device. It must be connected to a control interface.

For a Q-Channel control interface, the **QREQn** signal is an input and **QACCEPTn**, **QDENY**, and **QACTIVE** are outputs.

A Q-Channel interface that is described as a device interface

An interface where SSE-710 is the controller. It must be connected to a device interface.

For a Q-Channel device interface, the **QREQn** signal is an output and **QACCEPTn**, **QDENY**, and **QACTIVE** are inputs.


Note

Unless stated otherwise, all Q-Channel interfaces are single Q-Channel instances.

4.2 Host CPU-GIC socket interfaces

This section describes the interfaces that integrate a Host processor and Host GIC into the SSE-710.

4.2.1 Clock interface

The Host CPU Clock interface, **HOSTCPUCLKOUT**, is driven by the **HOSTCPUCLK**. The GIC Clock interface, **GICCLKOUT**, is driven by the **GICCLK**.

4.2.2 Reset interface

The Host CPU Reset interface includes the following reset signals:

- **HOSTCPUWARMRESETn[HOST_CPU_NUM_CORES-1:0]**: Warm reset for the CORE{x} power domain
- **HOSTCPUPORESETn[HOST_CPU_NUM_CORES-1:0]**: Power on reset for the CORE{x} power domain
- **HOSTCLUSTOPWARMRESETn**: Warm reset for the CLUSTOP power domain
- **HOSTCLUSTOPPORESETn**: Power on reset for the CLUSTOP power domain

The GIC reset signal, **GICRESETn**, is driven by **CLUSTOPWARMRESETn**.

4.2.3 Host CPU Memory (HOSTCPUMEM) interface

The HOSTCPUMEM interface is an AXI5 slave interface to be connected to a Host CPU AXI master interface.

Table 4-1: HOSTCPUMEM properties

Clock	Power domain	Reset	Additional details
HOSTCPUCLKOUT	CLUSTOP	**HOSTCLUSTOPWARMRESETn**	• AXI5 with Wakeup_Signals property set to True: • 40-bit address **Note:** Accesses outside the first 4GB cause SSE-710 to generate a DECERR. • 128-bit data

4.2.4 GIC Master (GICM) interface

The GICM interface is to be connected to the AXI slave interface of the Host CPU GIC.

This interface allows configuration access to the registers of the Host CPU GIC.

Table 4-2: GICM interface properties

Clock	Power domain	Reset	Additional details
GICCLKOUT	CLUSTOP	**CLUSTOPWARMRESETn**	• AXI5 with the Wakeup_Signals and Untranslated_Transactions properties set to True. • 32-bit address • 32-bit data • 2 user bits for AR and AW Channels: Indicate which Host CPU core generated the transaction

4.2.5 Host CPU Debug APB (HOSTCPUDBG) interface

The HOSTCPUDBG interface is an APB4 master interface to be connected to the APB4 slave debug interface of the Host CPU.

Table 4-3: HOSTCPUDBG interface parameters

Clock	Power domain	Reset	Additional details
HOSTCPUCLKOUT	CLUSTOP	**CLUSTOPPORESETn**	• APB4, including a **PWAKEUP** output • 32-bit address and data

4.2.6 Host CPU Trace (HOSTCPUTRACE) interface

The HOSTCPUTRACE interface is an ATB4 slave interface to be connected to the ATB4 master interfaces, and to be integrated to the Host CPU trace sources.

Table 4-4: HOSTCPUTRACE interface properties:

Clock	Power domain	Reset	Additional details
HOSTCPUCLKOUT	CLUSTOP	**CLUSTOPPORESETn**	• ATB4 • 32-bit data

4.2.7 Host CPU CTI Channel In (HOSTCPUCTICHIN) interface

The HOSTCPUCTICHIN interface is an input CTI channel interface. It connects the CTIs within the Host CPU cores to the CTI network provided by SSE-710.

Table 4-5: Interface properties

Clock	Power domain	Reset	Additional details
HOSTCPUCLKOUT	CLUSTOP	**CLUSTOPPORESETn**	CTI channel interface

4.2.8 Host CPU CTI Channel Out (HOSTCPUCTICHOUT) interface

The HOSTCPUCTICHOUT interface is an output CTI channel interface. It connects the CTIs of the Host CPU to the CTI network provided by the SSE-710.

Table 4-6: Interface properties

Clock	Power domain	Reset	Additional details
HOSTCPUCLKOUT	CLUSTOP	**CLUSTOPPORESETn**	CTI channel interface

4.2.9 Host CPU Debug Authentication (HOSTCPUDBGAUTH) interface

The HOSTCPUDBGAUTH interface has the standard Arm® debug authentication signals for use by the Host CPU.

Table 4-7: HOSTCPUDBGAUTH interface properties

Clock	Power domain	Reset	Signals
Asynchronous	CLUSTOP	**SEPORESETn**	**DBGEN** **NIDEN** **SPIDEN** **SPNIDEN**

4.2.10 Host CPU Power Control (HOSTCPUPWR) interface

The HOSTCPUPWR interface contains signals for power, clock, and reset control.

The signals in the interface are listed below. In some cases, the width of the signal depends on SSE-710 configuration options.

Table 4-8: HOSTCPUPWR interface details

Interface	Description	Signals
CPU Q-Channels	Power Q-Channel for the Host CPU in CORE{0-3} power domain	• CPUQREQn[HOST_CPU_NUM_CORES-1:0] • CPUQACCEPTn[HOST_CPU_NUM_CORES-1:0] • CPUQDENY[HOST_CPU_NUM_CORES-1:0] • CPUQACTIVE[HOST_CPU_NUM_CORES-1:0]
L2 Q-Channel	Power Q-Channel for the L2 cache RAM	• L2QREQn • L2QACCPETn • L2QDENY • L2QACTIVE
CLUSTOP Ingress Q-Channel	Device Q-Channel interface	• CLUSTOPINGRESSQREQn • CLUSTOPINGRESSQACCEPTn • CLUSTOPINGRESSQDENY • CLUSTOPINGRESSQACTIVE
CLUSTOP Ingress Q-Channel	Device Q-Channel interface	• CLUSTOPINGRESSQREQn • CLUSTOPINGRESSQACCEPTn • CLUSTOPINGRESSQDENY • CLUSTOPINGRESSQACTIVE
CLUSTOP Egress Q-Channel	Device Q-Channel interface	• CLUSTOPEGRESSQREQn • CLUSTOPEGRESSQACCEPTn • CLUSTOPEGRESSQDENY • CLUSTOPEGRESSQACTIVE
For details of these signals, see the Technical Reference Manual of the supported Cortex®-A processors.		SMPEN[HOST_CPU_NUM_CORES-1:0]
		STANDBYWFI[HOST_CPU_NUM_CORES-1:0]
		STANDBYWFIL2
		DBGPWRUPREQ[HOST_CPU_NUM_CORES-1:0]
		DBGNOPWRDWN[HOST_CPU_NUM_CORES-1:0]
		DBGPWRDUP[HOST_CPU_NUM_CORES-1:0]
		L2RSTDISABLE
		WARMRSTREQ[HOST_CPU_NUM_CORES-1:0]
		DBGRSTREQ[HOST_CPU_NUM_CORES-1:0]
		L2FLUSHREQ
		L2FLUSHDONE
		WAKEUPREQ[HOST_CPU_NUM_CORES-1:0]

4.2.11 Host CPU Configuration (HOSTCPUCFG) interface

The HOSTCPUCFG interface contains configuration signals for the Host CPU.

Table 4-9: HOSTCPUCFG interface details

Interface	Signal	Additional details
HOSTCPUCFG	**CFGEND[HOST_CPU_NUM_CORES-1:0]** **CFGTE[HOST_CPU_NUM_CORES-1:0]** **VINITHI[HOST_CPU_NUM_CORES-1:0]** **CRYPTODISABLE** **CP15SDISABLE[HOST_CPU_NUM_CORES-1:0]**	-

4.2.12 Host System Generic Timestamp Gray (HOSTCNTVALUEG) interface

The Host System Generic Timestamp Gray interface provides a gray-encoded timestamp input.

The interface must be converted to a binary value and connected to **CNTVALUEB** of the Host CPU.

Table 4-10: HOSTCNTVALUEG interface properties

Clock	Power domain	Reset	Additional details
HOSTCPUCLKOUT	CLUSTOP	**HOSTCLUSTOPWARMRESETn**	64-bit gray-encoded timestamp

4.2.13 GIC Configuration (GICCFG) interface

The GIC Configuration Interface is made up of the **CFGSDISABLE** signal which is driven by the HOST_SYS_LCTRL_ST.HOST_GIC_LOCK field.

Table 4-11: GICCFG interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	CLUSTOP	**GICRESETn**	-

4.2.14 GIC Interrupt (GICINT) interface

The GIC Interrupt interface provides the interrupt outputs, which must be connected to the Host System GIC, where the source is within SSE-710.

The GIC Interrupt interface provides the interrupt outputs where the source is within SSE-710. These outputs are connected to the Host System GIC. The interface is made up of the following signals:

- **GICINTDBGTOP**
- **GICINTMHUS**
- **GICINTMHUNS**
- **GICINTUART**
- **GICINTWDOGS**
- **GICINTWDOGNS**

The **GICINTDBGTOP** signal has the following bit assignments:

Table 4-12: GICINTDBGTOP bit assignments

Bit offset	Source	Power domain	Clock domain	GIC SPI number
0	Host STM Sync IRQ	DBGTOP	GICCLK	69
1	Host ETR Buffer IRQ	DBGTOP	DBGCLK	70
2	Host CATU IRQ	DBGTOP	DBGCLK	71
3	Host CTI Trigger Out 4	DBGTOP	GICCLK	72
4	Host CTI Trigger Out 5	DBGTOP	GICCLK	73

The **GICINTMHUS** signal has the following bit assignments:

Table 4-13: GICINTMHUS bit assignments

Bit offset	Source	Power domain	Clock domain	GIC SPI number	Notes
0	Host to Secure Enclave MHU0 Combined IRQ	SYSTOP	ACLK	41	-
1	Secure Enclave to Host MHU0 Combined IRQ	SYSTOP	ACLK	42	-
2	Host to EXTSYS0 MHU0 Combined IRQ	SYSTOP	ACLK	43	-
3	EXTSYS 0 to Host MHU0 Combined IRQ	SYSTOP	ACLK	44	
4	Host to EXTSYS 1 MHU0 Combined IRQ	SYSTOP	ACLK	45	-
5	EXTSYS 1 to Host MHU0 Combined IRQ	SYSTOP	ACLK	46	
6-9	-	-	ACLK	47-50	Present and driven LOW

The **GICINTMHUNS** signal has the following bit assignments:

Table 4-14: GICINTMHUNS bit assignments

Bit offset	Source	Power domain	Clock domain	GIC SPI number	Notes
0	Host to Secure Enclave MHU1 Combined IRQ	SYSTOP	ACLK	76	-
1	Secure Enclave to Host MHU1 Combined IRQ	SYSTOP	ACLK	77	-
2	Host to EXTSYS 0 MHU1 Combined IRQ	SYSTOP	ACLK	78	-
3	EXTSYS 0 to Host MHU1 Combined IRQ	SYSTOP	ACLK	79	
4	Host to EXTSYS 1 MHU1 Combined IRQ	SYSTOP	ACLK	80	-
5	EXTSYS 1 to Host MHU1 Combined IRQ	SYSTOP	ACLK	81	
6-9	-	-	ACLK	82-85	Present and driven LOW

The **GICINTUART** signal has the following bit assignments:

Table 4-15: GICINTUART bit assignments

Bit offset	Source	Power domain	Clock domain	GIC SPI number
0	UART0	AONTOP	UARTCLK	51
1	UART1	AONTOP	UARTCLK	52

The **GICINTWDOGS** signal has the following bit assignments:

Table 4-16: GICINTWDOGS bit assignments

Bit offset	Source	Power domain	Clock domain	GIC SPI number
0	Secure Watchdog WS0	AONTOP	REFCLK	32
1	Non-secure Watchdog WS1	AONTOP	REFCLK	33

The **GICINTWDOGNS** signal has the following bit assignment:

Table 4-17: GICINTWDOGNS bit assignment

Bit offset	Source	Power domain	Clock domain	GIC SPI number
0	Non-secure Watchdog WS0	AONTOP	REFCLK	64

4.2.15 GIC Shared Interrupt (GICSHDINT) interface

The GIC Shared Interrupt interface is driven by the ICI1 interface of the Interrupt router.

Table 4-18: GICSHDINT interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	AONTOP	**AONTOPWARMRESWETn**	Width is `NUM_EXP_SHD_INT` + 32 The bit assignment is as defined in Table 11-1: Interrupt Router interface assignment on page 183.

4.2.16 GIC Wakeup (GICWAKEUP) interface

This interface is the **GICWAKEUP** signal. This signal is an active-HIGH request for the Host System GIC to become operational.

Table 4-19: GICWAKEUP interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	AONTOP	**AONTOPWARMRESETn**	-

4.3 External System Harness interface

This section describes the interfaces of the External System Harness.

All the interfaces are prefixed with `EXTSYS{x}`, where x is the number of the External System, between 0 and 1. Each External System has its own version of each interface, which is associated with the clock, reset, and power domain of that External System.

4.3.1 External System {0-1} Clock interface

The clock inputs for this interface are used inside the External System harness.

The External System {0-1} Clock interface is made up of the following clock inputs:

- **EXTSYS{0-1}DBGCLKS**
- **EXTSYS{0-1}DBGCLKM**
- **EXTSYS{0-1}ATCLK**
- **EXTSYS{0-1}CTICLK**
- **EXTSYS{0-1}ACLK**
- **EXTSYS{0-1}MHUCLK**

4.3.2 External System {0-1} Reset Output interface

This interface comprises control signals for the External System:

- **EXTSYS{0-1}CPUWAIT**: Instruction execution halt for any master in the External System that executes instructions.
- **EXTSYS{0-1}PORESETn**: Power-on reset for all logic in the External System.


Note

EXTSYS{0-1}PORESETn is a reset output from SSE-710, and a reset input to the External System.

4.3.3 External System {0-1} Reset Input interface

The following reset signals for this interface are asynchronous assert and synchronous deassert.

They are used by the components and interfaces of the External System Harness:

- **EXTSYS{0-1}DBGPRESETSn**
- **EXTSYS{0-1}DBGPRESETMn**

- **EXTSYS{0-1}ATRESETn**
- **EXTSYS{0-1}CTIRESETn**
- **EXTSYS{0-1}ARESETn**
- **EXTSYS{0-1}MHURESETn**

4.3.4 External System {0-1} Memory (EXTSYS{0-1}MEM) interface

The External System {0-1} Memory interface is an AXI5 slave interface. It must be connected to an External System AXI master interface to provide access to the Host System.

Table 4-20: EXTSYS{0-1}MEM interface properties

Clock	Power domain	Reset	Additional details
EXTSYS{0-1}ACLK	EXTSYS{0-1}TOP	**EXTSYS{0-1}ARESETn**	• AXI5, with Wakeup_Signals set to True • 32-bit address • Configurable data width

4.3.5 External System {0-1} MHU (EXTSYS{0-1}MHU) interface

The External System {0-1} MHU access interface is an APB4 slave interface providing access to MHUs.

12.1.4 External System memory map on page 198 defines the memory map for this interface.

Table 4-21: EXTSYS{0-1}MHU interface properties

Clock	Power domain	Reset	Additional details
EXTSYS{0-1}MHUCLK	EXTSYS{0-1}TOP	**EXTSYS{0-1}MHURESETn**	• APB4 including a **PWAKEUP** input • 19-bit address. • 32-bit data

4.3.6 External System {0-1} MHU Interrupt (EXTSYS{0-1}MHUINT) interface

The External System {0-1} MHU Interrupt interface has eight interrupt signals.

Table 4-22: EXTSYS{0-1}MHUINT interface properties

Clock	Power domain	Reset	Additional detail
EXTSYS{0-1}MHUCLK	EXTSYS{0-1}TOP	**EXTSYS{0-1}MHURESETn**	All signals use the following name format {**x}EXTSYS{0-1}MHUINT**, where x is one of the following: • **HES{x}0**: Combined interrupt output from Receiver frame of Host to External System {0-1} MHU 0. • **ESH{x}0**: Combined interrupt output from Sender frame of External System {0-1} to Host MHU 0. Only implemented when **EXT_SYS{0-1}_TZ_SPT** is 1. • **HES{x}0**: Combined interrupt output from Receiver frame of Host to External System {0-1} MHU 1 • **ES{x}H1**: Combined interrupt output from Sender frame of External System {0-1} to Host MHU 1. Only implemented when **EXT_SYS{0-1}_TZ_SPT** is 1. • **SEES{x}0** • **ES{x}SE0**: Combined interrupt output from the Sender frame of External System {0-1} to Secure Enclave MHU 0.: Combined interrupt output from the Receiver frame of Secure Enclave to External System {0-1} MHU 0. • **SEES{x}1**: Combined interrupt output from Receiver frame of Secure Enclave to External System {0-1} MHU 1.Only implemented when **EXT_SYS{0-1}_TZ_SPT** is 1. • **ES{x}SE1**: Combined interrupt output from Sender frame of External System {0-1} to Secure Enclave MHU 1. Only implemented when **EXT_SYS{0-1}_TZ_SPT** is 1.

4.3.7 External System {0-1} Trace Expansion (EXTSYS{0-1}TRACEEXP) interface

The External System {0-1} Trace Expansion interface enables trace data to be routed into the SSE-710 debug infrastructure.

Table 4-23: EXTSYS{0-1}TRACEEXP interface properties

Clock	Power domain	Reset	Additional details
EXTSYS{0-1}ATCLK	EXTSYS{0-1}TOP	**EXTSYS{0-1}ATRESETn**	• ATB4 • 32-bit data

4.3.8 External System {0-1} Debug APB (EXTSYS{0-1}DBG) interface

The External System {0-1} Debug APB interface enables a debugger to access debug logic within the External System.

Table 4-24: EXTSYS{0-1}DBG interface details

Clock	Power domain	Reset	Additional detail
EXTSYS{0-1}DBGCLKM	EXTSYS{0-1}TOP	**EXTSYS{0-1}DBGPRESETMn**	• APB4, including: ◦ A DP Abort output ◦ A **PWAKEUP** output • 32-bit address and data

4.3.9 External System {0-1} External Debug APB (EXTSYS{0-1}EXTDBG) interface

The External System {0-1} External Debug APB slave interface enables the External System to access SSE-710 debug infrastructure.

12.1.2 External Debug Bus memory map on page 193 defines the memory map for this interface.

Table 4-25: EXTSYS{0-1}EXTDBG interface properties

Clock	Power domain	Reset	Additional details
EXTSYS{0-1}DBGCLKS	EXTSYS{0-1}TOP	**EXTSYS{0-1}DBGPRESETSn**	• APB4, including a **PWAKEUP** input • 32-bit address and data

4.3.10 External System {0-1} CTI Channel In (EXTSYS{0-1}CTICHIN) interface

The External System {0-1} CTI Channel In interface is an input CTI channel interface. The interface enables the External System to connect its local CTIs to the CTI network provided by the SSE-710. This connection lets the External System drive triggers into SSE-710.

Table 4-26: Interface properties

Signal	Power domain	Reset	Additional details
EXTSYS{0-1}CTICLK	EXTSYS{0-1}TOP	**EXTSYS{0-1}CTIRESETn**	CTI channel interface

4.3.11 External System {0-1} CTI Channel Out (EXTSYS{0-1}CTICHOUT) interface

The External System {0-1} CTI Channel Out interface is an output CTI channel interface. The interface enables the External System to connect its local CTIs to the CTI network provided by the SSE-710. This connection lets the External System receive triggers from the SSE-710.

Table 4-27: Interface properties

Signal	Power domain	Reset	Additional details
EXTSYS{0-1}CTICLK	EXTSYS{0-1}TOP	**EXTSYS{0-1}CTIRESETn**	CTI channel interface

4.3.12 External System {0-1} Shared Interrupt (EXTSYS{0-1}SHDINT) interface

The External System {0-1} Shared Interrupt interface is driven by ICI{2-3} of the Interrupt Router.

For more information on Shared Interrupt number assignment, see Table 11-1: Interrupt Router interface assignment on page 183.

Table 4-28: EXTSYS{0-1}SHDINT interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	AONTOP	NA	Width is `NUM_EXP_SHD_INT` + 32

4.3.13 External System {0-1} Memory Clock Q-Channel (EXTSYS{0-1}ACLKQ) interface

The External System {0-1} Memory Clock Q-Channel interface is a control Q-Channel interface that enables a clock controller in the External System to perform high-level clock gating of **EXTSYS{0-1}ACLK**.

Table 4-29: EXTSYS{0-1}ACLKQ interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	EXTSYS{0-1}TOP	**EXTSYS{0-1}ARESETn**	Q-Channel

4.3.14 External System {0-1} MHU Clock Q-Channel (EXTSYS{0-1}MHUCLKQ) interface

The External System {0-1} MHU Clock Q-Channel interface is a control Q-Channel interface that enables a clock controller in the External System to perform high-level clock gating of **EXTSYS{0-1}MHUCLK**.

Table 4-30: EXTSYS{0-1}MHUCLKQ interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	EXTSYS{0-1}TOP	**EXTSYS{0-1}MHURESETn**	Control Q-Channel

4.3.15 External System {0-1} Trace Clock Q-Channel (EXTSYS{0-1}ATCLKQ) interface

The External System {0-1} Trace Clock Q-Channel interface is a control Q-Channel interface that allows a clock controller in the External System to perform high-level clock gating of **EXTSYS{0-1}ATCLK**.

Table 4-31: EXTSYS{0-1}ATCLKQ interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	EXTSYS{0-1}TOP	**EXTSYS{0-1}ATRESETn**	Control Q-Channel

4.3.16 External System {0-1} Debug Master Clock Q-Channel (EXTSYS{0-1}DBGCLKMQ) interface

The External System {0-1} Debug Master Clock Q-Channel interface is a control Q-Channel interface that enables a clock controller in the External System to perform high-level clock gating of **EXTSYS{0-1}DBGCLKM**.

Table 4-32: EXTSYS{0-1}DBGCLKMQ interface properties01

Clock	Power domain	Reset	Additional details
Asynchronous	EXTSYS{0-1}TOP	**EXTSYS{0-1}DBGPRESETMn**	Control Q-Channel

4.3.17 External System {0-1} Debug Slave Clock Q-Channel (EXTSYS{0-1}DBGCLKSQ) interface

The External System {0-1} Debug Slave Clock Q-Channel interface is a control Q-Channel interface that enables a clock controller in the External System to perform high-level clock gating of **EXTSYS{0-1}DBGCLKS**.

Table 4-33: EXTSYS{0-1}DBGCLKSQ interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	EXTSYS{0-1}TOP	**EXTSYS{0-1}DBGPRESETSn**	Control Q-Channel

4.3.18 External System {0-1} CTI Clock Q-Channel (EXTSYS{0-1}CTICLKQ) interface

The External System {0-1} CTI Clock Q-Channel interface is a control Q-Channel interface that enables a clock controller in the External System to perform high-level clock gating of **EXTSYS{0-1}CTICLK**.

Table 4-34: EXTSYS{0-1}CTICLKQ interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	EXTSYS{0-1}TOP	**EXTSYS{0-1}CTIRESETn**	Control Q-Channel

4.3.19 External System {0-1} Memory Power Q-Channel (EXTSYS{0-1}MEMPWRQ) interface

The External System {0-1} Memory Power Q-Channel interface is a control Q-Channel interface. It enables the power controller in the External System to request an associated External System {0-1} Memory interface (**EXTSYS{0-1}MEM**) to leave or enter quiescent state for power control of the interface.

Table 4-35: EXTSYS{0-1}MEMPWRQ interface properties

Clock	Power domain	Reset	Additional data
Asynchronous	EXTSYS{0-1}TOP	**EXTSYS{0-1}ARESETn**	Control Q-Channel

4.3.20 External System {0-1} MHU Power Q-Channel (EXTSYS{0-1}MHUPWRQ) interface

The External System {0-1} MHU Power Q-Channel interface is a control Q-Channel interface. It enables the power controller in the External System to request an associated External System {0-1}

MHU interface (**EXTSYS{0-1}MHU**) to leave or enter quiescent state for power control of the interface.

Table 4-36: EXTSYS{0-1}MHUPWRQ interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	EXTSYS{0-1}TOP	**EXTSYS{0-1}MHURESETn**	Control Q-Channel

4.3.21 External System {0-1} MHU Power Request (EXTSYS{0-1}MHUPWRREQ) interface

The External System {0-1} MHU Power Request interface is a single signal, **EXTSYS{0-1}MHUPWRREQ**. This signal indicates when the Host System or Secure Enclave wants to send a message to the External System.


Note

This signal is a combination of signals from MHUs in the SYSTOP and SECENCTOP power domains.

Table 4-37: EXTSYS{0-1}MHUPWRREQ interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	AONTOP	**SYSTOPWARMRESETn** **SECENCWARMRESETn**	This signal is a combination of signals from MHUs in the SYSTOP and SECENCTOP power domains.

4.3.22 External System {0-1} Trace Expansion Power Q-Channel (EXTSYS{0-1}TRACEEXPPWRQ) interface

The External System {0-1} Trace Expansion Power Q-Channel interface is a control Q-Channel interface. It enables the power controller in the External System to request an associated External System {0-1} Trace Expansion interface (EXTSYS{0-1}TRACEEXP) to enter or leave the quiescent state for power control of the interface.

Table 4-38: EXTSYS{0-1}TRACEEXPPWRQ interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	EXTSYS{0-1}TOP	**EXTSYS{0-1}ATRESETn**	Control Q-Channel

4.3.23 External System {0-1} Debug APB Power Q-Channel (EXTSYS{0-1}DBGPWRQ) interface

The External System {0-1} Debug Power Q-Channel interface is a control Q-Channel interface. It enables the power controller in the External System to request an associated External System Debug interface (EXTSYS{0-1}DBG) to enter or leave quiescent state for power control of the interface.

Table 4-39: EXTSYS{0-1}DBGPWRQ interface properties

Clock	Power domain	Reset	Additional detail
Asynchronous	DBGTOP	**DBGTOPWARMRESETn**	Control Q-Channel

4.3.24 External System {0-1} External Debug APB Power Q-Channel (EXTSYS{0-1}EXTDBGPWRQ) interface

The External System {0-1} External Debug APB Power Q-Channel interface is a control Q-Channel interface. It enables the power controller in the External System to request an associated External System {0-1} External Debug APB interface (EXTSYS{0-1}EXTDBG) to enter or leave quiescent state for power control of the interface.

Table 4-40: EXTSYS{0-1}EXTDBGPWRQ interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	EXTSYS{0-1}TOP	**EXTSYS{0-1}DBGPRESETSn**	Control Q-Channel

4.3.25 External System {0-1} CTI In Power Q-Channel (EXTSYS{0-1}CTIINPWRQ) interface

The External System {0-1} CTI Power Q-Channel interface is a control Q-Channel interface. It enables the power controller in the External System to request an associated CTI Channel In (EXTSYS{0-1}CTICHIN) to enter or leave quiescent state for power control of the CTI Channel interface.

Table 4-41: EXTSYS{0-1}CTIINPWRQ interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	EXTSYS{0-1}TOP	**EXTSYS{0-1}CTIRESETn**	Control Q-Channel

4.3.26 External System {0-1} CTI Out Power Q-Channel (EXTSYS{0-1}CTIOUTPWRQ) Interface

The External System {0-1} CTI Power Q-Channel interface is a control Q-Channel interface. It allows the power controller in the External System to request an associated CTI Channel Out

(EXTSYS{0-1}CTICHOUT) to enter or leave quiescent state for power control of the CTI Channel interface.

Table 4-42: EXTSYS{0-1}CTIOUTPWRQ interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	DBGTOP	**DBGTOPWARMRESETn**	Control Q-Channel

4.3.27 External System {0-1} DBGTOP Q-Channel (EXTSYS{0-1}DBGTOPQ) interface

The External System {0-1} DBGTOP Q-Channel interface is a device Q-Channel interface that is controlled by a power controller for the DBGTOP domain.

Table 4-43: EXTSYS{0-1}DBGTOPQ interface properties

Clock	Power domain	Reset	Additional detail
Asynchronous	AONTOP	**AONTOPWARMRESETn**	Device Q-Channel

4.3.28 External System {0-1} SYSTOP Q-Channel (EXTSYS{0-1}SYSTOPQ) interface

The External System {0-1} SYSTOP Q-Channel interface is a device Q-Channel interface that is controlled by a power controller for the SYSTOP domain.

Table 4-44: EXTSYS{0-1}SYSTOPQ interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	AONTOP	**AONTOPWARMRESETn**	Device Q-Channel

4.3.29 External System {0-1} AONTOP Q-Channel (EXTSYS{0-1}AONTOPQ) interface

The External System {0-1} AONTOP Q-Channel interface is a device Q-Channel interface that is controlled by the Reset Controller.

Table 4-45: EXTSYS{0-1}AONTOPQ interfaces properties

Clock	Power domain	Reset	Additional details
Asynchronous	AONTOP	**AONTOPWARMRESETn**	Device Q-Channel

4.3.30 External System {0-1} Power Request (EXTSYS{0-1}PWRREQ) interface

The External System {0-1} Power Request interface is formed from two CoreSight™ **CDBGPWRUPREQ**/**ACK** handshakes. One of the **CDBGPWRUPREQ**/**ACK** handshakes is from the EXTDBG ROM and the other is from the HOST AXIAP ROM.

Table 4-46: EXTSYS{0-1}PWRREQ interfaces details

Clock	Power domain	Reset	Additional details
Asynchronous	DBGTOP	**DBGTOPWARMRESETn**	Formed of the following signals: • **EXTDBGROMCDBGPWRUPREQ**: Output • **EXTDBGROMCDBGPWRUPACK**: Input • **AXIAPROMCSYSPWRUPREQ**: Output • **AXIAPROMCSYSPWRUPACK**: Input

4.3.31 Reset Syndrome (EXTSYS{0-1}RSTSYN) interface

The Reset Syndrome interface is a 5-bit signal that indicates the cause of the last reset for the associated External System.


Note

Arm® strongly recommends that this interface is exposed to software running on the External System, in a reset syndrome register.

The following table shows the bit assignment:

Table 4-47: EXTSYS{0-1}RSTSYN bit assignment

Bit	Name	Description
0	POR	Indicates the last reset of the External System was caused by one of the following: • **PORESETn** pin being asserted • DP **CDBGRSTREQ** being asserted • SoC Watchdog reset request • Secure Enclave Crypto Accelerator Error reset request • Secure Enclave Watchdog reset request • SOC_RST_CTRL.RST_REQ bit set to `0b1` • Secure Enclave software reset request
1	nSRST	Indicates that the last reset of the External System was caused by either: • **nSRST** pin being asserted • DP ROM **CSYSRSTREQ** being asserted
2	Reserved	Tied to `0b0`
3	HOST	Indicates the last reset of the External System was caused by a Host System reset request

Bit	Name	Description
4	EXT	Indicates the last reset of the External System was caused by a request to reset this External System

4.4 Host System Interfaces

This section describes the interfaces of the Host System, excluding those of the CPU_GIC socket.

4.4.1 On-chip Volatile Memory (CVM) interface

There is a single AXI master expansion interface for on-chip volatile memory, for example SRAM.

The CVM interface has a Firewall Component associated with it. For more information, see 10.3.5 Host System firewall on page 169.

This interface includes the AMBA® AXI5 QoS Signals.

The CVM interface has an *Access Control Gate* (ACG), that moves SYSTOP to ON when there is an access to the On-chip Volatile Memory when the SYSTOP domain is not in the ON power-mode.

Table 4-48: CVM interface properties

Clock	Power domain	Reset	Additional details
ACLKOUT	SYSTOP	**SYSTOPWARMRESETn**	• AXI5 with the following properties set to True: ◦ Wakeup_Signals ◦ Untranslated_transactions • 32-bit address • Configurable data width

4.4.2 eXecute-in-place Non-volatile Memory (XNVM) interface

The *eXcute-in-Place* (XiP) Non-volatile Memory interface is an AXI Master interface. It provides access to an XiP Non-volatile memory, for example, Flash memory.

The XNVM interface has a Firewall Component associated with it.

This interface includes the AMBA® AXI5 QoS Signals.

Table 4-49: XNVM interface properties

Clock	Power domain	Reset	Additional details
ACLKOUT	SYSTOP	**SYSTOPWARMRESETn**	• AXI5 with the following properties set to True: ◦ Wakeup_Signals ◦ Untranslated_transactions • 32-bit address • Configurable data width

4.4.3 Off-chip Volatile Memory (OCVM) interface

The Off-Chip Volatile Memory interface is an AXI Master interface. It provides access to off-chip volatile memory, for example, DRAM.

The OCVM interface has a Firewall Component associated with it.

This interface includes the AMBA® AXI5 QoS Signals.

Table 4-50: OCVM interface properties

Clock	Power domain	Reset	Additional details
ACLKOUT	SYSTOP	**SYSTOPWARMRESETn**	• AXI5 with the following properties set to true: ◦ Wakeup_Signals ◦ Untranslated_transactions • 32-bit address • Configurable data width

4.4.4 Host Expansion Slave {0-1} (HOSTEXPSLV{0-1}) interfaces

There are two slave expansion interfaces of the Host System.

Each HOSTEXPSLV{x} interface has a Firewall Component associated with it.

This interface includes the AMBA® AXI5 QoS Signals.

Table 4-51: HOSTEXPSLV{0-1} interfaces properties

Clock	Power domain	Reset	Additional details
ACLKOUT	SYSTOP	**SYSTOPWARMRESETn**	• AXI5 with the following properties set to True: ◦ Wakeup_Signals ◦ Untranslated_transactions • 32-bit address • Configurable data width

4.4.5 Host Expansion Master {0-1} (HOSTEXPMST{0-1}) interfaces

There are two Host Expansion master expansion interfaces for the Host System.

Each HOSTEXPMST{x} interface has a Firewall Component associated with it.

This interface includes the AMBA® AXI5 QoS Signals.

Table 4-52: HOSTEXPMST{0-1} interface properties

Clock	Power domain	Reset	Additional details

ACLKOUT	SYSTOP	**SYSTOPWARMRESETn**	• AXI5 with the following properties set to True:
			  ◦ Wakeup_Signals
			  ◦ Untranslated_transactions
			• 32-bit address
			• Configurable data width

4.4.6 Host AON Expansion Master (HOSTAONEXPMST) interface

The HOSTAONEXPMST interface allows integrators to add their own peripherals to the AONTOP domain.

Table 4-53: HOSTAONEXPMST interface properties

Clock	Power domain	Reset	Additional details
HOSTAONEXPCLK	AONTOP	**AONTOPWARMRESETn**	• APB4, including a **PWAKEUP** output • 32-bit address and data

4.4.7 Host Debug APB Expansion (HOSTDBGEXP) interface

The HOSTDBGAPBEXP interface allows integrators to add their own debug components to the Host System.

Table 4-54: HOSTDBGEXP interface properties

Clock	Power domain	Reset	Additional details
DBGCLKOUT	DBGTOP	**DBGTOPWARMRESETn**	• APB4, including a **PWAKEUP** output • 32-bit address and data

4.4.8 Host Debug Trace Expansion (HOSTDBGTRACEEXP) interface

The HOSTDBGTRACEEXP interface allows integrators to add their own trace source components to the Host System.

Table 4-55: HOSTDBGTRACEEXP interface properties

Clock	Power domain	Reset	Additional details
DBGCLKOUT	DBGTOP	**DBGTOPWARMRESETn**	• ATB4 • 32-bit data

4.4.9 Host CTI Channel In Expansion (HOSTCTICHINEXP) interface

The Host CTI Channel In Expansion interface allows for an external CTI to connect and send trigger events to the internal CTI of the Host System.

Table 4-56: HOSTCTICHINEXP interface properties

Clock	Power domain	Reset	Additional details
DBGCLKOUT	DBGTOP	**DBGTOPWARMRESETn**	CTI Channel interface

4.4.10 Host CTI Channel Out Expansion (HOSTCTICHOUTEXP) interface

The Host CTI Channel Out Expansion interface allows for an external CTI to connect and receive triggers from the internal CTI of the Host System.

Table 4-57: HOSTCTICHOUTEXP interface details

Clock	Domain	Reset	Additional details
DBGCLKOUT	DBGTOP	**DBGTOPWARMRESETn**	CTI Channel interface

4.4.11 Host Debug Timestamp (HOSTTSVALUEB) interface

The CoreSight™ timestamp input interface is used by the CoreSight™ STM-500.

Table 4-58: HOSTTSVALUEB interface properties

Clock	Power domain	Reset	Additional details
REFCLK	DBGTOP	**DBGTOPWARMRESETn**	64-bit binary encoded timestamp

4.4.12 Host Debug Authentication (HOSTDBGAUTH) interface

The Host Debug Authentication interface provides the CoreSight™ authentication signals that must control any debug logic added to the Host System *Debug Authentication Zone* (DAZ).

The interface has the following properties:

Table 4-59: HOSTDBGAUTH interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	DBGTOP	**SEPORESETn**	4-bit signal made up of: • **DBGEN** • **NIDEN** • **SPIDEN** • **SPNIDEN**

4.4.13 Host STM DMA Peripheral Request (HOSTSTMDPR) interface

The Host STM DMA Peripheral Request interface enables communication between the STM and a DMA engine. It supports the DMA Peripheral Request interface, defined by CoreSight™ STM-500.

Table 4-60: HOSTSTMDPR interface properties

Clock	Power domain	Reset	Additional detail
DBGCLK	DBGTOP	**DBGTOPWARMRESETn**	CoreSight™ STM-500 DMA peripheral request interface. For more information on the signals and protocol, see the *Arm® CoreSight™ STM-500 System Trace Macrocell Technical Reference Manual*.

4.4.14 Host System REFCLK Generic Timestamp Gray (HOSTCNTVALUEG) interface

The Host System REFCLK Generic Timestamp Gray interface provides a Gray-encoded REFCLK application timestamp output. It can also be used by other components in the SoC added by the SoC integrator.

Table 4-61: HOSTCNTVALUEG interface properties

Clock	Power domain	Reset	Additional details
HOSTCNTCLKOUT	AONTOP	**AONTOPWARMRESETn**	64-bit Gray encoded timestamp

4.4.15 Host System REFCLK Generic Timestamp Binary (HOSTCNTVALUEB) interface

The Host System REFCLK Generic Timestamp binary interface provides an application timestamp value. This value is used for by the REFCLK timers {0-3} integrated in the SSE-710.

Table 4-62: HOSTCNTVALUEB interface properties

Clock	Power domain	Reset	Additional details
HOSTCNTCLKOUT	AONTOP	**AONTOPWARMRESETn**	Has a 64-bit binary encoded timestamp

4.4.16 Host System S32K Timestamp Gray (HOSTS32KCNTVALUEG) interface

The Host System 32K Generic Timestamp Gray interface provides a Gray-encoded 32K timestamp output. Connect this interface to the HOSTS32KCNTVALUEB interface to provide a timestamp for the S32K timers {0-1}. It can also be used by other components in the SoC that are added by the integrator.

Table 4-63: HOSTS32KCNTVALUEG interface properties

Clock	Power domain	Reset	Additional details
HOSTS32CNTCLKOUT	AONTOP	**AONTOPWARMRESETn**	64-bit Gray encoded timestamp

4.4.17 Host System S32K Timestamp Binary (HOSTS32KCNTVALUEB) interface

The Host System 32K Generic Timestamp binary interface provides the 32K timestamp, which is used by S32K timers {0-1}.

Table 4-64: HOSTS32KCNTVALUEB interface properties

Clock	Power domain	Reset	Additional details
HOSTS32CNTCLKOUT	AONTOP	**AONTOPWARMRESETn**	64-bit binary encoded timestamp

4.4.18 Host System Debug Power Request (HOSTDBGPWRREQ) interface

The Host System Debug Power Request interface allows you to add power domains to the Host System. The power domains are controlled by a debug agent.

Table 4-65: HOSTDBGPWRREQ interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	DBGTOP	**DBGTOPWARMRESETn**	16 CoreSight™ Debug Power Request (Request/Acknowledge) handshakes from the Host AXIAP ROM CSYSPWRUP

4.4.19 Host System UART {0,1} (HOSTUART{0,1}) interface

The following table shows the Host System UART 0 and 1 interfaces, which provide a standard UART interface, with flow control.

Table 4-66: HOSTUART{0,1} interfaces properties

Clock	Power domain	Reset	Additional details
Asynchronous	AONTOP	**AONTOPWARMRESETn**	Table 4-67: Host System UART interface signals on page 61

Table 4-67: Host System UART interface signals

Signal name	Input/Output	Description
HOSTUART{0,1}TX	Output	UART Transmit data
HOSTUART{0,1}RX	Input	UART Receive data
HOSTUART{0,1}RTSn	Output	UART Request to Send
HOSTUART{0,1}CTSn	Input	UART Clear to Send
HOSTUART{0,1}RIn	Input	UART Ring Indicator
HOSTUART{0,1}DCDn	Input	UART Data Carrier Detect
HOSTUART{0,1}DSRn	Input	UART Data Set Ready
HOSTUART{0,1}DTRn	Output	UART Data Terminal Ready
HOSTUART{0,1}OUT1n	Output	UART Out1
HOSTUART{0,1}OUT2n	Output	UART Out2

4.5 Secure Enclave interfaces

4.5.1 Crypto Accelerator socket interfaces

The Secure Enclave Crypto Accelerator socket interfaces are used to integrate a Crypto Accelerator into the SSE-710.

You can divide the Crypto Accelerator between the SECENCTOP and AONTOP power domains. The two pieces are called:

- Crypto Accelerator sockeT SECENTOP
- Crypto Accelerator *Always-On* socket AON

It is **IMPLEMENTATION DEFINED** whether the Crypto Accelerator is implemented as whole or separate.

4.5.1.1 Crypto Clock interface

The Crypto Clock interface **CRYPTOCLKOUT** is driven by the **SECENCCLK**. It is used only for the integration of the Crypto Accelerator into the Crypto Accelerator socket SECENTOP.

4.5.1.2 Crypto AON Clock (CRYPTOAONCLKOUT) Interface

The Crypto AON Clock interface **CRYPTOAONCLKOUT** is driven by **SECENCCLK**. It is used only for the integration of the Crypto Accelerator Always-On into the Crypto Accelerator socket AON.

4.5.1.3 Crypto Reset (CRYPTORESETn) interface

The Crypto Reset interface **CRYPTORESETn** is driven by **SECENCWARMRESETn**.

The Crypto Reset interface **CRYPTORESETn** is driven by **SECENCWARMRESETn**. It is used only for the integration of the Crypto Accelerator into the Crypto Accelerator Socket SECENCTOP.

Table 4-68: CRYPTORESETn interface properties

Clock	Power domain	Reset	Additional details
Active-LOW asynchronous assertion, synchronous de-assertion with respect to **CRYPTOCLKOUT**.	SECENTOP	**SECENCWARMRESETn**	-

4.5.1.4 Crypto AON Reset Interface

The Crypto Reset interface **CRYPTOAONRESETn** is driven by **SEPORESETn**. The reset signal is active-LOW asynchronous assertion, synchronous de-assertion with respect to **CRYPTOAONCLKOUT**.

It is used only for the integration of the Crypto Accelerator AON into the Crypto Accelerator Socket AON.

4.5.1.5 Crypto Accelerator DMA (CAD) interface

If required, use the Crypto Accelerator DMA (CAD) interface to connect the DMA interface of a Crypto Accelerator.

The Crypto Accelerator DMA interface is an AHB5 slave interface, with the following properties:

Table 4-69: Interface properties

Clock	Power domain	Reset	Details
CRYPTOCLKOUT	SECENCTOP	**CRYPTORESETn**	AHB5 Extended Memory Types :False Secure Transfers Endian: BES Stable between Clock: True Exclusive Transfers: False Multi Copy Atomicity: False User Signalling: False 32-bit address 32-bit data

The CAD interface allows access to the following sections of the Secure Enclave memory map:

- Secure Enclave ROM
- Secure Enclave RAM
- Host Access region

Any other regions of the Secure Enclave address map are treated as RAZ/WI and generate an error.

Access from the CAD interface to the Secure Enclave RAM and ROM ignores the security world of the transaction.


Note

If the CAD interface is not used, it must be tied off.

For more information on the CAM interface, see the *Arm® Corstone™ SSE-710 Subsystem Configuration and Integration Manual*.

4.5.1.6 Crypto Accelerator Configuration (CAC) interface

The Crypto Accelerator Configuration interface connects to theconfiguration interface of the Crypto Accelerator.


Note

All accesses to the CAC interface are always marked as secure because inside the Secure Enclave only a single security world is defined. Accesses in the address range `0x2F00_0000` to `0x2FFF_FFFF` are routed to the CAC interface.

Table 4-70: CAS interface proprties

Clock	Power domain	Reset	Additional details
CRYPTOCLKOUT	SECENCTOP	**CRYPTORESETn**	• AHB5 • Extended Memory Types :False • Secure Transfers • Endian: BES • Stable between Clock: True • Exclusive Transfers: False • Multi Copy Atomicity: False • User Signalling: False • 32-bit address • 32-bit data

4.5.1.7 Crypto Accelerator AON (CAAON) interface

The Crypto Accelerator AON interface is used for communication between the two halves of the **IMPLEMENTATION DEFINED** Crypto Accelerator.

The Crypto Accelerator AON interface is made up of two multi-bit signals, **CA2CAAON** and **CAAON2CA**, allowing the two halves of the Crypto Accelerator to communicate. The signals are used for:

- Information sent from the Crypto Accelerator to Crypto Accelerator Always-on. (**CA2CAAON**).
- Information sent from the Crypto Accelerator Always-on to Crypto Acceleratorr. (**CAAON2CA**).

The properties of the interface differ for each signal:

- **CA2CAAON**:
  - Width equal to CAAON2CA_WIDTH
  - **CRYPTOCLKOUT**
  - **CRYPTORESETn**
- **CAAON2CA**:

- Width equal to CA2CAAON_WIDTH
- **CRYPTOAONCLKOUT**
- **CRYPTOAONRESETn**

These signals are used only for the integration of the Crypto Accelerator into the SSE-710 subsystem.


Note

There are two instances of the CAAON interface crossing the power domain. In one instance the **CA2CAAON** and **CAAON2CA** signals are input and output respectively. In the other, the **CA2CAAON** and **CAAON2CA** signals are output and input respectively.

4.5.1.8 Crypto Accelerator lifecycle Control (CALC) interface

The Crypto Accelerator lifecycle Control interface is used to control the advance of the lifecycle of the SoC.

Table 4-71: CALC interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	AONTOP	NA	-

4.5.1.9 Crypto Accelerator Interrupt (CAINT) interface

The Crypto Accelerator Interrupt interface connects the interrupts from the Crypto Accelerator to the Cortex®-M0+ NVIC:

Table 4-72: CAINT interface properties

Clock	Power domain	Reset	Additional details
CRYPTOCLKOUT	SECENTOP	**CRYPTORESETn**	1 bit signal

4.5.1.10 Crypto Accelerator Power Q-Channel (CAPWRQ) interface

The Crypto Accelerator Power Q-Channel interface is used to control the power mode of the Crypto Accelerator.

Table 4-73: Interface properties

Clock	Power domain	Reset	Details
Asynchronous	AONTOP	**CRYPTOAONRESETn**	1x Q-Channel Device interface

The CAPWRQ interface is in Q_RUN state when the SECENCTOP PPU is in the ON power mode. The CAPWRQ interface is in the Q_STOPPED state when the SECENCTOP PPU is in the OFF, MEM_RET or WARM_RST power modes.

The **CAPWRQACTIVE** signal is OR combined into the PACTIVE[ON] to the SECENCTOP PPU.

4.5.1.11 Crypto Accelerator Error (CAE) interface

The Crypto Accelerator lifecycle control interface provides information about an error event of the Crypto Accelerator, which software cannot handle.

Table 4-74: CAE interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	AONTOP	**CRYPTORESETn**	1-bit input

After the CAE interface is asserted, it must remain asserted until the Internal Power-on-Reset (IPoR) has occurred. It is **UNPREDICTABLE** whether the IPoR is generated if you deassert the CAE interface before the IPoR sequence has completed.

For details of the IPoR sequence, see 7.5.2 Internal Power-on-Reset (IPoR) on page 115.

When the CAE interface is asserted and SE_RST_MSK.CA_ERR_MSK is set to `0b0` the following occurs:

- IPoR reset is generated asserting **AONTOPPORESETn**, **AONTOPWARMRESETn**, **SEPORESETn**, **EXTSYS{0-1}PORESETn** resets.
- Secure Enclave Base System SOC_RST_SYN indicates that the last reset was caused by the Secure Enclave.
- Secure Enclave System Control SE_RST_SYN indicates that the last reset was caused by an initialization error.
- Host Base System SOC_RST_SYNC indicates the last reset was caused by a POR.
- External System Reset Syndrome (EXTSYS{0-1}RSTSYN) interfaces indicates the last reset was caused by a POR.
- All SCBs are driven to `0b0`.
- Waits particular number of S32KCLK clock cycles for reset to be applied to the entire SoC (the number of cycles is set by the `SOC_RST_DLY` parameter).
- IPoR reset is released deasserting **AONOPPORESETn**, **AONTOPWARMRESETn**, **SEPORESETn**, **EXTSYS{0-1}PORESETn** resets.

Arm® strongly recommends that only those errors, which software cannot handle on the Secure Enclave or compromise the security of the SoC, assert the CAE interface.

4.5.2 Secure Enclave UART (SECENCUART) interface

The Secure Enclave UART has a standard interface with flow control.

The following table shows the SECENCUART interface signals:

Table 4-75: SECENCUART interface

Signal	Description	I/O
SECENCUARTTX	UART Transmit data	O
SECENCUARTRX	UART Receive data	I
SECENCUARTRTSn	UART Request to Send	O
SECENCUARTCTSn	UART Clear to Send	I
SECENCUARTRIn	UART Ring Indicator	I
SECENCUARTDCDn	UART Data Carrier Detect	I
SECENCUARTDSRn	UART Data Set Ready	I
SECENCUARTDTRn	UART Data Terminal Ready	O
SECENCUARTOUT1n	UART Out1	O
SECENCUARTOUT2n	UART Out2	O

Table 4-76: SECENCUART interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	AONTOP	**SEPORESETn**	-

4.5.3 Security Control Bits (SCB) interface

The Security Control Bits provide an input interface to control security features within the SSE-710 subsystem.

For more details on the SCB see 9.1.3 Security Control Bits (SCB) on page 150.

Table 4-77: SCB interface properties

Clock	Power domain	Reset	Additional details
Asynchronous	AONTOP	**CRYPTOAONRESETn**	Multi-bit signal, as defined in 9.1.3 Security Control Bits (SCB) on page 150

4.6 SSE-710 interfaces

This section describes the interfaces of the SSE-710.

Table 4-78: SSE-710 interfaces

Interface name	Clock	Power domain	Reset	Description
Clocks	See 5. Clocks on page 72			
Resets	See 7. Reset on page 108			
PLL Lock (PLLLOCK)	Both signals are treated as asynchronous	-	-	Indicates when a PLL is locked and safe to use by software. Uses two signals: • **CPUPLLLOCK** • **SYSPLLLOCK**

Interface name	Clock	Power domain	Reset	Description
Expansion Shared Interrupt (EXPSHDINT)	Asynchronous	AONTOP	**AONTOPWARMRESETn**	Drives the SII interface of the Interrupt Router.
JTAG/Serial Wire Debug (SWJ)	**SWCLKTCK**	AONTOP	**nTRST/PORESETn** **Note:** When either reset is asserted, the SWJ interface is reset.	Supports both JTAG and Serial Wire Debug protocols.
Trace Port Interface Unit (TPIU)	**TRACECLKOUT**	DBGTOP	**DBGTOPWARMRESETn**	A TPIU interface that enables trace data to be captured off-chip. The interface is formed of the following signals: • **TRACEDATA[31:0]** • **TRACECTL** • **TRACEMAXDATASIZE[4:0]** • **TPCTLVALID** The width of the TPIU interface can be configured to any number of bits, starting with **TRACEDATA[0]**, using the **TRACEMAXDATASIZE[4:0]** signal. For more information on the interface, see the Soc Identification interface (SOCID)
SoC Configuration (SOCCFG)	**S32KCLK**	AONTOP	**AONTOPWARMRESETn**	This is a static configuration interface. The interface is formed of the following signals: • **CVMSIZE[7:0]** • **XNVMSIZE[7:0]** • **OCVMSIZE[7:0]** (Only implemented when OCVM_EN is 1). For more information, see 10.3.5 Host System firewall on page 169
SoC Identification (SOCID)	-	AONTOP	-	Sets the SoC Identification information in the System ID registers. The signals are considered pseudo static.

Interface name	Clock	Power domain	Reset	Description
SoC Security Control (SOCSC)	Asynchronous	AONTOP	**SEPORESETn**	Contains miscellaneous control signals for controlling SoC features. The SOCSC interface is formed of the following output signals: • **DFTENABLE[1:0]**: ◦ Bit[0] is driven by bit[34] of the SCB interface. An active-HIGH signal, which indicates when Design for Test functionality can be enabled on the SoC, except for logic in the Secure Enclave. For example, scan or MBIST. ◦ Bit[1] is driven by bit[63] of the SCB interface. An active-HIGH signal, which indicates when Design for Test functionality can be enabled on the logic inside the Secure Enclave. For example, scan or MBIST. ◦ Both of these signals must not prevent access to the debug DP, in the SSE-710. • **SCBEXP[63:0]**: ◦ Driven by bits[127:64] of the SCB interface. ◦ Behavior of these signals is defined by the integrator.
SoC Lifecycle Control (SOCLCC)	Asynchronous	AONTOP	-	This interface is an input pin that controls the lifecycle state of the SoC.

Table 4-79: SoC Identification interface signals

Signal	Description
SOCPRTID[11:0]	SoC Product ID
SOCVAR[3:0]	SoC Variant
SOCREV[3:0]	SoC Revision
SOCIMPLID[10:0]	SoC Implementer JEP106 Code ID: • [6:0]: JEP106 identity code • [10:7]: JEP106 continuation code
DPROMPRTID[11:0]	DP ROM table Part ID
DPROMVAR[3:0]	DP ROM table Variant
DPROMREV[3:0]	DP ROM table Revision
DPROMIMPLID[10:0]	DP ROM table implementor JEP106 Code ID: • [6:0]: JEP106 identity code • [10:7]: JEP106 continuation code
EXTDBGROMPRTID[11:0]	EXTDBG ROM table Part ID
EXTDBGROMVAR[3:0]	EXTDBG ROM table Variant
EXTDBGROMREV[3:0]	EXTDBG ROM table Revision
EXTDBGROMIMPLID[10:0]	EXTDBG ROM table implementor JEP106 Code ID: • [6:0]: JEP106 identity code • [10:7]: JEP106 continuation code
HOSTROMPRTID[11:0]	Host ROM table Part ID

Signal	Description
HOSTROMVAR[3:0]	Host ROM table Variant
HOSTROMREV[3:0]	HOST ROM table Revision
HOSTROMIMPLID[10:0]	HOST ROM table implementor JEP106 Code ID: • [6:0]: JEP106 identity code • [10:7]: JEP106 continuation code
HOSTAXIAPROMPRTID[11:0]	Host AXIAP ROM table Part ID
HOSTAXIAPROMVAR[3:0]	Host AXIAP ROM table Variant
HOSTAXIAPROMREV[3:0]	Host AXIAP ROM table Revision
HOSTAXIAPROMIMPLID[10:0]	Host AXIAP ROM table implementor JEP106 Code ID: • [6:0]: JEP106 identity code • [10:7]: JEP106 continuation code

4.7 Clock Control interfaces

The SSE-710 provides several Q-Channel interfaces to support high-level clock gating of either clock inputs or outputs.

Table 4-80: Clock control interfaces

Interface name	Clock	Power domain	Reset	Description
REFCLK Q-Channel (REFCLKQ) interface	Asynchronous	AONTOP	**PORESETn**	Q-Channel control interface that allows high-level clock gating of the **REFCLK**. All signals have the format **REFCLK{x}**, where x is the standard Q-Channel protocol name.
ACLK Q-Channel (ACLKQ) interface	Asynchronous	SYSTOP	**SYSTOPWARMRESETn**	Formed of one or more device Q-Channels. There are (**NUM_ACLK_QCH + OCVM_EN**) device Q-Channels.
DBGCLK Q-Channel (DBGCLKQ) interface	Asynchronous	DBGTOP	**DBGTOPWARMRESETn**	Formed of one or more device Q-Channels. There are `NUM_DBGCLK_QCH` device Q-Channels

4.8 Power control interfaces

This section gives an overview of the Power Control interfaces.

Table 4-81: Power control interfaces

Interface name	Clock	Power domain	Reset	Description
SYSTOPQ	Asynchronous	AONTOP	**AONTOPWARMRESETn**	Enables expansion of the SYSTOP power domain. Formed of three device Q-Channels: • Bit 0 is Egress Q-Channel • Bit 1 is Internal Q-Channel • Bit 2 is Ingress Q-Channel
DBGTOPQ	Asynchronous	AONTOP	**AONTOPWARMRESETn**	Enables expansion of the DBGTOP power domain. Formed of three device Q-Channels: • Bit 0 is Egress Q-Channel • Bit 1 is Internal Q-Channel • Bit 2 is Ingress Q-Channel

5. Clocks

This chapter describes the primary and internal clocks within the SSE-710 subsystem.

5.1 Clock inputs

The SSE-710 has multiple clock inputs.

The SSE-710 has the following clock inputs:

Table 5-1: Clock inputs

Clock name	Description	Mandatory
REFCLK	**REFCLK** is the reference clock for the system, except for the Secure Enclave. It is used in many places and is the default option for all generated clocks. The REFCLK Q-Channel interface indicates when SSE-710 requires **REFCLK**. For more information, see 4.7 Clock Control interfaces on page 70. **REFCLK** can be gated, when the Q-Channel enters the Q_STOPPED state. When QACTIVE is asserted **REFCLK** must be ungated and the Q-Channel returned to the Q_RUN state.	Yes
S32KCLK	**S32KCLK** is used by the S32K counter, timers, and SoC watchdog in the system. It is also used by the wakeup logic when in the BSYS.SLEEP1 power state.	Yes
SECENCREFCLK	**SECENCREFCLK** is the reference clock for the Secure Enclave. It is only used to generate the clocks of the Secure Enclave.	Yes
CPUPLL	**CPUPLL** is used to generate the clocks used by the Host CPU. The **CPUPLL** input must be from a PLL, added by the integrator. There is also a **CPUPLLLOCK** signal. The software is responsible for switching to the **CPUPLL** when the PLL provides a clock and is locked to the frequency. **Note:** If the CPU PLL is present, but does not provide a lock signal. Then it is **IMPLEMENTATION DEFINED** how the software determines when the PLL is stable enough to be used.	Optional
SYSPLL	**SYSPLL** is the clock output of the system PLL. It generates all clocks in the system. Alongside the **SYSPLL** input, there is the **SYSPLLLOCK** signal. The software is responsible for switching to the **SYSPLL** only when the PLL provides a clock and is locked to the frequency. **Note:** If the System PLL does not provide a lock signal. Then it is **IMPLEMENTATION DEFINED** how the software determines when the PLL is stable enough to be used.	Yes
SWCLKTCK:	**SWCLKTCK** is the clock that drives the JTAG and *Serial Wire Debug* (SWD) interface.	Yes
TRACECLKIN:	**TRACECLKIN** is the clock that drives the TPIU. **TRACECLKIN** is required when the TPIU must export the trace. The integrator needs to ensure that **TRACECLKIN** is provided in the SoC.	Yes

Clock name	Description	Mandatory
UARTCLK:	The **UARTCLK** is the clock that generates the **HOSTUARTCLK** that is used by the Host System UARTs. The Host System UART has a programmable Baud rate. The maximum baud rate is max(**REFCLK**, **S32KCLK**, **UARTCLK**)/16. For more information, see the *PrimeCell UART (PL011) Technical Reference Manual*. The **UARTCLK** is always available when SSE-710 is in the BSYS.RUN and BSYS.SLEEP0 power states.SSE-710 supports both **UARTCLK** and **S32KCLK** as source for UARTs when SSE-710 is in SLEEP1 power state to enable wakeup from either of the Host UARTs. If the **UARTCLK** is not provided in the BSYS.SLEEP1 state and wakeup from either Host UART is required, software must select the **S32KCLK** as the source for the **HOSTUARTCLK**.	Yes
External System Clock Inputs	These are provided for logic inside the External System harness. See Table 5-2: External System Clock Inputs on page 73 Each of the external system clocks support high-level clock gating and a corresponding clock Q-Channel is provided. All the EXTSYS{0-1}{x} clocks must be provided whenever the clock Q-Channel QACTIVE is HIGH.	Yes

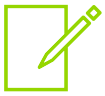
Note

You can drive multiple clock inputs from the same clock source, depending on SoC requirements.

Table 5-2: External System Clock Inputs

Clock name	Description
EXTSYS{0-1}DBGCLKS	Clock used for the EXTSYS{0-1}DBG interface
EXTSYS{0-1}DBGCLKM	Clock used for the EXTSYS{0-1}EXTDBG interface
EXTSYS{0-1}ATCLK	Clock used for the EXTSYS{0-1}TRACEEXP interface
EXTSYS{0-1}CTICLK	Clock used for the EXTSYS{0-1}CTICHIN/EXTSYS{0-1}CTICHOUT interfaces
EXTSYS{0-1}ACLK	Clock used for the EXTSYS{0-1}MEM interface
EXTSYS{0-1}MHUCLK	Clock used for the EXTSYS{0-1}MHU interface

5.2 Internal clocks

There are multiple internal clocks within the SSE-710:

The following figure shows the clock domain mapped over the SSE-710 topology diagram.

Figure 5-1: Clock domains of SSE-710

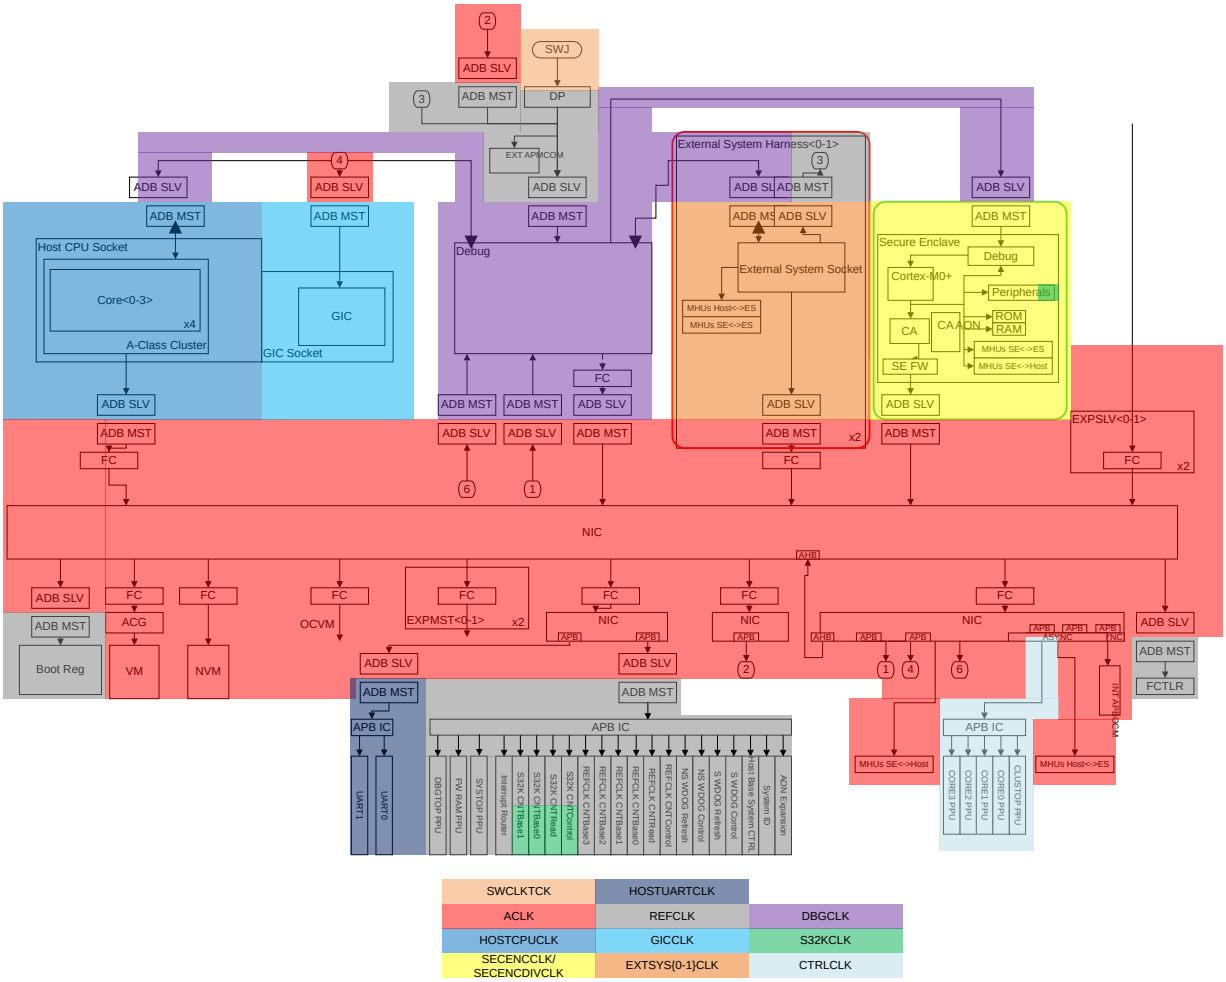

Note

For details of Clock inputs, see 5.1 Clock inputs on page 72.

5.2.1 ACLK

The following table describes the **ACLK** clock.

Table 5-3: ACLK summary table

Clock name	ACLK		
Sources	Name	Default	Divider support
	REFCLK	Yes	No
	SYSPLL	No	Yes

Clock name	ACLK
High-level clock gating support	Yes
Components	• NIC(Main) • NIC(AONPERIPH) • NIC(SYSPERIPH) • NIC(DBGPERIPH) • The following are Firewall Components of the Host System firewall: ◦ SYSPERIPH ◦ DBGPERIPH ◦ AONPERIPH ◦ XNVM ◦ CVM ◦ HOSTCPU ◦ EXTSYS{0-1} ◦ EXPSLV{0-1} ◦ EXPMST{0-1} ◦ OCVM • Host System side of MHUs with External Systems and Secure Enclave • INT APBCOM

5.2.2 GICCLK

The following table describes the **GICCLK** clock.

Table 5-4: GICCLK summary table

Clock name	GICCLK		
Sources	Name	Default	Divider Support
	REFCLK	Yes	No
	SYSPLL	No	Yes
High-level clock gating support	No		
Components	GIC		

5.2.3 DBGCLK

The following table describes the **DBGCLK** clock.

Table 5-5: DBGCLK summary table

Clock name	DBGCLK		
Sources	Name	Default	Divider support
	REFCLK	Yes	No

Clock name	DBGCLK		
	SYSPLL	No	Yes
High-level clock gating support	Yes		
Components	• Host ETR • Host CATU • Host STM • Host CTI • Host Funnel • Host Replicator • Host AXI AP • Host APB AP • Host AXIAP ROM • Host ROM • Host CTM • SoC CTM • SoC TPIU Replicator • SoC TPIU • SoC ETR • SoC CATU • Counter CTI • Debug Firewall Component of the Host System firewall • Channel Gates for the following DAZs: ◦ HOSTAUTH ◦ TPIUAUTH ◦ COUNTERAUTH		

5.2.4 SECENCCLK

SECENCCLK enables you to integrate the Crypto Accelerator into SSE-710 using **CRYPTOCLKOUT** and **CRYPTOAONCLKOUT**. They cannot be used for any other logic.

Table 5-6: SECENCCLK

Clock name	SECENCCLK		
Sources	Name	Default	Divider Support
	SECENCREFCLK	Yes	No
	SYSPLL	No	Yes
High-level clock gating support	No		
Components	Crypto Accelerator Secure Enclave firewall (FCTLR and FC1)		

Inside the Secure Enclave, there are several gated versions of **SECENCCLK** to reduce dynamic power.

5.2.5 SECENCDIVCLK

SECENCDIVCLK is an integer divided version of **SECENCCLK**. It can be configured to be either 1:1 or 1:2, depending on the programming of the Clock Divider Control register in the Secure Enclave System Control Unit. The default value is 2.

The following table summarizes **SECENCDIVCLK**.

Table 5-7: SECENCDIVCLK

Clock name	SECENCDIVCLK		
Sources	Name	Default	Divider Support
	SECENCCLK	Yes	Yes
High-level clock gating support	No		
Components	Secure Enclave Cortex®-M0+ Secure Enclave RAM Secure Enclave ROM Secure Enclave Watchdog Secure Enclave Timer 0 and 1 Sender frames of the SEH{0-1} and SEES{0-1}{0-1} MHUs Receiver frame of the HSE{0-1} and ES{0-1}SE{0-1} MHUs Secure Enclave System Control Secure Enclave Base System Control Secure Enclave UART Secure Enclave AHB AP Secure Enclave CS ROM Secure Enclave CTI Channel Gate for SECENCAUTH SECENC PPU		

5.2.6 HOSTCPUCLK

The following table describes the **HOSTCPUCLK** clock.

Table 5-8: HOSTCPUCLK summary table

Clock name	HOSTCPUCLK		
Sources	Name	Default	Divider support
	REFCLK	Yes	No
	SYSPLL	No	Yes
	CPUPLL	No	Yes
High-level clock gating support	No		
Components	Host CPU		

HOSTCPUCLK is provided as an output clock, **HOSTCPUCLKOUT**, which is only used within the CLUSTOP and CORE{x} power domains.

5.2.7 CTRLCLK

The following table describes the **CTRLCLK** clock.

Table 5-9: CTRLCLK summary table

Clock name	CTRLCLK		
Sources	Name	Default	Divider support
	REFCLK	Yes	No
	SYSPLL	No	Yes
High-level clock gating support	Yes		
Components	• CORE{0-3} PPU • CLUSTOP PPU		

5.2.8 REFCLK

The following table describes the **REFCLK** clock.

Table 5-10: REFCLK summary table

Clock name	REFCLK
Source	**REFCLK**
High-level clock gating support	Yes

Clock name	REFCLK
Components	• Interrupt Router • REFCLK CNTBase{0-3} • REFCLK CNTRead • REFCLK CNTControl • S32K CNTBase{0-1} • S32K CNTRead • S32K CNTControl • Non-secure Watchdog Refresh and Control • Secure Watchdog Refresh and Control • Host Base System Control • System ID • HOSTAONEXPMST • Firewall Controller of the Host System firewall • DBGTOP PPU • Firewall PPU • SYSTOP PPU • Boot Register

The **REFCLK** generates the **HOSTAONEXPCLK** and **HOSTCNTCLKOUT** output clocks.

5.2.9 S32KCLK

The following table describes the **S32KCLK** clock.

Table 5-11: S32KCLK summary table

Clock name	S32KCLK
Source	**S32KCLK**
High-level clock gating support	No
Components	• S32K CNTBase{0-1} • S32K CNTRead • S32K CNTControl • SoC Watchdog

The **S32KCLK** generates the **HOSTS32KCNTCLKOUT** output clock.

5.2.10 HOSTUARTCLK

The following table describes the **HOSTUARTCLK** clock.

Table 5-12: HOSTUARTCLK summary table

Clock name	HOSTUARTCLK		
Sources	Name	Default	Divider support
	REFCLK	Yes	No
	UARTCLK	No	Yes
	S32KCLK	No	No
High-level clock gating support	No		
Components	Host UART {0,1}		

5.3 Clock outputs

This section describes the clock outputs of SSE-710.

Table 5-13: Clock outputs

Clock name	Description
ACLKOUT	• Driven by **ACLK**. • Used only in SYSTOP power domain.
GICCLKOUT	• Driven by **GICCLK** • Used only in CLUSTOP power domain.
DBGCLKOUT	• Driven by **DBGCLK** • Used only in DBGTOP power domain.
CRYPTOCLKOUT	• Driven by **SECENCCLK** • Used only to integrate the Crypto Accelerator into the SECNECTOP power domain in the Secure Enclave.
CRYPTOAONCLKOUT	• Driven by **SECENCCLK** • Used only to integrate the Crypto Accelerator Always-on into the AONTOP power domain in the Secure Enclave.
HOSTAONEXPCLK	• Driven by **REFCLK** • Used only for logic connected to the **HOSTAONEXPMST**
HOSTCNTCLKOUT	• Driven by **REFCLK** • Used only for distribution of, and any additional timers added to, the **REFCLK** timestamp, provided by **HOSTTSVALUEG**.
HOSTCPUCLKOUT	• Driven by **HOSTCPUCLK** • Used only in CLUSTOP power domain.

Clock name	Description
HOSTS32KCNTCLKOUT	• Driven by **S32KCLK** • Used only for distribution of, and any additional timers added to, the S32K timestamp, provided by **HOSTS32KCNTVALUEG**
TRACECLKOUT:	• Driven by the **TPIU** • Half the frequency of **TRACECLKIN** • Used only for the TPIU interface

6. Power

This chapter describes the voltage and power domains within the SSE-710 subsystem.

6.1 Power overview

The SSE-710 supports multiple voltage and power domains.

Each voltage domain has at least one power domain, and each power domain supports several power modes.

6.2 PPU

Each power domain is controlled by a PCK-600 PPU that conforms to *Arm® Power Policy Unit Architecture Specification, version 1.1*. The PPU is responsible for controlling the power domain.

All PPUs support the following power modes: WARM_RST, ON, and OFF. Some PPUs support extra power modes.

PPUs have two modes of operation:

Static mode transitions

To calculate whether to remain in the current power mode or transition to the policy mode, the PPU uses:

- The policy in the PWR_POLICY register
- Hardware indications from components inside the power domain that use the **PACTIVE** or **QACTIVE** signals

In this mode of operation, the PPU can only transition to the programmed policy mode.

Dynamic mode transitions

The PPU uses the policy and hardware indications to calculate the lowest power mode the domain must be in. In this mode of operation, the PPU can transition to any of the required power modes autonomously without direct software intervention.

Arm strongly recommends that dynamic modes are enabled for all PPUs integrated within the SSE-710 subsystem by default. Software is not expected to require direct access to the PPU after performing an initial configuration at boot. Therefore, software is expected to use the Host CPU Cluster Power State register and Base System Power Request register to control the SSE-710 power modes and state power modes.

The PPU interrupts, excluding the SECENCTOP PPU, are combined to form a single interrupt called the Host PPU Combined interrupt. This interrupt is used by the Host CPU or Secure Enclave

depending on the settings of the Interrupt Router. Software uses the HOST_PPU_INT_ST register to identify which PPU has generated the interrupt. The interrupt from the SECENCTOP PPU is connected only to the Secure Enclave Cortex®-M0+ NVIC.

For more information on:

- The PCK-600 PPU, see the *Arm® CoreLink™ PCK-600 Power Control Kit Technical Reference Manual*.
- The PPU v1.1 architecture, see the *Arm® Power Policy Unit Architecture Specification, version 1.1*.
- The expected software usage of the PPU in the SSE-710, see 14.1 Power control on page 293.

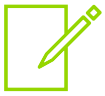
Note

Arm strongly recommends that software:

- Does not set the PPU to static mode transitions. This can lead to loss of functionality or deadlock of the SoC.
- Never programs the policy of the PPU to WARM_RST. This can lead to deadlock of the SoC.

PPU configurations

The PPUs within the SSE-710 are configured with the following:

- P-Channels for the device interfaces
- Default power mode of dynamic OFF
- All power modes are enabled for both static and dynamic transitions, unless stated otherwise in the following sections.
- No support for operating modes
- Support power mode entry delay, except Core{0-3} and CLUSTOP PPUs

PPU debug

The PPU includes the registers PPU_PWCR and PPU_PTCR for debugging power issues during SoC development.

Bit[35] of the *Security Control Bits* (SCB) interface controls whether these registers are read/write or are treated as WI and generate an error. When the SCB is HIGH, writes to the registers are allowed. For more information on the SCBs, see 9.1.3 Security Control Bits (SCB) on page 150


Note

Using PPU_PWCR or PPU_PTCR can lead to **UNPREDICTABLE** results depending on the power transitions being attempted. Arm strongly recommends that you use these registers for debugging purposes only.

6.3 Voltage domains

SSE-710 has the multiple voltage supplies.

Table 6-1: Voltage domains

Voltage domain name	Description
VSYS	• Voltage supply for AONTOP, SYSTOP, SECENCTOP, and DBGTOP power domains. • Voltage supply for firewall RAM domain. • Voltage supply is always available, unless SoC is in BSYS.OFF power state.
VCLUS	• Voltage supply for CLUSTOP and CORE{0-3} power domains • DVFS supported • Voltage supply is automatically enabled when domain requires it without any software intervention.
VEXTSYS{0-1}	• Voltage supply for EXTSYS{0-1}TOP power domain • DVFS supported • Voltage supply is automatically enabled when domain requires it without any software intervention.


Note

It is legal for an implementation to drive all voltages from a single voltage supply, VSYS.

6.4 Power domains

The SSE-710 has multiple power domains.

The power domains are mapped over a topology diagram of the SSE-710. The following figure shows the power domains that are supported by the SSE-710.

Figure 6-1: Power domains of SSE-710

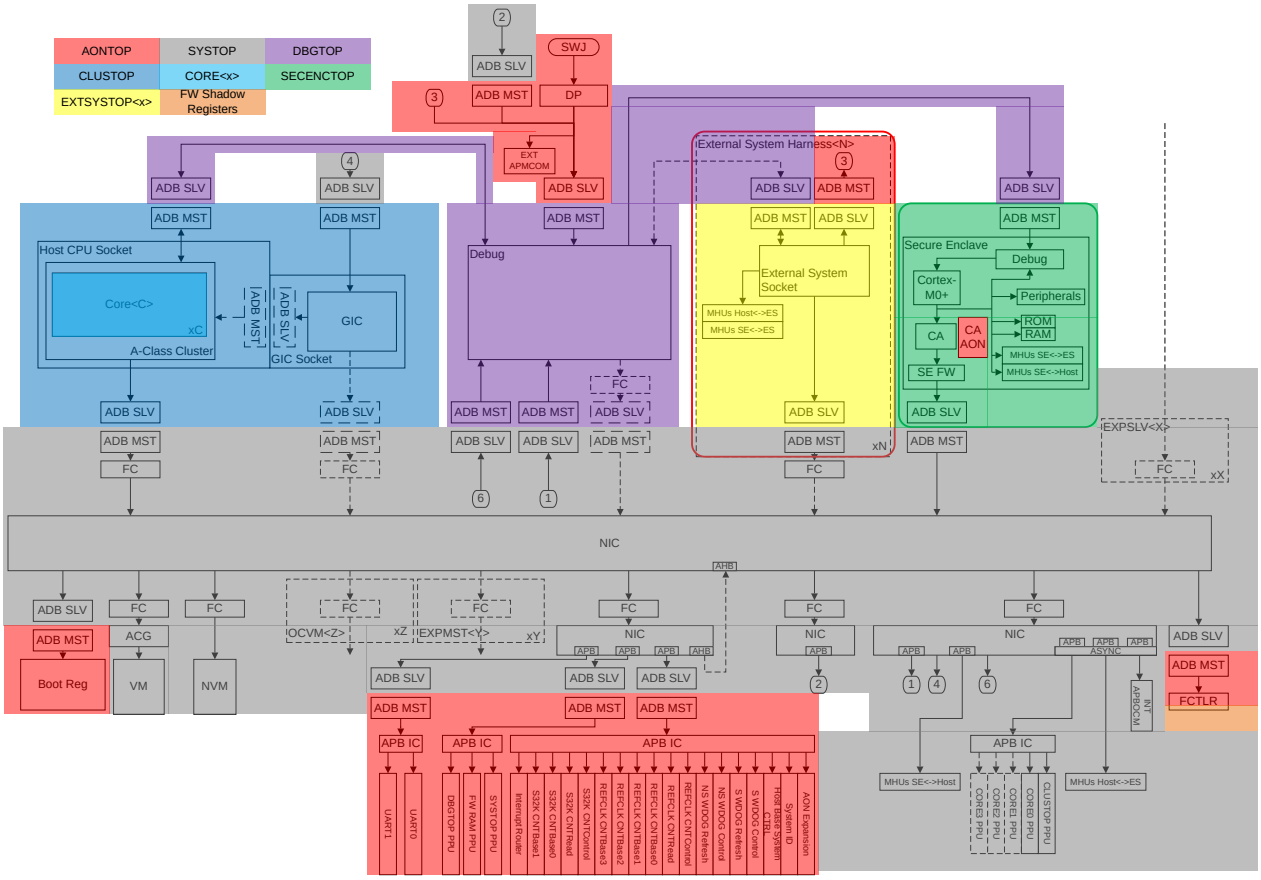

The following subsections provide information about:

- The components that are contained within each power domain
- The power modes that are supported by the domain
- The conditions on interfaces into and out of the domain, when in a specific power mode

6.4.1 AONTOP

AONTOP is the always-on power domain in the system.

AONTOP contains the following components:

- Generic counters for REFCLK and S32K time domains
- Generic timers for both REFCLK and 32K time domains
- Secure and Non-secure Watchdogs
- Interrupt Router
- Host Base System Control

- System ID
- Host System Firewall Controller
- CoreSight DP and ROM table
- EXT APBCOM
- GPIO Control
- Wakeup logic
- Firewall, SYSTOP, DBGTOP PPUs, and PCSMs.
- Firewall Controller
- Boot register
- PCSM of the CLUSTOP PPU


Note

This does not include the CLUSTOP PPU itself.

- SoC Watchdog
- Secure Enclave UART
- SECENCTOP PPU and PCSM
- Crypto Accelerator Always-on
- Secure Enclave System and Base System Control

The AONTOP domain is not controlled by a PPU. Instead, the logic is always powered whenever VSYS is applied.

6.4.1.1 Firewall RAM

The firewall PPU only controls the Host System Firewall shadow registers. These registers are implemented as RAM.

The firewall PPU supports the following power modes:

ON
: The Host System firewall shadow registers are powered.

OFF
: The Host System firewall shadow registers are not powered.

FUNC_RET
: The Host System firewall shadow registers are in a retention state.

WARM_RST
: It has no effect on the Host System firewall and is considered the same as ON.

The following applies when firewall PPU is in, or attempts to enter, the following modes:

Table 6-2: Firewall RAM power mode behavior

Power mode	Behavior
ON	Any access to the shadow registers of the Host System firewall is allowed.
OFF	The software must never request entry into the OFF power mode.
FUNC_RET	Any access to the shadow registers of the Host System firewall is stalled.
WARM_RST	The software must never request entry into the WARM_RST power mode.

6.4.2 SYSTOP

The SYSTOP power domain contains multiple components.

SYSTOP contains the following components:

- NIC-450 (Main/AONPERIPH/SYSPERIPH/DBGPERIPH)
- INT APBCOM
- Sender frames of the following MHUs, if they are present:
  - HSE{0-1}
  - HES{0-1}{0-1}
- Receiver frames of the following MHUs, if they are present:
  - SEH{0-1}
  - ES{0-1}H{0-1}
- The following Host System Firewall Components:
  - SYSPERIPH
  - DBGPERIPH
  - AONPERIPH
  - XNVM
  - CVM
  - HOSTCPU
  - EXTSYS{0-1}
  - EXPSLV{0-1}
  - EXPMST{0-1}
  - OCVM
- CLUSTOP and CORE{0-3} power control logic: PPU and PCSM except for CLUSTOP PCSM which is in AONTOP.


Note

The SYSTOP domain is further split into logic and memory subdomains. The memory domain only contains the on-chip volatile memory. All other logic is in the logic subdomain.

The SYSTOP power domain is controlled by the SYSTOP PPU, which supports the following power modes:

WARM_RST
SYSTOP is powered ON but all logic is held in reset.

ON
Logic and the volatile memory are powered.

FUNC_RET
Logic is powered, but the volatile memory is in retention.

MEM_RET
Logic is not powered, but the volatile memory is in retention.

The following applies when the SYSTOP domain is in the various modes:

Table 6-3: SYSTOP power mode behavior

Power mode	Behavior
OFF	• Any memory accesses from following masters or interfaces, listed below, to the Host System address space are stalled and a wakeup is generated to enter the FUNC_RET power mode: ◦ Host ETR and CATU ◦ SoC ETR and CATU ◦ AXI-AP ◦ External System {0-1} ◦ Secure Enclave • BSYS_PWR_ST.SYSTOP_PWR_ST field is set to `0b000` in both the Host and Secure Enclave Base System Control registers.
MEM_RET	• Any memory access from the following masters or interfaces, listed below, to the Host System address space is stalled and a wakeup is generated to enter the FUNC_RET power mode: ◦ Host ETR and CATU ◦ SoC ETR and CATU ◦ AXI-AP ◦ External System {0-1} ◦ Secure Enclave • BSYS_PWR_ST.SYSTOP_PWR_ST field is set to `0b001` in both the Host and Secure Enclave Base System Control registers.

Power mode	Behavior
FUNC_RET	• Any memory access to the On-chip Volatile Memory is stalled at the ACG on the CVM interface, and a request is generated to enter the ON power mode. • Any memory access from the following masters or interfaces to the Host System address space is allowed: ◦ Host CPU ◦ Host GIC ◦ Host ETR and CATU ◦ SoC ETR and CATU ◦ AXI-AP ◦ External System {0-1} ◦ Secure Enclave • BSYS_PWR_ST.SYSTOP_PWR_ST field is set to `0b010` in both the Host and Secure Enclave Base System Control registers.
ON	• Any memory access to the On-chip Volatile Memory is allowed through the ACG on the CVM interface. • Any memory access from the following masters or interfaces to the Host System address space is allowed: ◦ Host CPU ◦ Host GIC ◦ Host ETR and CATU ◦ SoC ETR and CATU ◦ AXI-AP ◦ External System {0-1} ◦ Secure Enclave • BSYS_PWR_ST.SYSTOP_PWR_ST field is set to `0b100` in both the Host and Secure Enclave Base System Control registers.
WARM_RST	• Software must never request entry into the WARM_RST power mode because it cannot cause the SYSTOP power domain to exit the mode. • Any memory access to the On-chip Volatile Memory is stalled at the ACG. • BSYS_PWR_ST.SYSTOP_PWR_ST field is set to `0b000` in both the Host and Secure Enclave Base System Control registers.

The SYSTOP power domain can be extended by an implementation, using the Q-Channels of the SYSTOPQ interface. For more information on the SYSTOPQ interface, see 4.8 Power control interfaces on page 70.

6.4.3 DBGTOP

The DBGTOP power domain contains multiple components

The DBGTOP power domain contains the following components:

- Host APB AP
- Host AXI AP
- Host AXIAP ROM
- Host ROM
- Host ETR
- Host CATU
- Host STM
- Host CTI
- Host Funnel
- Host Replicator
- Host CTM
- SoC TPIU Funnel
- SoC TPIU Replicator
- SoC TPIU
- SoC CTI
- SoC ETR
- SoC CATU
- SoC CTM
- Counter CTI
- EXTDBG ROM
- Debug Firewall Component, of the Host System Firewall
- Channel Gates for the following DAZs:
  - HOSTAUTH
  - TPIUAUTH
  - COUNTERAUTH

The DBGTOP power domain is controlled by the DBG PPU, which supports the following power modes:

ON
Debug logic is powered.

OFF
Debug logic is not powered.

WARM_RST

Reset all logic in the domain.

The following applies when the DBGTOP domain is in the various power modes:

Table 6-4: DBGTOP power mode behavior

Power mode	Behavior
ON	• Any access to the Host System Debug is allowed from any debug agent. • Any access to the External Debug Bus is allowed. • Any trace data from the following interfaces is accepted: ◦ HOSTCPUTRACE ◦ EXTSYS{0-1}TRACEEXP ◦ HOSTDBGTRACEEXP • Any CTI events are allowed on the following interfaces or components: ◦ HOSTCPUCTICH{IN,OUT} ◦ EXTSYS{0-1}CTICH{IN,OUT} ◦ HOSTCTICH{IN,OUT}EXP ◦ Secure Enclave CTI • BSYS_PWR_ST.DBGTOP_PWR_ST is `0b1` in both the Host and Secure Enclave Base System Control registers. • Accesses to the Host System memory map are allowed.

Power mode	Behavior
OFF	• Any access to the Host System Debug, from any debug agent, returns an error response. • Any access to the External Debug Bus, except for the DP ROM or EXT APBCOM, from any debug agent, returns an error response. • Any access to the CoreSight™ STM-500 Extended Stimulus Port is treated as RAZ/WI with no error response. • Any trace data from the following interfaces is ignored: ◦ HOSTCPUTRACE ◦ EXTSYS{0-1}TRACEEXP ◦ HOSTDBGTRACEEXP. • Any CTI events from the following interfaces or components are ignored: ◦ HOSTCPUCTICH{IN,OUT} ◦ EXTSYS{0-1}CTICH{IN,OUT} ◦ HOSTCTICH{IN,OUT}EXP ◦ Secure Enclave CTI • BSYS_PWR_ST.DBGTOP_PWR_ST is `0b0` in both the Host and Secure Enclave Base System Control registers.
WARM_RST	• Any access to the Host System Debug, from any debug agent, returns an error response. • Any access to the External Debug Bus, except for the DP ROM or EXT APBCOM, from any debug agent, returns an error response. • Any access to the CoreSight™ STM-500 Extended Stimulus Port is treated as RAZ/WI with no error response. • Any trace data from the following interfaces is ignored: ◦ HOSTCPUTRACE ◦ EXTSYS{0-1}TRACEEXP ◦ HOSTDBGTRACEEXP • Any CTI events are ignored from the following interfaces or components: ◦ HOSTCPUCTICH{IN,OUT} ◦ EXTSYS{0-1}CTICH{IN,OUT} ◦ HOSTCTICH{IN,OUT}EXP ◦ Secure Enclave CTI • BSYS_PWR_ST.DBGTOP_PWR_ST is `0b0` in both the Host and Secure Enclave Base System Control registers.

The debug agent is required to power on the DBGTOP domain, before attempting to configure or use the debug logic in the DBGTOP power domain, using one of the following methods:

- Setting the DP ROM **CDBGPWRUPREQ0** HIGH and waiting for **CDBGPWRUPACK0** to go HIGH
- Using a software **IMPLEMENTATION DEFINED** method
- Setting the BSYS_PWR_REQ.DBGTOP_PWR_REQ to `0b1` and waiting for BSYS_PWR_ST.DBGTOP_PWR_ST to become `0b1`


Note

This can be done in either the Host or Secure Enclave Base System Control registers.


Note

Arm recommends that:

- Debug agents using JTAG/SWD use the DP ROM **CDBGPWRUPREQ0/ACK0** method.
- All other debug agents use the BSYS_PWR_REQ and BSYS_PWR_ST registers, or the software **IMPLEMENTATION DEFINED** method.

The DBGTOP power domain can be extended by an implementation, using the Q-Channels of the DBGTOPQ interface. For more information on the DBGTOPQ interface, see 4.8 Power control interfaces on page 70.

6.4.4 CLUSTOP

The CLUSTOP power domain contains the Host CPU cluster logic and Host GIC, and supports multiple power modes.

The CLUSTOP power domain contains the following logic:

- Host CPU cluster logic and L2 cache: The Host CPU cluster logic includes other components. For more information, see the respective Technical Reference Manual for the Host CPU.
- Host GIC.

The CLUSTOP PPU is in the SYSTOP power domain. However, the PCSM associated with the PPU is in the AONTOP domain. The CLUSTOP PPU supports direct transitions from OFF to MEM_RET. This enables the PPU to be restored back to the power mode it was in before SYSTOP entered the OFF-power mode.

The CLUSTOP domain is controlled by CLUSTOP PPU, and supports the following power modes:

OFF

All logic and L2 cache is not powered.

MEM_RET

Enters L2 cache RAMs into retention. All other logic is not powered.

FUNC_RET

Enters L2 cache RAMs into retention. All other logic remains powered.

ON

All logic and L2 cache RAMs are powered

WARM_RST

WARM_RST power-mode is intended to be used for debugging only. It is not used during normal operation

All logic is powered, but the functional logic is held in reset. The debug logic within the Host CPU Cluster is available.

When CLUSTOP transitions into WARM_RST power mode, the logic within CORE{0-3} power domains are also reset.

SSE-710 only supports CLUSTOP transition into WARM_RST power mode if CORE{0-3} power domains are in OFF mode. Therefore, Arm® strongly recommends that the CORE{0-3} PPUs are in OFF power mode when you set the policy to WARM_RST for the CLUSTOP PPU.

WARM_RST can only be entered using static transitions. The software requesting these transitions must not be executing on the Host CPU, otherwise, a deadlock occurs.

The behavior depends on the power mode of the CLUSTOP domain.

Table 6-5: Power mode behavior

Power mode	Behavior
OFF	• An access on the HOSTCPUDBG interface gets an error response and does not cause a change in the power mode. • A CTI event on the HOSTCPUCTICHIN/OUT interfaces is ignored and does not cause a change in the power mode. • An access to the GICM interface is stalled and a request is generated to enter the FUNC_RET power mode. **Note:** SSE-710 supports the L2 cache flush using the **L2FLUSHREQ** and **L2FLUSHACK** signals when CLUSTOP transition into the OFF power mode. For more information on L2 Cache flush, see the Technical Reference Manual of one of the supported Cortex®-A processors. **Note:** The software must manually clean any dirty cache lines in the L2 for any locations that cannot be written using the same address that was used to fetch the initial data, for example, data initially fetched from a NAND flash device that has been modified.

Power mode	Behavior
MEM_RET	• An access on the HOSTCPUDBG interface generates an error response and does not cause a change in the power mode. • A CTI event on the HOSTCPUCTICHIN/OUT interfaces is ignored and does not cause a change in the power mode. • An access to the GICM interface is stalled and causes a request to enter the FUNC_RET power mode. • When CLUSTOP is exiting the MEM_RET mode the contents of the L2 cache is preserved in SSE-710.
FUNC_RET	• An access on the HOSTCPUDBG interface is allowed to the Host CPU external Debug registers. • A CTI event on the HOSTCPUCTICHIN/OUT interface is allowed. • An access to the GICM interface is allowed. • Access on the HOSTCPUMEM interfaces is allowed.
ON	• An access on the HOSTCPUDBG interface is allowed to the Host CPU external Debug registers. • A CTI event on the HOSTCPUCTICHIN/OUT interface is allowed. • Access to the GICM interface is allowed. • Access on the HOSTCPUMEM interfaces are allowed.
WARM_RST	• An access on the HOSTCPUDBG interface generates an error response and does not cause a change in the power mode. • A CTI event on the HOSTCPUCTICHIN/OUT interfaces is ignored and does not cause a change in the power mode. • An access to the GICM interface is stalled and causes a request to enter the FUNC_RET power mode.

6.4.5 CORE{0-3}

Each CORE{0-3} power domain contains a single Host CPU core.

The number of CORE power domains equals the number of cores in the Host CPU. Each CORE power domain is controlled by an associated CORE{0-3} PPU, which supports the following power modes:

ON

Resets the functional logic in the CORE domain.

OFF

All logic in the CORE domain is not powered.

FULL_RET

All logic in the CORE domain is in a retention state.

OFF_EMU

All logic in the CORE domain is powered. However, the functional logic is held in reset.

WARM_RST

Resets the functional logic in the CORE domain.

The CORE PPU and PCSMs are in the SYSTOP power domain.

The behavior depends on the power mode of the CORE{0-3} domain:

Table 6-6: Power mode behavior

Power mode	Behavior
ON	Any access to the debug logic within the Host CPU core {x} is allowed.
OFF	Any access to the debug logic within the Host CPU core {x} generates an error response.
FULL_RET	Any access to the debug logic within the Host CPU core {x} is allowed.
OFF_EMU	Any access to the debug logic within the Host CPU core {x} is allowed.
WARM_RST	Any access to the debug logic within the Host CPU core {x} is allowed.

WAKEUPREQ

When the **WAKEUPREQ[x]** signal is HIGH, there is a request to send an interrupt to the Host CPU core. The power control logic only responds to this signal in the OFF or OFF_EMU power modes. Otherwise it is ignored.

Core Primary Boot

The power control logic for CORE{0-3} power domains enables the selection of the Host CPU cores that automatically enter the ON power mode when the CLUSTOP power domain performs one of the following transitions:

- OFF to FUNC_RET/ON/WARM_RST.
- MEM_RET to FUNC_RET/ON/WARM_RST.

By default, CORE0 is selected, but software can configure this using the HOST_CPU_BOOT_MSK register in the Host Base System Control.


Note

Arm® recommends that only a single Host CPU core is enabled in the HOST_CPU_BOOT_MSK register at any given time.

6.4.6 SECENCTOP

This section describes the SECENCTOP power domain.

The SECENCTOP power domain contains the following components:

- Secure Enclave Cortex®-M0+
- Secure RAM and ROM
- Secure Enclave CMSDK Timer 0 and 1
- Secure Enclave Watchdog
- Sender frame of the SEH{0-1} and SEES{0-1}{0-1} MHUs
- Receiver frame of the HSE{0-1} and ES{0-1}SE{0-1} MHUs
- Secure Enclave Firewall
- Crypto Accelerator
- Secure Enclave AHB AP
- Secure Enclave CS ROM
- Secure Enclave CTI
- Channel Gate for SECENCAUTH DAZ

The SECENCTOP power domain supports the following power modes:

ON

All logic and Secure Enclave RAM is powered.

OFF

All logic, including the RAM is powered off.

MEM_RET

Secure Enclave RAM is in a retention state, all other logic is powered off.

WARM_RST

WARM_RST is supported by the SECENCTOP PPU, but Arm® strongly recommends that software never directs the PPU to WARM_RST.

The behavior depends on the power mode of the SECENCTOP domain.

Table 6-7: Power mode behavior

Power mode	Behavior
ON	• Any debug access to the Cortex®-M0+ is allowed. • A CTI event to or from the Secure Enclave CTI is allowed. • Any access to the Secure Enclave RAM is allowed.

Power mode	Behavior
OFF	• Any debug access to the Cortex®-M0+ generates an error. • A CTI event to or from the Secure Enclave CTI is ignored and does not cause a change in the power mode.
MEM_RET	• Any debug access to the Cortex®-M0+ generates an error. • A CTI event to or from the Secure Enclave CTI is ignored and does not cause a change in the power mode.
WARM_RST	• Software must never request entry into the WARM_RST power mode because it is unable to cause the SECENCTOP power domain to exit the WARM_RST power mode. • Any debug access to the Cortex®-M0+ generates an error. • A CTI event to or from the Secure Enclave CTI is ignored and does not cause a change in the power mode.


Note

Software is responsible for guaranteeing that the Crypto Accelerator is idle before entering into any power mode other than ON.

6.4.7 External System Power Domains

The section describes the external (EXTSYS{0-1}TOP) power domains.

Each External System has a power domain, referred to as EXTSYS{0-1}TOP. It is **IMPLEMENTATION DEFINED**, and includes the following logic:

- Sender frames of MHUs between the External and Host System. (ES{x}H{0-1}, where x is the External System number).
- Sender frames of MHUs between the External System and Secure Enclave. (ES{x}SE{0-1}, where x is the External System number).
- Receiver frames of MHUs between the Host and External System. (HES{x}{0-1}, where x is the External System number).
- Receiver frames of MHUs between the Secure Enclave and External System. (SEES{x}{0-1}, where x is the External System number).
- External System logic, which is **IMPLEMENTATION DEFINED**.
- Domain crossing logic for the following interfaces:
  - EXTSYS{0-1}MEM.
  - EXTSYS{0-1}MHU.
  - EXTSYS{0-1}TRACEEXP.

- EXTSYS{0-1}DBG.
- EXTSYS{0-1}EXTDBGAPB.
- EXTSYS{0-1}CTICHIN.
- EXTSYS{0-1}CTICHOUT.

An error is generated for an access to the External System {x} region of the External Debug Bus memory map. When the EXTSYS{x}TOP power domain is in a power mode where the EXTSYS{0-1}DBGQ interface is in the Q_STOPPED state.

For more information on the External Debug Bus memory map, see 12.1.2 External Debug Bus memory map on page 193.

6.5 Power domain relationship

There are different relationships between specific power domains.

The relationships between specific power domains in the SSE-710 are:

Table 6-8: Power domain relationships

Power domain	Relationship
AONTOP	Relatively always-on to all other power domains in the SoC. This includes any other power domains added by the SoC implementor.
SYSTOP	• Relatively always-on to CLUSTOP. CLUSTOP must be in OFF or MEM_RET power mode before SYSTOP enters the OFF or MEM_RET power mode. • Independent relationship with DBGTOP. • Independent relationship with EXTSYS{0-1}TOP.
DBGTOP	Independent relationship with CLUSTOP, SYSTOP, SECENCTOP, and EXTSYS{0-1}TOP.
SECENCTOP	• Independent relationship with SYSTOP and DBGTOP. • Independent relationship with EXTSYS{0-1}TOP.
CLUSTOP	Relatively always-on to CORE{0-3}. • When any CORE{0-3} is in any of the ON, FULL_RET, or OFF_EMU power modes, then CLUSTOP must be in ON or FUNC_RET power mode. • When all CORE{0-3} are in OFF power mode, then CLUSTOP can enter MEM_RET or OFF.
CORE{0-3}	Independent relationship between CORE{0-3} power domains.
EXTSYS{0-1}TOP	• Independent to SYSTOP, DBGTOP, and between EXTSYS{0-1} power domains. • Independent relationship with SECENCTOP.

6.6 Power states

The power states represent a software-visible abstraction of available power modes of each power domain.

A power state defines the wakeup capability, loss of context, and the power consumption. This section describes power states for theSSE-710, but the External Systems can define their own power states.

The power state of SSE-710 is controlled using the following registers:

- Host Base System Control registers:
  - HOST_CPU_CLUS_PWR_REQ
  - BSYS_PWR_REQ
- Secure Enclave Base System Control registers:
  - BSYS_PWR_REQ

The power state determines the behavior for shared resources, for example the SYSTOP and REFCLK. Alongside the shared resources, there are implemented behaviors for the dedicated resources, such as the CLUSTOP power domain of the Host System.

No system has overall control of the power state and shared resources. The Host System and Secure Enclave both have copies of the BSYS_PWR_REQ register to request the power state. The External System does not have control of the power state directly but can make software requests to the Host System or Secure Enclave. All copies of the BSYS_PWR_REQ register, plus other hardware indicators are used to determine the power state of the SoC.

Table 6-9: Power states of the SSE-710 subsystem

The following table defines the state of each power domain for each power state of the subsystem.

Power state	AON TOP	FW SR	SYS TOP	DBG TOP	CLUS TOP	SECEN CTOP	EXTSYS {0-1}TOP	REF CLK	S32K CLK	VSYS
BSYS.RUN	ON	ON/ FUNC _RET[a]	Any[b]	Any	Any[bcd]	Any	Any	Yes	Yes	ON
BSYS.SLEEP0	ON	ON/ FUNC _RET[a]	MEM_ RET / OFF	Any	MEM_ RET / OFF[cd]	Any	Any	Yes	Yes	ON
BSYS.SLEEP1	ON	ON/ FUNC _RET[a]	MEM_ RET / OFF	OFF	MEM_ RET / OFF[cd]	Any	Any	No	Yes	ON
BSYS.OFF	OFF	OFF[a]	OFF	OFF	OFF[cd]	OFF	OFF	No	No	OFF

[a] Host System Firewall RAM.
[b] SYSTOP and CLUSTOP cannot be in any of the following: SYSTOP and CLUSTOP in OFF; SYSTOP in OFF and CLUSTOP in MEM_RET; SYSTOP in MEM_RET and CLUSTOP in OFF; SYSTOP and CLUSTOP in MEM_RET.
[c] For the CLUSTOP to be in OFF or MEM_RET, all the CORE{x} power domains must be in OFF. The CORE{x} power domains can be in any legal power mode when CLUSTOP is in FUNC_RET or ON.
[d] Whenever the CLUSTOP domain is not in OFF the VCLUS, if implemented, must be ON.
[e] In BSYS.OFF, VSYS can be removed, but this is not required. Under some conditions, the system enters BSYS.OFF but VSYS is not removed, that is, **nSRST** is asserted.


Note

Cells with a value of “Any” mean that the power domain can enter any legal power mode for that domain.

In BSYS.SLEEP1 where **REFCLK** is marked as OFF, it is legal for **REFCLK** to remain running and be used by other components outside SSE-710, within the SoC. However, the **REFCLK** is internally gated by SSE-710. Any component in theSSE-710 using **REFCLK** no longer has a clock. For example, the REFCLK timestamp does not update and software must reprogram the time from the S32K counter on entry into BSYS.SLEEP0 or BSYS.RUN.

When transitioning between any power state, except all transitions to BSYS.OFF and BSYS.OFF-> BSYS.SLEEP1, **REFCLK** is required and is requested by SSE-710 using the REFCLKQ interface.

The following sections show the conditions required to transition to a power state.

Power state transition conditions for BSYS.RUN to BSYS.SLEEP0

This transition requires that all the following conditions are true:

- CLUSTOP power domain is in MEM_RET or OFF power mode.
- SYSTOP power domain is in MEM_RET or OFF power mode.
- Host or Secure Enclave Base System Control BSYS_PWR_REQ.REFCLK_REQ is `0b1`.
- Host and Secure Enclave Base System Control BSYS_PWR_REQ.SYSTOP_PWR_REQ is `0b000` or `0b001`.
- HOST_CPU_CLUS_PWR_REQ.PWR_REQ is `0b0`.
- When Host Base System Controller BSYS_PWR_REQ.WAKEUP_EN is 0b1, then the following conditions must be true:
  - GICWAKEUP is 0b0

  The following interrupts are not asserted:
  - Secure Watchdog WS0
  - Non-secure Watchdog WS0 or WS1
  - UART{0,1} UARTINTR.
  - HSE{0-1} Combined IR
  - SEH{0-1} Combined IRQ
  - HES{0-1}{0-1} Combined IRQ
  - ES{0-1}H{0-1} Combined IRQ
  - Host UART{0,1} UARTINTR

  The following interrupts are either not asserted or not routed to the Host GIC via the Interrupt Router:
  - REFCLK Timer {0-3} and S32KCLK Timer {0,1}

- CoreSight™ SDC-600
- Host System Firewall
- Host PPU Combined

When Host Base System Control BSYS_PWR_REQ.WAKEUP_EN is `0b0`, the following signals are considered to be `0b0`:

- GICWAKEUP
- Status of the following interrupts:
  - Secure Watchdog WS0
  - Non-secure Watchdog WS0 or WS1
  - HSE{0-1} Combined IRQ
  - SEH{0-1} Combined IRQ
  - HES{0-3}{0-1} Combined IRQ
  - ES{0-3}H{0-1} Combined IRQ
  - Host UART{0,1} UARTINTR
- Status of the following interrupts, independent of their routing via the Interrupt Router:
  - REFCLK Timer {0-3}
  - S32KCLK Timer {0,1}
  - CoreSight™ SDC-600
  - Host System firewall
  - Host PPU Combined

- Host ROM **CDBGPWRUPREQ0** is LOW.
- AXIAP ROM **CSYSPWRUPREQ[1:0]** are both LOW.
- There are no outstanding transactions from/on any of the following masters or interfaces to the Host System address space:
  - Host ETR or CATU
  - SoC ETR or CATU
  - AXI-AP
  - External System {0-1}
  - Secure Enclave and GICM
- CoreSight™ SDC-600 REMPUR is `0b0`.
- **CLUSTOPINGRESSQACTIVE** is LOW.
- **CLUSTOPEGRESSQACTIVE** is LOW.
- All SYSTOPQ interfaces **QACTIVE** is LOW.

Power state transition conditions for BSYS.SLEEP0 to BSYS.RUN

This transition is enabled when any of the following occurs:

- **GICWAKEUP** is HIGH and Host Base System Control BSYS_PWR_REQ.WAKEUP_EN is `0b1`.
- Any of the following interrupts are asserted and the Host Base System Control BSYS_PWR_REQ.WAKEUP_EN is `0b1`:
  - Secure Watchdog WS0.
  - Non-secure Watchdog WS1.
  - Non-secure Watchdog WS0.
  - HSE{0-1} Combined IRQ.
  - SEH{0,1} Combined IRQ.
  - HES{0-1}{0-1} Combined IRQ.
  - ES{0-1}H{0-1} Combined IRQ.
  - Host UART{0,1} UARTINTR.
- Any of the following interrupts are asserted, the Host Base System Control BSYS_PWR_REQ.WAKEUP_EN is `0b1` and the interrupt is routed to the Host GIC by the Interrupt Router:
  - REFCLK Timer {0-3}.
  - S32K Timer {0,1}.
  - Host System firewall.
  - Host PPU Combined.
  - CoreSight™ SDC-600.
- A memory access from the Secure Enclave or External Systems {0-1}.
- CoreSight™ SDC-600 REMPUR is `0b1`
- Host or Secure Enclave BSYS_PWR_REQ.SYSTOP_PWR_REQ is `0b010` or greater.
- Any of the following MHUs have ACCESS_REQUEST.ACC_REQ set to `0b1`:
  - SEH{0,1}.
  - ES{0-1}H{0,1}.
- Any SYSTOPQ interface QACTIVE is `0b1`.


Note

Wakeup sources which are from components which use the **REFCLK** can only generate wakeup events if the event occurred before the SSE-710 entered the SLEEP1 state. Software must not rely on these sources for wakeup from SLEEP1 and is expected to move any wakuep sources to components which have a clock or select a clock source for the clock which is running in that power state.

Power state transition conditions for BSYS.SLEEP0 to BSYS.SLEEP1

This transition is enabled when all the following conditions are true:

- All BSYS_PWR_REQ.SYSTOP_PWR_REQ are set to `0b00x`.
- CLUSTOP power domain is in MEM_RET or OFF power mode.
- SYSTOP power domain is in MEM_RET or OFF power mode.
- Both BSYS_PWR_REQ.REFCLK is set to `0b0`.
- DBGTOP power domain is in OFF power mode.
- DP **CDBGPWRUPREQ** is LOW.
- **GICWAKEUP** is LOW or Host Base System Control BSYS_PWR_REQ.WAKEUP_EN is 0b0.
- All of the following interrupts are not asserted or the Host Base System Control BSYS_PWR_REQ.WAKEUP_EN is `0b0`:
  - Secure Watchdog WS0.
  - Non-secure Watchdog WS1.
  - Non-secure Watchdog WS0.
  - HSE{0-1} Combined IRQ.
  - SEH{0,1} Combined IRQ.
  - HES{0-1}{0,1} Combined IRQ.
  - ES{0-1}H{0-1} Combined IRQ.
  - UART{0,1} UARTINTR.
- All of the following interrupts are not asserted or are not routed to the Host GIC using the Interrupt Router or the Host Base System Control BSYS_PWR_REQ.WAKEUP_EN is `0b0`:
  - REFCLK Timer {0-3}.
  - S32K Timer {0,1}.
  - Host System Firewall.
  - Host PPU Combined.
  - CoreSight™ SDC-600.
- No memory access from the Secure Enclave or any External Systems {0-1} to the Host System address space.
- Host or Secure Enclave BSYS_PWR_REQ.SYSTOP_PWR_REQ is `0b000` or `0b001`.
- All of the following MHUs have ACCESS_REQUEST.ACC_REQ set to `0b0`:
  - SEH{0,1}.
  - ES{0-1}H{0,1}.
- All SYSTOPQ interface **QACTIVE** is LOW.

Power state transition conditions for BSYS.SLEEP1 to BSYS.RUN

This transition is enabled when any of the following occurs:

- **GICWAKEUP** is HIGH and Host Base System Control BSYS_PWR_REQ.WAKEUP_EN is `0b1`.
- Any of the following interrupts are asserted and the Host Base System Control BSYS_PWR_REQ.WAKEUP_EN is `0b1`:

- Secure Watchdog WS0.
- Non-secure Watchdog WS1.
- Non-secure Watchdog WS0.
- HSE{0-1} Combined IRQ.
- SEH{0,1} Combined IRQ.
- HES{0-1}{0-1} Combined IRQ.
- ES{0-1}H{0-1} Combined IRQ.
- Host UART{0,1} UARTINTR.

- Any of the following interrupts are asserted, the Host Base System Control BSYS_PWR_REQ.WAKEUP_EN is `0b1` and the interrupt is routed to the Host GIC by the Interrupt Router:
  - REFCLK Timer {0-3}.
  - S32K Timer {0,1}.
  - Host System firewall.
  - Host PPU Combined.
  - CoreSight™ SDC-600.
- A memory access from the Secure Enclave or External Systems {0-1}.
- CoreSight™ SDC-600 REMPUR is `0b1`
- Host or Secure Enclave BSYS_PWR_REQ.SYSTOP_PWR_REQ is `0b010` or greater.
- Any of the following MHUs have ACCESS_REQUEST.ACC_REQ set to `0b1`:
  - SEH{0,1}.
  - ES{0-1}H{0,1}.
- Any SYSTOPQ interface QACTIVE is `0b1`.

Wakeup sources which are from components which use the **REFCLK** can only generate wakeup events if the event occurred before the SSE-710 entered the SLEEP1 state. Software must not rely on these sources for wakeup from SLEEP1 and is expected to move any wakeup sources to components which have a clock, or select a clock source for the clock which is running in that power state.

Power state transition conditions for BSYS.SLEEP1 to BSYS.SLEEP0

This transition is enabled when any of the following occur:

- DP **CDBGPWRUPREQ** is HIGH.
- Host or Secure Enclave Base System Control BSYS_PWR_REQ.DBGTOP_PWR_REQ field becomes `0b1`.
- Host or Secure Enclave Base System Control BSYS_PWR_REQ.REFCLK_REQ field becomes `0b1`.

Power state transition conditions for BSYS.OFF to BSYS.SLEEP1

This transition is enabled when VSYS is applied and **PORESETn** is deasserted.

The Power on Reset sequence causes SSE-710 to perform the following sequence: BSYS.OFF -> BSYS.SLEEP1 -> BSYS.SLEEP0.

Power state transition conditions for BSYS.{RUN/ SLEEP0/ SLEEP1} to BSYS.OFF

This transition is enabled when any of the following occur:

- SSE-710 **PORESETn** is asserted.
- SSE-710 **nSRST** input is asserted.
- Secure Enclave Watchdog reset request is asserted.
- SoC Watchdog reset request is asserted.
- DP ROM **CSYSRSTREQ** is HIGH.
- DP **CDBGRSTREQ** is HIGH.
- VSYS is removed.

Transition from BSYS.SLEEP1 to BSYS.RUN enters BSYS.SLEEP0 first, followed by a BSYS.SLEEP0 to BSYS.RUN transition.

When performing transitions from BSYS.RUN to BSYS.SLEEP0, when a change in the conditions during the transitions occurs, the transition completes first. Then, the conditions are re-evaluated and SSE-710 returns to BSYS.RUN. When performing transitions from BSYS.SLEEP0 to BSYS.SLEEP1, when a change in the conditions during the transitions occurs, whether or not the transition is completed depends on the timing of the arrival of the condition change:

- If a REFCLK request happens at the same time as REFCLKQ moves into Q_STOPPED, then SSE-710 enters into BSYS.SLEEP1.
- If a REFCLK request happens when REFCLKQ is in Q_REQUEST, then this request is denied, SSE-710 aborts the transition and returns to BSYS.SLEEP0.

6.6.1 Secure Enclave sleep states

The Secure Enclave has three sleep states: SLEEPING, SLEEPDEEP, and SLEPDEEP PG.

For details of Secure Enclave Sleep States, see the *Arm® Corstone™-1000 Cryptographic Extension Technical Reference Manual Addendum*.

6.6.2 External System power states

The power states of the External System are limited to the power domain EXTSYS{0-1}TOP and any other **IMPLEMENTATION DEFINED** power domains which are part of the External System.

The power states of the External System and the method by which a power state is selected is **IMPLEMENTATION DEFINED** as long as the following requirements are met:

- EXTSYS.OFF power state:
  - All power domains, within the External System are OFF.
  - The only exit condition from this state is via a power on reset of the External System.
  - If VEXT is implemented, then it can be turned OFF in this power state but must be able to be requested without software intervention.
- EXTSYS.RUN power state:
  - Must be able to perform its primary task, i.e. execute instructions if the **EXTSYS{0-1}CPUWAIT** signal is not asserted.
  - Debugger must be able to access the resources of the External System.
  - If VEXT is implemented, then it must be turned ON in this power state.
- Must be able to transition from EXTSYS.OFF to EXTSYS.RUN directly.
  - Part of the power on reset sequence for the External System.
- Must be able to transition from EXTSYS.RUN to EXTSYS.OFF directly.
  - In a controlled manner, using an **IMPLEMENTATION DEFINED** method. This is intended for a clean shutdown of the SSE-710.
  - Through any of the following:
    - **PORESETn** is asserted.
    - **nSRST** input is asserted.
    - Secure Enclave Watchdog reset request is asserted.
    - Secure Enclave Host System reset request is asserted.
    - DP ROM **CSYSRSTREQ** is asserted.
    - DP **CDBGRSTREQ** is asserted.
    - Removal of VSYS.
    - Removal of VEXT, if implemented.

Other power states can be implemented, between OFF and RUN, specific to the External System. For example, an EXTSYS.SLEEP can be defined where the External System enters a low-power state, with only minimal wakeup logic remaining powered, whilst the rest of the External System is not powered.

7. Reset

This chapter describes the reset requests, inputs, and outputs of the SSE-710 subsystem.

7.1 Reset overview

The SSE-710 reset strategy is built around a reset tree that incorporates a Reset Controller and the PPUs.

The Reset Controller handles top-level reset conditions, for example **PORSETn** or debug ROM table resets. The PPUs handle reset of the power domain that they control.

The following figure shows the SSE-710 reset tree.

Figure 7-1: Reset tree

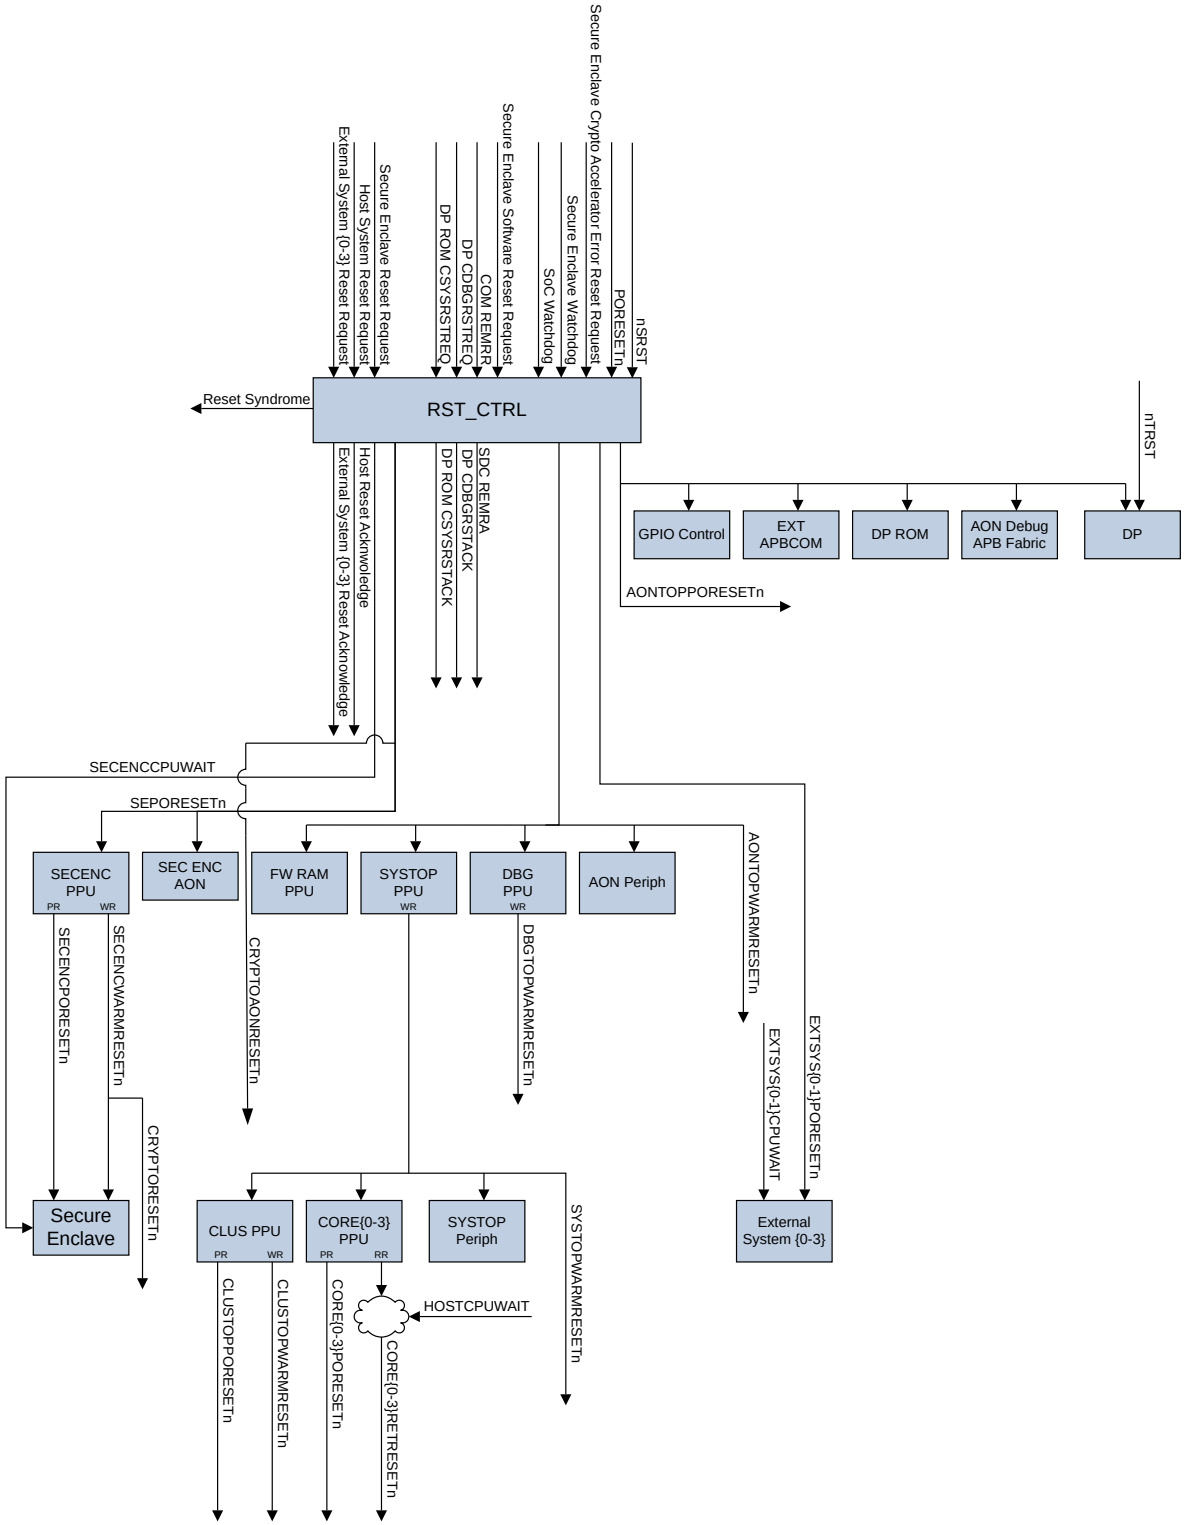

7.2 Reset inputs

All reset inputs to the SSE-710 subsystem are treated as asynchronous assert.

The **PORESETn** and **nTRST** resets are treated as asynchronous deassert, as the SSE-710 implements an internal reset synchronizer.

SSE-710 has the following reset inputs:

Table 7-1: Reset inputs

Reset name	Description
PORESETn	SoC Power-on reset
nTRST	Reset for the JTAG/Serial wire debug interface
EXTSYS{0-1}ARESETn	Reset for the EXTSYS{0-1}MEM interface
EXTSYS{0-1}MHURESETn	Reset for the EXTSYS{0-1}MHU interface
EXTSYS{0-1}DBGPRESETMn	Reset for the EXTSYS{0-1}DBG interface
EXTSYS{0-1}DBGPRESETSn	Reset for the EXTSYS{0-1}EXTDBG interface
EXTSYS{0-1}ATRESETn	Reset for the EXSYS{0-1}TRACEEXP interface
EXTSYS{0-1}CTIRESETn	Reset for the EXTSYS{0-1}CTICHIN and EXTSYS{0-1}CTICHOUT interfaces

7.3 Reset requests

Reset requests in SSE-710 provide different levels of reset.

Table 7-2: Reset Requests

Name	Type	Description
PORESETn	SSE-710 input	SoC Power-on reset. Top-level asynchronously asserted and deasserted active-LOW reset of SSE-710. The **PORESETn** input is synchronized within SSE-710 to the **S32KCLK** before being used, to reset the Reset Controller. When asserted, all logic in the SoC is reset. **Note:** **PORESETn** is used to reset the JTAG/SWD interface of the Debug port, however no synchronization of **PORESETn** takes place for the Debug port.
nTRST	SSE-710 input	JTAG/SWD interface reset. Top-level asynchronously asserted and deasserted active-LOW reset of SSE-710. The **nTRST** input is used to reset the JTAG/SWD interface of the Debug port.

Name	Type	Description
nSRST	SSE-710 input	SoC Warm Halted Reset. Top-level asynchronously assert and deasserted active-LOW reset request. The **nSRST** input signal is internally synchronized to the **S32KCLK**. When the input is asserted, all logic in the SoC is reset, except for the logic on the **AONTOPPORESETn**. After some cycles, which are configured through the SOC_RST_DLY parameter, the SoC is released from reset. However, the Secure Enclave Cortex®-M0+ is prevented from executing instructions while the **nSRST** signal is asserted LOW. This enables a debug agent to perform any initial setup of the debug logic of the SoC to facilitate debug from reset.
DP CDBGRSTREQ	SSE-710 internal	SoC Power-on reset request, from the Debug port. Equivalent to a **PORESETn**, except the Reset Controller is not reset.
DP CDBGRSTREQ	SSE-710 internal	DBGTOP reset. Debug reset request from the DP ROM. Causes a reset of the DBGTOP power domain.
DP ROM CSYSRSTREQ	SSE-710 internal	SoC Warm-halted reset. System reset request from DP ROM. Equivalent to **nSRST**. While **CSYSRSTREQ** is HIGH, the Secure Enclave Cortex®-M0+ is prevented from executing instructions.
Secure Enclave Crypto Accelerator Error Reset Request	SSE-710 internal	SoC Power-on reset. Reset request from the Secure Enclave indicating that a Crypto Accelerator Error is occurred. Equivalent to **PORESETn**, except the Reset Controller is not reset. For more information, see 4.5.1 Crypto Accelerator socket interfaces on page 61
Secure Enclave Watchdog Reset Request	SSE-710 internal	SoC Power-on reset Reset request from the Secure Enclave watchdog. Equivalent to **PORESETn**, except the Reset Controller is not reset. For more information on the use of the Secure Enclave Watchdog, see 14.5 Watchdog usage on page 304.
Secure Enclave Software Reset Request	SSE-710 internal	SoC Warm reset. Reset request from the Secure Enclave Cortex®-M0+ driven by the AIRCR.SYSRESETREQ field. For more information about **SYSRESETREQ** , see *Cortex®-M0+ Technical Reference Manual*. This reset is the equivalent to **nSRST**, except that the Secure Enclave Cortex®-M0+ is not prevented from executing instructions after the reset has been applied. For more information, see the *Arm®v6-M Architecture Reference Manual*.
SoC Watchdog Reset Request	SSE-710 internal	SoC Power-on reset Reset request from the SoC Watchdog. Equivalent to **PORESETn**, except the Reset Controller does not reset. For more information on the usage of the SoC Watchdog, see 14.5 Watchdog usage on page 304.

Name	Type	Description
Power domain PPU	SSE-710 internal	Power domain reset. As the PPU enter the different power modes, resets are applied to various parts of the logic. The WARM_RST power mode can be used on any power domain to reset certain logic within the domain. Typically, this causes the functional logic to be reset, while the debug logic is not reset. Arm® strongly recommends that you ensure that: • Software only sets the policy to WARM_RST when the domain is known to be idle. • When using WARM_RST policy, that it does not enter a deadlock condition because software cannot update the policy once in WARM_RST.
Software External System Reset Request	SSE-710 internal	In the Host Base System Control registers there is a register for each External System that is implemented to enable a software request to reset the associated External System. When software sets the EXT_SYS{0-1}_RST_CTRL.RST_REQ bit to `0b1`, the **EXTSYS{0-1}PORESETn** is asserted. This causes all logic within the External System to be reset. **Note:** It is not guaranteed that **EXTSYS{0-1}PORESETn** is asserted if the interfaces of the External System are not quiescent.
Software Host System Reset Request	SSE-710 internal	In the Secure Enclave Base System Control registers there is a register that enables software to request that the Host and all External Systems implemented are reset. When software sets the HOST_SYS_RST_CTRL.RST_REQ bit to `0b1`, the **AONTOPWARMRESETn** and **EXTSYS{0-1}PORESETn** are asserted. This causes all logic within the Host and External Systems to be reset. This also includes the debug logic of SSE-710, except for the External Debug Bus. **Note:** It is not guaranteed that the **AONTOPWARMRESETn** is asserted if the interfaces between the Secure Enclave and the Host System are not quiescent.
SoC Reset Request	SSE-710 internal	In the Secure Enclave Base System Control registers there is a register that enables software to request that the entire SoC is reset. When software sets the SOC_RST_CTRL.RST_REQ bit to `0b1`, all logic within the SoC is reset. This is equivalent to **PORESETn** being asserted, except that the Reset Controller is not reset.
DBGRSTREQ	Host CPU_GIC Socket input	Host CPU core reset. Request to reset the functional logic of the Host CPU core. Arm® strongly recommends that this reset is only used when the Host CPU core is known to be idle. Requesting a reset when the Host CPU core is not idle can lead to **UNPREDICTABLE** results.
WARMRSTREQ	Host CPU_GIC Socket input	Host CPU core reset. Request to reset the function logic of the Host CPU core. Software must execute a WFI before the reset of the Host CPU core is applied.

Name	Type	Description
External System Reset Inputs The External System Harness defines the following reset input signals, which reset interfaces of the External System: **EXTSYS{0-1}ARESETn** **EXTSYS{0-1}MHURESETn** **EXTSYS{0-1}DBGPRESETMn** **EXTSYS{0-1}DBGPRESETSn** **EXTSYS{0-1}ATRESETn** **EXTSYS{0-1}CTIRESETn**	External System Harness inputs	These resets are generated by power control logic of the External System and must be either directly or indirectly asserted when **EXTSYS{0-1}PORESETn** is asserted.

The difference between a SoC Warm reset and a SoC Warm-halted reset is as follows: during a SoC Warm-halted reset, the SoC is prevented from executing any instructions, while the conditions that caused it, or any other condition which can cause it, remain asserted.

7.4 Reset outputs

SSE-710 has several reset outputs. Each reset output hasr a specific function.

All reset outputs are asynchronously asserted and asynchronously deasserted. The resets are driven either:

- Directly from the **DEVPORESETn**, **DEVRETRESETn**, and **DEVWARMRESETn** outputs of the PPU
- From the Reset Controller

Arm® strongly recommends that reset outputs are only used as described in the following table.

Table 7-3: SSE-710 reset outputs

Reset Output	Usage
AONTOPPORESETn	This reset provides resets for additional components that are implemented within the AONTOP power domain. This components are not affected by the **nRSRT** reset.

Reset Output	Usage
AONTOPWARMRESETn	Resets extended logic implemented by partners within the AONTOP power domain, which is affected by the **nSRST** functionality.
CRYTPORESETn	Resets the Crypto Accelerator.
CRYPTOAONRESETn	Resets the Crypto Accelerator Always-on.
SYSTOPWARMRESETn	Resets any extended logic implemented by partners within the SYSTOP power domain.
DBGTOPWARMRESETn	Resets any logic within the DBGTOP power domain.
EXTSYS{0-1}PORESETn	Either directly or indirectly resets all logic within the External System, including reset control logic of the External System. An External System reset only reset logic on the **EXTSYS{0-1}PORESETn**. An External System reset is caused by software setting the EXT_SYS{0-1}_RST_CTRL.RST_REQ field to `0b1`. The logic within the External System can be reset using an internal request, for example a watchdog, but the interfaces with the logic outside the External System must be quiescent before performing the reset.
HOSTCLUSTOPPORESETn	Resets logic, within the CLUSTOP power domain, which must not be affected by a WARM_RST.
HOSTCLUSTOPWARMRESETn	Resets logic, within the CLUSTOP power domain, which must be affected by a WARM_RST.
HOSTCPUPORESETn[x:0]	Resets debug logic of the Host CPU Core{0-3}. **Note:** This is a superset reset of **HOSTCPUWARMRESETn[x:0]**.
HOSTCPUWARMRESETn[x:0]	Resets functional logic of the Host CPU core{0-3}.


Note

For **HOSTCPUPORESETn** and **HOSTCPUWARMRESETn**, x is equal to `HOST_CPU_NUM_CORES`-1.

7.5 Reset types

The SSE-710 subsystem has multiple types of resets.

7.5.1 External Power-on-Reset (EPoR)

External Power on Reset is triggered by the **PORESETn** input and resets all logic within the SoC including the Reset Controller, with exception of the logic on the **nTRST** input.

Example 7-1: Reset sequence for External Power-on-Resets

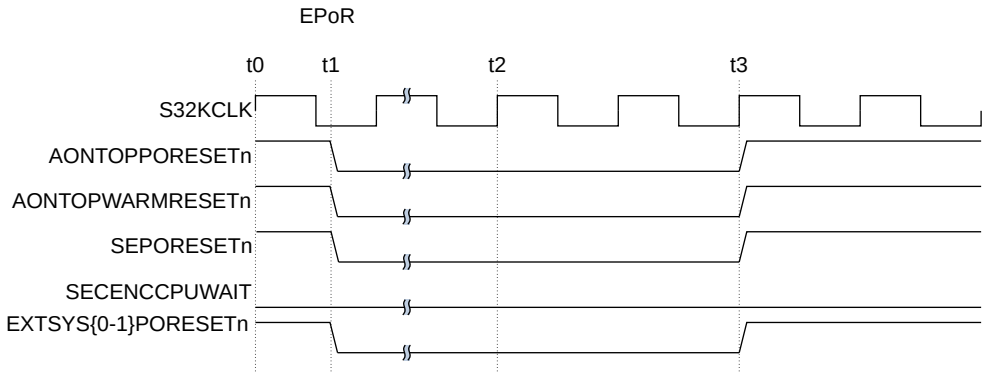

Note

The diagram only shows the reset outputs of the Reset Controller.

1. At t1 the **PORESETn** pin is asserted. This causes a reset of the Reset Controller and all logic in the SoC.
2. Between t1 and t2 an undefined number of S32KCLK cycles can pass before **PORESETn** is deasserted.
3. At t3 the Reset Controller deasserts the **AONTOPPORESETn**, **AONTOPWARMRESETn**, **SEPORESETn** and **EXTSYS{0-1}PORESETn** outputs.

7.5.2 Internal Power-on-Reset (IPoR)

This reset is similar to an External Power on Reset, except that the Reset Controller is not reset. All other logic is reset, except the logic on the **nTRST** reset.

An IPoR is triggered, using one of the following reset requests:

- SoC Watchdog
- Secure Enclave Watchdog
- Secure Enclave Crypto Accelerator error reset request
- **DP CDBGRSTREQ**
- SoC reset request

Figure 7-3: Example of a reset sequence for Internal Power-on-Resets

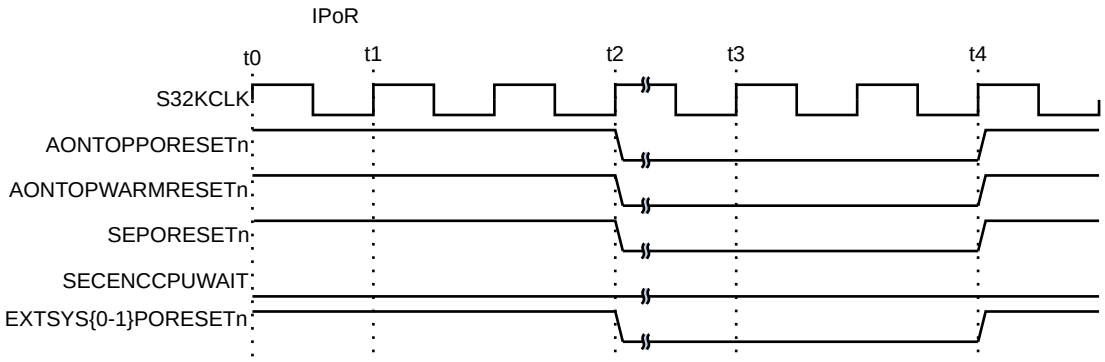

Note

The diagram only shows the reset outputs of the Reset Controller.

1. Between t0 and t1 the request to perform an IPoR is asserted to the Reset Controller.
2. The Reset Controller is required to synchronizer the request to the **S32KCLK** before performing the reset. This occurs between t1 and t2.

   Note

   In this example, there are no other reset requests at this point.
3. At t2 the Reset Controller asserts the **AONTOPPORESETn**, **AONTOPWARMRESETn**, **SEPORESETn** and **EXTSYS{0-1}PORESETn**.
4. Between t2 and t3, an undefined number of S32KCLK cycles can pass before the reset request is deasserted. For all reset request other than the **DP CDBGRSTREQ**, the reset request is deasserted because the source is reset as part of this sequence.
5. Between t3 and t4, the Reset Controller is required to synchronizer the request deassertion to the S32KCLK before deasserting the reset outputs.
6. At t4 the Reset Controller deasserts the **AONTOPPORESETn**, **AONTOPWARMRESET**, **SEPORSETn**, and **EXTSYS{0-1}PORESETn** output.

7.5.3 Debug Reset (DBGRST)

A debug reset resets all logic in the SoC except for the logic on the **nTRST** and **AONTOPORESETn**.

This reset can be triggered by the following requests:

- **nSRST**

- DP ROM CSYSRSTREQ
- Secure Enclave software reset request

There are two types of debug reset, in both cases when the request is received the Reset Controller applies the reset and then automatically releases the logic from reset:

- If the request is from the **nSRST** or DP ROM CSYSRSTREQ, the request also prevents the Secure Enclave Cortex®-M0+ from executing instructions while the request is asserted.
- If the request is from a Secure Enclave software reset request, the Secure Enclave Cortex®-M0+ is not prevented from executing instructions.

Note

This type of reset allows a debugger to remain connected to parts of the debug logic which reset using the **AONTOPPORESETn**.

Example 7-3: Reset sequence for debug reset request from nSRST or DP ROM CSYSRSTREQ,

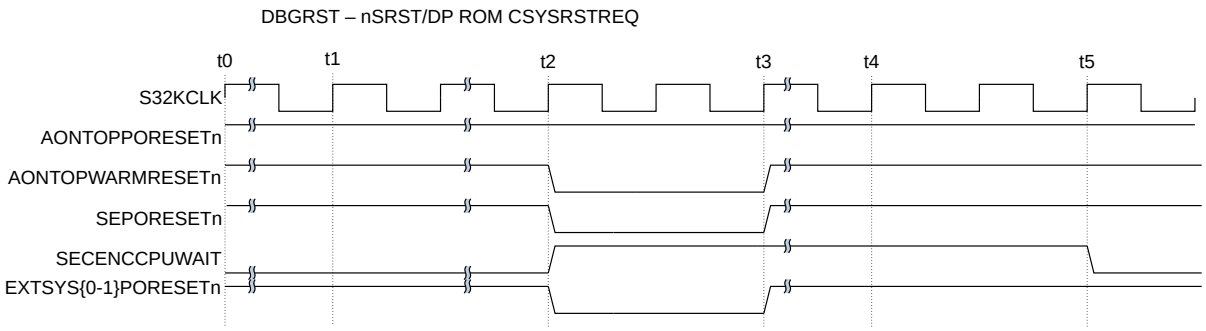

Note

The diagram only shows the reset outputs of the Reset Controller.

1. Between t0 and t1, the request to perform a DBGRST is asserted to the Reset Controller.
2. Between t1 and t2, the Reset Controller synchronizes the request to the S32KCLK before performing the reset. Any actions to handle quiescent handshakes also occur after the request is synchronized, but before t2.

   Note

   In this example, there are no other reset requests at this point and all quiescent handshakes have been accepted.

3. At t2 the Reset Controller asserts **AONTOPWARMRESETn**, **SEPORESETn**, and **EXTSYS{0-1}PORESETn**. The Reset Controller also asserts the **SECENCCPUWAIT** signal.
4. At t3 the Reset Controller deasserts the **AONTOPWARMRESETn**, **SEPORESETn** and **EXTSYS{0-1}PORESETn**. The **SECENCCPUWAIT** signal remains asserted preventing the Secure Enclave Cortex®-M0+ from executing instructions.
5. Between t3 and t4 an undefined number of S32KCLK cycles can occur until the request to perform a DBGRST is de-asserted. The request must be deasserted by the debug agent by either deasserting **nSRST** input or setting the DP ROM CSYSRSTREQ to `0b0`.
6. Between t4 and t5, the Reset Controller is required to synchronizer the request deassertion to the S32KCLK before deasserting the reset outputs.
7. At t5, the Reset Controller deasserts the **SECENCCPUWAIT** and the Secure Enclave Cortex®-M0+ starts to execute instructions.

Example 7-4: Reset sequence for debug reset caused by Secure Enclave software reset

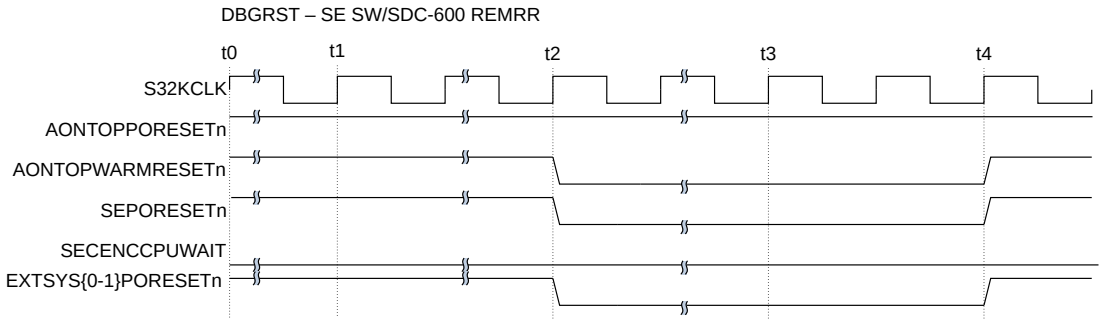

Note

The diagram only shows the reset outputs of the Reset Controller. It does not show any other outputs of the Reset Controller

1. Between t0 and t1 the request to perform a DBGRST is asserted to the Reset Controller.
2. Between t1 and t2, any actions to handle quiescent handshakes also occur after the request is synchronized but before t2.

Note

In this example, there are no other reset requests at this point and all quiescent handshakes have been accepted.

3. At t2 the Reset Controller asserts **AONTOPWARMRESETn**, **SEPORESETn**, and **EXTSYS{0-1}PORESETn**. In this case the **SECENCCPUWAIT** is not asserted.
4. Between t2 and t3, an undefined number of S32KCLK cycles can occur until the request to perform a DBGRST is deasserted. In both cases this occurs as part of the reset sequence.

5. Between t3 and t4, the Reset Controller is required to synchronizer the request deassertion to the S32KCLK before deasserting the reset outputs.
6. At t4, the Reset Controller deasserts the **AONTOPWARMRESETn**, **SEPORESETn**, and **EXTSYS{0-1}PORESETn**.

7.5.4 Host system reset

A Host System reset resets all the logic on the **AONTOPWARMRESETn** and **EXTSYS{0-1}PORESETn**.

A Host System reset is triggered via the Secure Enclave setting the HOST_SYS_RST_CTRL.RST_REQ to `0b1` in the Secure Enclave Base System Control registers.

Example 7-5: Reset sequence for Host system reset

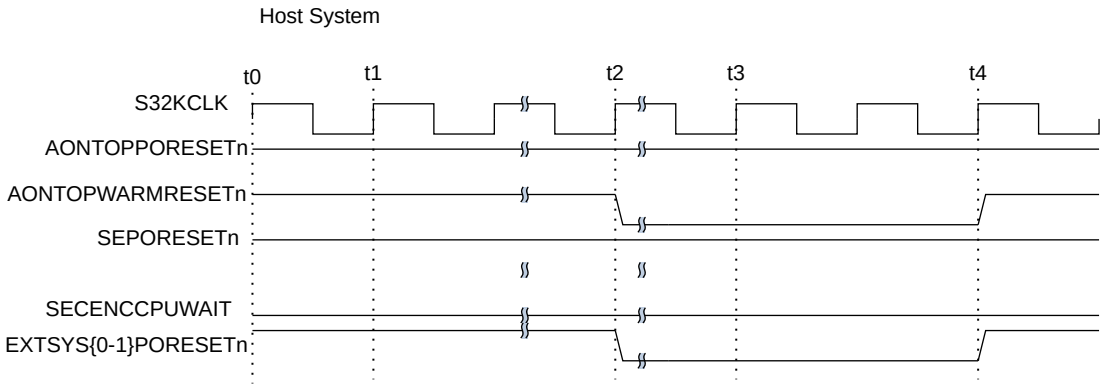

Note

The diagram only shows the reset outputs of the Reset Controller.

1. Between t0 and t1, a request to perform a Host System reset is asserted to the Reset Controller.
2. Between t1 and t2, the Reset Controller is required to synchronizer the request to the S32KCLK before performing the reset. Any actions to handle quiescent handshakes also occur after the request is synchronizer but before t2.

Note

In this example, there are no other reset requests at this point and all quiescent handshakes have been accepted.

3. At t2 the Reset Controller asserts **AONTOPWARMRESETn** and **EXTSYS{0-1}PORESETn**.

Note

This only occurs if the Reset Controller was able to quiescent interfaces between Host or External Systems and either the Secure Enclave or **AONTOPPORESETn** reset domain.

4. Between t2 and t3, an undefined number of S32KCLK cycles can occur before the reset request is deasserted.
5. Between t3 and t4, the Reset Controller is required to synchronizer the request deassertion to the S32KCLK before deasserting the reset outputs.
6. At t4 the Reset Controller deasserts **AONTOPWARMRESETn** and **EXTSYS{0-1}PORESETn**.

7.5.5 External system reset

An External System reset only resets logic on the **EXTSYS{0-1}PORESETn**. An External System reset is caused by software setting the EXT_SYS{0-1}_RST_CTRL.RST_REQ field to `0b1`.

Note

The logic within the External System can be reset from an internal request, for example an External System watchdog reset request, but the interfaces with the logic outside the External System must be quiescent before performing the reset.

Example 7-6: Reset requence for External system reset

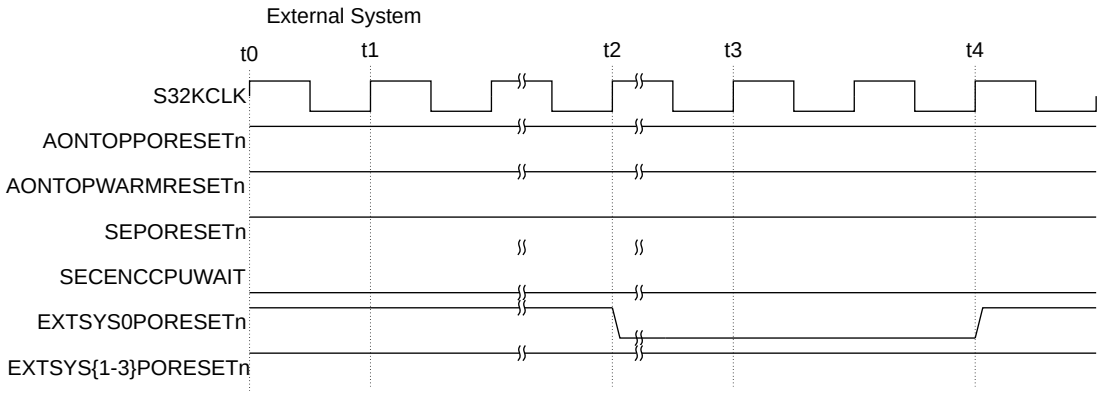

Note

The diagram only shows the reset outputs of the Reset Controller. It does not show any other outputs of the Reset Controller.

1. Between t0 and t1, a request to perform an External System reset is asserted to the Reset Controller.


Note

In this example, only External System 0 is being reset. However, the same sequence applies for any External System.

2. The Reset Controller is required to synchronize the request to the S32KCLK before performing the reset. This occurs between t1 and t2. Any actions to handle quiescent handshakes, also occur after the request is synchronized but before t2.


Note

In this example, there are no other reset requests at this point and all quiescent handshakes have been accepted.

3. At t2 the Reset Controller asserts **EXTSYS0PORESETn** and leaves **EXTSYS{1-3}PORESETn** deasserted..


Note

This only occurs if the Reset Controller was able to quiescent interfaces of the External System.

4. Between t2 and t3, an undefined number of S32KCLK cycles can occur before the reset request is deasserted.
5. Between t3 and t4, the Reset Controller is required to synchronizer the request deassertion to the S32KCLK before de-asserting the reset outputs.
6. At t4, the Reset Controller deasserts **EXTSYS0PORESETn**.

7.5.6 Reset controller behavior

If multiple changes in the reset requests occur at the same time the request with the highest priority.

The priority is defined as follows (from highest to lowest):

1. **PORESETn** – Highest priority
2. Secure Enclave Crypto Accelerator Error reset request
3. Secure Enclave Watchdog
4. **DP CDBGRSTREQ**
5. SoC reset request, from the Secure Enclave Base System Control register
6. Secure Enclave software reset request
7. **nSRST**
8. **DP ROM CSYSRSTREQ**

9. Host System reset request
10. External System {0-3} reset request – Lowest priority

Arm® strongly recommends:

- That the debugging agent does not request a reset using **DP ROM CSYSRSTREQ** at the same time.
- When the Secure Enclave is unable to handle the refreshing the watchdogs it disables the watchdogs. For example, when debugging the Secure Enclave firmware.

The Reset Controller has number of Q-Channels to manage the quiescent of interfaces between logic which is being reset and logic which is not being reset. The Q-Channels are used to control domain bridges or access control gates on these interfaces. The handshake on the Q-Channel is performed before starting the reset sequence and only if all Q-Channel requests are accepted.

After the Reset Controller starts one of the sequences, any removal of the reset request is ignored. If the Q-Channel request is denied, depending on the reset request being attempted by the Reset Controller the behavior is as follow:

- **PORESETn** is always guaranteed to be applied as this is the Power on Reset for the entire SoC including the Reset Controller.
- For SoC Watchdog reset request, SoC reset request, Secure Enclave Watchdog reset request, and **DP CDBGRSTREQ** the reset is always guaranteed to be applied. The reset is guaranteed because all logic is reset within the SoC, except for the Reset Controller. The only exception is if **PORESETn** is asserted.
- For a Secure Enclave software reset request, **nSRST**, and **DP ROM CSYSRSTREQ** reset request, the behaviour is as follows:
  - If reset request has been removed and no other reset request is asserted, the Reset Controller returns to the idle state.
  - If reset request has been removed and another reset request is asserted, the Reset Controller handles the new reset request.
  - If reset request remains and is still the highest priority reset request asserted, the Reset Controller repeats the sequence.
  - If reset request remain and is no longer the highest priority reset request asserted, the Reset Controller handles the new reset request as it is higher priority.
- For the Host and External System reset requests the Reset Controller repeats the request a number of times.


Note

The Reset Controller always responds to a higher priority reset request being asserted.

If the reset request is removed the behavior is the same as the behavior for the Secure Enclave software reset request, **nSRST**, and **DP ROM CSYSRSTREQ** reset requests.

7.6 Power-on reset

When the **PORESETn** signal is asserted, SSE-710 is in a *Power on Reset* (PoR) state.

During a PoR state, the **REFCLKQACTIVE** is asserted. To exit the PoR state, the following conditions must be true before the **PORESETn** signal is deasserted:

- **S32KCLK** must be running.
- **SECENCREFCLK** must be running.
- VSYS voltage supply must be available and stable.
- If the VCLUS voltage supply is implemented, then it must either be:
  - Available and stable.
  - Requestable without the use of software running on any processor in the SoC when a debugger connects to the SoC through the SWJ interface. The method by which this is achieved is **IMPLEMENTATION DEFINED**.
- If VEXTSYS{0-1} voltage supplies are implemented, then they must either be:
  - Available and stable
  - Requestable without the use of software running on any processor in the SoC when a debugger connects to the SoC through the SWJ interface. The method by which this is achieved is **IMPLEMENTATION DEFINED**.

When exiting PoR state, SSE-710 automatically performs the following sequence:

- Transition to BSYS.SLEEP1 on deasserted of **PORESETn**. The SECENCTOP power domain automatically enters the ON power mode and the Secure Enclave Cortex®-M0+starts to execute instructions.
- When the REFCLK Q-Channel enters the Q_RUN state, SSE-710 transitions to BSYS.SLEEP0.

SSE-710 enters the BSYS.RUN power state only when a request for SYSTOP to exit the OFF power mode is detected.

For details of REFCLK Q-channel, see 4.7 Clock Control interfaces on page 70.

8. Debug

This chapter describes the SSE-710 subsystem debug infrastructure.

8.1 Debug overview

SSE-710 provides a debug infrastructure for the Secure Enclave and Host System and provides interfaces for the debug infrastructure of the External Systems. It is compliant with the *Arm® Debug Interface Architecture Specification ADIv6.0*.

SSE-710 debug infrastructure provides the following features:

- Self-hosted debug of Host or External Systems
- External debug of Secure Enclave, Host, or External System by an off-chip debugger
- External debug of Secure Enclave, Host, or External System from one of the other systems in the SSE-710
- Full cross trigger support
- Trace support
- Granular power control
- Debug through power down, for the Host and External Systems
- Debug from reset, for all systems
- Fine grain control of debug privileges provided to the debug agent
- Support for single or multi-system debug
- Certificate injection using SDC-600
- Support for Serial Wire Debug
- Support for JTAG debug

In this section, the following terms are used:

Target system

The system which is being debugged.

The target system can be an individual system or multiple systems within the SSE-710.

Debug agent

The entity performing the debug on the target system.

This can be one of:

- Software running on the target system, for example, software running on a Host CPU core debugging itself or another Host CPU core in the Host System.

- Software running on another system within SSE-710. For example, an application on the Host CPU debugging an External System. This also includes where the software is running a USB or WiFi stack to perform debug over functional IO.
- Off-chip debugger, for example, a JTAG or Serial Wire debugger (SWJ).

Self-host debug agent

The debug agent is software running on a processor which is being debugged.

For example, this is where a Host CPU core uses the System registers to access its own debug logic.

External debug agent

The debug agent is one of the following:

- Software running on another processor to the one being debugged. This includes when the processors are operating in a multiple processor configuration, all running the same operating system.
- Software running on another system.
- Off-chip debugger.

The debug infrastructure in SSE-710 is constructed from multiple separate blocks:

- External Debug Bus
- Host System Debug
- Secure Enclave Debug
- SoC Debug
- AXI AP Debug
- External System Debug

Some of the blocks provide debug features to a single system, other blocks provide features to all systems within SSE-710.

8.2 Authentication

SSE-710 lets you control debug authentication and authorization.

SSE-710 uses *Debug Authentication Zones* (DAZ) and *Debug Authorization Access Control Gate* (DAACG) to:

- Control the types of debug:
  - Non-secure/secure
  - Invasive/non-invasive
- Control which debug agents can debug a specific target system.:
  - Self-hosted

- External: It can control which debug agents can be the external debuggers for a target system

8.2.1 Debug Authentication Zone (DAZ)

DAZ controls which debug features are enabled for the debug logic within the zone.

There are two types of DAZ:

- Controls the level of debug within a system, for example HOSTAUTH
- Controls of external debug access to the system, for example HOSTEXTAUTH

A DAZ is controlled by a unique set of the Arm® CoreSight™ authentication signals: **DBGEN**, **NIDEN**, **SPIDEN**, and **SPNIDEN**. The values of the authentication signals are driven by the SCB of the Secure Enclave. Not all DAZ support all the authentication signals.

Table 8-1: DAZ Signals in the SSE-710

DAZ name	Components	Extendable	DBGEN	NIDEN	SPIDEN	SPNIDEN
DPAUTH	DP ROM EXTDBG ROM	No	Y	Y	Y	Y
COMAUTH The **DBGEN** signal is referred to as the PEN and NIDEN is referred to as the RRDIS. These signals then drive the **CFG_PEN** and **CFG_RRDIS** inputs of EXT APBCOM.	SDC-600	No	Y	Y	N	N
HOSTAXIAUTH	Host AXI AP Host AXIAP ROM	No	Y	Y	Y	Y
HOSTEXTAUTH There is no difference between invasive and non-invasive debug and therefore: • **DBGEN** and **NIDEN** and referred to as NS and set to the same value. • **SPIDEN** and **SPNIDEN** and referred to as S and set to the same value. For example, the HOSTEXTAUTH DAZ has two signals: NS and S. The NS signal sets the value of **DBGEN** and **NIDEN** of the Host APB AP. The S signal sets the value of **SPIDEN** and **SPNIDEN** of the Host APB AP.	Host APB AP	No	Y	Y	Y	Y

DAZ name	Components	Extendable	DBGEN	NIDEN	SPIDEN	SPNIDEN
HOSTAUTH	Host ROM Host ETR Host STM Host CTI Host Replicator Host CPU Cluster The Host CPU Cluster includes a number of debug components. For more information, see the *Arm® Cortex®-A32 Processor Technical Reference Manual*, *Arm® Cortex®-A35 Processor Technical Reference Manual*, or *Arm® Cortex®-A53 MPCore Processor Technical Reference Manual*. Host Expansion Host CATU	Yes	Y	Y	Y	Y
TPIUAUTH	SoC TPIU SoC TPIU Funnel SoC TPIU Replicator SoC ETR SoC CATU SoC CTI	No	Y	Y	Y	Y
COUNTERAUTH	Counter CTI	No	Y	Y	Y	Y
SECENCAUTH Secure Enclave has only a single security world. Any debug components which support both Secure and Non Secure debug have the same debug privileges enabled in either security world. This is achieved by driving the **DBGEN** and **SPIDEN** signals and the **NIDEN** and **SPNIDEN** signals to the same value.	Secure Enclave AHB-AP The ap_en and ap_secure_en are both driven by the logical OR of the **DBGEN** and **NIDEN** from the SECENCAUTH DAZ.	No	Y	Y	N	N
	Secure Enclave CTI					
	Secure Enclave Cortex®-M0+ The Secure Enclave Cortex®-M0+ includes a number of debug components. For more information, see *Cortex®-M0+ Technical Reference Manual*.					
	Secure Enclave CS ROM					

8.2.2 Debug Authorization Access Control Gate (DAACG)

The *Debug Authorization Access Control Gate* (DAACG) is a type of *Access Control Gate* (ACG) which grants access from master(s) to downstream components.

ACG status is controlled by the SCB of the Secure Enclave. The ACG is either in:

Open state

DBGEN is HIGH. Accesses are allowed through the ACG.

Closed state

DBGEN is LOW. Accesses are terminated by the DAACG with an error.

SSE-710 includes an instance of the DAACG for each master interface to the External Debug Bus. DAACGs enable the software to configure which external debug agents are able to access the components on the External Debug Bus.

Arm® strongly recommends that when **DBGEN** is changed, no transactions are outstanding on the slave interface. It is **UNPREDICTABLE** whether the transaction sees the new or old value of **DBGEN**.

You can enable debug from more than one external debug agent. However, you cannot grant different external debug agents different debug privileges to the same target system or to different target systems. Arm® strongly recommends that only a single external debug agent is enabled at any time.

You can have one target system performing self-hosted debug while another system is being debugged by an external debug agent. This also includes where the debug agent is debugging another CPU within the same system. For example, one Host CPU core debugging another Host CPU core.

8.3 Debug blocks

The following sections describe the blocks that form the debug infrastructure of the SSE-710.

8.3.1 External Debug Bus

SSE-710 provides an External Debug bus.

The following figure shows the External Debug bus, which provides access to any debug agent that is performing external debug on a target system.

Figure 8-1: External Debug Bus

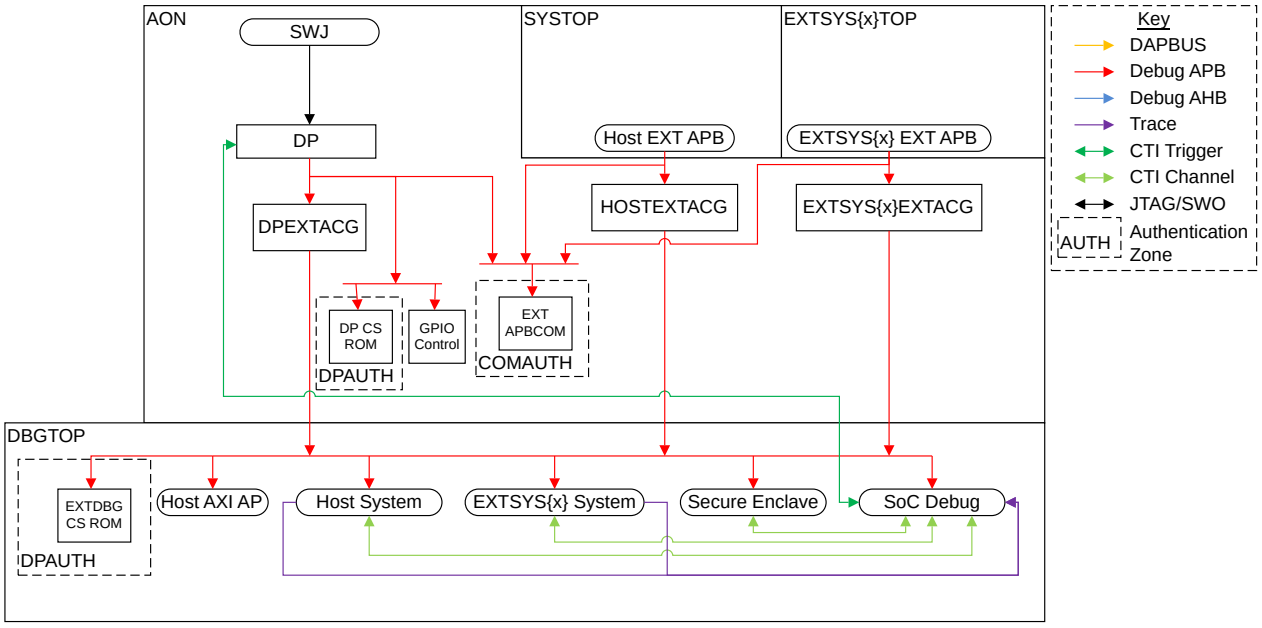

Note

This figure shows a single EXTSYS{x} EXT APB interface and EXTSYS{x} system. There is one for each External System in SSE-710.

The External Debug bus contains the following:

Debug Port (DP)

Provides access to an off-chip JTAG or a Serial Wire Debug Agent

Host External APB (Host EXT APB)

Host External APB (Host EXT APB): Provides access for the Host System to be an external debug agent to any system in the SSE-710.

External System {0-1}

External APB (EXTSYS{x} EXT APB): Provides access for the External System {0-1} to be an external debug agent to any system in the SSE-710.

EXT APBCOM

Provides the ability to insert a certificate to enable debug access for an external debug agent

Note

When the debug agent is software running on a system within the SoC, Arm® strongly recommends an **IMPLEMENTATION DEFINED** method is used to request access to the External Debug Bus, instead of using the EXT APBCOM. For more information on the EXT APBCOM, see 11.7 CoreSight SDC-600 on page 187.

Debug Port ROM (DP ROM)

ROM table for DP, pointing to EXT APBCOM, EXTDBG ROM and GPIO Control.

External Debug ROM (EXTDBG ROM)

ROM table for all components on the External Debug Bus, except for EXT APBCOM and GPIO Control.

GPIO Control (GPIO Control)

ROM table for all components on the External Debug Bus, except for EXT APBCOM and GPIO Control.

GPIO Control (GPIO Control)

Debug Component to be able to control general purpose inputs and outputs. In SSE-710 the GPIO Control component supports a single GPO and no GPI. The GPO is used to drive the CALC interface, along with the SOCLCC interface. For more information see 8.10 GPIO control on page 146.

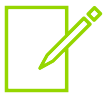

The GPIO Control block has its **DBGEN** input tied HIGH in SSE-710.

Debug APB ports

Debug APB ports to:

- Host System Debug (Host System)
- Host System (Host AXI AP)
- External System Debug (EXTSYS{x} System)
- Secure Enclave Debug (Secure Enclave)
- SoC Debug logic (SoC Debug)

DAZ

Two DAZ:

- DPAUTH: Controls the level of debug functionality of the DP and EXTDBG ROM.
- COMAUTH: Controls the level of functionality provided by the EXT APBCOM. **DBGEN** drives the **CFG_PEN** input and **NIDEN** drives the **CFG_RRDIS** input, through an inverter.

DAAC

The following DAACGs:

- DPEXTACG: Prevents access to all debug logic other than the DP ROM, EXT APBCOM and GPIO Control from the DP.
- HOSTEXTACG: Prevents access to all debug logic other than EXT APBCOM, from the Host System using the Host EXT APB interface.
- EXTSYS{x}EXTACG: Prevents access to all debug logic other than EXT APBCOM, from the External System using the EXTSYS{x}EXTAPB interface.

12.1.2 External Debug Bus memory map on page 193 shows the memory map for the External Debug Bus.

8.3.2 Host AXI AP debug

The AXI AP and the associated AXI ROM table provide direct access into the Host System memory map, as the following figure shows.

Figure 8-2: AXI AP

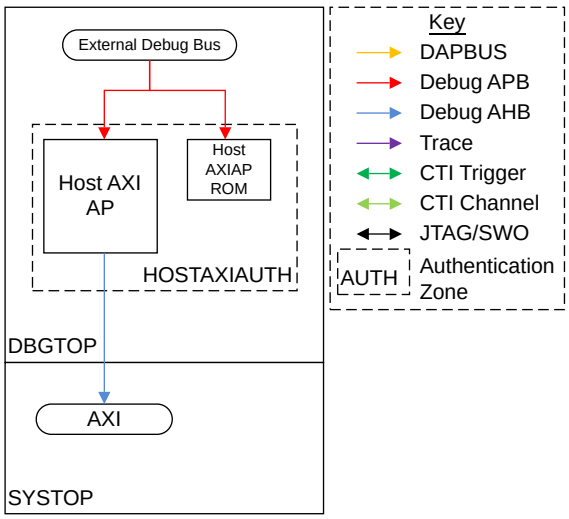

For information on the Host System memory map, see 12.1.1 Host System memory map on page 188. Both the AXI AP and ROM are part of the HOSTAXIAUTH authentication zone. The AXI AP ROM provides power control for the power domains of the SoC which can be accessed through the Host AXI AP. For more information, see8.7 ROM tables on page 142.

8.3.3 Secure Enclave debug

The Secure Enclave provides access to the Secure Enclave debug logic through an AHB AP.

The Secure Enclave debug logic provides the following:

- Secure Enclave AHB AP
- Secure Enclave CS ROM
- Secure Enclave CTI
- Secure Enclave Channel Gate
- Secure Enclave Cortex®-M0+ Debug

Figure 8-3: Secure Enclave Debug Access

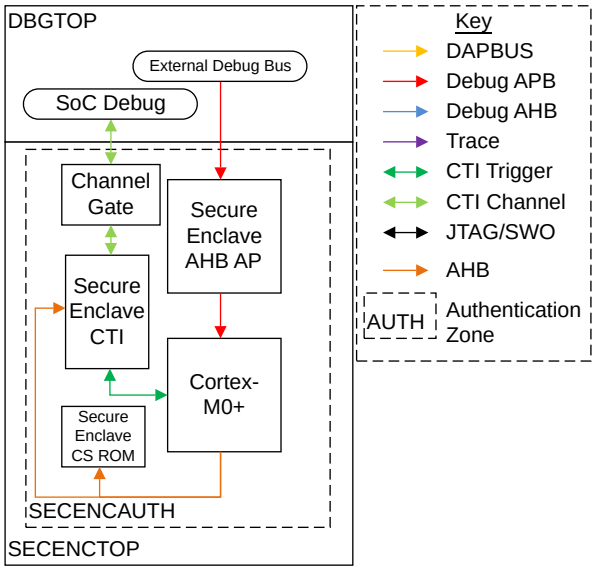

The Secure Enclave AHB-AP, and all debug logic within the Secure Enclave, is part of the SECENCAUTH authentication zone.

8.3.4 Host System debug

This section describes the Host system debug.

The following figure shows the Host System debug block.

Figure 8-4: Host System debug block

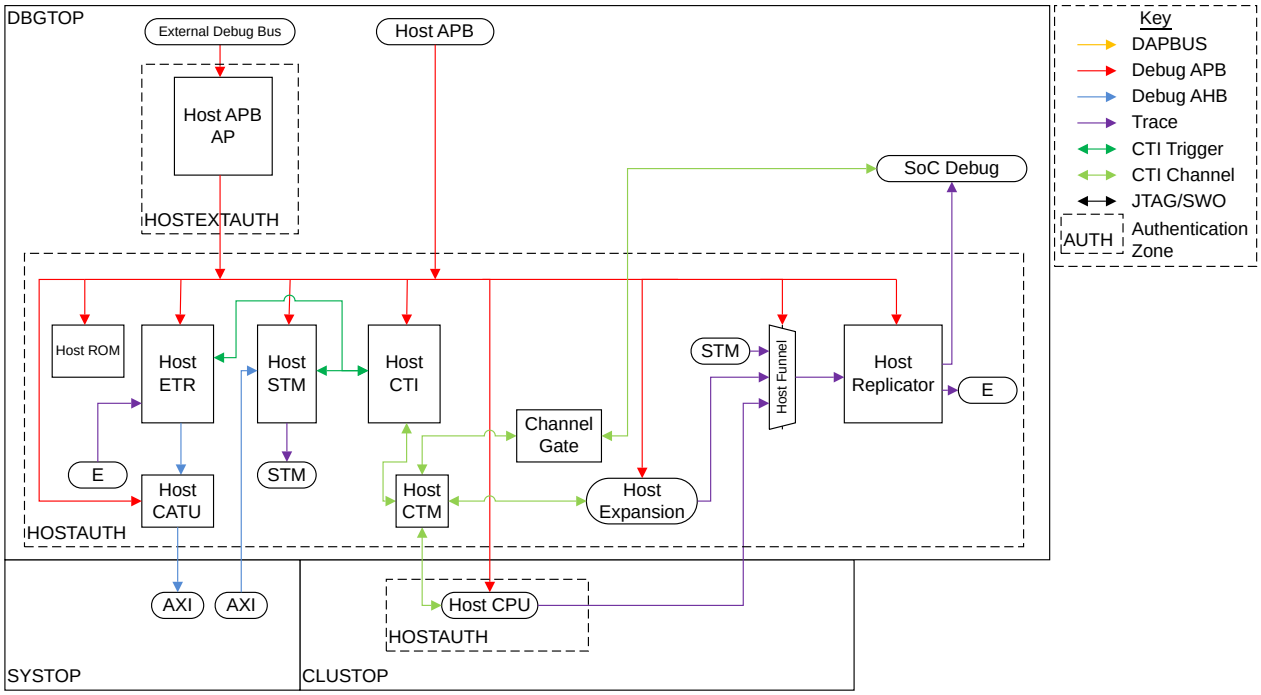

The Host System debug block provides the following features:

- Host APB-AP: Provides access for an external debug agent
- Host ROM: CoreSight™ ROM table with GPR functionality
- Host ETR: Ability to send only Host System trace to memory
- Host CATU: Ability to perform address translation for the Host ETR AXI transactions
- Host STM: Ability to perform instrumented software trace
- Host CTI: Trigger capabilities for Host ETR and STM
- Host CTM and Channel gate
- Host Funnel: Combines trace from Host STM, HOSTCPUTRACE, and HOSTTRACEEXP interfaces
- Host Replicator: Replicates combined trace from Host Funnel, to SoC Debug block or Host ETR
- Host APB interface: Allows access from the Host System memory map
- Host CPU APB interface (HOSTCPUDBG): Access to external debug registers of the Host CPU
- Host CPU Trace interface (HOSTCPUTRACE): Trace interface for Host CPU
- Host CTI Channel interface (HOSTCPUCTICHIN/HOSTCPUCTICHOUT): Connection with CTIs within the Host CPU.
- Two DAZ:
  - HOSTAUTH: Controls the level of debug allowed in the Host System, for either self-host or external debug agents. This DAZ is exposed on the HOSTDBGAUTH interface and is connected to all debug logic in the Host System.

- HOSTEXTAUTH: Controls the level of debug access by an external debug agent. Used only for the Host APB-AP.

The Host CPU debug logic is not part of SSE-710. SSE-710 provides an interface to integrate the debug logic, provided by the Host CPU cluster.SSE-710 requires that the ETM is enabled in each Host CPU core.

The Host System debug provides expansion interfaces to allow for expansions and to integrate Host CPU. The expansion interfaces are:

- Host CPU expansion (Host CPU in the preceeding figure):
  - APB, with 16MB of allocated memory space (HOSTCPUDBG)
  - CTI channel interface (HOSTCPUCTICHIN/HOSTCPUCTICHOUT)
  - ATB interface (HOSTCPUTRACE). If there is more than a single CPU Core in the Host CPU cluster, an ATB funnel is added.
- Host System expansion (Host Expansion in the preceeding figure):
  - APB, with 16MB of allocated memory space (HOSTDBGAPBEXP)
  - CTI Channel interface (HOSTCTICHINEXP/HOSTCTICHOUTEXP)
  - ATB interface (HOSTDBGTRACEEXP)

12.1.1.1.1 Host System Debug on page 189 describes the memory map of the Host debug block.

8.3.5 External System {0-1} debug

The following figure shows the External System Debug, implemented by the External System harness.

Figure 8-5: External System debug

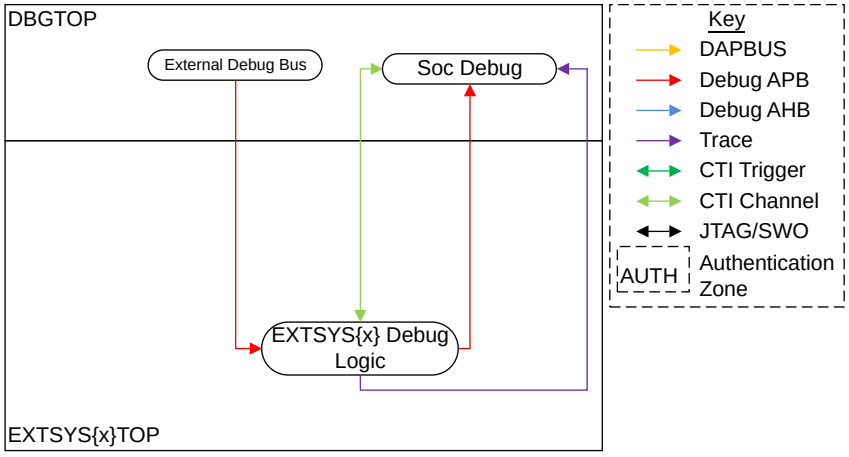

The External System debug block is part of the External System harness. The External System debug block includes the following:

- External System Debug APB (EXTSYS{0-1}DBG) interface: Provides the ability for the External System to integrate its own debug logic.
- External System External Debug Access (EXTSYS{0-1}EXTDBG): Provides the ability for the External System {0-1} to be an external debug agent for another system in the SoC.
- CTI Channel interface (EXTSYS{0-1}CTICHIN/EXTSYS{0-1}CTICHOUT): Expansion of the SSE-710 CTI network into the External System {0-1}.
- Trace Interface: Expansion of the SSE-710 trace network into the External System {0-1}.

The debug logic that is provided in the External System {0-1} (EXTSYS{x} Debug Logic) is **IMPLEMENTATION DEFINED**. No debug logic is also a valid configuration.

8.3.6 SoC debug

SSE-710 has some debug logic in the SoC Debug block, which is shared among all systems.

The following figure shows the SoC debug block:

Figure 8-6: SoC debug logic

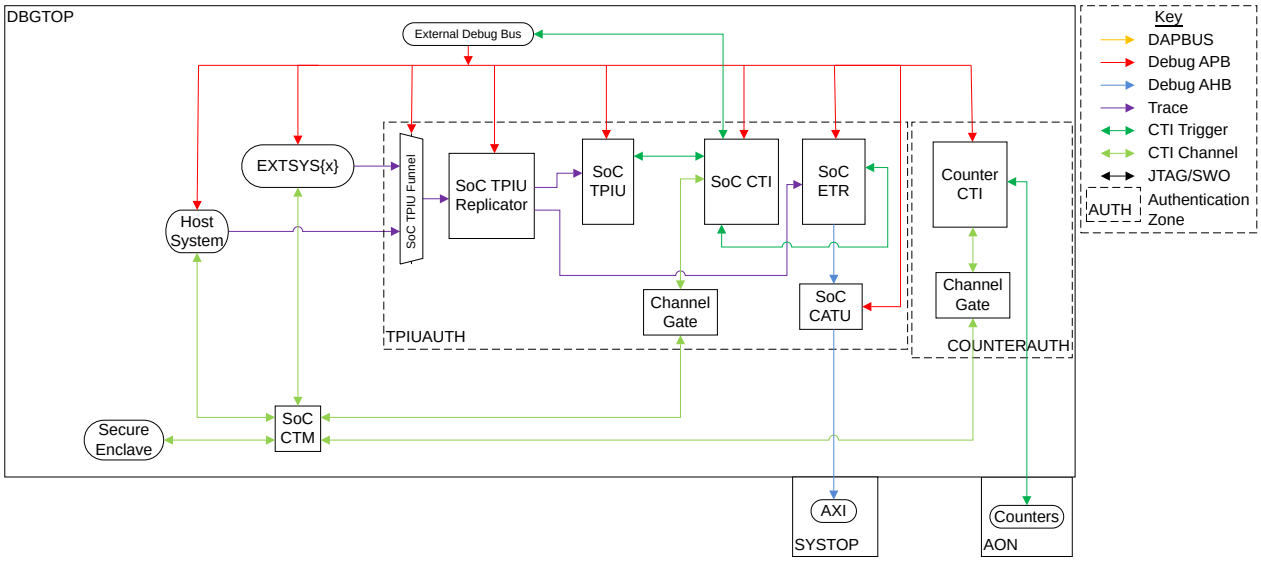

The SoC debug block contains the following:

- SoC TPIU: Provides trace off-chip
- SoC TPIU Funnel: Combines trace from all External Systems and the Host System
- SoC TPIU Replicator: Replicates the trace from the SoC TPIU Funnel to be sent
- SoC CTI: Provides triggers to and from the SoC TPIU, ETR, and DP

- SoC ETR: Allows trace to be routed to a memory location in the Host System
- SoC CATU: Allows for address translation of the SoC ETR AXI transactions
- SoC CTM: Combines CTI Channel interfaces from all External Systems and Host Systems
- Counter and TPIU Channel Gates
- Counter CTI: Provides triggers to and from the REFCLK and S32K counters.
- Two DAZ:
  - TPIUAUTH: Controls the level of debug that is provided by all logic in the SoC Debug, except for the Counter CTI
  - COUNTERAUTH: Controls the level of debug of the Counter CTI

12.1.2 External Debug Bus memory map on page 193 describes the memory map of the SoC Debug block.

8.4 Cross Trigger Infrastructure

The Cross Trigger Interfaces (CTIs) and Cross Trigger Interfaces (CTMs) are interconnected.

The following figure shows the interconnect between CTIs and CTMs.

Figure 8-7: SSE-710 Cross Trigger Infrastructure

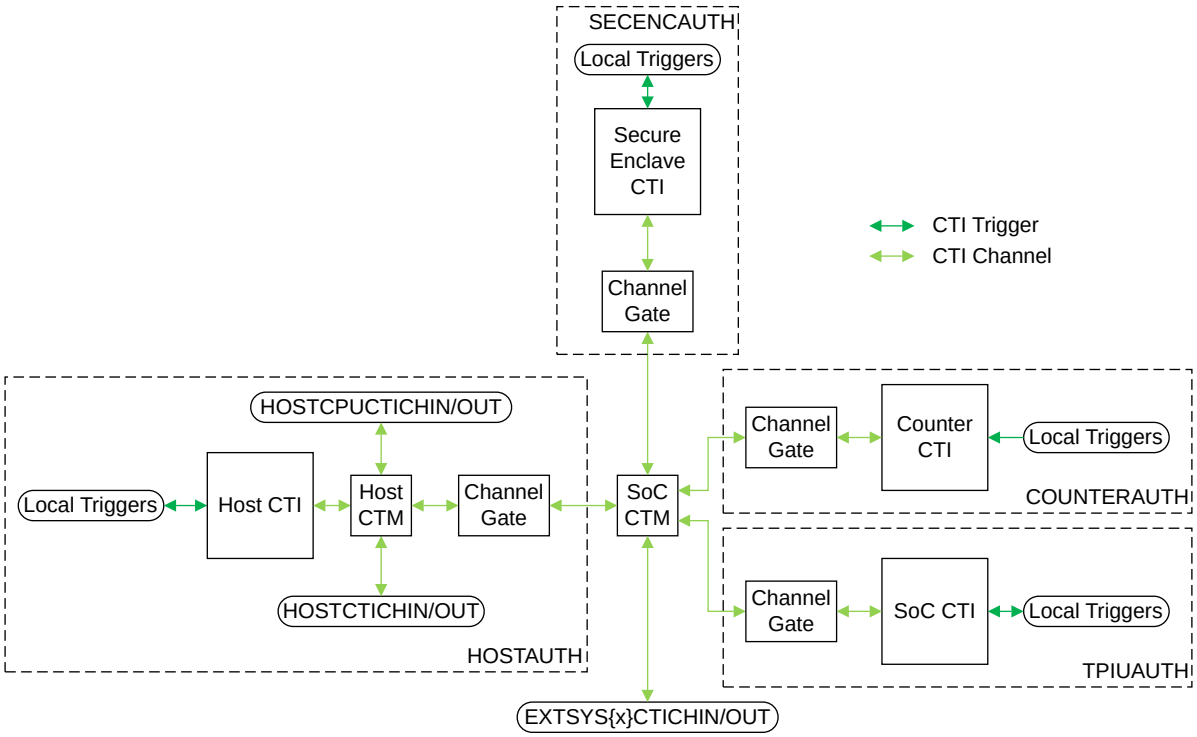

In SSE-710 there are the following CTIs and CTMs:

- Host CTI: Connected to Host ETR and Host STM
- Host CTM: Connects the Host CTI and SoC CTI. Also provides the following interfaces:
  - HOSTCPUCTICHIN and HOSTCPUCTICHOUT interfaces for the Host CPU
  - HOSTCTICHINEXP and HOSTCTICHOUTEXP interfaces for expansion of the Host debug logic
- SoC CTI: Connected to SoC TPIU, SoC ETR, and DP
- Counter CTI: Connected to the REFCLK and S32K Counter halt interfaces
- SoC CTM: Connects the SoC, Counter, Host, Secure Enclave CTIs and also provides the EXTSYS{0-1}CTICHIN and EXTSYS{0-1}CTICHOUT interfaces


Note

When a system exposes an event, it is globally broadcast to all systems. The debug agent is responsible for only exposing events which are safe to be exposed with all other systems.

SSE-710 also provides Channel gates between the SoC CTM and the following components:

- Host CTM
- Secure Enclave CTI
- Counter CTI
- SoC CTI

The Channel Gate is placed on the Channel interfaces between CTI components to control whether events can be sent between CTI components. Each Channel Gates is controlled by a Channel Enable (CHEN) signal:

- HIGH: the Channel Gate is enabled and the events can be forwarded.
- LOW: the Channel Gate is disabled and the events can not be forwarded.

The CHEN is controlled by SCB from a Crypto Accelerator. For more information on SCB control, see 9.1.3 Security Control Bits (SCB) on page 150.

8.4.1 Cross Trigger Interface (CTI)

The following tables show the assignment of trigger inputs and outputs for each CTI.

Table 8-2: Host CTI Trigger In

CTI Trigger In offset	Source component	Event type	Notes
0	Host STM TRIGOUTSPTE	Pulse	For more information on STM triggers, see the *Arm® CoreSight™ STM-500 System Trace Macrocell Technical Reference Manual*.
1	Host STM TRIGOUTSW		

CTI Trigger In offset	Source component	Event type	Notes
2	Host STM TRIGOUTHETE		
3	Host STM ASYNCOUT		
4	Host ETR FULL	Level	For more information ETR triggers, see the *Arm® CoreSight™ System-on-Chip SoC-600 Technical Reference Manual*.
5	Host ETR ACQCOMP		
6	Host ETR FLUSHCOMP	Pulse	

Table 8-3: Host CTI Trigger Out

CTI Trigger Out offset	Sink component	Software handshake	Notes
0	Host STM HWEVENTS[0]	No	For more information on STM triggers, see the *Arm® CoreSight™ STM-500 System Trace Macrocell Technical Reference Manual*.
1	Host STM HWEVENTS[2]		
2	Host ETR TRIGIN		For more information ETR triggers, see the *Arm® CoreSight™ System-on-Chip SoC-600 Technical Reference Manual*.
3	Host ETR FLUSHIN		
4	Host GIC Interrupt 72		For more information on the connection to the Host GIC, see 12.2.1 Host CPU interrupt map on page 199.
5	Host GIC Interrupt 73		

Table 8-4: SoC CTI Trigger In

CTI Trigger In offset	Source component	Event type	Notes
0	SoC TPIU FLUSHCOMP	Pulse	For more information on TPIU triggers, see the *Arm® CoreSight™ System-on-Chip SoC-600 Technical Reference Manual*.
1	SoC ETR FULL	Level	For more information ETR triggers, see the *Arm® CoreSight™ System-on-Chip SoC-600 Technical Reference Manual*.
2	SoC ETR ACQCOMP		
3	SoC ETR FLUSHCOMP	Pulse	

Table 8-5: SoC CTI Trigger Out

CTI Trigger Out offset	Sink component	Software handshake	Notes
0	TPIU TRIGIN	No	For more information on TPIU triggers, see the *Arm® CoreSight™ System-on-Chip SoC-600 Technical Reference Manual*.
1	TPIU FLUSHIN		
2	DP Event Status	Yes	For more information on DP triggers, see the *Arm® CoreSight™ System-on-Chip SoC-600 Technical Reference Manual*.

CTI Trigger Out offset	Sink component	Software handshake	Notes
3	SoC ETR TRIGIN	No	For more information ETR triggers, see the *Arm® CoreSight™ System-on-Chip SoC-600 Technical Reference Manual*.
4	SoC ETR FLUSHIN		

Table 8-6: Counter CTI Trigger Out

CTI Trigger Out offset	Source component	Software handshake	Notes
0	**REFCLK** Counter Halt	No	-
1	**REFCLK** Counter Restart		
2	32K Counter Halt		
3	32K Counter Restart		

Table 8-7: Secure Enclave CTI Trigger In

CTI Trigger In offset	Source component	Event type	Notes
0	Secure Enclave Cortex®-M0+ Halted	Pulse	-

Table 8-8: Secure Enclave CTI Trigger Out

CTI Trigger Out offset	Source component	Software handshake	Notes
0	Secure Enclave Cortex®-M0+ External Debug Request	Yes	-
1	Reserved	-	
2	Secure Enclave Cortex®-M0+ NVIC Interrupt 11	No	
3	Secure Enclave Cortex®-M0+ NVIC Interrupt 12		
4	Reserved	-	
5			
6			
7	Secure Enclave Cortex®-M0+ Debug Restart	No	

8.4.2 Cross Trigger Matrix (CTM)

SSE-710 provides an SoC and Host CTM, which connects all CTIs together.

The Host CTM connects the Host CTI to the SoC CTM and provides the following expansion interfaces:

- HOSTCTICHINEXP and HOSTCTICHOUTEXP: For expansion of the Host System debug
- HOSTCPUCTICHIN and HOSTCPUCTICHOUT: For connecting to the CTIs within the Host CPU

The SoC CTM connects the Host CTM, SoC CTI, and Counter Secure Enclave CTI together and provides the following expansion interface:

- EXTSYS{0-1}CTICHIN and EXTSYS{0-1}CTICHOUT: For connecting any CTIs and CTMs provided by the External System

8.4.3 CTI expansion

SSE-710 provides the following interfaces for expansion of the CTI:

- HOSTCTICHINEXP/HOSTCTICHOUTEXP for expansion of CTI within the Host System
- EXTSYS{0-1}CTICHIN/EXTSYS{0-1}CTICHOUT for expansion of CTI within the External System. Arm® strongly recommends that if an External System implements CTI, it must provide a Channel Gate controlled by one of the expansion bits of the SCBs.
- HOSTCPUCTICHIN/HOSTCPUCTICHOUT for connection to a CTM within the Host CPU. Arm® strongly recommends that the integrator of the SSE-710 uses these interfaces to integrate the CTI provided by the Host CPU.

8.5 Trace

SSE-710 provides a trace infrastructure to collect trace data from sources and transport it to the trace sinks.

SSE-710 provides the following trace components and interfaces :

- Sources:
  - *System Trace Macrocell* (STM)
  - Host CPU trace: Connected to the *Embedded Trace Macrocells* (ETMs) of the Host CPU
  - Host Debug block trace expansion: To be used to add other trace sources to the Host System
  - External System {0-1} trace expansion: To be used to add trace sources to the External System, for example, Cortex®-M core ETM or ITM
- Funnels:
  - Host Funnel
  - SoC TPIU Funnel
- Replicators:
  - Host Replicator
  - SoC TPIU Replicator
- Sinks:
  - Host Embedded Trace Router
  - SoC Embedded Trace Router
  - SoC TPIU

SSE-710 requires an ETM for each Host CPU core implemented.

The following table shows the connection of the trace component input and output ports.

Table 8-9: Trace connectivity

Component	Input port	Source	Output port	Sink
STM	N/A	N/A	0	Host Funnel ATB port 0
Host CPU Funnel	0	Core 0 ETM	0	Host Funnel ATB port 1
	1	Core 1 ETM		
	2	Core 2 ETM		
	3	Core 3 ETM		
Host Funnel	0	System Trace Macrocell	0	Host Replicator
	1	Host CPU Funnel		
	2	HOSTDBGTRACEEXP		
Host Replicator	0	Host Funnel	0	SoC TPIU Funnel ATB port 0
			1	Host ETR
Host ETR	0	Host Replicator port 1	N/A	
SoC TPIU Funnel	0	Host Replicator	0	SoC TPIU Replicator
	1	EXTSYS0TRACEEXP		
	2	EXTSYS1TRACEEXP		
SoC TPIU Replicator	0	SoC TPIU Funnel	0	SoC TPIU
			1	SoC ETR
SoC TPIU	0	SoC TPIU Replicator port 0	N/A	-
SoC ETR	0	SoC TPIU Replicator port 1		

8.6 System Trace Macrocell

The *System Trace Macrocell* (STM) is a high-bandwidth trace source for software instrumentation and enables hardware events to generate trace data.

The STM has an Extended Stimulus interface, that occupies 16MB in the Host System memory map, and a configuration interface that occupies 4KB in the Host System memory map.

For more information on the Host System memory map, see 12.1.1 Host System memory map on page 188.

For more information on the System Trace Macrocell, see the *Arm® CoreSight™ STM-500 System Trace Macrocell Technical Reference Manual*.

Depending on the security of the master that generated the access to the Extended Stimulus interface, the STM generates trace packets with a different STPv2 MasterID. The following table shows the mapping of different masters in the SSE-710 to STPv2 MasterID.

Table 8-10: SSE-710 Master ID to STM STPv2 Master ID mapping

Master ID	Logical system master	STPv2 Master ID for secure accesses	STPv2 Master ID for non-secure accesses
0	Secure Enclave	0	64
1	Host CPU	1	65

Master ID	Logical system master	STPv2 Master ID for secure accesses	STPv2 Master ID for non-secure accesses
4	AXI AP	4	68
32-63	**IMPLEMENTATION DEFINED**	32-63	96-127
Others	Default master	3	67


Note

Any STPv2 MasterID values that are not listed in the above table are Reserved.

The MasterID used by the STM is taken from the StreamID used by the Firewalls.

Access to the Extended Stimulus port of the STM when the DBGTOP power domain is in the OFF or WARM_RST power mode is treated as RAZ/WI.

Using its hardware event interface the STM can also be used to generate trace packets, based on hardware events. The following table shows the STM hardware events, sensitivity, and source.

Table 8-11: STM hardware events

STM event input	Edge/level	Source
0	Edge	Rising edge of Host CTI Trigger Out 0
1	Edge	Falling edge of Host CTI Trigger Out 0
2	Edge	Rising edge of Host CTI Trigger Out 1
3	Edge	Falling edge of Host CTI Trigger Out 1
4-31	Edge	Reserved
32-64	Level	Reserved

The **NSGUAREN** input of the STM is tied HIGH. This means that Secure and Non-secure accesses to the extended stimulus port behave the same.

For more details of **NSGUAREN**, see *Arm® CoreSight™ STM-500 System Trace Macrocell Technical Reference Manual*.

8.7 ROM tables

SSE-710 has a CoreSight™ ROM table structure.

This structure is compliant with ADIv6, as the following figure shows:

Figure 8-8: ROM table structure

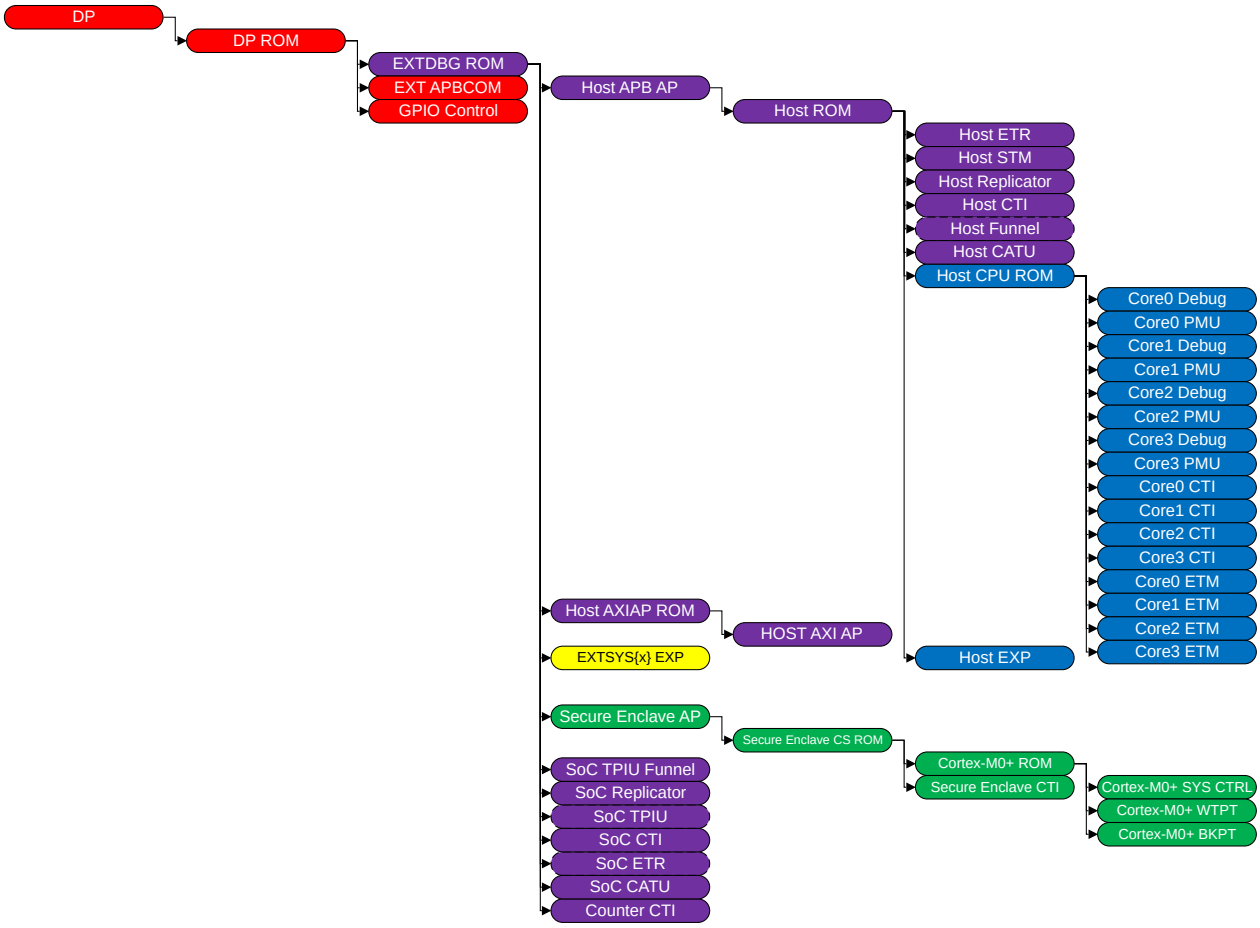

The colors represent the power domain that the component resides in:

Table 8-12: Key

Color	Power domain
Red	AONTOP
Purple	DBGTOP
Yellow	EXTSYS{0-1}TOP
Blue	CLUSTOP
Green	SECENCTOP

For any components inside the SSE-710 which support both an internal and external view, all ROM table entries use the address of the external view.

SSE-710 conforms to the ADIv6 and includes a mix of Class 1 and Class 9 ROM tables. The following ROM tables, implemented by the SSE-710, are Class 9 ROM tables:

- DP ROM
- EXTDBG ROM
- Host ROM
- Host AXIAP ROM

All other ROM tables are implemented as Class 1.

SSE-710 also requires that the integrator adds additional ROM tables for the:

- External System, if it includes more than one debug component in the External System {x} Debug region of the External Debug bus memory map
- Host CPU: SSE-710 supports the supported Cortex®-A processors, which already includes a ROM table.
- Host EXP ROM, if additional debug logic is added to the Host System

The additional ROM tables can be either Class 1 or Class 9 ROM tables, depending on the implementation. Arm® strongly recommends that Class 9 ROM tables are used.

8.8 Granular Power Requestor (GPR)

SSE-710 within the system, uses the *Granular Power Requestor* (GPR) functionality of the ROM tables and the DP CTRL/STATUS register.

The following ROM tables provided by SSE-710 are Class 9 ROM tables with GPR functionality:

- DP ROM
- EXTDBG
- Host ROM
- Host AXIAP ROM

All other ROM tables are implemented as Class 1.

Not all power and reset request signals from SSE-710 ROM tables and DP are used. The following table shows the function of each signal of the component.

Table 8-13: ROM table GPR connectivity

Component name	Signal name	Function
DP	**CDBGPWRUPREQ**	Request for **REFCLK** to be enabled
	CDBGPWRUPACK	Driven HIGH when REFCLK Q-Channel is in the Q_RUN state
	CDBGRSTREQ	Request reset of entire SoC. Signal is connected to Reset Controller.
	CDBGRSTACK	Acknowledge when SoC has been reset. Driven by the Reset Controller.
	CSYSPWRUPREQ	Reserved. Connected to **CSYSPWRUPACK** of the DP.
	CSYSPWRUPACK	Reserved. Connected to **CSYSPWRUPREQ** of the DP.

Component name	Signal name	Function
DP ROM	**CDBGPWRUPREQ0**	Power request for DBGTOP
	CDBGPWRUPACK0	Driven HIGH when DBGTOP is in the ON power mode
	CDBGPWRUPREQ[31:1]	Reserved
	CDBGPWRUPACK[31:1]	
	CSYSPWRUPREQ[31:0]	
	CSYSPWRUPACK[31:0]	
	CDBGRSTREQ	Request reset of DBGTOP
	CDBGRSTACK	Driven HIGH when DBGTOP is in WARM_RESET power mode
	CSYSRSTREQ	Request **nSRST**. Connected to Reset Controller.
	CSYSRSTACK	Acknowledge when SoC has been reset and is held before Secure Enclave starts to boot. Driven by Reset Controller.
SDC-600	**REMPUR**	Power request for SYSTOP
	REMPUA	Driven HIGH when SYSTOP is in the ON or FUNC_RET power mode
	REMRR	Connected to CoreSight™ SDC-600 **REMRA**
	REMRA	Connected to CoreSight™ SDC-600 **REMRR**
EXTDBG ROM	**CDBGPWRUPREQ0**	Power request for SECENCTOP
	CDBGPWRUPACK0	Driven HIGH when the SECENCTOP is in ON
	CDBGPWRUPREQ[2:1]	Power request for External System {0-1}. This signal is part of the External System harness, EXTSYS{0-1}PWRREQ interface. When the associated External System is not implemented, then the signal is Reserved – Connected to the respective **CDBGPWRUPACK**.
	CDBGPWRUPACK[2:1]	Acknowledge when the External System has entered a power mode where accesses to the debug components is allowed. This signal is part of the External System Harness EXTSYS{0-1}PWRREQ interface. When the associated External System is not implemented, then the signal is Reserved: Connected to respective **CDBGPWRUPREQ**.
	CDBGPWRUPREQ[31:3]	Reserved
	CDBGPWRUPACK[31:3]	
	CSYSPWRUPREQ[31:0]	
	CSYSPWRUPACK[31:0]	
	CDBGRSTREQ	Reserved. Connected to **CDBGRSTACK**.
	CDBGRSTACK	Reserved. Connected to **CDBGRSTREQ**.
	CSYSRSTREQ	Reserved. Connected to **CSYSRSTACK**.
	CSYSRSTACK	Reserved. Connected to **CSYSRSTREQ**.
Host ROM	**CDBGPWRUPREQ0**	Power request for Host CPU power domain. This signal is connected to the CLUSTOP PPU.
	CDBGPWRUPACK0	Acknowledge when the Host CPU is accessible. This signal is driven by the CLUSTOP PPU, when the domain is in the ON or FUNC_RET power modes.
	CDBGPWRUPREQ[31:1]	Reserved
	CDBGPWRUPACK[31:1]	
	CSYSPWRUPREQ[31:0]	
	CSYSPWRUPACK[31:0]	
	CDBGRSTREQ	Reserved. Connected to **CDBGRSTACK**.
	CDBGRSTACK	Reserved. Connected to **CDBGRSTREQ**.
	CSYSRSTREQ	Reserved. Connected to **CSYSRSTACK**.

Component name	Signal name	Function
	CSYSRSTACK	Reserved. Connected to **CSYSRSTREQ**.
Host AXIAP ROM	**CDBGPWRUPREQ[31:0]**	Reserved
	CDBGPWRUPACK[31:0]	
	CSYSPWRUPREQ0	Power request for SYSTOP to be in ON or FUNC_RET. This signal is connected to the SYSTOP PPU.
	CSYSPWRUPACK0	Acknowledge when SYSTOP is in the ON or FUNC_RET power mode
	CSYSPWRUPREQ1	Power request for CLUSTOP to be in ON or FUNC_RET. This signal is connected to the CLUSTOP PPU.
	CSYSPWRUPACK1	Acknowledge when CLUSTOP is in the ON or FUNC_RET power mode
	CSYSPWRUPREQ[3:2]	Power request for EXTSYS{0-1}TOP power domain. When the associated External System is not implemented, then the signal is Reserved, connected to respective **CDBGPWRUPACK**.
	CSYSPWRUPACK[3:2]	Acknowledge for EXTSYS{0-1}TOP power domain is in a power mode to be able to process transactions. When the associated External System is not implemented, then the signal is Reserved, connected to respective **CDBGPWRUPREQ**.
	CSYSPWRUPREQ[15:4]	Reserved
	CSYSPWRUPACK[15:4]	
	CSYSPWRUPREQ[31:16]	**IMPLEMENTATION DEFINED**. Exposed in the HOSTDBGPWRREQ interface. The integrator uses these signals for any other power domains, which can be accessed by the Host AXI AP.
	CSYSPWRUPACK[31:16]	**IMPLEMENTATION DEFINED**. Exposed in the HOSTDBGPWRREQ interface. The integrator uses these signals to indicate when the power domain powered.
	CDBGRSTREQ	Reserved. Connected to **CDBGRSTACK**.
	CDBGRSTACK	Reserved. Connected to **CDBGRSTREQ**.
	CSYSRSTREQ	Reserved. Connected to **CSYSRSTACK**.
	CSYSRSTACK	Reserved. Connected to **CSYSRSTREQ**.

For the **CSYSPWRUPREQ/ACK** and **CDBGPWRUPREQ/ACK** interfaces of the ROM tables which are described as Reserved in the table above, the **REQ** is looped back to the **ACK**. The value in the ROM table DEVID.NUMREQ field reflects the number of **REQ/ACK** signals implemented by the ROM table.

8.9 CoreSight timestamp

SSE-710 supports a CoreSight™ timestamp. However, its control and operation frequency is **IMPLEMENTATION DEFINED**.

8.10 GPIO control

The GPIO Control is a peripheral on the External Debug bus which is only accessible through the DP. Access from any other masters treats the location as Reserved.

The GPIO Control enables a debugger to control the CALC interface.

For more details, see 4.5.1.8 Crypto Accelerator lifecycle Control (CALC) interface on page 65.

9. Secure Enclave

This chapter describes the Secure Enclave, which is deigned to be Root of Trust within the BSA and considers all other logic outside the Enclave as less trustworthy than itself.

9.1 Secure Enclave components

This section describes Secure Enclave components.

9.1.1 Cryptographic Accelerator

The Secure Enclave defines a Crypto Accelerator Socket to allow for integration of any cryptographic accelerator and lifecycle management logic to be added to the SoC.

The Crypto Accelerator Socket provides the ability to implement a Crypto Accelerator split across the SECENCTOP and AONTOP power domains. This specification refers to the parts of the Crypto Accelerator as Crypto Accelerator and Crypto Accelerator AON respectively. Alongside this it uses the term Crypto Accelerator Socket SECENCTOP and Crypto Accelerator Socket AON to refer to the two parts of the Crypto Accelerator Socket to accept the two parts of the Crypto Accelerator.

The following requirements must be obeyed of any Crypto Accelerator which is implemented into the Crypto Accelerator Socket:

- The cryptographic algorithms which are implemented by the Crypto Accelerator are **IMPLEMENTATION DEFINED**.
- The Crypto Accelerator must provide a method to store the current lifecycle state of the SoC in a way which is only modifiable to memory accesses via the CAC interface.
- The Crypto Accelerator must prevent the lifecycle state from returning to a previous lifecycle state.
- The Crypto Accelerator must store the lifecycle state in non-volatile storage which is only accessible by the Crypto Accelerator.
- When the Crypto Accelerator is released from reset it must read the lifecycle state from the non-volatile storage location and use the value to set the default values on the SCBs.
- The Crypto Accelerator must drive the SCB based on the rules defined in 9.1.3 Security Control Bits (SCB) on page 150 section and the current lifecycle state of the SoC.
- When integrating the Crypto Accelerator into the Crypto Accelerator Socket the integrator must tie-off any un-used interface to prevent any deadlock from occuring or any weaking of the security features implemented by the Secure Enclave.

9.1.2 Lifecycle States (LCS)

Using the **IMPLEMENTATION DEFINED** cryptographic engine, the Secure Enclave provides the *LifeCycle State* (LCS) of the SoC.

The lifecycle state is always one of the following:

Table 9-1: Lifecycle States

LCS	Description
Chip Manufacture	Initial state after manufacture
Device Manufacture	Used during device development
Secure Enable	Used when the device is deployed
Return Merchandise Authorization	Used when the device has reached end-of-life. At this point only Secure Enclave can boot, no other system can boot . It is **IMPLEMENTATION DEFINED** whether the Secure Enclave firmware is able to request the Host or External System to be able to boot by updating the corresponding Security Control Bit (SCB) value.

The lifecycle runs in a linear direction from Chip Manufacture → Device Manufacture → Secure Enable → Return Merchandise Authorization.

To change the lifecycle state of the SoC, software executing on the Secure Enclave Cortex®-M0+ must perform an **IMPLEMENTATION DEFINED** sequence.

The SSE-710 has the *Crypto Accelerator Lifecycle Control* (CALC) interface. This interface provides a physical signal to prevent advancement of the lifecycle state by accidental or malicious means. The CALC interface is `0b1` under one of the following conditions:

- *SoC Lifecycle Control* (SOCLCC) interface of SSE-710 is `0b1`. The Agent performing the sequence must assert the SOCLCC interface before releasing **PORESETn** to the SoC and keep it asserted until the transition is complete.
- GPO0 of the GPIO Control is `0b1`. The debug agent must follow these steps:
  1. Use the DP to program the GPIO Control to assert the GPO0 output
  2. Cause a debug reset by either:
     - Asserting the **nSRST** input
     - Setting the DP ROM CSYSRSTREQ register to `0b1` and wait for the DP ROM CSYSRSTACK to become `0b1`
  3. Allow the Secure Enclave Cortex®-M0+ to perform the software sequence by either:
     - De-asserting the **nSRST** input
     - Setting the DP ROM CSYSRSTREQ to `0b0` and waiting for the DP ROM CSYSRSTACK to become `0b0`


Note

The sequence either uses the **nSRST** input or the DP ROM CSYSRSTREQ/ACK handshake to cause the reset.

9.1.3 Security Control Bits (SCB)

The Secure Enclave uses the Security Control Bits (SCB) to control features within the SSE-710 subsystem.

The following table shows the bit assignment of the SCB interface as defined in the 4.5.3 Security Control Bits (SCB) interface on page 67. Any Reserved bits are tied LOW.

Table 9-2: SCB Interface bit assignment

SCB offset	Name	Description
0-1	Reserved	
2	SECENCAUTH_DBGEN	Controls the SECENCAUTH DAZ
3	SECENCAUTH_NIDEN	
4	SECENCAUTH_CHEN	Controls the Channel Gate in the SECENCAUTH DAZ
5	DPAUTH_DBGEN	Controls the DPAUTH DAZ
6	DPAUTH_NIDEN	
7	DPAUTH_SPIDEN	
8	DPAUTH_SPNIDEN	
9	COMAUTH_PEN	Controls the COMAUTH DAZ
10	COMAUTH_RRDIS	
11-12	Reserved	-
13	TPIUAUTH_DBGEN	Controls the TPIUAUTH DAZ
14	TPIUAUTH_NIDEN	
15	TPIUAUTH_SPIDEN	
16	TPIUAUTH_SPNIDEN	
17	TPIUATH_CHEN	Controls the Channel Gate in the TPIUAUTH DAZ
18	COUNTERAUTH_DBGEN	Controls the COUNTERAUTH DAZ
19	COUNTERAUTH_NIDEN	
20	COUNTERAUTH_SPIDEN	
21	COUNTERAUTH_SPNIDEN	-
22	COUNTERAUTH_CHEN	Controls the Channel Gate in the COUNTERAUTH DAZ
23	HOSTEXTAUTH_NS	Controls HOSTEXTAUTH DAZ
24	HOSTEXTAUTH_S	
25	HOSTAXIAUTH_DBGEN	Controls HOSTAXIAUTH DAZ
26	HOSTAXIAUTH_NIDEN	

SCB offset	Name	Description
27	HOSTAXIAUTH_SPIDEN	
28	HOSTAXIAUTH_SPNIDEN	
29	HOSTAUTH_DBGEN	Controls the HOSTAUTH DAZ, and drives the HOSTCPUDBGAUTH and HOSTDBGAUTH interfaces.
30	HOSTAUTH_NIDEN	
31	HOSTAUTH_SPIDEN	
32	HOSTAUTH_SPNIDEN	
33	HOSTAUTH_CHEN	Controls the Channel Gate in the HOSTAUTH DAZ
34	SOC_DFTENABLE	Drives the **DFTENABLE[0]** signal of the SOCSC interface
35	PPU_DBGEN	Controls whether the PPU debug functionality is enabled
36	SECENC_FW_BYPASS	Drives the Bypass interfaces of the Secure Enclave Firewall
37	HOST_FW_BYPASS	Drives the Bypass interfaces of the Host System Firewall
38	DPEXTACG	Controls the DAACG on the DP port of the External Debug Bus
39	HOSTEXTACG	Controls the DAACG on the Host System port of the External Debug Bus
40	EXTSYS0EXTACG	Controls the DAACG on the External System 0 port of the External Debug Bus
41	EXTSYS1EXTACG	Controls the DAACG on the External System 1 port of the External Debug Bus
42-47	Reserved	-
48	HOST_CPUWAIT_WEN	Controls whether the HOST_SYS_RST_CTRL.CPUWAIT field is writeable
49	EXT_SYS0_CPUWAIT_WEN	Controls whether the EXT_SYS0_RST_CTRL.CPUWAIT field is writeable
50	EXT_SYS1_CPUWAIT_WEN	Controls whether the EXT_SYS1_RST_CTRL.CPUWAIT field is writeable
51-62	Reserved	-
63	SECENC_DFTENABLE	Drives the **DFTENABLE[1]** signal of the SOCSC interface
64-127	SoC Expansion	Used for **IMPLEMENTATION DEFINED** use-cases within the SoC, for example, controlling the debug privileges of the External System

The SSE-710 subsystem defines the following rules for the SCB:

- The default values of an SCB are dependent on the lifecycle state of the SoC. Until the lifecycle state is known all SCB must have a value of `0b0`.
- The values of the SCBs must not change to the lifecycle dependent value until the lifecycle has been determined and the LCS interface is stable.
- In all lifecycle states the **COMAUTH – PEN** is `0b1` by default in all lifecycle states.
- In all lifecycle states the **COMAUTH – RRDIS** is **IMPLEMENTATION DEFINED** in all lifecycle states.
- In the Chip Manufacture lifecycle state:
  - The value of all SCBs, other than **COMAUTH** signals, is **IMPLEMENTATION DEFINED**.
  - It is **IMPLEMENTATION DEFINED** whether the SCB can be updated by either:
    - Software running on the Secure Enclave. Arm® recommends that this is done as part of a certificate authentication.
    - Debug accesses using the Secure Enclave Cortex®-M0+


Note

Arm® recommends that in the Chip Manufacture state, the minimum enabled features allow for a debug agent to be able to transition the SoC to the Device Manufacture state.

- In the Device Manufacture lifecycle state:
  - The value of all SCBs, other than **COMAUTH** signals, is **IMPLEMENTATION DEFINED**.
  - It is **IMPLEMENTATION DEFINED** whether the SCB can be updated by either:
    - Software running on the Secure Enclave. Arm® recommends that this is done as part of a certificate authentication.
    - Debug accesses using the Secure Enclave Cortex®-M0+


Note

Arm® recommends that in the Device Manufacture state, the minimum enabled features allow for a debug agent to be able to debug the SoC. This can be by the injection of a certificate using the CoreSight™ SDC-600, or having debug access enabled by default

- In the Secure Enclave lifecycle state:
  - The values of all SCBs, other than **COMAUTH**, **HOST_CPUWAIT_WEN**, and **EXT_SYS{0-1}_CPUWAIT_WEN** signals, must be 0.
  - The **HOST_CPUWAIT_WEN** and **EXT_SYS{0-1}_CPUWAIT_WEN** signals must be 1.
  - The following SCBs must be updatable:
    - DPAUTH_{x}
    - TPIUAUTH_{x}
    - COUNTERAUTH_{x}
    - HOTEXTAUTH_{x}
    - HOSTAUTH_{x}
    - HOSTAXIAUTH_{x}
    - HOSTEXTACG
    - EXTSYS{0-1}EXTACG

    by either:
    - Software running on the Secure Enclave. Arm® recommends that this is done as part of a certificate authentication.
    - Debug accesses using the Secure Enclave Cortex®-M0+


Note

Arm® strongly recommends that both methods are implemented.

  - For all other SCBs it is **IMPLEMENTATION DEFINED** whether the bits are updatable.

- In the Return Merchandise Authorization lifecycle state:
  - The values of all SCBs, other than **COMAUTH – PEN**, is **IMPLEMENTATION DEFINED**.
  - It is **IMPLEMENTATION DEFINED** whether the SCB can be updated by either:
    - Software running on the Secure Enclave. Arm® recommends that this is done as part of a certificate authentication.
    - Debug access using the Secure Enclave Cortex®-M0+

  Arm® recommends that for:

  Security requirements of the SoC in the lifecycle

  To prevent the Secure Enclave keys being extracted by an attacker when enabling debug, the keys must be destroyed on entry into the Return Merchandise Authorization state.

  Debuggability requirements of the SoC in the lifecycle

  Discovering the issues with returned devices is hampered if a debug agent is not enabled, or there is no way to re-enable debug in the Return Merchandise Authorization state.

9.1.4 Secure Enclave Cortex-M0+

The Secure Enclave contains a Cortex®-M0+ processor, configured as follows:

- Little-endian data
- NVIC with support for 32 interrupts
- MPU with 8 regions
- Two data watchpoints
- Four breakpoints
- Halted debug support
- Privileged and Unprivileged support
- Systick Timer
- Vector Table Offset Register support
- Disables support for individual interrupts. The value is `0xEB5F_8010`.
- Architectural Clock Gating support

The SysTick timer of the Secure Enclave Cortex®-M0+ uses the following clocks:

- **SECENCDIVCLK**
- **S32KCLK**

For the definitions of Clock Domains, see 5. Clocks on page 72

9.1.5 Secure Enclave reset

The Secure Enclave can request a reset of itself, or other systems in the SoC, using any of the following mechanisms:

- A Secure Enclave or SoC watchdog reset request. When either of these events occur, the entire SoC is reset.
- A Secure Enclave software reset request triggered using the Cortex®-M0+ AIRCR.SYSRESETREQ field. The entire SoC is reset, except for all logic on **AONTOPPORESETn**.

  For more information on Reset Domains, see 7. Reset on page 108.
  For more information on Cortex®-M0+ AIRCR.SYSRESETREQ, see the *Cortex®-M0+ Technical Reference Manual*.
- Using the Secure Enclave Base System Control registers:
  - Initiate reset of the Host and External System, using the HOST_SYS_RST_CTRL.RST_REQ field.
  - Initiate reset of the SoC, using the SOC_RST_CTRL.RST_REQ field.

9.1.6 Secure Enclave peripherals

This section describes peripheral devices attached to the Secure Enclave.

9.1.6.1 ROM and RAM

The Secure Enclave has its own dedicated ROM and RAM.

For more information on the Secure Enclave ROM and RAM, see the 12.1.3.1 Secure Enclave ROM region on page 195 and 12.1.3.2 Secure Enclave RAM region on page 196.

9.1.6.2 Interrupt collator

The Secure Enclave has an interrupt collator enabling more than 32 interrupts to be handled by the Secure Enclave Cortex-M0+ processor.

It is controlled by the Secure Enclave Base System Control registers. For more information on these registers, see the 12.3.2.1 Secure Enclave Base System Control register summary on page 236.

9.1.6.3 Timers

The Secure Enclave has two CMSDK timers, Timer 0 and 1. These are in the SECENCTOP power domain and use the **SECENCDIVCLK** clock.

For more information, see 5.2.5 SECENCDIVCLK on page 77.

Both of the Secure Enclave CMSDK timers support being halted when the Secure Enclave Cortex®-M0+ iis halted by a debugger. The timer is halted when the Secure Enclave Cortex®-M0+ is halted if software sets bit 1 of the CTRL register of the timer to `0b1`. For more information on the Timers see the *Arm® Cortex®-M System Design Kit Technical Reference Manual*.

9.1.6.4 Watchdogs

The Secure Enclave includes two CMSDK watchdogs:

- The Secure Enclave Watchdog is in the SECENCTOP power domain and uses **SECENCDIVCLK**
- The SoC Watchdog is in the AONTOP power domain and uses **S32KCLK**

An access to the SoC Watchdog takes 6-7 **S32KCLK** clock cycles. Accessing it too often might lead to software performance issues.

For more information on the CMSDK Watchdog, see the *Arm® Cortex®-M System Design Kit Technical Reference Manual*.

On the first expiry of the counter, both watchdogs generate interrupts to the Secure Enclave Cortex®-M0+. On the second expiry, they generate a reset request to the Reset Controller of SSE-710 and cause a reset of the entire SoC. For more details, see 7. Reset on page 108.

When the Cortex®-M0+ is halted by a debugger, both Secure Enclave watchdogs are halted.

9.1.6.5 Secure Enclave MHUs

The MHUs between the Host System and the Secure Enclave are split across the SECENCTOP and SYSTOP power domains. The MHUs between the External Systems and the Secure Enclave are split across the EXTSYS{0-1}TOP and SECENCTOP power domain.

For more details of MHUs, see 11.5 MHU on page 184.

9.1.6.6 System Control registers

The System Control registers are in the AONTOP power domain and control various aspects of the Secure Enclave.

For more information see the 12.3.2.2 Secure Enclave System Control register summary on page 249.

9.1.6.7 Base System Control registers

The Base System Control registers are in the Always ON power domain. They control the features of the Host and External Systems in the integration of the Secure Enclave System into the SSE-710 subsystem.

For more information, see the 12.3.2.1 Secure Enclave Base System Control register summary on page 236.

9.1.6.8 UART

The Secure Enclave has a PL011 UART located in the AONTOP power domain.

The Secure Enclave supports hardware-based flow control, and uses the **SECENCDIVCLK** for both the **PCLK** and **UARTCLK** inputs of PL011. The Secure Enclave uses only the combined interrupt from the UART.

For more information on the PL011 UART see the *PrimeCell UART (PL011) Technical Reference Manual*.


Note

The SSE-710 subsystem only supports the PL011 as an RS232-compliant UART.

4.5.2 Secure Enclave UART (SECENCUART) interface on page 66 describes the Secure Enclave UART interface, SECENCUART.

9.1.6.9 SECENCTOP PPU

The Secure Enclave has a *Power Policy Unit* (PPU) controlling the SECENCTOP power domain.

For more information on the PPU and SECENCTOP, see the *Arm® CoreLink™ PCK-600 Power Control Kit Technical Reference Manual* and 6.4.6 SECENCTOP on page 96.

9.1.6.10 Secure Enclave Firewall

The Secure Enclave has a dedicated firewall instance.


Note

The SSE-710 subsystem has two firewalls. This section only describes the Secure Enclave firewall. For more details, see Firewall appendix.

This firewall is configured as follows:

- Lockdown Extension level 0

- Save and Restore Extension level 0
- Security Extension level 1

The Secure Enclave firewall has the same firewall interfaces as defined in the Firewall appendix , except that it does not have Lockdown and Tamper Interrupt interfaces.

The Secure Enclave firewall is in the SECENCTOP power domain and includes two Firewall Components (FC0 – FCTLR and FC1). FC1 is on the memory path to the Host System.

The Protection Size of the Secure Enclave Firewall Component 1 is set to 4GB. 10.3.3 IMPLEMENTATION DEFINED behavior on page 164 defines the firewall **IMPLEMENTATION DEFINED** behavior. All of this behavior applies to the Secure Enclave firewall.

The following table describes the Firewall Component configuration in:

Table 9-3: Firewall component configuration

Firewall component	PE_LVL	ME_LVL	TE_LVL	RSE_LVL	NUM_RGN	MNRS	MXRS	NUM_MPE	SINGLE_MS
FCTLR	1	0	0	1	3	7	21	1	1
FC1	2	2	2	0	8	7	32	1	1

The following configurable values apply to all Firewall Components:

- MST_ID_WIDTH is 1 bit. All transactions have a fixed StreamID of 0 for the Secure Enclave firewall.
- SEC_SPT is `0b1`
- MA_SPT is `0b1`
- SH_SPT is `0b0`
- INST_SPT is `0b1`
- PRIV_SPT is `0b1`

For Firewall Components that use PE.1, the following regions are predefined:

- Firewall Controller:
  - As defined in the Firewall appendix.
  - The Configuration Master is set to the ID of the Secure Enclave.

When the Secure Enclave firewall terminates a read transaction, or detects a read transaction marked as either an AMBA AXI5 SLVERR or DECERR, and the Firewall's monitor logic is enabled, the read data value is set to `0xDEAD_DEAD`.

For more information on the SSE-710 subsystem Firewall, see the Firewall appendix.

10. Interconnect

This chapter describes the Host System Interconnect, which connects masters and slaves of the Host System.

10.1 NIC

SSE-710 contains an internal fabric responsible for connecting all master and slaves within SSE-710. This internal fabric is created using CoreLink™ NIC-400 switches.

SSE-710 contains the following switches:

- NIC (Main)
- NIC (AONPERIPH)
- NIC (SYSPERIPH)
- NIC (DBGPERIPH)

The following figure shows the connection between the NIC switches.

Figure 10-1: Interconnect interfaces

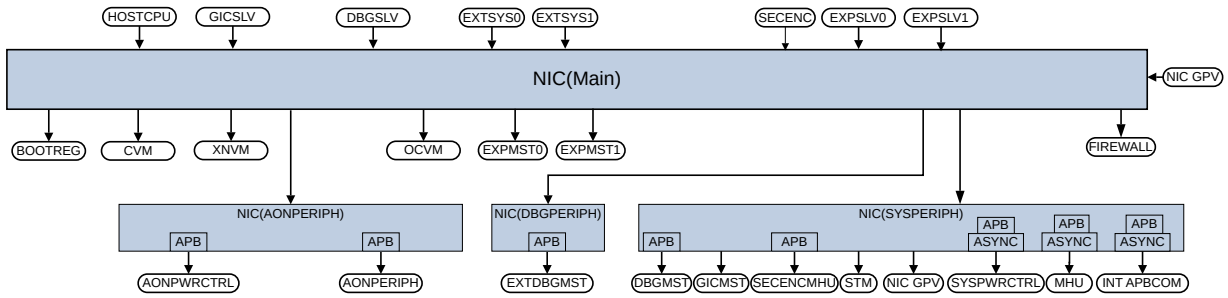

The Main NIC is the main switch connecting all the masters, both internal and external to SSE-710, to the slaves of the SoC. For some slaves, access is via one of the peripheral NICs (AONPERIPH / SYSPERIPH / DBGPERIPH).

The peripheral NICs (AONPERIPH/ SYSPERIPH/ DBGPERIPH) provide access to peripherals in the AONTOP, SYSTOP, and DBGTOP power domains. Some of the ports on these NICs provide protocol conversion to APB and implement clock domain crossing.

NIC(Main) includes a *Global Programmers View* (GPV).

The NIC interconnect is not a full crossbar and certain masters can only access certain slaves as the following table shows:

Table 10-1: Slave access to master interfaces of interconnect

Master interface	Address ranges	Slave interface					
		HOST CPU	GIC SLV	DBG SLV	EXT SYS {0-1}	SEC ENC	EXP SLV {0-1}
BOOTREG	0x0000_0000 – 0x0000_0FFF	Y	N	Y	N	Y	N
CVM	0x0200_0000 – 0x03FF_FFFF		Y	Y	Y	Y	Y
XNVM	0x0800_0000 – 0x0FFF_FFFF		Y	Y	Y	Y	Y
DBGMST	0x1000_0000 – 0x17FF_FFFF		N	N	N	Y	N
EXTDBGMST	0x1800_0000 – 0x19FF_FFFF		N	N	N	Y	N
AONPWRCTRL	0x1A02_0000 – 0x1A04_FFFF		N	Y	N	Y	N
AONPERIPH	0x1A00_0000 – 0x1A01_FFFF 0x1A20_0000 – 0x1A6F_FFFF		N	Y	Y	Y	Y
FIREWALL	0x1A80_0000 – 0x1A9F_FFFF		N	Y	N	Y	N
MHU	0x1B00_0000 – 0x1B0F_FFFF		N	Y	N	Y	N
SECENCMHU	0x1B80_0000 – 0x1B83_FFFF		N	Y	Y	Y	N
INT APBCOM	0x1B90_0000 – 0x1B90_FFFF		N	Y	N	Y	N
SYSPWRCTRL	0x1BC0_0000 – 0x1BC4_FFFF		N	Y	N	Y	N
GICMST	0x1C00_0000 – 0x1CFF_FFFF		N	Y	Y	Y	Y
STM	0x1D00_0000 – 0x1DFF_FFFF		N	Y	N	Y	Y
NIC GPV	0x1E00_0000 – 0x1E0F_FFFF		N	Y	N	Y	N
EXPMST{0-1}	0x4000_0000 – 0x7FFF_FFFF		Y	Y	Y	Y	Y
OCVM	0x8000_0000 – 0xFFFF_FFFF		Y	Y	Y	Y	Y

Attributes of the AXI master interface

The following table lists AXI ID width and read/write issuing capability of the SSE-710 master AXI interfaces.

Table 10-2: AXI master interface attributes

Interface	Read issuing capability	Write issuing capability	AXI ID width
CVM	32	32	Depends on the interface ID widths of the AXI slave. Calculate according to the following formula: Max(AXI slave interface ID width) + 1 + ceil(log2(number of incoming ports)) Max(AXI slave interface ID width) is 8 or maximum of `EXPSLV{0-1}_ID_WIDTH` or `EXTSYS{0-1}_ID_WIDTH` if any is set to bigger than 8 Number of incoming ports = 3 + `NUM_EXP_SLV` + `NUM_EXT_SYS` **Note:** `NUM_EXP_SLV` and `NUM_EXT_SYS` are fixed to 2 in SSE-710.
XNVM			
OCVM			
EXPMST{0-1}			

Attributes of the AXI slave interface

The following table lists the AXI ID width and the read/write issuing capability of the SSE-710 slave AXI interfaces.

Table 10-3: AXI slave interface attributes

Interface	Read issuing capability	Write issuing capability	AXI ID width
HOST CPU	32	16	8
EXTSYS 0	8	8	EXT_SYS0_ID_WIDTH
EXTSYS 1	8	8	EXT_SYS1_ID_WIDTH
EXPSLV{0-1}	32	16	Configurable

10.2 Quality of Service (QoS)

The SSE-710 Main CoreLink™ NIC-400 implementation supports a programmable QoS scheme for the Host CPU, External Systems {0-1}, Debug AXI, and Secure Enclave interfaces.

Also, the Expansion Slaves {0-1} provide QoS values from the protocol interface ("From Master"):

- "Programmable" for Host CPU, External Systems {0,1}, Debug AXI and Secure Enclave:
  - The interface does not include the **ARQOS** and **AWQOS** signals.
  - The QoS setting is configured via the GPV.
- "From master" for the Expansion Slave {0,1}: In this case, the interface has the **ARQOS** and **AWQOS** signals.

The following interfaces have the **ARQOS** and **AWQOS** signals included:

- CVM
- XNVM
- OCVM
- HOSTEXPMST{0-1}

To support the programmable option for QoS, the *Global Programmers View*(GPV) has been enabled for the respective slave interfaces. The address offsets from the CoreLink™ NIC-400 Main GPV base are listed below:

- `0x42000-0x42FFF`: HOSTCPU
- `0x43000-0x43FFF`: EXTSYS0
- `0x44000-0x44FFF`: EXTSYS1
- `0x45000-0x45FFF`: SECENC
- `0x46000-0x46FFF`: DBGSLV

10.3 Firewall

The Firewall has been implemented within the subsystem to provide both monitoring and protection of the address space, between the many different entities within the SSE-710.

This section describes the SSE-710 Host System Firewall configurations and the **IMPLEMENTATION DEFINED** behavior of the Firewall. For the full set of generic Firewall features, see C. Firewall on page 346.

10.3.1 Firewall Component interfaces

The Firewall Component includes clock, reset, bypass, DFT control, AXI4 Stream configuration, AXI5 Master and Slave, and Q-Channel power and clock control interfaces.

The following figure shows the *Firewall Component* (FC) interfaces.

Figure 10-2: Firewall Component interfaces

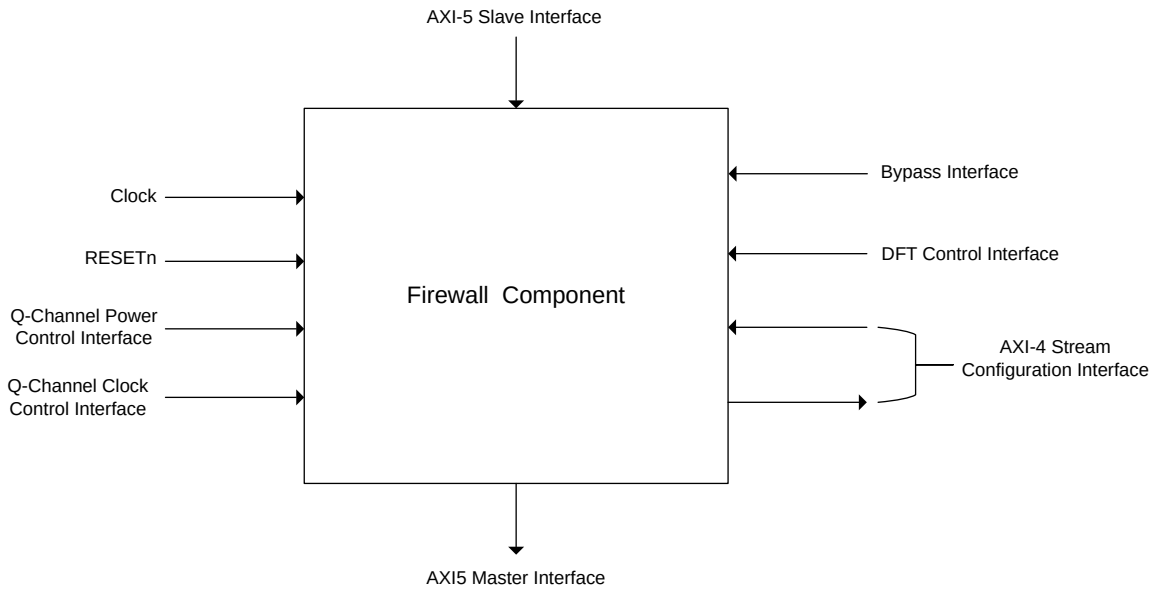

10.3.1.1 AXI5 Slave and Master interfaces

The slave interface is the ingress port where a transaction enters the FC, and the master interface is the egress where a transaction exits.

The interfaces have the following properties:

- AMBA AXI5 protocol with the following properties set to true:
  - Wakeup_Signal
  - Untranslated_Transaction

For more details, see *AMBA® AXI and ACE Protocol Specification*.

10.3.1.2 AXI4 Stream Configuration interface

This configuration interface is bidirectional, and is used for communication between the *Firewall Controller* (FCTLR) and the FCs.

The interface is implemented based on the *AMBA® 4 AXI4-Stream Protocol Specification*, in which messages can be initiated from either side.

10.3.1.3 Low Power interfaces

The Firewall Component has two low power Q-Channel interfaces.

The low power interfaces are:

- A Q-channel for clock control
- A Q-channel for power gating

The low power states implemented in the Firewall are based on the states defined in the *Arm® Power Control System Architecture Specification*.

10.3.1.4 Bypass interface

The Bypass interface serves to determine the behavior of protection logic checks that are applied by the Firewall on transactions.

In the SSE-710, the Bypass interface of the Host System Firewall is driven by an SCB bit.

For more information on the Bypass interface, see C. Firewall on page 346.

10.3.2 Firewall Controller interfaces

The Firewall Controller includes clock, reset, interrupt, bypass, lockdown, protection, DFT control, AXI4 Stream configuration, AXI5 programming, P-Channel Power Control, and Q-Channel Clock control interfaces.

The following figure shows the *Firewall Controller* (FCTLR) interfaces.

Figure 10-3: FCTLR interfaces

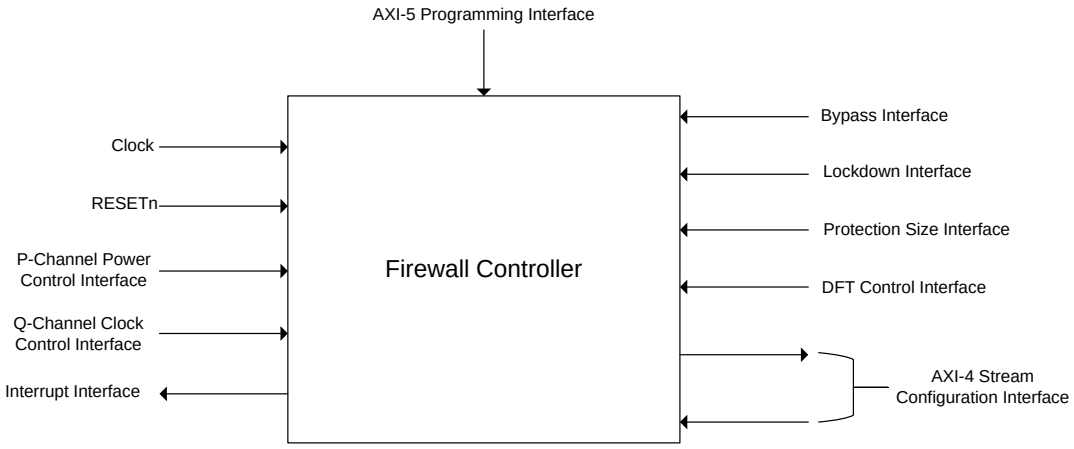

AXI5 Programming interface

This interface is used to program the Firewall configuration registers inside the FCTLR and FCs.

The interface has the following properties:

AMBA AXI5 protocol with the following properties set to true:

- Wakeup_Signal
- Untranslated_Transaction

For more information on supported configuration access through this interface to the Firewall Controller, see 10.3.3 IMPLEMENTATION DEFINED behavior on page 164.

AXI4 Stream Configuration interface

This configuration interface is bidirectional, and is used for communication between the FCTLR and the FCs.

For more information, see 10.3.1.2 AXI4 Stream Configuration interface on page 162.

Lockdown interface

The lockdown input controls the lockdown extension functionality of the Firewall.

For more information on the lockdown behavior, see C. Firewall on page 346.

Interrupt interface

The Firewall Controller implements two interrupt interfaces.

The two interrupt interfaces are for:

- Firewall interrupt: A single interrupt signal for the entire Firewall

- Tamper interrupt: A separate interrupt interface to generate tamper interrupt to the system

Protection Size interface

The Protection Size interface is an input vector which is organized as a single dimensional array in which each 8 bits is allocated to one FC.

In the SSE-710, the sampled values on the protection interface of the Host FCs are defined in 10.3.5 Host System firewall on page 169.

Low Power interfaces

The Firewall Controller has two low power interfaces.

These low power interfaces are:

- A Q-Channel interface for clock control
- A P-Channel interface for power gating. The P-Channel supports the following power modes: OFF, FUNC_RET, ON.

The low power states implemented in Firewall are based on the states defined in *Arm® Power Control System Architecture Specification*.

Bypass interface

The Bypass interface serves to determine the behavior of protection logic checks that are applied by the Firewall on transactions.

The FCTLR Bypass interface behaves in a same way as FC Bypass interface. For more information, see 10.3.1.4 Bypass interface on page 162.

10.3.3 IMPLEMENTATION DEFINED behavior

This section applies to both the Host System and Secure Enclave Firewalls, unless otherwise stated.

.C. Firewall on page 346 describes behaviors that are **IMPLEMENTATION DEFINED**.


Note

This section must be read with C. Firewall on page 346.

The following **IMPLEMENTATION DEFINED** behaviors that are visible to software or impact on the SSE-710:

- The firewall always occupies 2MB with any unused 64KB pages, allocated to firewall components, being marked as Reserved.
- Bus slave, Bus master, and Programming interfaces:

    - Burst based with address that is reported in the Fault Entries and Error Detection Reports using the starting address of the transaction.
    - There is one Bus slave and master interface per Firewall Component which supports the same address, data, and transaction properties.
    - Any transaction that is received on the Bus slave interface is forwarded to the Bus master interface if it passes the protection logic checks or the Firewall Bypass interface is asserted.
- Transaction processing, for protection and monitoring:
    - Firewall components process transactions for read and write transactions separately.
    - There are no guarantees about the order in which Fault Entries or Error Detection Reports are generated:
        - Between read and write transactions
        - Between read or write transactions which have different AXI IDs
- Firewall Controller only implements region 0 to 2 as required by the specification.
- Shadow Registers, when Save and Restore Extension level 1 (SRE.1) is implemented:
    - Implemented as SRAM
    - Support memory retention
    - Applies compression
- When a firewall component (including the Firewall Controller), which supports Protection Extension level 1 (PE.1) or greater, terminates a transaction it:
    - Generates an AMBA AXI5 DECERR, if PE_ST.ERR is set to `0b1`.
    - Generate StreamID specific read data responses, if PE_ST.RAZ is set to `0b0`, as defined by the `FC_ERR_RESP_DEF` or `FIREWALL_F0_CFG_GLOBAL_SSE710_ERR_RESP_PER_MST_ID{x}_VAL` at design time, where x is the StreamID value.
    - Mark the response as being generated by the Firewall using the RUSER[0] or BUSER[0] bits.
- When a firewall component, which supports Monitor Extension LEVEL 2 (ME.2), detects an error transaction it:
    - Generate StreamID specific read data responses, if ME_ST.RDUM is set to `0b0`, as defined by the `FC_ERR_RESP_DEF` or `FIREWALL_F0_CFG_GLOBAL_SSE710_ERR_RESP_PER_MST_ID{x}_VAL` at design time, where x is the StreamID value.
    - Mark the response as being generated by the Firewall using the RUSER[0] or BUSER[0] bits.
- The firewall controller only supports configuration accesses with the following properties, otherwise a Configuration Access Error is generated:
    - 32-bit word-aligned access only.
    - Memory type must be either Device-nRnGnE or Device-nRnGE. See Table 10-4: Firewall memory attribute to AMBA AXI5 AxCACHE mapping on page 166
- The Firewall Controller does not support any exclusive accesses and treats the access as normal.

- The Firewall Controller responds to a configuration access, which generates a Configuration Access Error by:
  - Generating an AMBA AXI5 SLVERR, if FW_ST.ERR is set to `0b1`.
  - Generate StreamID specific read data responses, if FW_ST.RAZ is set to `0b0`, as defined by the `FC_ERR_RESP_DEF` or `FIREWALL_F0_CFG_GLOBAL_SSE710_ERR_RESP_PER_MST_ID{x}_VAL` at design time, where x is the StreamID value.
- Reset value of the PE_CTRL.FLT_CFG is `0b10`.
- Reset value of the PE_CTRL.RAZ is `0b0`.
- Reset value of the PE_CTRL.ERR is `0b1`.
- All firewall components, in the SSE-710 use the AMBA AXI5 bus protocol.
- The firewall:
  - Ignores the transient property.
  - Treats any memory which is not Write-Back cacheable at both inner and outer as Non-cacheable.

The following table shows a summary of the mapping between the firewall, RGN_TCFG2.MA field, and the AMBA AXI5 memory types. The table also shows how the firewall handles the AxLOCK signal on the incoming transaction.

Table 10-4: Firewall memory attribute to AMBA AXI5 AxCACHE mapping

Firewall memory attribute	AMBA AXI5 AxCACHE	AxLOCK	Note
Device-nGnRnE	Device Non-bufferable	As incoming transaction	-
Device-nGnRE Device-nGRE Device-GRE	Device Bufferable		
Normal-iNC-o{NC,WT,WB} Normal-iWT-o{NC,WT,WB} Normal-iWB-o{NC,WT}	Normal Non-Cacheable Bufferable		

Firewall memory attribute	AMBA AXI5 AxCACHE	AxLOCK	Note
Normal-iWB-oWB	Write-Back No Allocate Write-Back Read-Allocate Write-Back Write-Allocate Write-Back Read and Write-Allocate	Set to 0	The value that is selected depends on the read and write allocation policy that is defined in the RGN_TCFG2.MA field for the inner domain.

Where:

- i: Inner
- o: Outer
- NC: Non-cacheable
- WT: Write-Through cacheable
- WB: Write-Back cacheable

The following table shows a summary of the mapping between the AMBA AXI5 memory types and the firewall memory attribute types.

Table 10-5: AMBA AXI5 AxCACHE to Firewall memory attribute mapping

AMBA AXI5 AxCACHE	Firewall memory attribute	Notes
Device Non-bufferable	Device-nGnRnE	-
Device Bufferable	Device-nGnRE	
Normal Non-Cacheable Bufferable Normal Non-Cacheable Non-Bufferable Write-Through No Allocate Write-Through Read-Allocate Write-Through Write-Allocate Write-Through Read and Write-Allocate	Normal-iNC-oNC	

AMBA AXI5 AxCACHE	Firewall memory attribute	Notes
Write-Back No Allocate Write-Back Read-Allocate Write-Back Write-Allocate Write-Back Read and Write-Allocate	Normal-iWB-oWB	The read and write allocation policy, for inner and outer, is set to the same value as the AXI5 AxCACHE fields.

For more information about the memory types, see *Arm® Architecture Reference Manual Armv8, for Armv8-A architecture profile*.

10.3.4 Firewall read response value

This section describes how to configure the firewall to set the read data response to a value based on the StreamID.

The firewall can be configured to set the read data response to a value based on the StreamID when one of the following has occurred:

- A transaction is terminated by the protection logic of the firewall.
- A bus error is detected by the monitoring logic of the firewall.
- A Configuration Access Error has been generated.

The read data response value that is returned is a design time configuration, set using the following parameters:

- `FC_ERR_RESP_DEF`
- `FIREWALL_F0_CFG_GLOBAL_SSE710_FC_MST_ID{x}_VAL`
- `FIREWALL_F0_CFG_GLOBAL_SSE710_ERR_RESP_PER_MST_ID{x}_VAL`

The firewall determines which value is returned by performing the following:

- If SINGLE_MST is set to 1, then return FC_ERR_RESP_DEF.
- Looking up the MasterID from the StreamID of the transaction against the values of `FIREWALL_F0_CFG_GLOBAL_SSE710_FC_MST_ID{x}_VAL` configuration option.
  - If the value is found, then the corresponding value in `FIREWALL_F0_CFG_GLOBAL_SSE710_ERR_RESP_PER_MST_ID{x}_VAL` is returned.

    Note

    `FIREWALL_F0_CFG_GLOBAL_SSE710_FC_MST_ID{x}_VAL` and `FIREWALL_F0_CFG_GLOBAL_SSE710_ERR_RESP_PER_MST_ID{x}_VAL` are matched based on the value of x.

  - If the value is not found, then the value that is defined by `FC_ERR_RESP_DEF` is returned.

10.3.5 Host System firewall

The Host System firewall supports multiple configurations.

The following configurations are supported:

- Lockdown Extension level 2 (LE.2)
- Save and Restore Extension level 1(SRE.1)
- Security Extension level 1(SE.2)


Note

SSE-710 contains two firewalls, in the Secure Enclave and the Host System. This section applies only to the Host System firewall. For more information, see 9.1.6.10 Secure Enclave Firewall on page 156.

The Host System firewall is distributed across multiple clock and power domains, with the Firewall Controller located in AONTOP.

All transactions issued on the slave interfaces of the interconnect, must pass through at least one firewall component of the Host System Firewall before being issued on the master interfaces of the interconnect. The location of the firewall components within SSE-710 is shown in the following figure.

Figure 10-4: Location of firewall components

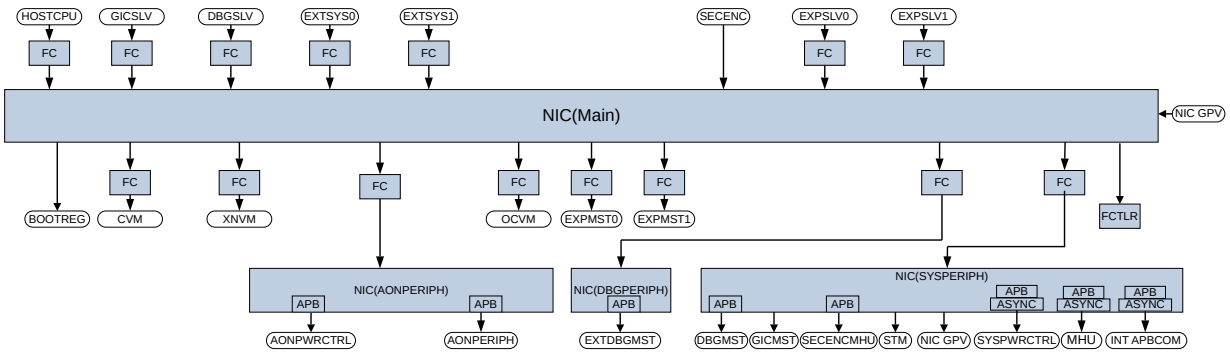

The following table shows the location of the firewall components and the mapping to the firewall component IDs. The following table shows location names.

Table 10-6: Host System firewall component locations

Firewall component	Location	Firewall Component ID
Firewall Controller (FCTLR)	Firewall	0
SYSPERIPH	Between NIC(Main) and NIC(SYSPERIPH)	1
DBGPERIPH	Between NIC(Main) and NIC(DBGPERIPH)	2
AONPERIPH	Between NIC(Main) and NIC(AONPERIPH)	3

Firewall component	Location	Firewall Component ID
XNVM	XNVM	4
CVM	CVM	5
HOSTCPU	HOSTCPU	6
EXTSYS0	EXTSYS0	7
EXTSYS1	EXTSYS1	8
EXPSLV0	EXPSLV0	9
EXPSLV1	EXPSLV1	10
EXPMST0	EXPMST0	11
EXPMST1	EXPMST1	12
OCVM	OCVM	13 (if OCVM_EN = 1)
Debug	DBGSLV	13 (if OCVM_EN =) 14 (if OCVM_EN = 1)

SSE-710 allows an implementation to configure certain aspects of the firewall components of the Host System firewall.

The following table shows the configuration values of the firewall components. Where the value is fixed, an explicit value is provided. Where the value is configurable, a range of values is provided.

Table 10-7: Host System firewall component configuration

Firewall component	PE_LV	ME_LVL	TE_LVL	RSE_LVL	NUM_RGN	MNRS	MXRS	NUM_MPE	SINGLE_MST
FCTLR	1	0	0	1	3	7	21	1	0
SYSPERIPH	1	0	0	0	22	7	29	1	0
DBGPERIPH	1	0	0	0	1	7	29	1	0
AONPERIPH	1	0	0	0	40	7	23	1	0
XNVM	2	0	0	0/1	16/32/48/64	7	27	4	0
CVM	2	0	0	0/1	16/32/48/64	3	25	4	0
HOSTCPU	0	2	0	0	NA	NA	NA	NA	1
DBG	2	2	2	1	4/8	7	32	2	0
EXTSYS{0-1}	2	2	2	1	8/16	7	32	1	1
EXPSLV{0-1}	2	2	2	1	8/16/32	7	32	4	0
EXPMST{0-1}*	2	0	0	0/1	8/16/32	3	29	4	0
	1	0	0	0	1-64	7	29	4	0
OCVM	2	0	0	0/1	16/32/64	7	31	4	0

* The allowed values depend on the level of Protection Extension implemented. EXMPST0 can be configured to support either PE.1 or PE.2. EXPMST1 only supports PE.2

The following configurable values always apply for all firewall components:

- MST_ID_WIDTH is 8
- SEC_SPT is `0b1`

- MA_SPT is `0b1`
- SH_SPT is `0b0`
- INST_SPT is `0b1`
- PRIV_SPT is `0b1`

1 fault entries and 1 error detection report are implemented by each Firewall .

For firewall components that implement PE.1, the following regions are implemented:

- Firewall controller:
  - As defined in C. Firewall on page 346
  - Configuration Master is set to the StreamID of the Secure Enclave.
- SYSPERIPH:
  - 1 region for each Host External System MHU
  - 1 region for CLUSTOP PPU
  - 1 region for each CORE{0-3} PPU
  - 1 region for INT APBCOM
  - 1 region for each MHU in the Secure Enclave
  - 1 region for STM Extended Stimulus port
  - 1 region for NIC GPV
  - 1 region for the GIC
  - 1 region Host Debug address space
- DBGPERIPH:
  - 1 region External Host Debug address space
- AONPERIPH:
  - 1 region for FW RAM PPU
  - 1 region for SYSTOP PPU
  - 1 region for DBGTOP PPU
  - 1 region for Interrupt Router
  - 1 region for each S32K CNTBase{0-1}
  - 1 region for S32K CNTRead
  - 1 region for S32K CNTControl
  - 1 region for S32K CNTCTL
  - 1 region for each **REFCLK** CNTBase{0-3}
  - 1 region for **REFCLK** CNTRead
  - 1 region for **REFCLK** CNTControl
  - 1 region for **REFCLK** CNTCTL

  - 1 region for Non-secure WDOG Refresh
  - 1 region for Non-secure WDOG Control
  - 1 region for Secure WDOG Refresh
  - 1 region for Secure WDOG Control
  - 1 region for Host Base System Control
  - 1 region for System ID
  - 1 region for each UART{0-1}
  - 16 regions allocated to the AON Expansion. The size and base address of these regions is **IMPLEMENTATION DEFINED**, but must be allocated exclusively within the AON Expansion address space.
- EXPMST0:
  - **IMPLEMENTATION DEFINED** when the EXPMST0 firewall component is configured with a PE_LVL of 1.

The **SOCCFG** interface signals, drive the Protection Size interface of the following firewall components:

- **CVMSIZE** is connected to the CVM firewall component.
  - Legal values are: 256KB, 512KB, 1MB, 2MB, 4MB, 8MB, 16MB, 32MB.
  - Any values less than 256KB are treated as if 256KB was set.
  - Any values greater than 32MB are treated as if 32MB was set.
- **XNVMSIZE** is connected to the XNVM firewall component.
  - Legal values are: 4MB, 8MB, 16MB, 32MB, 64MB, 128MB.
  - Any values less than 4MB are treated as if 4MB was set.
  - Any values greater than 128MB are treated as if 128MB was set.
- **OCVMSIZE** is connected to the OCVM firewall component.
  - Legal values are: 0B, 2MB, 4MB, 8MB, 16MB, 32MB, 64MB, 128MB, 256MB, 512MB, 1GB, 2GB.
  - Any values less than 2MB, other than 0B, are treated as if 2MB was set.
  - Any values greater than 2GB are treated as if 2GB was set.


Note

See C. Firewall on page 346 for more information on the Protection Size interface.

For all other firewall components with the Protection Size interface, the interface is tied off internal to SSE-710 to match the address width of the firewall component.

10.3.6 Host System firewall regions

A region defines an area of memory where certain masters can perform memory operations.

AONPERIPH, DBGPERIPH, and SYSPERIPH firewall components have pre-defined regions.

10.3.6.1 AONPERIPH firewall component regions

The following table shows the AONPERIPH firewall component regions.

Table 10-8: AONPERIPH Firewall Component regions

Region number	Peripheral	Field	Value
0	System ID	Base Address	0x1A00_000
		Size	0x0C
		MULnPO2	0b0
1	Host Base System Control	Base Address	0x1A01_0000
		Size	0x0C
		MULnPO2	0b0
2	FW PPU	Base Address	0x1A02_0000
		Size	0x0C
		MULnPO2	0b0
3	SYSTOP PPU	Base Address	0x1A03_0000
		Size	0x0C
		MULnPO2	0b0
4	DBGTOP PPU	Base Address	0x1A04_0000
		Size	0x0C
		MULnPO2	0b0
5	REFCLK CNTControl	Base Address	0x1A20_0000
		Size	0x0C
		MULnPO2	0b0
6	REFCLK CNTRead	Base Address	0x1A21_0000
		Size	0x0C
		MULnPO2	0b0
7	REFCLK CNTCTL	Base Address	0x1A22_0000
		Size	0x0C
		MULnPO2	0b0

Region number	Peripheral	Field	Value
8	REFCLK CNTBase0	Base Address	0x1A23_0000
		Size	0x0C
		MULnPO2	0b0
9	REFCLK CNTBase1	Base Address	0x1A24_0000
		Size	0x0C
		MULnPO2	0b0
10	REFCLK CNTBase2	Base Address	0x1A25_0000
		Size	0x0C
		MULnPO2	0b0
11	REFCLK CNTBase3	Base Address	0x1A26_0000
		Size	0x0C
		MULnPO2	0b0
12	NS WDOG CTRL	Base Address	0x1A30_0000
		Size	0x0C
		MULnPO2	0b0
13	NS WDOG Refresh	Base Address	0x1A31_0000
		Size	0x0C
		MULnPO2	0b0
14	S WDOG CTRL	Base Address	0x1A32_0000
		Size	0x0C
		MULnPO2	0b0
15	S WDOG Refresh	Base Address	0x1A33_0000
		Size	0x0C
		MULnPO2	0b0
16	S32K CNTControl	Base Address	0x1A40_0000
		Size	0x0C
		MULnPO2	0b0
17	S32K CNTRead	Base Address	0x1A41_0000
		Size	0x0C
		MULnPO2	0b0
18	S32K CNTCTL	Base Address	0x1A42_0000
		Size	0x0C

Region number	Peripheral	Field	Value
		MULnPO2	0b0
19	S32K CNTBase0	Base Address	0x1A43_0000
		Size	0x0C
		MULnPO2	0b0
20	S32K CNTBase1	Base Address	0x1A44_0000
		Size	0x0C
		MULnPO2	0b0
21	Interrupt Router	Base Address	0x1A50_0000
		Size	0x0C
		MULnPO2	0b0
22	UART 0	Base Address	0x1A51_0000
		Size	0x0C
		MULnPO2	0b0
23	UART 1	Base Address	0x1A52_0000
		Size	0x0C
		MULnPO2	0b0
24-39	AON Expansion 0-15	Base Address	0x1A60_0000 - 0x1A6F_FFFF, defined by FIREWALL_F0_CFG_SSE710_AONPERIPH_FC_RGN{x}_BASE_ADDR
		Size	0x00, 0x0C-0x14, defined by FIREWALL_F0_CFG_SSE710_AONPERIPH_FC_RGN{x}_SIZE
		MULnPO2	0b0

10.3.6.2 SYSPERIPH firewall component regions

The regions of the SYSPERIPH firewall component vary depending on the SSE-710 configuration.

The following table shows the SYSPERIPH Firewall Component regions.

Table 10-9: SYSPERIPH Firewall Component regions

Region number	Peripheral	Field	Value
0	Host System Debug	Base Address	0x1000_0000
		Size	0x1B
		MULnPO2	0b0
1	HSE MHU0	Base Address	0x1B80_0000
		Size	0x0C

Region number	Peripheral	Field	Value
		MULnPO2	0b0
2	SEH MHU0	Base Address	0x1B81_0000
		Size	0x0C
		MULnPO2	0b0
3	HSE MHU1	Base Address	0x1B82_0000
		Size	0x0C
		MULnPO2	0b0
4	SEH MHU1	Base Address	0x1B83_0000
		Size	0x0C
		MULnPO2	0b0
5	CoreSight™ SDC-600	Base Address	0x1B90_0000
		Size	0x0C
		MULnPO2	0b0
		Size	0x0C
		MULnPO2	0b0
6	CLUSTOP PPU	Base Address	0x1BC0_0000
		Size	0x0C
		MULnPO2	0b0
7	CORE 0 PPU	Base Address	0x1BC1_0000
		Size	0x0C
		MULnPO2	0b0
8	CORE 1 PPU	Base Address	0x1BC2_0000, only implemented when HOST_CPU_NUM_CORES > 1
		Size	0x0C
		MULnPO2	0b0
9	CORE 2 PPU	Base Address	0x1BC3_0000, only implemented when HOST_CPU_NUM_CORES > 2
		Size	0x0C
		MULnPO2	0b0
10	CORE 3 PPU	Base Address	0x1BC4_0000, only implemented when HOST_CPU_NUM_CORES > 3
		Size	0x0C
		MULnPO2	0b0
7 + HOST_CPU_NUM_CORES	HES0 MHU0	Base Address	0x1B00_0000
		Size	0x0C
		MULnPO2	0b0

Region number	Peripheral	Field	Value
8 + HOST_CPU_NUM_CORES	ES0H MHU0	Base Address	0x1B01_0000
		Size	0x0C
		MULnPO2	0b0
9 + HOST_CPU_NUM_CORES	HES0 MHU1	Base Address	0x1B02_0000, only implemented when EXT_SYS0_TZ_SPT = 1
		Size	0x0C
		MULnPO2	0b0
10 + HOST_CPU_NUM_CORES	ES0H MHU1	Base Address	0x1B03_0000 ,only implemented when EXT_SYS0_TZ_SPT = 1
		Size	0x0C
		MULnPO2	0b0
9 + HOST_CPU_NUM_CORES + EXT_SYS0_TZ_SPT * 2	HES1 MHU0	Base Address	0x1B04_0000
		Size	0x0C
		MULnPO2	0b0
10 + HOST_CPU_NUM_CORES + EXT_SYS0_TZ_SPT * 2	ES1H MHU0	Base Address	0x1B05_0000
		Size	0x0C
		MULnPO2	0b0
11 + HOST_CPU_NUM_CORES + EXT_SYS0_TZ_SPT * 2	HES1 MHU1	Base Address	0x1B06_0000, only implemented when EXT_SYS1_TZ_SPT = 1
		Size	0x0C
		MULnPO2	0b0
12 + HOST_CPU_NUM_CORES + EXT_SYS0_TZ_SPT * 2	ES1H MHU1	Base Address	0x1B07_0000, only implemented when EXT_SYS1_TZ_SPT = 1
		Size	0x0C
		MULnPO2	0b0
7 + HOST_CPU_NUM_CORES + 2 * (EXT_SYS0_TZ_SPT + EXT_SYS1_TZ_SPT) + 4	CoreLink™ GIC-400 Distributor	Base Address	0x1C00_0000
		Size	0x13
		MULnPO2	0b0
7 + HOST_CPU_NUM_CORES + 2 * (EXT_SYS0_TZ_SPT + EXT_SYS1_TZ_SPT) +	CoreSight STM-500 Extend Stimulus Port	Base Address	0x1D00_0000
		Size	0x18
		MULnPO2	0b0
7 + HOST_CPU_NUM_CORES + 2 * (EXT_SYS0_TZ_SPT + EXT_SYS1_TZ_SPT) + 6	NIC(Main) GPV	Base Address	0x1E00_000
		Size	0x14

10.3.6.3 DBGPERIPH firewall component regions

This section details the firewall component region of External debug bus.

Table 10-10: DBGPERIPH Firewall Component regions

Region Number	Peripheral	Field	Value
0	External Debug Bus	Base Address	0x1800_0000
		Size	0x19
		MULnPO2	0b0

10.3.6.4 EXPMST{0,1} firewall component regions

EXPMST{0,1}_PE_LVL is configurable, if it is 2 then regions are software configurable, if it is 1 then the regions are user configurable with parameter.

The EXPMST0 firewall components are described in the following table.

Table 10-11: EXPMST{0-1} Firewall Component Regions

Region number	Peripheral	Field	Value
The number of regions implemented is set by the EXPMST0_NUM_RGN configuration option.	The peripheral or peripherals protected by the region is **IMPLEMENTATION DEFINED**.	Base Address Defined by FIREWALL_F0_CFG_SSE710_EXPMST0_FC_RGN{x}_BASE_ADDR	The Base address of the region must be within the address range defined for that firewall component. For EXPMST0 firewall component: • The Base and Upper Address of a region must be within 0x4000_0000 – 0x5FFF_FFFF. • The size of the region must be one of the following: 0x0 – region is disabled; 0x0C to 0x1C.
		Size Defined by FIREWALL_F0_CFG_SSE710_EXPMST0_FC_RGN{x}_SIZE	
The number of regions that are implemented is set by the EXPMST1_NUM_RGN configuration option.	The peripheral or peripherals protected by the region is **IMPLEMENTATION DEFINED**.	Base address Defined by FIREWALL_F0_CFG_SSE710_EXPMST1_FC_RGN{x}_BASE_ADDR	The Base address of the region must be within the address range defined for that firewall component. For EXPMST1 firewall component: • The Base and Upper Address of a region must be within 0x6000_0000 – 0x7FFF_FFFF. • The size of the region must be one of the following: ◦ 0x0 – Region is disabled ◦ 0x0C to 0x1
		Size Defined by FIREWALL_F0_CFG_SSE710_EXPMST1_FC_RGN{x}_SIZE	

10.4 StreamID and CPUID

Each master or group of masters has a unique StreamID, which is used to identify the transactions issued by that master.

SSE-710 implements a firewall, as defined in C. Firewall on page 346. SSE-710 assigns a StreamID value to all the internal masters and provides inputs and outputs on the CVM, XNVM, OCVM, HOSTEXPSLV{0-1}, and HOSTEXPMST{0-1} interfaces.

The width of the inputs and outputs is 8.


Note

The interfaces support the AMBA AXI5 Untranslated_Transaction property including the **AxMMUSID** signals as part of the AR and AW channels. These signals are used to pass the StreamID associated with the transaction.

The following table shows the StreamIDs allocated to masters within SSE-710 and those available for use by the integrator.

Table 10-12: StreamID assignment

StreamID	Master
0	Secure Enclave
1	Host CPU
2	Reserved
3	Host ETR
4	Host AXI-AP
5	SoC ETR
6-15	Reserved
16	External System 0
17	External System 1
18-31	Reserved
32-255	Expansion

The StreamID is passed through SSE-710, unaltered from the slave to master interface, using the **AxMMUSID** signal.

For certain interfaces, as well as the StreamID, there is also a CPUID transported over ARUSER[1:0] and AWUSER[1:0], which identifies which Host CPU core generated the transaction. For information on the interfaces, see 4.1 Interfaces overview on page 36.

The Host CPU core number is binary encoded on the user signals. For transactions generated by another master, those signals are set to one of the following values:

- `0b00` by the Secure Enclave, EXTSYS{0-1}
- Any two-bit value given by HOSTEXPSLV{0-1} through the user signals of the Host Expansion slave Interface

10.5 Reserved address space and error responses

The reserved address space can be in the Host System address space or in the expansion regions.

The reserved address space in an SSE-710 design can be:

- Address space marked as Reserved in Host System address space, defined in 12.1.1 Host System memory map on page 188.

  When any read access (either instruction fetch or data access) is issued to the reserved address space defined in Host System address space, the error response is returned along with a read data from the firewall.

  For internal masters of SSE-710, the read data value is defined as the following table shows. For masters added to SSE-710, for example the External Systems and masters connecting to expansion slave ports, the read data value is defined by the integrator.
- Any unused address space within expansion regions, for example the unused regions of the volatile, non-volatile and off-chip volatile memory regions.

  When any read access is issued to the reserved address space within expansion regions, the response is determined by an integrator. Arm® has defined rules for reserved address space the integrator should obey when memory and peripherals are integrated into SSE-710.

The read data response is generated by the Host System firewall and is configured using the

`HOST_FC_ERR_RESP_DEF` or `FIREWALL_F0_CFG_GLOBAL_SSE710_ERR_RESP_PER_MST_ID{x}_VAL`

configuration option. SSE-710 defines the values of the error response read data for masters within SSE-710, including the Host CPU. For masters added to SSE-710, the integrator must set the appropriate value for each master.

The following table shows the value for the error response read data for masters which are part of SSE-710.

Table 10-13: Error response read data for internal masters of SSE-710

Master	Read data value
Secure Enclave	`0xDEAD_DEAD`
Host CPU	`0xEC00EC00`
Host ETR	`0x0000_0000`
Host AXI-AP	`0x0000_0000`
SoC ETR	`0x0000_0000`

11. Host System peripherals

This chapter describes the peripherals, which are part of the SSE-710 Host System.

11.1 Counters and timers

SSE-710 has two time-domains, **REFCLK** and **S32K**.

Both time domains are based on Generic Time, as defined by the *Arm® Architecture Reference Manual Armv8, for Armv8-A architecture profile*.

REFCLK time domain

The **REFCLK** time domain increments by 1 on each **REFCLK** cycle.

The **REFCLK** time domain includes:

- A memory-mapped counter, **REFCLK** counter. For more information on the **REFCLK** counter register, see 12.3.4 REFCLK Counter CNTControl register summary on page 264
- Four memory-mapped timers, **REFCLK** Timer {0-3} as defined by *Arm® Architecture Reference Manual Armv8, for Armv8-A architecture profile*
- Each Host CPU core implements an Arm® Generic Timer, as defined by *Arm® Architecture Reference Manual Armv8, for Armv8-A architecture profile*
- Two Generic Watchdogs, Secure and Non-secure Watchdogs, as defined by the *Arm® Server Base System Architecture 5.0*

The **REFCLK** time domain can be halted during debug, using the Counter CTI.

The **REFCLK** time domain only runs when SSE-710 is in BSYS.RUN or BSYS.SLEEP0 power states. After exiting BSYS.SLEEP1 or BSYS.OFF, software is responsible to restore the counter.

The **REFCLK** time domain is distributed using a 64-bit gray-encoded timestamp value. SSE-710 provides the following interfaces for the **REFCLK** time domain:

- **HOSTCNTVALUEG**: Output from the **REFCLK** counter (gray-encoded timestamp value).
- **HOSTCNTVALUEB**: Input for the **REFCLK** Timers, Secure and Non-secure Watchdogs (binary-encoded timestamp value).

S32K time domain

The **S32K** time domain increments by 1 every S32KCLK cycle and includes:

- A memory-mapped counter, S32K counter, as defined by the *Arm® Architecture Reference Manual Armv8, for Armv8-A architecture profile*
- Two memory-mapped timers, S32K Timer 0 and 1, as defined by the *Arm® Architecture Reference Manual Armv8, for Armv8-A architecture profile*

The **S32K** time domain can be halted during debug using the Counter CTI.

The **S32K** time domain runs in all SSE-710 power states, except for BSYS.OFF. It is software's responsibility for programming the S32K counter when the SSE-710 exits the BSYS.OFF power state.

The S32K time domain is distributed using a 64-bit gray-encoded timestamp value. SSE-710 provides the following interfaces:

- **HOSTS32KCNTVALUEG**: Output from the S32K counter (gray-encoded timestamp value)
- **HOSTS32KCNTVALUEB**: Input for the S32K Timer 0 and 1 (binary-encoded timestamp value)

11.2 Watchdog

The Host System includes two Generic Watchdogs: Secure and Non-secure.

See *Arm® Server Base System Architecture* for more information about the Generic Watchdog.

Each Watchdog has a control and refresh frame that is located in the Host System memory map. Both watchdogs use the **REFCLK** timestamp to generate the interrupt events. The two watchdog signals, **WS0** and **WS1**, of the Secure and Non-secure watchdogs are routed as interrupts to:

- Non-secure watchdog **WS0** and **WS1** are both routed to the GIC as interrupts 64 and 33 respectively.
- Secure watchdog **WS0** is routed to the GIC as interrupt 32.
- Secure watchdog **WS1** is routed to the Secure Enclave as interrupt 9.

11.3 Host Base System Control

The Host Base System Control registers allow software to configure features of the Host System:

- Host CPU static configurations
- Host System clock control
- Host CPU power control
- Host CPU boot mask
- Reset Syndrome status
- External System Reset control and status
- Power Control of the Base System

For more information on the Host Base System Control registers, see 12.3.1 Host Base System Control register summary on page 204.

11.4 Interrupt Router

SSE-710 includes an Interrupt Router, which routes the interrupts signals of the EXPSHDINT interface and some internal SSE-710 interrupts between the:

- GICSHDINT interface
- Secure Enclave Cortex®-M0+NVIC
- EXTSYS{0-1}SHDINT interfaces

The Interrupt Router is configured with:

- NUM_ICI between four ICI interfaces:
  - ICI0: Secure Enclave
  - ICI1: Host System GIC
  - ICI2: External System 0
  - ICI3: External System 1
- NUM_SHD_INT is set to 32 plus NUM_EXP_SHD_INT.
- SI{x}_ICI_MSK and SI{x}_DEF_ICI are as defined in the table below.
- Level 2 of the Lockdown Extension
- Secure Enclave is the Security Monitor for the Interrupt Router in SSE-710.

The following table shows the allowed routing for Interrupt Router:

- "Y" indicates that the interrupt can be routed to that ICI{x}
- "N" indicates it cannot be routed to that ICI{x}

Table 11-1: Interrupt Router interface assignment

Interrupt Source	SII	ICI0 (Secure Enclave)	ICI1 (Host System)	ICI{2-3} (External System)
Host System Firewall Interrupt	0	Y[a]	Y	N
SDC-600	1	Y[a]	Y	N
Host PPU Combined	2	Y	Y[a]	N
REFCLK Timer 0	3	Y	Y[a]	Y
REFCLK Timer 1	4	Y	Y[a]	Y
REFCLK Timer 2	5	Y	Y[a]	Y
REFCLK Timer 3	6	Y	Y[a]	Y
S32K Timer 0	7	Y	Y[a]	Y
S32K Timer 1	8	Y	Y[a]	Y

[a] Default routing for this interrupt.

Interrupt Source	SII	ICI0 (Secure Enclave)	ICI1 (Host System)	ICI{2-3} (External System)
SoC ETR	9	N	Y	Y
SoC CATU	10	N	Y	Y
Reserved	11-31	N	N	N
EXPSHDINT[n][b]	32+n	[c] [d]	[d]	

Note

For **EXPSHDINT[n]**, n is between 0 and NUM_EXP_SHD_INT.

IMPLEMENTATION DEFINED behavior

This section describes the **IMPLEMENTATION DEFINED** behaviors that are visible or impact on the SSE-710.

Note

This section must be read with the C. Firewall on page 346.

- Both level and edge-based shared interrupts are supported.
- Lockdown interface is implemented.
- Tamper Interrupt interface is implemented. Tamper Interrupt is reported to Secure Enclave.

Note

While the Interrupt Router supports both level and edge-based interrupts, the destination of that interrupt must be considered. For example, the Secure Enclave interrupts only support level-based interrupts.

B. Interrupt Router on page 340 defines the generic behavior of the Interrupt Router.

11.5 MHU

SSE-710 provides *Message Handling Units* (MHUs).

The MHUs are compliant to the MHUv2.1 specification, defined in A. Message Handling Unit on page 311.

[b] When interrupt source is not implemented, this bit of the SII is tied 0 and the SI{x}_ICI_DST is set to 0.
[c] Only the first 32 **EXPSHDINT[n]** can be routed to the Secure Enclave.
[d] Dependent on the value of the SI{x}_ICI_DST and SI{x}_DEF_ICI.

An MHU has two halves:

- A Sender frame, or Sender
- A Receiver frame, or Receiver

The MHU facilitates communication between two systems in SSE-710, with one system having the Sender and the other having the Receiver of a single MHU.

MHUs are implemented in pairs to allow for full-duplex communications, with Sender and Receiver being reversed between the two MHUs in the pair. For example, between the Host System and Secure Enclave: one of the MHUs has the Sender in the Host System and the Receiver in the Secure Enclave, while the other MHU has the Sender in the Secure Enclave and the Receiver in the Host System.

SSE-710 subsystem supports Arm® TrustZone® for each External System, two pairs of MHUs are implemented. This enables software to use one pair for Secure communication and the other for Non-secure communication.


Note

In this document, the MHUs are either referred to by the MHU Name or Short Name, defined in the table below. When referring to either the Sender or Receiver frame, the Short Name is followed by either Sender or Receiver.

The MHU is designed to be split across clock, reset, and power boundaries. The software is responsible for requesting the Receiver portion of the MHU is powered. The software does this by using the Ready to Send protocol defined in A. Message Handling Unit on page 311. Software is also responsible for making sure the Sender frame remains powered until the message has been received by the Receiver. How software knows when the Receiver has received the message, depends on the Transport Protocol being used. For more information on Transport Protocols see A.3 Transport protocols on page 334.

The following table shows the SSE-710 MHUs.

Table 11-2: MHUs

MHU name	Short name	Sender system	Receiver system	Notes
Host to Secure Enclave MHU0	HSE0	Host	Secure Enclave	
Secure Enclave to Host MHU0	SEH0	Secure Enclave	Host	
Host to Secure Enclave MHU1	HSE1	Host	Secure Enclave	
Secure Enclave to Host MHU1	SEH1	Secure Enclave	Host	
Secure Enclave to External System 0 MHU 0	SEES00	Secure Enclave	External System 0	
External System 0 to Secure Enclave MHU 0	ES0SE0	External System 0	Secure Enclave	
Secure Enclave to External System 0 MHU 1	SEES01	Secure Enclave	External System 0	Implemented when EXT_SYS0_TZ_SPT is 1

MHU name	Short name	Sender system	Receiver system	Notes
External System 0 to Secure Enclave MHU 1	ES0SE1	External System 0	Secure Enclave	
Secure Enclave to External System 1 MHU 0	SEES10	Secure Enclave	External System 1	
External System 1 to Secure Enclave MHU 0	ES1SE0	External System 1	Secure Enclave	
Secure Enclave to External System 1 MHU 1	SEES11	Secure Enclave	External System 1	Implemented when EXT_SYS1_TZ_SPT is 1
External System 1 to SSecure Enclave MHU 1	ES1SE1	External System 1	Secure Enclave	
Host to External System 0 MHU 0	HES00	Host	External System 0	
External System 0 to Host MHU 0	ES0H0	External System 0	Host	
Host to External System 0 MHU 1	HES01	Host	External System 0	Implemented when EXT_SYS1_TZ_SPT is 1
External System 0 to Host MHU 1	ES0H1	External System 0	Host	
Host to External System 1 MHU 0	HES10	Host	External System 1	
External System 1 to Host MHU 0	ES1H0	External System 1	Host	
Host to External System 1 MHU 1	HES11	Host	External System 1	Implemented when EXT_SYS1_TZ_SPT is 1
External System 1 to Host MHU 1	ES1H1	External System 1	Host	

Using the MHU_{x}_NUM_CH parameter (where x is the MHU short name), all MHUs in SSE-710, can be configured with 1 to 32 channels.

11.6 Boot Register

The Boot Register is a write-once set of registers providing the initial instructions to the Host CPU.

The Boot Register is in the AONTOP power domain and uses **REFCLK**. This enables the Boot Register to retain its values in all power states.

The Boot Register is only accessible through:

- Secure read
- Secure writes from the Secure Enclave only

Any other accesses generate an error response.

The method by which the Boot Register identifies the master which issued the read or write operation, uses the StreamID. For more information, see 10.4 StreamID and CPUID on page 178.

For more information on the register assignment of the Boot Register, see 12.3.6 Boot register summary on page 273.

11.7 CoreSight SDC-600

To enable a debug agent to pass data to the target system,SSE-710 includes a CoreSight™ SDC-600 component. An example would be the passing of a debug authentication certificate.

The CoreSight™ SDC-600 in SSE-710 has two parts:

- An External APBCOM, compliant with the ADIv6 specification, in the AONTOP power domain and in the **REFCLK** domain. It is connected to the External Debug Bus.
- An Internal APBCOM in the SYSTOP power domain and in the **ACLK** domain, connected to the NIC. The interrupt is routed, via the Interrupt Router, to the Host System GIC or the NVIC of the Secure Enclave.

- In this manual, the External and Internal APBCOM are referred to as EXT APBCOM and INT APBCOM respectively.
- SSE-710 does not support the usage of the REMRR to generate a reset request and the debug agent must use the **nSRST** or the **CSYSRSTREQ** of the DP ROM.

For more information on the CoreSight™ SDC-600, see the*Arm® CoreSight™ SDC-600 Secure Debug Channel Technical Reference Manual*.

11.8 UART

SSE-710 includes two PL011 UARTs in the AONTOP domain, Host UART{0,1}.

The UART supports the hardware-based follow control. Host UART{0,1} uses the HOSTUARTCLK for both the PCLK and UARTCLK inputs of PL011. The Host System only uses the combined interrupt from the UART.

For more information on the UARTs, see the *PrimeCell UART (PL011) Technical Reference Manual*.

SSE-710 only supports the use of the PL011 as an RS232-compliant UART.

12. Programmers model

This chapter describes the SSE-710 subsystem memory and interrupt maps, and provides detailed information for all programmable registers.

12.1 Memory map

This section describes the SSE-710 memory maps.

12.1.1 Host System memory map

Peripherals and devices in the Host System are allocated memory in units of 64KB pages. This allows systems to use any granular page size defined by the table page formats.

Any unoccupied regions of its allocation are Reserved. Reserved regions are treated as RAZ/WI and generate an error if accessed.

Some regions of memory are marked as Secure. Secure access can access these regions, while Non-Secure access is blocked and generates an error.

- Issuing master
- Security
- Privilege type
- Access type (Read/Write/Execute)

The Firewall applies additional constraints, and it cannot remove any constraints already specified. A secure access region cannot be made accessible to non-secure access attempts.

At the highest level, the Host System Memory Map is divided into a few key regions as the table below shows.


Note

In the following table the omission of the Security column, or an entry of "-", indicates that the region or peripheral is accessed from any security world.

The entire Host System memory map is protected by the Host System Firewall, see C. Firewall on page 346. It enables software to limit access to areas of the memory map, based on the following properties:

Table 12-1: Host System memory map

Offset	Size	Security	Region	Notes
0x00_0000_0000	4KB	Secure	Boot Register	12.3.6 Boot register summary on page 273
0x00_0000_1000	1020KB	-	Reserved	-

Offset	Size	Security	Region	Notes
0x00_0010_0000	15MB	-	Reserved	-
0x00_0100_0000	16MB	-	Reserved	-
0x00_0200_0000	32MB	-	Volatile Memory	4.4.1 On-chip Volatile Memory (CVM) interface on page 55
0x00_0400_0000	64MB	-	Reserved	-
0x00_0800_0000	128MB	-	eXecute-in-place Non-Volatile Memory	4.4.2 eXecute-in-place Non-volatile Memory (XNVM) interface on page 55
0x00_1000_0000	160MB	-	Debug	12.1.1.1 Debug region on page 189
0x00_1A00_0000	608MB	-	Host Peripherals	12.1.1.2 Host Peripheral Region on page 190
0x00_4000_0000	1GB	-	Host Master Expansion	12.1.1.4 Host Master Expansion Region on page 193
0x00_8000_0000	2GB	-	Off-chip Volatile Memory	4.4.3 Off-chip Volatile Memory (OCVM) interface on page 56
0x01_0000_0000	1020GB	-	Reserved	-

12.1.1.1 Debug region

The following table shows the Debug region of the Host System memory map.

Table 12-2: Debug region memory map

Offset	Size	Component	Notes
0x00_1000_0000	128MB	Host System Debug	See 12.1.1.1.1 Host System Debug on page 189
0x00_1800_0000	32MB	External Debug Bus	See 12.1.2 External Debug Bus memory map on page 193

The Host Debug area provides access to the debug infrastructure of the Host System. The External Debug Bus region provides access to all debug logic, for example, the debug infrastructure of an External System.

12.1.1.1.1 Host System Debug

The Host System Debug area provides access to the debug infrastructure of the Host System.

This memory is accessed via the:

- Host System processor and Secure Enclave, via the Host System memory map at 0x1000_0000.
- Any External Debug Bus master, using the Host APB AP in the External Debug Memory Map, for example, the DP. See 12.1.2 External Debug Bus memory map on page 193.

The following table shows the Host Debug memory map.

Table 12-3: Host Debug memory map

Offset	Size	Component	Notes
0x0000_0000	64KB	Host Debug ROM	-

Offset	Size	Component	Notes
0x0001_0000	576KB	Reserved	-
0x000A_0000	64KB	Host Funnel	-
0x000B_0000	384KB	Reserved	-
0x0011_0000	64KB	Host Replicator	-
0x0012_0000	64KB	Host ETR	-
0x0013_0000	64KB	Host CATU	-
0x0014_0000	64KB	Host CTI	-
0x0015_0000	128KB	Reserved	-
0x0017_0000	64KB	STM APB	-
0x0018_0000	512KB	Reserved	-
0x0020_0000	14MB	Reserved	-
0x0100_0000	16MB	Expansion	Maps to Host Debug APB expansion Interface (HOSTDBGEXP)
0x0200_0000	16MB	Host CPU Debug	Maps to Host CPU Debug APB Interface (HOSTCPUDBG)
0x0300_0000	16MB	Host CPU Debug Internal View	Maps to Host CPU Debug APB Interface (HOSTCPUDBG) with the **PADDR31** signal tied LOW
0x0400_0000	64MB	Reserved	-
0x0800_0000	896MB	Reserved	Not directly accessible in the Host System memory map
0x4000_0000	3GB	Reserved	Not directly accessible in the Host System memory map

12.1.1.2 Host Peripheral Region

The Host Peripheral Region gives access to all peripherals within the Host System.

This memory is accessed via the system interconnect by the:

- Host Processor
- Secure Enclave
- Host Slave Expansion interfaces (HOSTEXPSLV{0-1})
- External Systems

The following table shows the Host Peripheral Region peripherals and memory area.

Table 12-4: Host Peripheral Memory Map

Offset	Size	Security	Component	Notes
0x1A00_0000	64KB	-	System ID	See 12.3.6 Boot register summary on page 273
0x1A01_0000	64KB	-	Host Base System Control	See 12.1.1 Host System memory map on page 188
0x1A02_0000	64KB	Secure	Firewall PPU	See *Arm® CoreLink™ PCK-600 Power Control Kit Technical Reference Manual*
0x1A03_0000	64KB	Secure	SYSTOP PPU	See *Arm® CoreLink™ PCK-600 Power Control Kit Technical Reference Manual*
0x1A04_0000	64KB	Secure	DBGTOP PPU	See *Arm® CoreLink™ PCK-600 Power Control Kit Technical Reference Manual*

Offset	Size	Security	Component	Notes
0x1A05_0000	1728KB	-	Reserved	-
0x1A20_0000	64KB	Secure	REFCLK CNTControl	See 11.1 Counters and timers on page 181
0x1A21_0000	64KB	-	REFCLK CNTRead	
0x1A22_0000	64KB	-	REFCLK CNTCTL	
0x1A23_0000	64KB	-	REFCLK CNTBase0	
0x1A24_0000	64KB	-	REFCLK CNTBase1	
0x1A25_0000	64KB	-	REFCLK CNTBase2	
0x1A26_0000	64KB	-	REFCLK CNTBase3	
0x1A27_0000	640KB	-	Reserved	
0x1A30_0000	64KB	-	NS WDOG CTRL	
0x1A31_0000	64KB	-	NS WDOG Refresh	
0x1A32_0000	64KB	Secure	Secure WDOG CTRL	
0x1A33_0000	64KB	Secure	Secure WDOG Refresh	
0x1A34_0000	768KB	-	Reserved	
0x1A40_0000	64KB	Secure	S32K CNTControl	
0x1A41_0000	64KB	-	S32K CNTRead	
0x1A42_0000	64KB	-	S32K CNTCTL	
0x1A43_0000	64KB	-	S32K CNTBase0	
0x1A44_0000	64KB	-	S32K CNTBase1	
0x1A45_0000	704KB	-	Reserved	-
0x1A50_0000	64KB	-	Interrupt Router	See B. Interrupt Router on page 340
0x1A51_0000	64KB	-	UART0	See 11.8 UART on page 187, B. Interrupt Router on page 340
0x1A52_0000	64KB	-	UART1	See 11.8 UART on page 187, B. Interrupt Router on page 340
0x1A53_0000	832KB	-	Reserved	-
0x1A60_0000	1MB	-	AON Expansion	-
0x1A70_0000	1MB	-	Reserved	-
0x1A80_0000	2MB	-	Host System Firewall	See 10.3.5 Host System firewall on page 169
0x1AA0_0000	6MB	-	Reserved	-
0x1B00_0000	64KB	-	Host to External System 0 MHU0	See 11.5 MHU on page 184
0x1B01_0000	64KB	-	External System 0 to Host MHU0	
0x1B02_0000	64KB	-	Host to External System 0 MHU1	See 11.5 MHU on page 184
0x1B03_0000	64KB	-	External System 0 to Host MHU1	Reserved if EXT_SYS0_TZ_SPT = 0
0x1B04_0000	64KB	-	Host to External System 1 MHU0	See 11.5 MHU on page 184
0x1B05_0000	64KB	-	External System 1 to Host MHU0	
0x1B06_0000	64KB	-	Host to External System 1 MHU1	See 11.5 MHU on page 184
0x1B07_0000	64KB	-	External System 1 to Host MHU1	Reserved if EXT_SYS1_TZ_SPT = 0
0x1B08_0000	512KB	-	Reserved	-
0x1B10_0000	7MB	-	Reserved	-

Offset	Size	Security	Component	Notes
0x1B80_0000	64KB	-	Host to Secure Enclave MHU0	See 11.5 MHU on page 184
0x1B81_0000	64KB	-	Secure Enclave to Host MHU0	
0x1B82_0000	64KB	-	Host to Secure Enclave MHU1	
0x1B83_0000	64KB	-	Secure Enclave to Host MHU1	
0x1B84_0000	768KB	-	Reserved	-
0x1B90_0000	64KB	-	INT APBCOM	See 11.7 CoreSight SDC-600 on page 187
0x1B91_0000	64KB	-	Reserved	-
0x1B92_0000	64KB	-	Reserved	-
0x1B93_0000	832KB	-	Reserved	-
0x1BA0_0000	6MB	-	Reserved	-
0x1BC0_0000	64KB	Secure	CLUSTOP PPU	See *Arm® CoreLink™ PCK-600 Power Control Kit Technical Reference Manual*
0x1BC1_0000	64KB	Secure	CORE 0 PPU	See *Arm® CoreLink™ PCK-600 Power Control Kit Technical Reference Manual*
0x1BC2_0000	64KB	Secure	CORE 1 PPU	Reserved if HOST_NUM_CPU_CORE < 2 See *Arm® CoreLink™ PCK-600 Power Control Kit Technical Reference Manual*
0x1BC3_0000	64KB	Secure	CORE 2 PPU	Reserved if HOST_NUM_CPU_CORE < 3 See *Arm® CoreLink™ PCK-600 Power Control Kit Technical Reference Manual*
0x1BC4_0000	64KB	Secure	CORE 3 PPU	Reserved if HOST_NUM_CPU_CORE < 4. See *Arm® CoreLink™ PCK-600 Power Control Kit Technical Reference Manual*
0x1BC5_0000	3778KB	-	Reserved	-
0x1C00_0000	16MB	-	GIC Region	See 12.1.1.3 GIC Region on page 192
0x1D00_0000	16MB	-	CoreSight STM-500 Extended Stimulus Port	See 8.6 System Trace Macrocell on page 141
0x1E00_0000	1MB	Secure	NIC GPV	See *Arm® CoreLink™ NIC-450 Network Interconnect Technical Overview*
0x1E10_0000	543MB	-	Reserved	-

12.1.1.3 GIC Region

SSE-710 subsystem supports CoreLink™ GIC-400.

The following table shows the CoreLink™ GIC-400 memory region allocation.

Table 12-5: GIC region allocation

Offset	Size	Component
0x000_0000	64KB	Reserved
0x001_0000	4KB	GIC Distributor
0x001_1000	120KB	Reserved
0x002_F000	8KB	GIC CPU Interface

Offset	Size	Component
0x003_1000	120KB	Reserved
0x004_F000	4KB	GIC Virtual Interface Control
0x005_0000	4KB	GIC Virtual Interface Control Alias
0x005_1000	120KB	Reserved
0x006_F000	8KB	GIC Virtual CPU Interface
0x007_1000	15932KB	Reserved


Note

For more information on the programmers model of GIC see the *CoreLink™ GIC-400 Generic Interrupt Controller Technical Reference Manual*.

12.1.1.4 Host Master Expansion Region

The Host Master Expansion Region contains a 1GB area of memory routed to Host Expansion Master {0-1} (HOSTEXPMST{0-1}) Interfaces. Each interface is assigned 512MB.

The table below shows the Host Master Expansion Region.

Table 12-6: Host Master Expansion Region

Offset	Size	Security	Component	Notes
0x00_4000_0000	512MB	-	EXPMST0	-
0x00_6000_0000	512MB	-	EXPMST1	-

By adding use-case specific peripherals the integrator can expand the SSE-710 subsystem.

Access to the Host Expansion Master {0-1} interfaces are filtered by the EXPMST{0-1} Firewall Components.

For more information on the Host System Firewall, see 10.3.5 Host System firewall on page 169.

12.1.2 External Debug Bus memory map

The External Debug Bus memory map provides access to all debug logic in the SSE-710 subsystem, including the SoC Debug infrastructure.

Access is via different masters using different offsets:

- Host System processor and Secure Enclave, via the Host System memory map, at 0x1800_0000
- JTAG/Serial Wire Debug, via the DP at 0x0000_0000
- Each associated External System External Debug Interface (EXTSYS{0-1}EXTDBG) of the External System, at an **IMPLEMENTATION DEFINED** address

The table below shows the External Debug Bus memory map.

External Debug Memory Map access is controlled via the Debug Authorization Access Control Gate.

Table 12-7: External Debug Bus memory map

Offset	Size	Component	Notes
0x0000_0000	64KB	DP ROM	Only accessible by the DP. When accessed from any other master this location is Reserved.
0x0001_0000	64KB	GPIO Control	
0x0002_0000	64KB	EXT APBCOM	-
0x0003_0000	64KB	EXTDBG ROM	-
0x0004_0000	64KB	Secure Enclave AP	-
0x0005_0000	64KB	Host APB AP	-
0x0006_0000	64KB	Host AXIAP ROM	-
0x0007_0000	64KB	Host AXI AP	-
0x0008_0000	512KB	Reserved	-
0x0010_0000	64KB	SoC TPIU Funnel	-
0x0011_0000	64KB	SoC TPIU Replicator	-
0x0012_0000	64KB	SoC TPIU	-
0x0013_0000	64KB	SoC CTI	-
0x0014_0000	64KB	SoC ETR	-
0x0015_0000	64KB	SoC CATU	-
0x0016_0000	64KB	Counter CTI	-
0x0017_0000	576KB	Reserved	-
0x0020_0000	1MB	External System 0 Debug	-
0x0030_0000	1MB	External System 1 Debug	-
0x0040_0000	1MB	Reserved	-
0x0050_0000	1MB	Reserved	-
0x0060_0000	1018MB	Reserved	-
0x4000_0000	3GB	Reserved	-

Note

Reserved locations on the External Debug Bus are treated as RAZ/WI and generate an error when accessed.

External System {0-1} Debug

The External System{0-1} Debug regions are provided to allow the External System to implement its own debug logic.

Arm® recommends using only CoreSight™ ROM tables and APs in the External System {0-1} Debug regions. Accessing unused address space returns an error.

12.1.3 Secure Enclave memory map

The Secure Enclave has a 32-bit address space.

All unallocated memory regions are Reserved. Access to Reserved regions results in a RAZ/WI and a bus error response, except for the Reserved locations defined by the Cortex®-M0+ processor.

The following table shows the Secure Enclave memory map.

Table 12-8: Memory Map

Base address	Size	Region	Notes
0x0000_0000	128KB	Secure Enclave ROM Region	Provides access to the Secure Enclave ROM. See Secure Enclave ROM Region.
	96KB	Reserved	-
0x0002_0000	896KB	Reserved	
0x0010_0000	765MB		
0x2F00_0000	16MB	Crypto Accelerator socket	Exposed on the CAC interface. For more information on the CAC interface, see 4.5.1.5 Crypto Accelerator DMA (CAD) interface on page 62.
0x3000_0000	1MB	Secure Enclave RAM Region	Secure Enclave RAM. See section 12.1.3.2 Secure Enclave RAM region on page 196.
0x3010_0000	511MB	Reserved	-
0x5000_0000	256MB	Secure Enclave Peripheral Region	See 12.1.3.3 Secure Enclave Peripheral region on page 196
0x6000_0000	2GB	Host Access Region	See 12.1.3.4 Host Access region on page 197
0xE000_0000	1MB	Private Peripheral Bus	See Private Peripheral Bus (PPB)
0xE010_0000	255MB	Reserved	-
0xF000_0000	4KB	Secure Enclave CS ROM	-
0xF000_1000	4KB	Secure Enclave CTI	-
0xF000_2000	1016KB	Reserved	-
0xF010_0000	255MB	Reserved	-


Note

The Secure Enclave Memory Map applies to the Cortex®-M0+. The CAD interface only has visibility of the Host Access Region, the Secure Enclave ROM, and the Secure Enclave RAM.

12.1.3.1 Secure Enclave ROM region

The Secure Enclave ROM region provides a 128KB area of memory starting at address `0x0000_0000`. The Processor of the Secure Enclave core boots from this address.

The actual size of the ROM is defined by the SEC_ENC_ROM_SIZE configuration. All other unmapped locations in the Secure Enclave ROM Region is Reserved and is treated as RAZ/WI and generates an error. Any attempt to write to the Secure Enclave ROM generates is treated as WI and generates an error.

12.1.3.2 Secure Enclave RAM region

The Secure Enclave RAM region provides a 1MB area for accessing private on-chip SRAMs. This starts at address `0x3000_0000`.

The actual amount SRAM that is implemented is defined by the `SEC_ENC_RAM_SIZE` configuration and starts at address `0x3000_0000`. All other unmapped locations in the Secure Enclave RAM Region is Reserved and is treated as RAZ/WI and generates an error.

12.1.3.3 Secure Enclave Peripheral region

The Secure Enclave Peripheral region provides an area for accessing all internal Secure Enclave peripherals. The following table shows the memory map:

Table 12-9: Secure Enclave Peripheral region

Base address	Size	Region	Notes
`0x5000_0000`	`4KB`	Timer 0	CMSDK Timers. For more information, see the *Arm® Cortex®-M System Design Kit Technical Reference Manual*.
`0x5000_1000`	4KB	Timer 1	
`0x5000_2000`	4KB	Reserved	-
`0x5000_3000`	4KB	SEH 0 Sender	MHUs between Secure Enclave and Host System. For more details of these MHUs, see 11.5 MHU on page 184.
`0x5000_4000`	4KB	HSE 0 Receiver	
`0x5000_5000`	4KB	SEH 1 Sender	
`0x5000_6000`	4KB	HSE 1 Receiver	
`0x5000_7000`	36KB	Reserved	-
`0x5001_0000`	4KB	SEES0 0 Sender	MHUs between Secure Enclave and External System 0. For more details of these MHUs, see 11.5 MHU on page 184.
`0x5001_1000`	4KB	ES0SE 0 Receiver	
`0x5001_2000`	4KB	SEES0 1 Sender	Reserved when `EXT_SYS0_TZ_SPT` = 0
`0x5001_3000`	4KB	ES0SE 1 Receiver	MHUs between Secure Enclave and External System 0. For more details of these MHUs, see 11.5 MHU on page 184.
`0x5001_4000`	4KB	SEES1 0 Sender	MHUs between Secure Enclave and External System 1. For more details of these MHUs, see 11.5 MHU on page 184.
`0x5001_5000`	4KB	ES1SE 0 Receiver	

Base address	Size	Region	Notes
0x5001_6000	4KB	SEES1 1 Sender	Reserved when EXT_SYS1_TZ_SPT = 0 MHUs between Secure Enclave and External System 1. For more details of these MHUs, see 11.5 MHU on page 184.
0x5001_7000	4KB	ES1SE 1 Receiver	-
0x5001_8000	32KB	Reserved	-
0x5002_0000	96KB	Reserved	-
0x5008_0000	4KB	Secure Enclave System Control Register	Secure Enclave System Control Register. See section 12.3.2.2 Secure Enclave System Control register summary on page 249
0x5008_1000	4KB	Watchdog Timer	CSMDK Watchdog. See the *Arm® Cortex®-M System Design Kit Technical Reference Manual*.
0x5008_2000	44KB	Reserved	-
0x5008_D000	4KB	SECENCTOP PPU	See 9.1.6.9 SECENCTOP PPU on page 156
0x5008_E000	4KB	Secure Enclave Base System Control Register	See 12.3.2.1 Secure Enclave Base System Control register summary on page 236.
0x5008_F000	4KB	SoC Watchdog	See the *Arm® Cortex®-M System Design Kit Technical Reference Manual*.
0x5009_0000	4KB	UART	See the *PrimeCell UART (PL011) Technical Reference Manual*.
0x5009_1000	1084KB	Reserved	-
0x5020_0000	2MB	Secure Enclave Firewall	See 9.1.6.10 Secure Enclave Firewall on page 156
0x5040_0000	252MB	Reserved	-

12.1.3.4 Host Access region

The Host Access Region is a 2GB area starting at 0x6000_0000. It gives the Secure Enclave access to the Host System address space.

All access passes through FC1 of the Secure Enclave Firewall. The mapping between the Secure Enclave and Host System address space is controlled using the regions within the Secure Enclave Firewall and the Translation Extension. For more information on Firewall translation, see C. Firewall on page 346.


Note

The security of access to the Host System address space is controlled by the Firewall translation extension.

12.1.3.5 Cortex®-M0+ Private Peripheral Bus (PPB) region

The Private Peripheral Bus region is available only to the Cortex-M0+ core. For more details, see the *Arm®v6-M Architecture Reference Manual*.

12.1.3.6 Timers and Watchdog timers

Timers 0 and 1 and the Secure Enclave Watchdog timer are located in the SECENCTOP power domain.

The SoC Watchdog resides within the AONTOP domain of the Secure Enclave. For more information on these timers, see the *Arm® Cortex®-M System Design Kit Technical Reference Manual*.

12.1.3.7 MHUs

The Secure Enclave supports six pairs of MHUs. Each pair provides full-duplex communication between the Secure Enclave and either the Host System or an External System.

For more information on the MHUs, see A. Message Handling Unit on page 311.

12.1.4 External System memory map

The EXTSYS{0-1}MHU interface permits access to the MHUs. The following table shows the External System memory map.


Note

The full memory map of the External System is **IMPLEMENTATION DEFINED**.

Table 12-10: Memory map of External System MHU interface

Offset	Size	Security	Component	Notes
0x0_0000	4KB	Secure	HES{x}0 Receiver	When EXT_SYS{x}_TZ_SPT = 1
		-		When EXT_SYS{x}_TZ_SPT = 0
0x0_1000	60KB	-	Reserved	RAZ/WI and slave error is returned
0x1_0000	4KB	Secure	ES{x}H0 Sender	When EXT_SYS{x}_TZ_SPT = 1
		-		When EXT_SYS{x}_TZ_SPT = 0
0x1_1000	60KB	-	Reserved	RAZ/WI and slave error is returned
0x2_0000	4KB	-	HES{x}1 Receiver	Reserved when EXT_SYS{x}_TZ_SPT =0
0x2_1000	60KB	-	Reserved	RAZ/WI and slave error is returned
0x3_0000	4KB	-	ES{x}H1 Sender	Reserved when EXT_SYS{x}_TZ_SPT =0
0x3_1000	60KB	-	Reserved	RAZ/WI and slave error is returned
0x4_0000	4KB	Secure	SEES{x}0 Receiver	When EXT_SYS{x}_TZ_SPT = 1
		-		When EXT_SYS{x}_TZ_SPT = 0

Offset	Size	Security	Component	Notes
0x4_1000	60KB	-	Reserved	RAZ/WI and slave error is returned
0x5_0000	4KB	Secure	ES{x}SE0 Sender	When EXT_SYS{x}_TZ_SPT = 1
		-		When EXT_SYS{x}_TZ_SPT = 0
0x5_1000	60KB	-	Reserved	RAZ/WI and slave error is returned
0x6_0000	4KB	-	SEES{x}1 Receiver	Reserved when EXT_SYS{x}_TZ_SPT =0
0x6_1000	60KB	-	Reserved	RAZ/WI and slave error is returned
0x7_0000	4KB	-	ES{x}SE1 Sender	Reserved when EXT_SYS{x}_TZ_SPT =0
0x7_1000	60KB	-	Reserved	RAZ/WI and slave error is returned

Note

{x} is the External System number (0-1).

12.2 Interrupt map

This section covers the interrupt map of the Host Processor and the Secure Enclave interrupt map.

12.2.1 Host CPU interrupt map

The Host processor interrupt map depends on the configuration of the SSE-710.

Table 12-11: Host CPU interrupt map

Interrupt number	Interrupt source	Type	Level/edge	Notes
0-15	Reserved	SGI	Edge	-
16-24	Reserved	PPI	-	-
25	Virtual CPU Interface Maintenance	PPI	Level	-
26	Hypervisor timer	PPI	Level	-
27	Virtual timer	PPI	Level	-
28	Reserved	PPI	Level	-
29	Secure physical timer	PPI	Level	-
30	Non-secure physical timer	PPI	Level	-
31	Reserved	PPI	-	-
32	Secure WDOG WS0	SPI	Level	-
33	Non-secure WDOG WS1	SPI	Level	-
34	**REFCLK** Timer 0	SPI	Level	-
35	S32K Timer 0	SPI	Level	-
36	Host System Firewall	SPI	Level	-

Interrupt number	Interrupt source	Type	Level/edge	Notes
37	Host PPU Combined	SPI	Level	-
38	CoreSight™ SDC-600	SPI	Level	-
39	CPU **nEXTERRIRQ**	SPI	Level	See the relevant supported Cortex®-A processors TRM for more information.
40	CPU **nINTERRIRQ**	SPI	Level	
41	Host to Secure Enclave MHU0 Sender Combined IRQ	SPI	Level	-
42	Secure Enclave to Host MHU0 Receiver Combined IRQ	SPI	Level	-
43	Host to EXTSYS0 MHU0 Sender Combined IRQ	SPI	Level	-
44	EXTSYS 0 to Host MHU0 Receiver Combined IRQ	SPI	Level	-
45	Host to EXTSYS 1 MHU0 Sender Combined IRQ	SPI	Level	-
46	EXTSYS 1 to Host MHU0 Receiver Combined IRQ	SPI	Level	-
47-50	Reserved	SPI	-	-
51	Host UART0 UARTINTR	SPI	Level	-
52	Host UART1 UARTINTR	SPI	Level	-
53-63	SoC Expansion	SPI	**IMPLEMENTATION DEFINED**	The interrupts could be from GICSHDINT or the dedicated interrupts which do not pass through the Interrupt Router
64	Non-secure WDOG WS0	SPI	Level	-
65	**REFCLK** Timer 1	SPI	Level	-
66	**REFCLK** Timer 2	SPI	Level	-
67	**REFCLK** Timer 3	SPI	Level	-
68	S32K Timer 1	SPI	Level	-
69	Host STM Sync IRQ	SPI	Edge	-
70	Host ETR Buffer IRQ	SPI	Level	-
71	Host CATU IRQ	SPI	Level	-
72	Host CTI Trigger Out 4	SPI	Edge	-
73	Host CTI Trigger Out 5	SPI	Edge	-
74	SoC ETR IRQ	SPI	Level	-
75	SoC CATU IRQ	SPI	Level	-
76	Host to Secure Enclave MHU1 Sender Combined IRQ	SPI	Level	-
77	Secure Enclave to Host MHU1 Receiver Combined IRQ	SPI	Level	-
78	Host to EXTSYS 0 MHU1 Sender Combined IRQ	SPI	Level	Only implemented when `EXT_SYS0_TZ_SPT` = 1, otherwise Reserved
79	EXTSYS 0 to Host MHU1 Receiver Combined IRQ	SPI	Level	

Interrupt number	Interrupt source	Type	Level/edge	Notes
80	Host to EXTSYS 1 MHU1 Sender Combined IRQ	SPI	Level	Only implemented when `EXT_SYS1_TZ_SPT` = 1, otherwise Reserved
81	EXTSYS 1 to Host MHU1 Receiver Combined IRQ	SPI	Level	
82-85	Reserved	SPI	-	-
86	Host CPU0 Debug Comm Channel	SPI	Level	-
87	Reserved	SPI	-	-
88	Host CPU0 PMU Counter Overflow	SPI	Level	-
89	Host CPU0 CTI Trigger	SPI	Edge	-
90	Reserved	SPI	-	-
91	Host CPU1 Debug Comm Channel	SPI	Level	Only implemented when HOST_CPU_NUM_CORES > 1. Otherwise Reserved.
92	Reserved	SPI	-	-
93	Host CPU1 PMU Counter Overflow	SPI	Level	Only implemented when HOST_CPU_NUM_CORES > 1. Otherwise Reserved.
94	Host CPU1 CTI Trigger	SPI	Edge	
95	Reserved	SPI	-	-
96	Host CPU2 Debug Comm Channel	SPI	Level	Only implemented when HOST_CPU_NUM_CORES > 2. Otherwise Reserved.
97	Reserved	SPI	-	-
98	Host CPU2 PMU Counter Overflow	SPI	Level	Only implemented when HOST_CPU_NUM_CORES > 2. Otherwise Reserved.
99	Host CPU2 CTI Trigger	SPI	Edge	
100	Reserved	SPI	-	-
101	Host CPU3 Debug Comm Channel	SPI	Level	Only implemented when HOST_CPU_NUM_CORES > 3. Otherwise Reserved.
102	Reserved	SPI	-	-
103	Host CPU3 PMU Counter Overflow	SPI	Level	Only implemented when HOST_CPU_NUM_CORES > 3. Otherwise Reserved.
104	Host CPU3 CTI Trigger	SPI	Edge	Only implemented HOST_CPU_NUM_CORES > 3. Otherwise Reserved.
105-127	Reserved	SPI	-	-
>=128	SoC Expansion	SPI	**IMPLEMENTATION DEFINED**	The interrupts could be from GICSHDINT or the dedicated interrupts which do not pass through the Interrupt Router


Note

Reserved interrupts are tied LOW when not implemented.

Interrupt IDs 32-63 are lockable when they are configured as secure interrupts. For more information, see *Arm® Generic Interrupt Controller Architecture Specification, architecture version 2.0*.

12.2.2 Secure Enclave interrupt map

The following table summarizes the Secure Enclave interrupt map.

Table 12-12: Interrupt Map

Interrupt number	Interrupt source	Level/edge	Notes
NMI	SoC Watchdog Timer	Level	-
0	Secure Enclave Interrupt Expansion	Level	See 12.2.3 Secure Enclave interrupt expansion on page 202
1	Crypto Accelerator Interrupt 0	Level	If unused then must be tied to LOW
2	Crypto Accelerator Interrupt 1	Level	
3	Secure Enclave Watchdog Timer	Level	-
4	Reserved	-	-
5	CMSDK Timer 0	Level	-
6	CMSDK Timer 1	Level	-
7	Host System Firewall Tamper Interrupt	Level	-
8	Interrupt Router Tamper Interrupt	Level	-
9	Secure Watchdog WS1	Level	-
10	SECENCTOP PPU	Level	-
11	UART UARTINTR	Level	-
12	Secure Enclave Firewall Interrupt	Level	-
13	Secure Enclave CTI Trigger Out 2	Level	-
14	Secure Enclave CTI Trigger Out 3	Level	-
15-20	Reserved	-	-
21	SEH0 MHU Sender Combined interrupt	Level	-
22	Reserved	-	-
23	HSE0 MHU Receiver Combined interrupt	Level	-
24	Reserved	-	-
25	Reserved	-	-
26	SEH1 MHU Sender Combined interrupt	Level	-
27	Reserved	-	-
28	HSE1 MHU Receiver Combined interrupt	Level	-
29-31	Reserved	-	-

12.2.3 Secure Enclave interrupt expansion

The Secure Enclave interrupt collator provides 128 interrupts in addition to the interrupts provided by the Cortex®-M0+ NVIC.

The output drives interrupt 0 in the Secure Enclave Cortex®-M0+. If an interrupt is unmasked and Secure Enclave Base System Control **BSYS_PWR_REQ.WAKEUP_EN** is HIGH, these interrupts force the Secure Enclave to exit the OFF or MEM_RET power modes.

Software uses the SEC_ENC_INT_COL_ST{0-3} and SEC_ENC_INT_COL_MSK{0-3} registers to monitor the status, and mask the interrupt respectively. All interrupts routed to the Secure Enclave must be level-based as there is no logic to capture which source caused the interrupt. The Secure Enclave Expansion interrupts are referred to as SEEI{x}, where x is the interrupt number starting from 0.

The following table summarizes the Secure Enclave interrupt expansion.

Table 12-13: Secure Enclave interrupt expansion

Interrupt number (SEEI{x})	Interrupt source	Level/ edge	Notes
0	Host System Firewall Interrupt	Level	All interrupt sources are routed through the Interrupt Router. For more information about the Interrupt Router, see B. Interrupt Router on page 340.
1	CoreSight™ SDC-600	Level	
2	Host PPU Combined Interrupt	Level	
3	**REFCLK** Timer 0	Level	
4	**REFCLK** Timer 1	Level	
5	**REFCLK** Timer 2	Level	
6	**REFCLK** Timer 3	Level	
7	**S32K** Timer 0	Level	
8	**S32K** Timer 1	Level	
9-31	Reserved	-	-
32-63	Shared Interrupt 32-63	Level	All interrupt sources are routed through the Interrupt Router. The source of the interrupt is **IMPLEMENTATION DEFINED**.
64	SEES00 MHU Sender Combined Interrupt	Level	-
65	ES0SE0 MHU Receiver Combined Interrupt	Level	-
66	SEES01 MHU Sender Combined Interrupt	Level	Only implemented when `EXT_SYS0_TZ_SPT` = 1, otherwise Reserved.
67	ES0SE1 MHU Receiver Combined Interrupt	Level	
68	SEES10 MHU Sender Combined Interrupt	Level	
69	ES1SE0 MHU Receiver Combined Interrupt	Level	
70	SEES11 MHU Sender Combined Interrupt	Level	Only implemented when `EXT_SYS1_TZ_SPT` = 1, otherwise Reserved.
71	ES1SE1 MHU Receiver Combined Interrupt	Level	
72-127	Reserved	-	-

12.3 Register descriptions

This chapter describes the registers for SSE-710 components.

For the following standalone components, see their individual references:

- PPU
- Coresight IP components
- CoreSight™ SDC-600
- supported Cortex®-A processors
- NIC-400 GPV
- UART

For detailed register descriptions of the Firewall and MHU, see C. Firewall on page 346 and A. Message Handling Unit on page 311

12.3.1 Host Base System Control register summary

This section summarizes the Host Base System Control registers.

The Host Base System Control has registers to control the clocks, power, and reset for SSE-710. It resides within AONTOP power domain. The registers can be accessed at offset `0x1A01_0000` in the Host System memory map.

Table 12-14: Host Base System Control registers summary

Offset	Short name	Access	Name
`0x000`	CLUSTER_CONFIG	RW	Cluster Static Config 12.3.1.1 Cluster Static Config (CLUS_CFG) register on page 207
`0x004 – 0x00C`	-	RO	Reserved
`0x010`	PE0_CONFIG	RW	Processing Element 0 Static Config 12.3.1.2 Processing Element {0-3} Static Config (PE{0-3}_CFG) register on page 208
`0x014`	PE0_RVBARADDR_LW	RW	Processing Element 0 Reset Vector Base Address Lower 12.3.1.3 Processing Element {0-3} Reset Vector Base Address Lower and Upper (PE{0-3}_RVBAR_LW) register on page 209
`0x018`	PE0_RVBARADDR_UP	RO	Processing Element 0 Reset Vector Base Address Upper 12.3.1.4 Processing Element {0-3} Reset Vector Base Address Upper (PE{0-3}_RVBAR_UP) register on page 209
`0x01C`	-	RO	Reserved
`0x020`	PE1_CONFIG	RW	Processing Element 1 Static Config 12.3.1.2 Processing Element {0-3} Static Config (PE{0-3}_CFG) register on page 208
`0x024`	PE1_RVBARADDR_LW	RW	Processing Element 1 Reset Vector Base Address Lower 12.3.1.3 Processing Element {0-3} Reset Vector Base Address Lower and Upper (PE{0-3}_RVBAR_LW) register on page 209
`0x028`	PE1_RVBARADDR_UP	RO	Processing Element 1 Reset Vector Base Address Upper 12.3.1.4 Processing Element {0-3} Reset Vector Base Address Upper (PE{0-3}_RVBAR_UP) register on page 209
`0x02C`	-	RO	Reserved

Offset	Short name	Access	Name
0x030	PE2_CONFIG	RW	Processing Element 2 Static Config 12.3.1.2 Processing Element {0-3} Static Config (PE{0-3}_CFG) register on page 208
0x034	PE2_RVBARADDR_LW	RW	Processing Element 2 Reset Vector Base Address Lower 12.3.1.3 Processing Element {0-3} Reset Vector Base Address Lower and Upper (PE{0-3}_RVBAR_LW) register on page 209
0x038	PE2_RVBARADDR_UP	RO	Processing Element 2 Reset Vector Base Address Upper 12.3.1.4 Processing Element {0-3} Reset Vector Base Address Upper (PE{0-3}_RVBAR_UP) register on page 209
0x03C	-	RO	Reserved
0x040	PE3_CONFIG	RW	Processing Element 3 Static Config 12.3.1.2 Processing Element {0-3} Static Config (PE{0-3}_CFG) register on page 208
0x044	PE3_RVBARADDR_LW	RW	Processing Element 3 Reset Vector Base Address Lower 12.3.1.3 Processing Element {0-3} Reset Vector Base Address Lower and Upper (PE{0-3}_RVBAR_LW) register on page 209
0x048	PE3_RVBARADDR_UP	RO	Processing Element 3 Reset Vector Base Address Upper 12.3.1.4 Processing Element {0-3} Reset Vector Base Address Upper (PE{0-3}_RVBAR_UP) register on page 209
0x04C	-	RO	Reserved
0x050 – 0x1FC	-	RO	Reserved
0x200	HOST_RST_SYN	RO	Host Reset Syndrome 12.3.1.5 Host Reset Syndrome (HOST_RST_SYN) register on page 210
0x204 – 0x2FC	-	RO	Reserved
0x300	HOST_CPU_BOOT_MSK	RW	Host CPU Boot Mask 12.3.1.6 Host CPU Boot Mask (HOST_CPU_BOOT_MSK) register on page 211
0x304	HOST_CPU_CLUS_PWR_REQ	RW	Host CPU Cluster Power Request 12.3.1.7 Host CPU Cluster Power Request (HOST_CPU_CLUS_PWR_REQ) register on page 211
0x308	HOST_CPU_WAKEUP	RW	Host CPU Wakeup 12.3.1.8 Host CPU Wakeup (HOST_CPU_WAKEUP) register on page 212
0x30C	-	RO	Reserved
0x310	EXT_SYS0_RST_CTRL	RW	External System 0 Reset Control 12.3.1.9 External System {0-1} Reset Control (EXT_SYS{0-1}_RST_CTRL) register on page 212
0x314	EXT_SYS0_RST_ST	RO	External System 0 Reset Status 12.3.1.10 External System {0-1} Reset Status (EXT_SYS{0-1}_RST_ST) register on page 213
0x318	EXT_SYS1_RST_CTRL	RW	External System 1 Reset Control 12.3.1.9 External System {0-1} Reset Control (EXT_SYS{0-1}_RST_CTRL) register on page 212
0x31C	EXT_SYS1_RST_ST	RO	External System 1 Reset Status 12.3.1.10 External System {0-1} Reset Status (EXT_SYS{0-1}_RST_ST) register on page 213
0x320 – 0x3FC	-	RO	Reserved
0x400	BSYS_PWR_REQ	RW	Base System Power Request 12.3.1.11 Base System Power Request (BSYS_PWR_REQ) register on page 213
0x404	BSYS_PWR_ST	RO	Base System Power Status 12.3.1.12 Base System Power Status (BSYS_PWR_ST) register on page 214
0x408 – 0x4FC	-	RO	Reserved
0x500	HOST_SYS_LCTRL_ST	RO	Host System Lock Control Status 12.3.1.13 Host System Lock Control Status (HOST_SYS_LCTRL_ST) register on page 215

Offset	Short name	Access	Name
0x504	HOST_SYS_LCTRL_SET	WO	Host System Lock Control Set 12.3.1.14 Host System Lock Control Set (HOST_SYS_LCTRL_SET) register on page 217
0x508	HOST_SYS_LCTRL_CLR	WO	Host System Lock Control Clear 12.3.1.15 Host System Lock Control Clear (HOST_SYS_LCTRL_CLR) register on page 218
0x50C0x – 0x7FC	-	RO	Reserved
0x800	HOSTCPUCLK_CTRL	RW	Host CPU Clock Control 12.3.1.16 Host CPU Clock Control (HOSTCPUCLK_CTRL) register on page 220
0x804	HOSTCPUCLK_DIV0	RW	Host CPU Clock Divider 0 12.3.1.17 Host CPU Clock Divider 0 (HOSTCPUCLK_DIV0) register on page 220
0x808	HOSTCPUCLK_DIV1	RW	Host CPU Clock Divider 1 12.3.1.18 Host CPU Clock Divider 1 (HOSTCPUCLK_DIV1) register on page 221
0x80C	-	RO	Reserved
0x810	GICCLK_CTRL	RW	GIC Clock Control 12.3.1.19 GIC Clock Control (GICCLK_CTRL) register on page 222
0x814	GICCLK_DIV0	RW	GIC Clock Divider 0 12.3.1.20 GIC Clock Divider 0 (GICCLK_DIV0) register on page 222
0x818 – 0x81C	-	RO	Reserved
0x820	ACLK_CTRL	RW	AXI Clock Control 12.3.1.21 AXI Clock Control (ACLK_CTRL) register on page 223
0x824	ACLK_DIV0	RW	AXI Clock Divider 0 12.3.1.22 AXI Clock Divider 0 (ACLK_DIV0) register on page 224
0x828 – 0x82C	-	RO	Reserved
0x830	CTRLCLK_CTRL	RW	Control Clock Control 12.3.1.23 Control Clock Control (CTRLCLK_CTRL) register on page 225
0x834	CTRLCLK_DIV0	RW	Control Clock Divider 0 12.3.1.24 Control Clock Divider 0 (CTRLCLK_DIV0) register on page 225
0x838 – 0x83C	-	RO	Reserved
0x840	DBGCLK_CTRL	RW	Debug Clock Control 12.3.1.25 Debug Clock Control (DBGCLK_CTRL) register on page 226
0x844	DBGCLK_DIV0	RW	Debug Clock Divider 0 12.3.1.26 Debug Clock Divider 0 (DBGCLK_DIV0) register on page 227
0x848 – 0x84C	-	RO	Reserved
0x850	HOSTUARTCLK_CTRL	RW	Host UART Clock Control 12.3.1.27 Host UART Clock Control (HOSTUARTCLK_CTRL) register on page 228
0x854	HOSTUARTCLK_DIV0	RW	Host UART Clock Divider 0 12.3.1.28 Host UART Clock Divider 0 (HOSTUARTCLK_DIV0) register on page 228
0x858 – 0x85C	-	RO	Reserved
0x860	REFCLK_CTRL	RW	REFCLK Clock Control 12.3.1.29 REFCLK Clock Control register on page 229
0x864 – 0x9FC	-	RO	Reserved

Offset	Short name	Access	Name
0xA00	CLKFORCE_ST	RO	Clock Force Status 12.3.1.30 Clock Force Status (CLKFORCE_ST) register on page 230
0xA04	CLKFORCE_SET	WO	Clock Force Set 12.3.1.31 Clock Force Set (CLKFORCE_SET) register on page 230
0xA08	CLKFORCE_CLR	WO	Clock Force Clear 12.3.1.32 Clock Force Clear (CLKFORCE_CLR) register on page 231
0xA0C	-	RO	Reserved
0xA10	PLL_ST	RO	PLL Status 12.3.1.33 PLL Status (PLL_ST) register on page 232
0xA140x – 0xAFC	-	RO	Reserved
0xB00	HOST_PPU_INT_ST	RO	Host PPU Interrupt Status 12.3.1.34 Host PPU Interrupt Status (HOST_PPU_INT_ST) register on page 232
0xB04 – 0xFCC	-	RO	Reserved
0xFD0	PID4	RO	Peripheral ID4 12.3.1.35 Peripheral ID 4 (PID4) register on page 233
0xFD4	PID5	RO	Peripheral ID5 12.3.1.36 Peripheral ID 5 (PID5) register on page 234
0xFD8	PID6	RO	Peripheral ID6 12.3.1.37 Peripheral ID 6 (PID6) register on page 234
0xFDC	PID7	RO	Peripheral ID7 12.3.1.38 Peripheral ID 7 (PID7) register on page 234
0xFE0	PID0	RO	Peripheral ID0 12.3.1.39 Peripheral ID 0 (PID0) register on page 234
0xFE4	PID1	RO	Peripheral ID1 12.3.1.40 Peripheral ID 1 (PID1) register on page 235
0xFE8	PID2	RO	Peripheral ID2 12.3.1.41 Peripheral ID 2 (PID2) register on page 235
0xFEC	PID3	RO	Peripheral ID3 12.3.1.42 Peripheral ID 3 (PID3) register on page 235
0xFF0	CID0	RO	Component ID0 12.3.1.43 Component ID 0 (CID0) register on page 235
0xFF4	CID1	RO	Component ID1 12.3.1.44 Component ID 1 (CID1) register on page 236
0xFF8	CID2	RO	Component ID2 12.3.1.45 Component ID 2 (CID2) register on page 236
0xFFC	CID3	RO	Component ID3 12.3.1.46 Component ID 3 (CID3) register on page 236

The Host Base System Control registers support 32-bit word aligned access. Any other bit size or unaligned access results in an error response and RAZ/WI.

The Host Base System Control registers have the following behaviors:

- Any read to a write-only register is treated as RAZ.
- Any write to a read-only register is treated as WI.

In both cases, no error is generated.

12.3.1.1 Cluster Static Config (CLUS_CFG) register

The following table gives a bit-level description of the Cluster Static Config register.

Table 12-15: CLUS_CFG register

Bits	Name	Description	Type	Reset
[31:1]	-	Reserved.	RO	0x0000_000

Bits	Name	Description	Type	Reset
[0]	CRYPTODISABLE	Disable the cryptographic extensions of the Host processor Cores. Possible field values are: 0b0 – Crypto is enabled. 0b1 – Crypto is disabled. Changes only take effect at power-on reset of the core.	RW	0b0


Note

A power-on reset of a core is when the CORE{0-3} power domain performs an OFF to ON transition.

When HOST_SYS_LCTRL.HOST_LOCK is set to 0b1, writing to this register does not update it and generates an error.

12.3.1.2 Processing Element {0-3} Static Config (PE{0-3}_CFG) register

The following table gives a bit-level description of the PE{0-3}_CFG register.

Table 12-16: PE{0-3}_CFG register

Bits	Name	Description	Type	Reset
[31:4]	-	Reserved	RO	0x0000000
[3]	AA64nAA32	Select initial register width state. This field is only implemented if the HOST_CPU_TYPE is 2 or 3, otherwise it is Reserved and treated as RAZ/WI. • 0b0 – AArch32 • 0b1 – AArch64. Changes in this field only take effect at a power on reset of the core	RW	0b0
[2]	VINITHI	Locations of the exception vectors at reset. Sets the initial value of SCTLR.V. Possible field values are: • 0b0: Exception vector start at 0x000000000 • 0b1: Exception vector start at 0xFFFF00000 Changes in this field only take effect at a reset of the core.	RW	0b0
[1]	CFGTE	Enabling T32 exceptions. Sets the initial value of the SCTLR.TE. Possible field values are: • 0b0: Exceptions taken in Arm(v7) or AArch32(v8) state • 0b1: Exceptions taken in Thumb32 state Changes in this field only take effect at a reset of the core.	RW	0b0

Bits	Name	Description	Type	Reset
[0]	CFGEND	Endianness configuration at reset. Sets the initial value of the SCTLR_EL3.EE and SCTR_S.EE bits. Possible field values are: • `0b0`: Little-endian • `0b1`: Big-endian Changes in this field only take effect at a reset of the core.	RW	`0b0`


Note

A power-on reset of a core is when the CORE{0-3} power domain performs an OFF to ON transition. A reset of the core is when the CORE{0-3} power domain performs any of the following transitions:

- OFF to ON
- OFF_EMU to ON
- WARM_RST to ON

When HOST_SYS_LCTRL_ST.HOST_LOCK is set to `0b1` a write to this register does not update the register and generates an error response.

12.3.1.3 Processing Element {0-3} Reset Vector Base Address Lower and Upper (PE{0-3}_RVBAR_LW) register

The following table gives a bit-level description of the Reset Vector Lower register.

Table 12-17: PE{0-3}_RVBAR_LW register

Bits	Name	Description	Type	Reset
[31:2]	RVBAR31_2	Reset vector address bits [31:2]. Sets the initial value of the RVBAR_EL3 register. This field is only implemented when CPU_TYPE is 2 or 3, otherwise it is Reserved and treated as RAZ/WI. **Note:** Changes in this field only take effect at a power on reset of the core.	RW	`0x00000000`
[11:0]	-	Reserved	RO	`0b00`


Note

When HOST_SYS_LCTRL_ST.HOST_LOCK is set to `0b1` a write to this register does not update the register and generates an error response.

12.3.1.4 Processing Element {0-3} Reset Vector Base Address Upper (PE{0-3}_RVBAR_UP) register

The following table gives a bit-level description of the reset vector lower register.

Table 12-18: PE{0-3}_RVBAR_UP register

Bits	Name	Description	Type	Reset
[31:12]	-	Reserved	RO	0x00000
[11:0]	RVBAR43_32	Reset vector address bits [43:32]. Sets the initial value of RVBAR_EL3 register. As SSE-700 is a 32-bit system, this field is read-only and reads as all zeros.	RO	0x000


Note

When HOST_SYS_LCTRL_ST.HOST_LOCK, is set to 0b1 a write to this register does not update the register and generates an error response.

12.3.1.5 Host Reset Syndrome (HOST_RST_SYN) register

The following table gives a bit-level description of the Reset Syndrome register.

Table 12-19: HOST_RST_SYN register

Bits	Name	Description	Type	Reset
[31:4]	-	Reserved.	RO	0x0000000
[3]	HOST	Indicates the last reset of the Host System was caused by the HOST_SYS_RST_CTRL.RST_REQ bit being set to 0b1	RO	UNKNOWN
[2]	-	Reserved	RO	0b0
[1]	**nSRST**	Indicates that the last reset of the Host System was caused by either: • **nSRST** pin being asserted • **DP ROM CSYSRSTREQ** being asserted	RO	UNKNOWN
[0]	POR	Indicates that the last reset of the Host System was caused by one of the following: • **PORESETn** pin being asserted • **DP CDBGRSTREQ** being asserted • Secure Enclave Watchdog reset request • SoC Watchdog reset request • SOC_RST_CTRL.RST_REQ bit set to 0b1 • Secure Enclave software reset request • Secure Enclave Crypto Accelerator Error reset request	RO	UNKNOWN


Note

This register enables software to read the Reset Syndrome output of the Reset Controller. Therefore, the reset value of this register depends on the cause of the last reset of the Host System.

12.3.1.6 Host CPU Boot Mask (HOST_CPU_BOOT_MSK) register

The following table gives a bit-level description of the Boot Mask register.

Table 12-20: HOST_CPU_BOOT_MSK register

Bits	Name	Description	Type	Reset
[31:4]	-	Reserved.	RO	0x0000000
[3:0]	BOOT_MSK	Configures which Host processor Cores automatically boot when the CLUSTOP domain exits OFF or MEM_RET. Each bit in this field corresponds to a Host processor Core, with bit 0 corresponding to Host processor Core 0 and bit 1 corresponding to Host processor Core 1. Bit values: • 0: Host processor Core x does not automatically boot on CLUSTOP domain exit from OFF or MEM_RET. • 1: Host processor Core x does automatically boot on CLUSTOP domain exit from OFF or MEM_RET. Changes in this field only take effect at the next CLUSTOP transitions from OFF or MEM_RET. The legal values depend on the number of Host processor cores used, see the table below. All other values are Reserved and generate an error if software attempts to write to the register.	RW	0x1

Note

When HOST_SYS_LCTRL_ST.HOST_LOCK is set to 1, a write to this register does not update the register and generates an error.

Table 12-21: Legal values of the BOOT MSK register

HOST_CPU_NUM_CORE	BOOT_MSK Legal Values
1	0x1
2	0x1 – 0x3
3	0x1 – 0x7
4	0x1 – 0xF

12.3.1.7 Host CPU Cluster Power Request (HOST_CPU_CLUS_PWR_REQ) register

The following table gives a bit-level description of the Cluster Power State Request register.

Table 12-22: HOST_CPU_CLUS_PWR_REQ register

Bits	Name	Description	Type	Reset
[31:2]	-	Reserved.	RO	0x00000000

Bits	Name	Description	Type	Reset
[1]	MEM_RET_REQ	Possible field values are: 0: When entering a low power mode, the last level cache of the Host processor is not required to be retained 1: When entering a low power mode, the last level cache of the Host processor is required to be retained	RW	0b0
[0]	PWR_REQ	Possible field values are: 0: CLUSTOP can enter a low power mode, when all CORE{0-3} domains are OFF and the GIC is idle 1: CLUSTOP is required to be in a power mode greater than or equal to FUNC_RET even when all CORE{0-3} domains are OFF	RW	0b0

12.3.1.8 Host CPU Wakeup (HOST_CPU_WAKEUP) register

The following table gives a bit-level description of the Host CPU Wakeup register.

Table 12-23: HOST_CPU_WAKEUP register

Bits	Name	Description	Type	Reset
[31:4]	-	Reserved	RO	0x0000000
[3]	CORE3_WAKEUP	Wakeup Core 3. This field is only implemented when HOST_CPU_NUM_CORES is 3 otherwise it is Reserved and treated as RAZ/WI.	RW	0b0
[2]	CORE2_WAKEUP	Wakeup Core 2. This field is only implemented when `HOST_CPU_NUM_CORES` > 2, otherwise it is Reserved and treated as RAZ/WI.	RW	0b0
[1]	CORE1_WAKEUP	Wakeup Core 1. This field is only implemented when `HOST_CPU_NUM_CORES` > 1, otherwise it is Reserved and treated as RAZ/WI.	RW	0b0
[0]	CORE0_WAKEUP	Wakeup Core 0	RW	0b0

Arm recommends, that software only writes to this register to wake a Host processor core for the first time after the CLUSTOP power domain has exited one of the following power modes: OFF, MEM_RET or WARM_RST. In all other cases, Arm recommends using interrupts through the Host processor GIC.

12.3.1.9 External System {0-1} Reset Control (EXT_SYS{0-1}_RST_CTRL) register

The following table gives a bit-level description of the Reset Control register.

Table 12-24: EXT_SYS{0-1}_RST_CTRL register

Bits	Name	Description	Type	Reset
[31:2]	-	Reserved	RO	0x00000000

Bits	Name	Description	Type	Reset
[1]	RST_REQ	Reset request for External System. Possible field values are: • 0: No reset requested • 1: Reset requested	RW	0b0
[0]	CPUWAIT	CPU Wait control Possible field values are: • 0: CPUWAIT signal of the External System is de-asserted: • 1: CPUWAIT signal of the External System is asserted. When CPUWAIT becomes 0b0 any attempt to revert to 0b1 is ignored. This field only returns to 0b1 when any of the following occur: • External Power-on reset. • Internal Power-on reset. • Debug reset. • Host System reset. • Reset of the associated External System. This field becomes RO when associated EXT_SYS{0-1}_CPUWAIT_WEN is 0b0. Any attempt to set this field to 0b0 by software is ignored.	RW	0b1

12.3.1.10 External System {0-1} Reset Status (EXT_SYS{0-1}_RST_ST) register

The following table gives a bit-level description of the Reset Status register.

Table 12-25: EXT_SYS{0-1}_RST_ST register

Bits	Name	Description	Type	Reset
[31:2]	-	Reserved	RO	0x000000
[2:1]	RST_ACK	Status of reset request Possible field values are: 0b00: No reset requested 0b01: Reset request unable to complete 0b10: Reset request complete 0b11: Reserved	RO	0b00
[0]	-	Reserved	RO	0b0

12.3.1.11 Base System Power Request (BSYS_PWR_REQ) register

The following table gives a bit-level description of the Base Sytem Power Request register.

Table 12-26: BSYS_PWR_REQ register

Bits	Name	Description	Type	Reset
[31:6]	-	Reserved	RO	0x0000000
[5:3]	SYSTOP_PWR_REQ	Selects SYSTOP power domain behavior when no activity in the domain. Possible field values are: 0b000: No request for logic or volatile memory to be powered 0b001: No request for logic to be powered, but volatile memory must be retained 0b01x: Request for logic to be powered, but volatile memory can be either powered or retained 0b1xx: Request for logic and volatile memory to be powered	RW	0b000
[2]	DBGTOP_PWR_REQ	Selects DBGTOP power domain behavior when no activity in the domain. Possible field values are: 0b0: No request for DBGTOP to be powered 0b1: Request for DBGTOP to be powered	RW	0b0
[1]	REFCLK_REQ	Request **REFCLK** Possible field values are: 0b0: No request for **REFCLK** to be supplied 0b1: Request for **REFCLK** to be supplied	RW	0b0
[0]	WAKEUP_EN	Host System wakeup enable. Possible field values are: 0b0: Wakeup for Host System is disabled 0b1: Wakeup for Host System is enabled	RW	0b0

12.3.1.12 Base System Power Status (BSYS_PWR_ST) register

The following table gives a bit-level description of the Bit level Base System Power Status register.

Table 12-27: BSYS_PWR_ST register

Bits	Name	Description	Type	Reset
[31:6]	-	Reserved	RO	0x0000000

Bits	Name	Description	Type	Reset
[5:3]	SYSTOP_PWR_ST	SYSTOP power domain status. Possible field values are: `0b000`: SYSTOP is in the OFF or WARM_RST power mode `0b001`: SYSTOP is in the MEM_RET power mode `0b010`: SYSTOP is in the FUNC_RET power mode `0b100`: SYSTOP is in the ON power mode All other values are Reserved.	RO	UNKNOWN
[2]	DBGTOP_PWR_ST	DBGTOP power domain status. Possible field values are: `0b0`: DBGTOP is in the OFF or WARM_RST power mode `0b1`: DBGTOP is in the ON-power mode	RO	UNKNOWN
[1]	-	Reserved	RO	`0b0`
[0]	-	Reserved	RO	`0b0`

Note

The values in this register are driven from the PPUHWSTAT outputs of the respective PPU, and are only valid if the respective PPU is not making a transition.

12.3.1.13 Host System Lock Control Status (HOST_SYS_LCTRL_ST) register

The following table gives a bit-level description of the Lock Control register.

Table 12-28: HOST_SYS_LCTRL_ST register

Bits	Name	Description	Type	Reset
[31]	LOCK_CLR_DIS	Controls the behavior of the HOST_SYS_LCTRL_CLR register `0b0`: Writes to the HOST_SYS_LCTRL_CLR register take effect `0b1`: Writes to the HOST_SYS_LCTRL_CLR register are ignored	RO	`0b0`
[30:8]	-	Reserved	RO	`0x000000`
[7]	HOST_LOCK	Controls whether registers in the Host Base System Control, which are lockable, are locked or unlocked. `0b0`: Registers are unlocked `0b1`: Registers are locked	RO	`0b0`

Bits	Name	Description	Type	Reset
[6]	HOST_GIC_LOCK	Drives the **CFGSDISABLE** signal of the GICCFG interface. `0b0`: Signal is de-asserted. `0b1`: Signal is asserted. This field is set to `0b0` when CLUSTOP power domain enters on of the following: • OFF • MEM_RET • WARM_RST	RO	`0b0`
[5]	HOST_CPU3_LOCK	Drives the **CP15SDISABLE[3]** signal of the HOSTCPUCFG interface. `0b0`: Signal is de-asserted. `0b1`: Signal is asserted. This field is set to `0b0` when: • CORE3 power domain enters OFF, OFF_EMU or WARM_RST. • CLUSTOP power domain enters OFF, MEM_RET or WARM_RST. This field is Reserved and treated as RAZ/WI when HOST_CPU_NUM_CORE < 4.	RO	`0b0`
[4]	HOST_CPU2_LOCK	Drives the **CP15SDISABLE[2]** signal of the HOSTCPUCFG interface. `0b0`: Signal is de-asserted. `0b1`: Signal is asserted. This field is set to `0b0` when: • CORE2 power domain enters OFF, OFF_EMU, or WARM_RST • CLUSTOP power domain enters OFF, MEM_RET, or WARM_RST This field is Reserved and treated as RAZ/WI when HOST_CPU_NUM_CORE < 3.	RO	`0b0`
[3]	HOST_CPU1_LOCK	Drives the **CP15SDISABLE[1]** signal of the HOSTCPUCFG interface. `0b0`: Signal is de-asserted. `0b1`: Signal is asserted. This field is set to `0b0` when: • CORE1 power domain enters OFF, OFF_EMU, or WARM_RST • CLUSTOP power domain enters OFF, MEM_RET, or WARM_RST This field is Reserved and treated as RAZ/WI when HOST_CPU_NUM_CORE < 2.	RO	`0b0`

Bits	Name	Description	Type	Reset
[2]	HOST_CPU0_LOCK	Drives the **CP15SDISABLE[0]** signal of the HOSTCPUCFG interface. `0b0`: Signal is de-asserted. `0b1`: Signal is asserted. This field is set to `0b0` when: • CORE0 power domain enters OFF, OFF_EMU, or WARM_RST • CLUSTOP power domain enters OFF, MEM_RET, or WARM_RST	RO	`0b0`
[1]	INT_RTR_LOCK	Controls the Interrupt Router Lockdown interface. `0b0`: Interrupt Router Lockdown interface is de-asserted `0b1`: Interrupt Router Lockdown interface is asserted	RO	`0b0`
[0]	HOST_FW_LOCK	Controls the Host Firewall Lockdown interface. `0b0`: Host Firewall Lockdown interface is de-asserted `0b1`: Host Firewall Lockdown interface is asserted	RO	`0b0`


Note

If there is a simultaneous set and clear event, where the clear event can be either a write to the HOST_SYS_LCTRL_CLR register or a condition described in the above table, then the clear event always takes precedence. For information on how software should use the HOST_SYS_LCTRL_{ST/SET/CLR} registers, see 14. Software sequences on page 293

12.3.1.14 Host System Lock Control Set (HOST_SYS_LCTRL_SET) register

The following table gives a bit-level description of the Lock Control Set register.

Table 12-29: HOST_SYS_LCTRL_SET register

Bits	Name	Description	Type	Reset
[31]	LOCK_CLR_DIS	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.LOCK_CLR_DIS field to 1. Writing 0 to this field has no effect on the value of HOST_SYS_LCTRL_ST.LOCK_CLR_DIS field. This field always reads as 0.	WO	`0b0`
[30:8]	-	Reserved	RO	`0x000000`
[7]	HOST_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.HOST_LOCK field to 1. Writing 0 to this field has no effect on the value of HOST_SYS_LCTRL_ST.HOST_LOCK field. This field always reads as 0.	WO	`0b0`

Bits	Name	Description	Type	Reset
[6]	HOST_GIC_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.HOST_GIC_LOCK field to 1. Writing 0 to this field has no effect on the value of HOST_SYS_LCTRL_ST.HOST_GIC_LOCK field. This field always reads as 0.	WO	0b0
[5]	HOST_CPU3_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.HOST_CPU3_LOCK field to 1. Writing 0 to this field has no effect on the value of HOST_SYS_LCTRL_ST.HOST_CPU3_LOCK field. This field always reads as 0.	WO	0b0
[4]	HOST_CPU2_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.HOST_CPU2_LOCK field to 1. Writing 0 to this field has no effect on the value of HOST_SYS_LCTRL_ST.HOST_CPU2_LOCK field. This field always reads as 0.	WO	0b0
[3]	HOST_CPU1_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.HOST_CPU1_LOCK field to 1. Writing 0 to this field has no effect on the value of HOST_SYS_LCTRL_ST.HOST_CPU1_LOCK field. This field always reads as 0.	WO	0b0
[2]	HOST_CPU0_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.HOST_CPU0_LOCK field to 1. Writing 0 to this field has no effect on the value of HOST_SYS_LCTRL_ST.HOST_CPU0_LOCK field. This field always reads as 0.	WO	0b0
[1]	INT_RTR_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.INT_RTR_LOCK field to 1. Writing 0 to this field has no effect on the value of HOST_SYS_LCTRL_ST.INT_RTR_LOCK field. This field always reads as 0.	WO	0b0
[0]	HOST_FW_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.HOST_FW_LOCK field to 1. Writing 0 to this field has no effect on the value of HOST_SYS_LCTRL_ST.HOST_FW_LOCK field. This field always reads as 0.	WO	0b0

12.3.1.15 Host System Lock Control Clear (HOST_SYS_LCTRL_CLR) register

The following table gives a bit-level description of the Lock Control Clear register.

Table 12-30: HOST_SYS_LCTRL_CLR register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x000000
[7]	HOST_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.HOST_LOCK field to 0. Writing 0 to this field has no effect on the value of the HOST_SYS_LCTRL_ST.HOST_LOCK field. This field always reads as 0.	WO	0b0
[6]	HOST_GIC_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.HOST_GIC_LOCK field to 0. Writing 0 to this field has no effect on the value of the HOST_SYS_LCTRL_ST.HOST_GIC_LOCK field. This field always reads as 0.	WO	0b0
[5]	HOST_CPU3_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.HOST_CPU3_LOCK field to 0. Writing 0 to this field has no effect on the value of the HOST_SYS_LCTRL_ST.HOST_CPU3_LOCK field. This field always reads as 0.	WO	0b0
[4]	HOST_CPU2_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.HOST_CPU2_LOCK field to 0. Writing 0 to this field has no effect on the value of the HOST_SYS_LCTRL_ST.HOST_CPU2_LOCK field. This field always reads as 0.	WO	0b0
[3]	HOST_CPU1_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.HOST_CPU1_LOCK field to 0. Writing 0 to this field has no effect on the value of the HOST_SYS_LCTRL_ST.HOST_CPU1_LOCK field. This field always reads as 0.	WO	0b0
[2]	HOST_CPU0_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.HOST_CPU0_LOCK field to 0. Writing 0 to this field has no effect on the value of the HOST_SYS_LCTRL_ST.HOST_CPU0_LOCK field. This field always reads as 0.	WO	0b0
[1]	INT_RTR_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.INT_RTR_LOCK field to 0. Writing 0 to this field has no effect on the value of the HOST_SYS_LCTRL_ST.INT_RTR_LOCK field. This field always reads as 0.	WO	0b0

Bits	Name	Description	Type	Reset
[0]	HOST_FW_LOCK	Writing 1 to this field sets the HOST_SYS_LCTRL_ST.HOST_FW_LOCK field to 0. Writing 0b0 to this field has no effect on the value of the HOST_SYS_LCTRL_ST.HOST_FW_LOCK field. This field always reads as 0.	WO	0b0


Note

Writes to the HOST_SYS_LCTRL_CLR register are ignored when the HOST_SYS_LCTRL_ST.LOCK_CLR_DIS field is set to 1.

12.3.1.16 Host CPU Clock Control (HOSTCPUCLK_CTRL) register

The following table gives a bit-level description of the Host CPU Clock Control register.

Table 12-31: HOSTCPUCLK_CTRL register

Bits	Name	Description	Type	Reset
[31:16]	-	Reserved	RO	0x0000
[15:8]	CLKSELECT_CUR	Currently selected clock source for **HOSTCPUCLK**. Possible field values are: 0x00: Clock gate 0x01: **REFCLK** 0x02: **SYSPLL** 0x04: **CPUPLL** All other values are Reserved.	RO	UNKNOWN
[7:0]	CLKSELECT	Select the clock source for **HOSTCPUCLK**. Possible field values are: 0x00: Clock gate 0x01: **REFCLK** 0x02: **SYSPLL** 0x04: **CPUPLL** All other values are Reserved. Selecting a value which is Reserved can cause a deadlock.	RW	0x01

12.3.1.17 Host CPU Clock Divider 0 (HOSTCPUCLK_DIV0) register

The following table gives a bit-level description of the Host CPU Clock Divider 0 register.

Table 12-32: HOSTCPUCLK_DIV0 register

Bits	Name	Description	Type	Reset
[31:21]	-	Reserved	RO	0x000
[20:16]	CLKDIV_CUR	Current value of integer divider applied to **SYSPLL**. Possible field values are: 0x00: Divide by 1 0x01: Divide by 2 ... 0x0F: Divide by 32	RO	UNKNOWN
[15:5]	-	Reserved	RO	0x000
[4:0]	CLKDIV	Select the value of the integer divider applied to **SYSPLL**. Possible field values are: 0x00: Divide by 1 0x01: Divide by 2 ... 0x1F: Divide by 32	RW	0x00

12.3.1.18 Host CPU Clock Divider 1 (HOSTCPUCLK_DIV1) register

The following table gives a bit-level description of the Host CPU Clock Divider 1 register.

Table 12-33: HOSTCPUCLK_DIV1 register

Bits	Name	Description	Type	Reset
[31:21]	-	Reserved	RO	0x000
[20:16]	CLKDIV_CUR	Current value of integer divider applied to **CPUPLL**. Possible field values are: 0x00: Divide by 1 0x01: Divide by 2 ... 0x0F: Divide by 32	RO	UNKNOWN

Bits	Name	Description	Type	Reset
[15:5]	-	Reserved	RO	0x000
[4:0]	CLKDIV	Select the value of the integer divider applied to **CPUPLL**. Possible field values are: 0x00: Divide by 1 0x01: Divide by 2 ... 0x1F: Divide by 32	RW	0x00

12.3.1.19 GIC Clock Control (GICCLK_CTRL) register

The following table gives a bit-level description of the GIC Clock Control register.

Table 12-34: GICCLK_CTRL register

Bits	Name	Description	Type	Reset
[31:24]	ENTRY_DELAY	Reserved and treated as RAZ/WI	RW	0x00
[23:16]	-	Reserved	RO	0x00
[15:8]	CLKSELECT_CUR	Currently selected clock source for **GICCLK**. Possible field values are: 0x00: Clock gate 0x01: **REFCLK** 0x02: **SYSPLL** All other values are Reserved.	RO	UNKNOWN
[7:0]	CLKSELECT	Select the clock source for **GICCLK**. Possible field values are: 0x000: Clock gate 0x001: **REFCLK** 0x002: **SYSPLL** All other values are Reserved. Selecting a value which is Reserved can cause a deadlock.	RW	0x01

12.3.1.20 GIC Clock Divider 0 (GICCLK_DIV0) register

The following table gives a bit-level description of the GIC Clock Divider 0 register.

Table 12-35: GICCLK_DIV0 register

Bits	Name	Description	Type	Reset
[31:21]	-	Reserved.	RO	0x000
[20:16]	CLKDIV_CUR	Current value of integer divider applied to **SYSPLL**. Possible field values are: 0x00: Divide by 1 0x01: Divide by 2 ... 0x1F: Divide by 32	RO	UNKNOWN
[15:5]	-	Reserved	RO	0x000
[4:0]	CLKDIV	Select the value of the integer divider applied to **SYSPLL**. Possible field values are: 0x00: Divide by 1 0x01: Divide by 2 ... 0x1F: Divide by 32	RW	0x00

12.3.1.21 AXI Clock Control (ACLK_CTRL) register

The following table gives a bit-level description of the AXI Clock Control register.

Table 12-36: ACLK_CTRL register

Bits	Name	Description	Type	Reset
[31:24]	ENTRY_DELAY	Configure number of idle clock cycles before clock is gated	RW	0x00
[23:16]	-	Reserved	RO	0x00

Bits	Name	Description	Type	Reset
[15:8]	CLKSELECT_CUR	Currently selected clock source for **ACLK**. Possible field values are: `0x00`: Clock gate `0x01`: **REFCLK** `0x02`: **SYSPLL** All other values are Reserved.	RO	UNKNOWN
[7:0]	CLKSELECT	Select the clock source for **ACLK**. Possible field values are: `0x00`: Clock gate `0x01`: **REFCLK** `0x02`: **SYSPLL** All other values are Reserved. Selecting a value which is Reserved can cause a deadlock.	RW	`0x01`

12.3.1.22 AXI Clock Divider 0 (ACLK_DIV0) register

The following table gives a bit-level description of the AXI Clock Divider 0 register description.

Table 12-37: ACLK_DIV0 register

Bits	Name	Description	Type	Reset
[31:21]	-	Reserved	RO	`0x000`
[20:16]	CLKDIV_CUR	Current value of integer divider applied to **SYSPLL**. Possible field values are: `0x00`: Divide by 1. `0x01`: Divide by 2. ... `0x1F`: Divide by 32.	RO	UNKNOWN
[15:5]	-	Reserved	RO	`0x000`

Bits	Name	Description	Type	Reset
[4:0]	CLKDIV	Select the value of the integer divider applied to **SYSPLL**. Possible field values are: `0x00`: Divide by 1 `0x01`: Divide by 2 ... `0x1F`: Divide by 32	RW	`0x00`

12.3.1.23 Control Clock Control (CTRLCLK_CTRL) register

The following table gives a bit-level description of the Control Clock Control register.

Table 12-38: CTRLCLK_CTRL register

Bits	Name	Description	Type	Reset
[31:24]	ENTRY_DELAY	Configure number of idle clock cycles before clock is gated	RW	`0x00`
[23:16]	-	Reserved	RO	`0x00`
[15:8]	CLKSELECT_CUR	Currently selected clock source for **CTRLCLK**. Possible field values are: `0x00`: Clock gate `0x01`: **REFCLK** `0x02`: **SYSPLL** All other values are Reserved.	RO	UNKNOWN
[7:0]	CLKSELECT	Select the clock source for **CTRLCLK**. Possible field values are: `0x00`: Clock gate `0x01`: **REFCLK** `0x02`: **SYSPLL** All other values are Reserved. Selecting a value which is Reserved can cause a deadlock.	RW	`0x01`

12.3.1.24 Control Clock Divider 0 (CTRLCLK_DIV0) register

The following table gives a bit-level description of the Control Clock Divider 0 register.

Table 12-39: CTRLCLK_DIV0 register

Bits	Name	Description	Type	Reset
[31:21]	-	Reserved	RO	Reserved
[20:16]	CLKDIV_CUR	Current value of integer divider applied to **SYSPLL**. Possible field values are: `0x00`: Divide by 1 `0x01`: Divide by 2 ... `0x1F`: Divide by 32	RO	UNKNOWN
[15:5]	-	Reserved	RO	`0x000`
[4:0]	CLKDIV	Select the value of the integer divider applied to **SYSPLL**. Possible field values are: `0x00`: Divide by 1 `0x01`: Divide by 2 ... `0x1F`: Divide by 32	RW	`0x00`

12.3.1.25 Debug Clock Control (DBGCLK_CTRL) register

The following table gives a bit-level description of the Debug Clock Control register.

Table 12-40: DBGCLK_CTRL register

Bits	Name	Description	Type	Reset
[31:24]	ENTRY_DELAY	Configure number of idle clock cycles before clock is gated	RW	`0x00`
[23:16]	-	Reserved	RO	`0x00`

Bits	Name	Description	Type	Reset
[15:8]	CLKSELECT_CUR	Currently selected clock source for **DBGCLK**. Possible field values are: `0x00`: Clock gate `0x01`: **REFCLK** `0x02`: **SYSPLL** All other values are Reserved.	RO	UNKNOWN
[7:0]	CLKSELECT	Select the clock source for **DBGCLK**. Possible field values are: `0x00`: Clock gate `0x01`: **REFCLK** `0x02`: **SYSPLL** All other values are Reserved. Selecting a value which is Reserved can cause a deadlock.	RW	`0x01`

12.3.1.26 Debug Clock Divider 0 (DBGCLK_DIV0) register

The following table gives a bit-level description of the Debug Clock Divider 0 register.

Table 12-41: DBGCLK_DIV0 register

Bits	Name	Description	Type	Reset
[31:21]	-	Reserved	RO	`0x000`
[20:16]	CLKDIV_CUR	Current value of integer divider applied to **SYSPLL**. Possible field values are: `0x00`: Divide by 1 `0x01`: Divide by 2 ... `0x1F`: Divide by 32	RO	UNKNOWN
[15:5]	-	Reserved	RO	`0x000`

Bits	Name	Description	Type	Reset
[4:0]	CLKDIV	Select the value of the integer divider applied to **SYSPLL**. Possible field values are: `0x00`: Divide by 1. `0x01`: Divide by 2. ... `0x1F`: Divide by 32.	RW	`0x00`

12.3.1.27 Host UART Clock Control (HOSTUARTCLK_CTRL) register

The following table gives a bit-level description of the Host UART Clock Control register.

Table 12-42: HOSTUARTCLK_CTRL register

Bits	Name	Description	Type	Reset
[31:16]	-	Reserved	RO	`0x0000`
[15:8]	CLKSELECT_CUR	Currently selected clock source for **HOSTUARTCLK**. Possible field values are: `0x00`: Clock gate `0x01`: **REFCLK** `0x02`: **UARTCLK** `0x04`: **S32KCLK** All other values are Reserved.	RO	UNKNOWN
[7:0]	CLKSELECT	Select the clock source for **HOSTUARTCLK**. Possible field values are: `0x00`: Clock gate `0x01`: **REFCLK** `0x02`: **UARTCLK** `0x04`: **S32KCLK** All other values are Reserved. Selecting a value which is Reserved can cause a deadlock.	RW	`0x01`

12.3.1.28 Host UART Clock Divider 0 (HOSTUARTCLK_DIV0) register

The following table gives a bit-level description of the Host UART Clock Divider 0 register.

Table 12-43: HOSTUARTCLK_DIV0 register

Bits	Name	Description	Type	Reset
[31:21]	-	Reserved	RO	0x000
[20:16]	CLKDIV_CUR	Current value of integer divider applied to **UARTCLK**. Possible field values are: 0x00: Divide by 1 0x01: Divide by 2 ... 0x1F: Divide by 32	RO	UNKNOWN
[15:5]	-	Reserved	RO	0x000
[4:0]	CLKDIV	Select the value of the integer divider applied to **UARTCLK**. Possible field values are: 0x00: Divide by 1 0x01: Divide by 2 ... 0x1F: Divide by 32	RW	0x00

12.3.1.29 REFCLK Clock Control register

The following table gives a bit-level description of the REFCLK Clock Control register.

Table 12-44: REFLCK register

Bits	Name	Description	Type	Reset
[31:24]	ENTRY_DELAY	Configure number of idle clock cycles before clock is gated	RW	0x00
[23:16]	-	Reserved	RO	0x00
[15:8]	CLKSELECT_CUR	Currently selected clock source for **DBGCLK**. 0x01: **REFCLK** All other values are Reserved.	RO	UNKNOWN
[7:0]	CLKSELECT	Select the clock source for **DBGCLK:** 0x01: **REFCLK** All other values are Reserved.	RO	0x01

12.3.1.30 Clock Force Status (CLKFORCE_ST) register

The following table gives a bit-level description of the Clock Force Status register.

Table 12-45: CLKFORCE_ST register

Bits	Name	Description	Type	Reset
[31:7]	-	Reserved	RO	0x0000000
[6]	REFCLK_FORCE_ST	Status of **REFCLK** clock force. 0b0: High-level clock gating is enabled. 0b1: High-level clock gating is disabled.	RO	0b0
[5]	-	Reserved	RO	0b0
[4]	DBGCLK_FORCE_ST	Status of **DBGCLK** clock force. Possible field values are: 0b0: High-level clock gating is enabled. 0b1 – High-level clock gating is disabled.	RO	0b0
[3]	CTRLCLK_FORCE_ST	Status of **CTRLCLK** clock force. Possible field values are: 0b0: High-level clock gating is enabled. 0b1: High-level clock gating is disabled.	RO	0b0
[2]	ACLK_FORCE_ST	Status of **ACLK** clock force. Possible field values are: 0b0: High-level clock gating is enabled. 0b1: High-level clock gating is disabled.	RO	0b0
[1]	GICCLK_FORCE_ST	Reserved and treated as RAZ/WI	RO	0b0
[0]	-	Reserved	RO	0b0

12.3.1.31 Clock Force Set (CLKFORCE_SET) register

The following table gives a bit-level description of the Clock Force Set register.

Table 12-46: CLKFORCE_SET register

Bits	Name	Description	Type	Reset
[31:7]	-	Reserved	RO	0x0000000

Bits	Name	Description	Type	Reset
[6]	REFCLK_FORCE_SET	Set REFCLKCLK_FORCE_ST in CLKFORCE_ST register. Writing a value of 0b0 has no effect. This field always reads as 0b0.	WO	0b0
[5]	-	Reserved	RO	0b0
[4]	DBGCLK_FORCE_SET	Set DBGCLK_FORCE_ST in CLKFORCE_ST register. Writing a value of0b 0 has no effect. This field always reads as0b 0.	WO	0b0
[3]	CTRLCLK_FORCE_SET	Set CTRLCLK_FORCE_ST in CLKFORCE_ST register. Writing a value of 0b0 has no effect. This field always reads as 0b0.	WO	0b0
[2]	ACLK_FORCE_SET	Set ACLK_FORCE_ST in CLKFORCE_ST register. Writing a value of 0b0 has no effect. This field always reads as 0b0.	WO	0b0
[1]	GICCLK_FORCE_SET	Reserved and treated as RAZ/WI	WO	0b0
[0]	-	Reserved	WO	0b0

12.3.1.32 Clock Force Clear (CLKFORCE_CLR) register

The following table gives a bit-level description of the Clock Force Clear register description.

Table 12-47: CLKFORCE_CLR register

Bits	Name	Description	Type	Reset
[31:7]	-	Reserved	RO	0x0000000
[6]	REFCLK_FORCE_CLR	Clear REFCLK_FORCE_ST in CLKFORCE_ST register. Writing a value of 0b0 has no effect. This field always reads as 0b0.	WO	0b0
[5]	-	Reserved	RO	0b0
[4]	DBGCLK_FORCE_CLR	Clear DBGCLK_FORCE_ST in CLKFORCE_ST register. Writing a value of 0b0 has no effect. This field always reads as 0b0.	WO	0b0
[3]	CTRLCLK_FORCE_CLR	Clear CTRLCLK_FORCE_ST in CLKFORCE_ST register. Writing a value of 0b0 has no effect. This field always reads as 0b0.	WO	0b0

Bits	Name	Description	Type	Reset
[2]	ACLK_FORCE_CLR	Clear ACLK_FORCE_ST in CLKFORCE_ST register. Writing a value of `0b0` has no effect. This field always reads as `0b0`.	WO	`0b0`
[1]	GICCLK_FORCE_CLR	Reserved and treated as RAZ/WI	WO	`0b0`
[0]	-	Reserved	RO	`0b0`

12.3.1.33 PLL Status (PLL_ST) register

The following table gives a bit-level description of the Phase Locked Loop Status register.

Table 12-48: PLL_ST register

Bits	Name	Description	Type	Reset
[31:2]	-	Reserved	RO	`0x0000000`
[1]	CPUPLLLOCK_ST	Status of the **CPUPLLLOCK** input. Possible field values are: 0: **PLL** is not locked. 1: **PLL** is locked.	RO	UNKNOWN
[0]	SYSPLLLOCK_ST	Status of the **SYSPLLLOCK** input. Possible field values are: 0: **PLL** is not locked. 1: **PLL** is locked.	RO	UNKNOWN

12.3.1.34 Host PPU Interrupt Status (HOST_PPU_INT_ST) register

The following table gives a bit-level description of the PPU Interrupt Status register.

Table 12-49: HOST_PPU_INT_ST register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved.	RO	`0x000000`
[7]	CORE3_INT_ST	Status of CORE3 PPU Interrupt. Possible field values are: • 0: Interrupt is deasserted • 1: Interrupt is asserted This field is Reserved and treated as RAZ/WI when HOST_CPU_NUM_CORE < 4.	RO	UNKNOWN

Bits	Name	Description	Type	Reset
[6]	CORE2_INT_ST	Status of CORE2 PPU Interrupt. Possible field values are: • 0: Interrupt is deasserted • 1: Interrupt is asserted This field is Reserved and treated as RAZ/WI when HOST_CPU_NUM_CORE < 3.	RO	UNKNOWN
[5]	CORE1_INT_ST	Status of CORE1 PPU Interrupt. Possible field values are: • 0: Interrupt is deasserted • 1: Interrupt is asserted This field is Reserved and treated as RAZ/WI when HOST_CPU_NUM_CORE < 2.	RO	UNKNOWN
[4]	CORE0_INT_ST	Status of CORE0 PPU Interrupt. Possible field values are: • 0: Interrupt is deasserted • 1: Interrupt is asserted	RO	UNKNOWN
[3]	CLUSTOP_INT_ST	Status of CLUSTOP PPU Interrupt. Possible field values are: • 0: Interrupt is deasserted • 1: Interrupt is asserted	RO	UNKNOWN
[2]	SYSTOP_INT_ST	Status of SYSTOP PPU Interrupt. Possible field values are: • 0: Interrupt is deasserted • 1: Interrupt is asserted	RO	UNKNOWN
[1]	DBGTOP_INT_ST	Status of the DBGTOP PPU Interrupt. Possible field values are: • 0: Interrupt is deasserted • 1: Interrupt is asserted	RO	UNKNOWN
[0]	FW_INT_ST	Status of the Firewall PPU Interrupt. Possible field values are: • 0: Interrupt is de-asserted • 1: Interrupt is asserted	RO	UNKNOWN

When any bit in this register is 1, the PPU Combined interrupt is asserted.

12.3.1.35 Peripheral ID 4 (PID4) register

The following table gives a bit-level description of the Peripheral ID 4 register.

Table 12-50: PID4 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x000000
[7:4]	Size	Number of 4KB blocks occupied by the System ID block. This field is deprecated.	RO	0x0
[3:0]	DES_2	JEP Continuation	RO	0x4

12.3.1.36 Peripheral ID 5 (PID5) register

The following table gives a bit-level description of the Peripheral ID 5 register.

Table 12-51: PID5 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x00000000

12.3.1.37 Peripheral ID 6 (PID6) register

The following table gives a bit-level description of the Peripheral ID 6 register.

Table 12-52: PID6 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x00000000

12.3.1.38 Peripheral ID 7 (PID7) register

The following table gives a bit-level description of the Peripheral ID 7 register.

Table 12-53: PID7 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x00000000

12.3.1.39 Peripheral ID 0 (PID0) register

The following table gives a bit-level description of the Peripheral ID 0 register.

Table 12-54: PID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x000000
[7:0]	PART_0	Bits [7:0] of part ID	RO	0x78

12.3.1.40 Peripheral ID 1 (PID1) register

The following table gives a bit-level description of the Peripheral ID 1 register.

Table 12-55: PID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x000000
[7:4]	DES_0	Bits [3:0] of JEP 106 Identity	RO	0xB
[3:0]	PART_1	Bits [11:8] of part ID	RO	0x0

12.3.1.41 Peripheral ID 2 (PID2) register

The following table gives a bit-level description of the Peripheral ID 2 register.

Table 12-56: PID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x000000
[7:4]	REVISION	Major revision of the System ID block	RO	0x0b0
[3]	JEDEC	Indicates the use of JEDEC JEP106 identification scheme	RO	1
[2:0]	DES_1	Bits [6:4] of JEP 106 Identity	RO	0b011

12.3.1.42 Peripheral ID 3 (PID3) register

The following table gives a bit-level description of the Peripheral ID 3 register.

Table 12-57: PID3 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x000000
[7:4]	REVAND	Minor revision of the System ID block	RO	0x0
[3:0]	CMOD	Customer modification field	RO	0x0

12.3.1.43 Component ID 0 (CID0) register

The following table gives a bit-level description of the Component ID 0 register.

Table 12-58: CID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x000000
[7:0]	PRMBL_0	Preamble 0	RO	0x0D

12.3.1.44 Component ID 1 (CID1) register

The following table gives a bit-level description of the Component ID 1 register.

Table 12-59: CID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x000000
[7:4]	CLASS	Class of the component	RO	0xF
[3:0]	PRMBL_1	Preamble 1	RO	0x0

12.3.1.45 Component ID 2 (CID2) register

The following table gives a bit-level description of the Component ID 2 register.

Table 12-60: CID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x000000
[7:0]	PRMBL_2	Preamble 2	RO	0x05

12.3.1.46 Component ID 3 (CID3) register

The following table gives a bit-level description of the Component ID 3 register.

Table 12-61: CID3 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x000000
[7:0]	PRMBL_3	Preamble 3	RO	0xB1

12.3.2 Secure Enclave Registers

This section summarizes the Secure Enclave Base and System Control registers.

12.3.2.1 Secure Enclave Base System Control register summary

This section summarizes the Secure Enclave Base System Control registers.

The Secure Enclave Base System Control registers are in the Secure Enclave Peripheral Map. They are only accessible by the Secure Enclave Cortex®-M0+ processor. The Secure Enclave Base System Control registers are accessed at offset `0x0x5008_E000` in the Secure Enclave memory map. These registers are in the Always ON power domain.

Table 12-62: Secure Enclave Base System Control register summary

Offset	Short name	Access	Name
0x000	HOST_SYS_RST_CTRL	RW	Host System Reset Control
0x004	HOST_SYS_RST_ST	RO	Host System Reset Status
0x008	SOC_RST_CTRL	RW	SoC Reset Control
0x00C	SOC_RST_SYN	RO	SoC Reset Syndrome
0x010	SEC_ENC_INT_COL_ST0	RO	Secure Enclave Interrupt Collator Status 0
0x014	SEC_ENC_INT_COL_ST1	RO	Secure Enclave Interrupt Collator Status 1
0x018	SEC_ENC_INT_COL_ST2	RO	Secure Enclave Interrupt Collator Status 2
0x01C	SEC_ENC_INT_COL_ST3	RO	Secure Enclave Interrupt Collator Status 3
0x020	SEC_ENC_INT_COL_MSK0	RW	Secure Enclave Interrupt Collator Mask 0
0x024	SEC_ENC_INT_COL_MSK1	RW	Secure Enclave Interrupt Collator Mask 1
0x028	SEC_ENC_INT_COL_MSK2	RW	Secure Enclave Interrupt Collator Mask 2
0x02C	SEC_ENC_INT_COL_MSK3	RO	Secure Enclave Interrupt Collator Mask 3
0x030 – 0x3FC	-	RO	Reserved
0x400	BSYS_PWR_REQ	RW	Base System Power Request
0x404	BSYS_PWR_ST	RO	Base System Power Status
0x408 – 0x7FC	-	RO	Reserved
0x800	SECENCCLK_CTRL	RW	**SECENCCLK** Control
0x804	SECENCCLK_DIV	RW	**SECENCCLK** Divider
0x808 – 0x9FC	-	RO	Reserved
0xA00	CLKFORCE_ST	RO	Clock Force Status
0xA04	CLKFORCE_SET	WO	Clock Force Set
0xA08	CLKFORCE_CLR	WO	Clock Force Clear
0xA10	SEC_ENC_PLL_ST	RO	PLL Status
0xA14 – 0xFCC	-	RO	Reserved
0xFD0	PID4	RO	Peripheral ID4
0xFD4	PID5	RO	Peripheral ID5
0xFD8	PID6	RO	Peripheral ID6
0xFDC	PID7	RO	Peripheral ID7
0xFE0	PID0	RO	Peripheral ID0
0xFE4	PID1	RO	Peripheral ID1
0xFE8	PID2	RO	Peripheral ID2
0xFEC	PID3	RO	Peripheral ID3

Offset	Short name	Access	Name
0xFF0	CID0	RO	Component ID0
0xFF4	CID1	RO	Component ID1
0xFF8	CID2	RO	Component ID2
0xFFC	CID3	RO	Component ID3

The Secure Enclave Base System Control registers support only 32-bit word aligned accesses. Access by other bit sizes or unaligned accesses are treated as 32-bit word aligned access without generating an error response.

The Secure Enclave Base System Control registers have the following behavior:

- Any read to a write-only register is treated as RAZ.
- Any write to a read-only register is treated as WI.

In both cases no error is generated.

12.3.2.1.1 Host System Reset Control (HOST_SYS_RST_CTRL) register

The following table gives a bit-level description of the HOST_SYS_RST_CTRL register.

Table 12-63: HOST_SYS_RST_CTRL register

Bits	Name	Description	Type	Reset
[31:2]	-	Reserved	RO	0x0000_0000
[1]	RST_REQ	Reset request for Host System 0b0: No reset requested 0b1: Reset requested	RW	0b0
[0]	CPUWAIT	CPU Wait control 0b0: Host System's CPUWAIT signal is de-asserted. 0b1: Host System's CPUWAIT signal is asserted. When CPUWAIT becomes 0b0 attempts to set it back to 0b1 are ignored. This field only returns to 0b1 when any of the following occur: • External Power on Reset • Internal Power on Reset • Debug reset • Host System reset A request to set this field 0b0 at the same time as the field is to be set to 0b1 results in the field being set to 0b1. This field becomes RO when HOST_CPUWAIT_WEN is LOW. Any attempt to set this field to 0b0 by software is ignored.	RW	0b1

12.3.2.1.2 Host System Reset Status (HOST_SYS_RST_ST) register

The following table gives a bit-level description of the HOST_SYS_RST_ST register.

Table 12-64: HOST_SYS_RST_ST register

Bits	Name	Description	Type	Reset
[31:3]	-	Reserved	RO	0x00_0000
[2:1]	RST_ACK	Status of reset request 0b00 – No reset requested 0b01 – Reset request unable to complete 0b10 – Reset request complete 0b11 – Reserved	RO	0b00
[0]	-	Reserved	RO	0b0

12.3.2.1.3 SoC Reset Control (SOC_RST_CTRL) register

The following table gives a bit-level description of the SOC_RST_CTRL register.

Table 12-65: SOC_RST_CTRL register

Bits	Name	Description	Type	Reset
[31:2]	-	Reserved	RO	0x0000_0000
[1]	RST_REQ	Reset request for SoC. 0b0: No reset requested 0b1: Reset requested	RW	0b0
[0]	-	Reserved	RO	0b0

12.3.2.1.4 Secure Enclave Interrupt Collator Status 0 (SEC_ENC_INT_COL_ST0) register

The following table gives a bit-level description of the SEC_ENC_INT_COL_ST0 register.

Table 12-66: SEC_ENC_INT_COL_ST0 register

Bits	Name	Description	Type	Reset
[31:0]	SSI_ST	Status of the Secure Enclave Expansion Interrupt (SEEI) {0-31} after the respective mask field from SEC_ENC_INT_COL_MSK0 register has been applied. SEEI0 assigned to bit[0]. 0b0: Interrupt is de-asserted 0b1: Interrupt is asserted	RO	UNKNOWN

12.3.2.1.5 Secure Enclave Interrupt Collator Status 1 (SEC_ENC_INT_COL_ST1) register

The following table gives a bit-level description of the SEC_ENC_INT_COL_ST1 register.

Table 12-67: SEC_ENC_INT_COL_ST1 register

Bits	Name	Description	Type	Reset
[31:0]	SSI_ST	Status of the Secure Enclave Expansion Interrupt (SEEI) {32-63} after the respective mask field from SEC_ENC_INT_COL_MSK1 register has been applied. SEEI32 assigned to bit[0]: `0b0`: Interrupt is de-asserted `0b1`: Interrupt is asserted	RO	UNKNOWN

12.3.2.1.6 Secure Enclave Interrupt Collator Status 2 (SEC_ENC_INT_COL_ST2) register

The following table gives a bit-level description of the SEC_ENC_INT_COL_ST2 register.

Table 12-68: SEC_ENC_INT_COL_ST2 register

Bits	Name	Description	Type	Reset
[31:0]	SSI_ST	Status of the Secure Enclave Expansion Interrupt (SEEI) {64-95} after the respective mask field from SEC_ENC_INT_COL_MSK2 register has been applied. SEEI64 assigned to bit[0]: `0b0`: Interrupt is de-asserted `0b1`: Interrupt is asserted	RO	UNKNOWN

12.3.2.1.7 Secure Enclave Interrupt Collator Status 3 (SEC_ENC_INT_COL_ST3) register

The following table gives a bit-level description of the SEC_ENC_INT_COL_ST3 register.

Table 12-69: SEC_ENC_INT_COL_ST3 Register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	`0x0000_0000`

12.3.2.1.8 Secure Enclave Interrupt Collator Mask 0 (SEC_ENC_INT_COL_MSK0) register

The following table gives a bit-level description of the SEC_ENC_INT_COL_MSK0 register.

Table 12-70: SEC_ENC_INT_COL_MSK0 register

Bits	Name	Description	Type	Reset
[31:0]	SSI_MSK	Configures whether Secure Enclave Expansion Interrupt (SEEI) {0-31} generates an interrupt to the Secure Enclave Cortex®-M0+ core, with SEEI0 assigned to bit[0]: `0b0`: Interrupt is unmasked `0b1`: Interrupt is masked	RW	`0x0000_0000`

12.3.2.1.9 Secure Enclave Interrupt Collator Mask 1 (SEC_ENC_INT_COL_MSK1) register

The following table gives a bit-level description of the SEC_ENC_INT_COL_MSK1 register.

Table 12-71: SEC_ENC_INT_COL_MSK1 register

Bits	Name	Description	Type	Reset
[31:0]	SSI_MSK	Configures whether Secure Enclave Expansion Interrupt (SEEI) {32-63} generates an interrupt to the Secure Enclave Cortex®-M0+ core, with SEEI32 assigned to bit[0]: `0b0`: Interrupt is unmasked `0b1`: Interrupt is masked	RW	`0x0000_0000`

12.3.2.1.10 Secure Enclave Interrupt Collator Mask 2 (SEC_ENC_INT_COL_MSK2) register

The following table gives a bit-level description of the SEC_ENC_INT_COL_MSK2 register.

Table 12-72: SEC_ENC_INT_COL_MSK2 register

Bits	Name	Description	Type	Reset
[31:0]	SSI_MSK	Configures whether Secure Enclave Expansion Interrupt (SEEI) {64-95} generates an interrupt to the Secure Enclave Cortex®-M0+ core, with SEEI64 assigned to bit[0]: `0b0`: Interrupt is unmasked `0b1`: Interrupt is masked	RW	`0x0000_0000`

12.3.2.1.11 Secure Enclave Interrupt Collator Mask 3 (SEC_ENC_INT_COL_MSK3) register

The following table gives a bit-level description of the SEC_ENC_INT_COL_MSK3 register

Table 12-73: SEC_ENC_INT_COL_MSK3 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	`0x0000_0000`

12.3.2.1.12 Base System Power Request (BSYS_PWR_REQ) register

The following table gives a bit-level description of the BSYS_PWR_REQ register.

Table 12-74: BSYS_PWR_REQ register

Bits	Name	Description	Type	Reset
[31:6]	-	Reserved	RO	`0x0000_0000`

Bits	Name	Description	Type	Reset
[5:3]	SYSTOP_PWR_REQ	Select SYSTOP power domain behaviour when no activity in the domain: `0b000` – No request for logic or volatile memory to be powered. `0b001` – No request for logic to be powered, but volatile memory must be retained. `0b01x` – Request for logic to be powered, but volatile memory can be either powered or retained. `0b1xx` – Request for logic and volatile memory to be powered.	RW	`0b000`
[2]	DBGTOP_PWR_REQ	Select DBGTOP power domain behaviour when no activity in the domain: `0b0` – No request for DBGTOP to be powered `0b1` – Request for DBGTOP to be powered	RW	`0b0`
[1]	REFCLK_REQ	Request **REFCLK**: `0b0` – No request for **REFCLK** to be supplied `0b1` – Request for **REFCLK** to be supplied	RW	`0b0`
[0]	WAKEUP_EN	Secure Enclave wakeup enable: `0b0` – Wakeup for Secure Enclave is disabled `0b1` – Wakeup for Secure Enclave is enabled	RW	`0b0`

12.3.2.1.13 Base System Power Status (BSYS_PWR_ST) register

The following table gives a bit-level description of the BSYS_PWR_ST register.

Table 12-75: BSYS_PWR_ST register

Bits	Name	Description	Type	Reset
[31:6]	-	Reserved	RO	`0x0000_0000`
[5:3]	SYSTOP_PWR_ST	SYSTOP power domain status: `0b000`: SYSTOP is in the OFF or WARM_RST power mode `0b001`: SYSTOP is in the MEM_RET power mode `0b010`: SYSTOP is in the FUNC_RET power mode `0b100`: SYSTOP is in the ON power mode All other values are Reserved	RO	UNKNOWN

Bits	Name	Description	Type	Reset
[2]	DBGTOP_PWR_ST	DBGTOP power domain status: 0b0: DBGTOP is in the OFF or WARM_RST power mode 0b1: DBGTOP is in the ON-power mode	RO	UNKNOWN
[1]	-	Reserved	RO	0b0
[0]	-	Reserved	RO	0b0

The values in these registers are driven from the **PPUHWSTAT** outputs of the respective PPU. The values are only valid if the respective PPU is not making a transition. For more information on the PPU, see the *Arm® Power Policy Unit Architecture Specification, version 1.1*.

12.3.2.1.14 SECENCCLK Clock Control (SECENCCLK_CTRL) register

The following table gives a bit-level description of the SECENCCLK_CTRL register.

Table 12-76: SECENCCLK_CTRL register

Bits	Name	Description	Type	Reset
[31:24]	ENTRY_DELAY	Configure number of idle clock cycles before clock is gated	RW	0x00
[23:16]	-	Reserved	RO	0x00
[15:8]	CLKSELECT_CUR	Currently selected clock source for **SECENCCLK**: 0x00: Clock gate 0x01: **SECENCREFCLK** 0x02: **SYSPLL** All other values are Reserved.	RO	UNKNOWN
[7:0]	**CLKSELECT**	Select the clock source for **SECENCCLK**: 0x00: Clock gate 0x01: **SECENCREFCLK** 0x02: **SYSPLL** All other values are Reserved. Selecting a value which is Reserved can cause a deadlock.	RW	0x01

12.3.2.1.15 SECENCCLK Clock Divider (SECENCCLK_DIV) register

The following table gives a bit-level description of the SECENCCLK_DIV register.

Table 12-77: SECENCCLK_DIV register

Bits	Name	Description	Type	Reset
[31:21]	-	Reserved	RO	0x000
[20:16]	CLKDIV_CUR	Current value of integer divider applied to **SYSPLL**: 0x00: Divide by 1 0x01: Divide by 2 ... 0x1F: Divide by 32	RO	UNKNOWN
[15:5]	-	Reserved	RO	0x000
[4:0]	CLKDIV	Select the value of the integer divider applied to **SYSPLL**: 0x00: Divide by 1 0x01: Divide by 2 ... 0x1F: Divide by 32	RW	0x00

12.3.2.1.16 Clock Force Status (CLKFORCE_ST) register

The following table gives a bit-level description of the CLKFORCE_ST register.

Table 12-78: CLKFORCE_ST register

Bits	Name	Description	Type	Default
[31:1]	-	Reserved	RO	0x0000_0000
[0]	SECENCCLK_FORCE_ST	Status of **SECENCCLK** clock force: 0b0: High-level clock gating is enabled 0b1: High-level clock gating is disabled	RO	0b0

The SECENCCLK_FORCE_ST field applies to both high-level clock gating on **SECENCCLK** and **SECENCDIVCLK**.

12.3.2.1.17 Clock Force Set (CLKFORCE_SET) register

The following table gives a bit-level description of the CLKFORCE_SET register.

Table 12-79: CLKFORCE_SET register

Bits	Name	Description	Type	Reset
[31:1]	-	Reserved	RO	0x0000_0000
[0]	SECENCCLK_FORCE_SET	Set SECENC_FORCE_ST in CLKFORCE_ST register. Writing a value of 0b0 has no effect. This field always reads as 0b0.	WO	0b0

12.3.2.1.18 Clock Force Clear (CLKFORCE_CLR) register

The following table gives a bit-level description of the CLKFORCE_CLR register.

Table 12-80: CLKFORCE_CLR register

Bits	Name	Description	Type	Reset
[31:1]	-	Reserved	RO	0x0000_0000
[0]	SECENCCLK_FORCE_CLR	Clear SECENC_FORCE_ST in CLKFORCE_ST register. Writing a value of 0b0 has no effect. This field always reads as 0b0.	WO	0b0

12.3.2.1.19 Secure Enclave Phase Locked Loop (PLL) Status (SEC_ENC_PLL_ST) register

The following table gives a bit-level description of the SEC_ENC_PLL_ST register.

Table 12-81: SEC_ENC_PLL_ST register

Bits	Name	Description	Type	Reset
[31:1]	-	Reserved	RO	0x0000_0000
[0]	**SYSPLLLOCK_ST**	Status of the **SYSPLLLOCK** input. 0b0: PLL is not locked 0b1: PLL is locked	RO	UNKNOWN

12.3.2.1.20 Peripheral ID 4 (PID4) register

The following table gives a bit-level description of the Peripheral ID 4 register.

Table 12-82: PID4 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	Size	Number of 4KB occupied by the System ID block. This field is deprecated.	RO	0x0
[3:0]	DES_2	JEP Continuation	RO	0x4

12.3.2.1.21 Peripheral ID 5 (PID5) register

The following table gives a bit-level description of the Peripheral ID 5 register.

Table 12-83: PID5 fegister

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.2.1.22 Peripheral ID 6 (PID6) register

The following table gives a bit-level description of the Peripheral ID 6 register.

Table 12-84: PID6 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.2.1.23 Peripheral ID 7 (PID7) register

The following table gives a bit-level description of the Peripheral ID 7 register.

Table 12-85: PID7 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.2.1.24 Peripheral ID 0 (PID0) register

The following table gives a bit-level description of the Peripheral ID 0 register.

Table 12-86: PID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000

Bits	Name	Description	Type	Reset
[7:0]	PART_0	Bits [7:0] of part ID	RO	0x77

12.3.2.1.25 Peripheral ID 1 (PID1) register

The following table gives a bit-level description of the Peripheral ID 1 register.

Table 12-87: PID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	DES_0	Bits [3:0] of JEP 106 Identity	RO	0xB
[3:0]	PART_1	Bits [11:8] of part ID	RO	0x0

12.3.2.1.26 Peripheral ID 2 (PID2) register

The following table gives a bit-level description of the Peripheral ID 2 register.

Table 12-88: PID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVISION	Major revision of the System ID block.	RO	0x0
[3]	JEDEC	Indicates the use of JEDEC JEP106 identification scheme.	RO	0b1
[2:0]	DES_1	Bits [6:4] of JEP 106 Identity	RO	0b011

12.3.2.1.27 Peripheral ID 3 (PID3) register

The following table gives a bit-level description of the Peripheral ID 3 register.

Table 12-89: PID3 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVAND	Minor revision of the System ID block	RO	0x0
[3:0]	CMOD	Customer modification field	RO	0x0

12.3.2.1.28 Component ID 0 (CID0) register

The following table gives a bit-level description of the Component ID 0 register.

Table 12-90: CID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000

Bits	Name	Description	Type	Reset
[7:0]	PRMBL_0	Preamble 0	RO	0x0D

12.3.2.1.29 Component ID 1 (CID1) register

The following table gives a bit-level description of the Component ID 1 register.

Table 12-91: CID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	CLASS	Class of the component	RO	0xF
[3:0]	PRMBL_1	Preamble 1	RO	0x0

12.3.2.1.30 Component ID 2 (CID2) register

The following table gives a bit-level description of the Component ID 2 register.

Table 12-92: CID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_2	Preamble 2	RO	0x05

12.3.2.1.31 Component ID 3 (CID3) register

The following table gives a bit-level description of the Component ID 3 register.

Table 12-93: CID3 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_3	Preamble 3	RO	0xB1

12.3.2.1.32 SoC Reset Syndrome (SOC_RST_SYN) register

The following table gives a bit-level description of the SOC_RST_SYN register.

Table 12-94: SOC_RST_SYN register

Bits	Name	Description	Type	Reset
[31]	SEC_ENC	Indicates the source of the last reset is captured in the SE_RST_SYN register in the Secure Enclave System Control registers	RO	UNKNOWN
[30:4]	-	Reserved	RO	0x000_0000
[3]	SOC_WDOG	Indicates the last reset of the SoC was caused by the SoC Watchdog	RO	UNKNOWN
[2]	-	Reserved	RO	0x000_0000

Bits	Name	Description	Type	Reset
[1]	nSRST	Indicates that the last reset of the SoC was caused by either: • **nSRST** pin being asserted • DP ROM **CSYSRSTREQ** being asserted	RO	UNKNOWN
[0]	POR	Indicates that the last reset of the SoC was caused by either: • **PORESETn** pin being asserted • DP **CDBGRSTREQ** being asserted • SOC_RST_CTRL.RST_REQ bit set to `0b1`	RO	UNKNOWN

12.3.2.2 Secure Enclave System Control register summary

This section summarizes the Secure Enclave System Control registers.

The Secure Enclave System Control Registers are in the Secure Enclave Memory Map. These registers are only accessible to the Secure Enclave Cortex®-M0+processor. The Secure EnclaveSystem Control registers are accessed at offset `0x5008_0000` in the Secure Enclave memory map. These registers are in the Always ON power domain.

Table 12-95: Secure Enclave System Control register summary

Offset	Short name	Access	Name
`0x000`	SE_RST_SYN	RO	Secure Enclave Reset Syndrome
`0x004`	12.3.2.2.2 Secure Enclave Reset Mask (SE_RST_MSK) register on page 251	RW	Secure Enclave Reset Mask
`0x008`	SE_PWR_CTRL	RW	Secure Enclave Power Control
`0x00C – 0x010`	-	RO	Reserved
`0x014`	SE_GP0	RW	Secure Enclave General Purpose 0
`0x018`	SE_GP1	RW	Secure Enclave General Purpose 1
`0x01C`	SE_GP2	RW	Secure Enclave General Purpose 2
`0x020`	SE_GP3	RW	Secure Enclave General Purpose 3
`0x024 -0x02C`	-	RO	Reserved
`0x030`	SE_CLK_DIV	RW	Secure Enclave Clock Divider Control
`0x034 -0x0F4`	-	RO	Reserved
`0x0F8`	SE_BLD_CFG	RO	Secure Enclave Build Configuration
`0x0FC - 0xFCC`	-	RO	Reserved
`0xFD0`	12.3.2.1.20 Peripheral ID 4 (PID4) register on page 245	RO	Peripheral ID4
`0xFD4`	12.3.2.1.21 Peripheral ID 5 (PID5) register on page 246	RO	Peripheral ID5
`0xFD8`	12.3.2.1.22 Peripheral ID 6 (PID6) register on page 246	RO	Peripheral ID6
`0xFDC`	12.3.2.1.23 Peripheral ID 7 (PID7) register on page 246	RO	Peripheral ID7
`0xFE0`	12.3.2.1.24 Peripheral ID 0 (PID0) register on page 246	RO	Peripheral ID0
`0xFE4`	12.3.2.1.25 Peripheral ID 1 (PID1) register on page 247	RO	Peripheral ID1

Offset	Short name	Access	Name
0xFE8	12.3.2.1.26 Peripheral ID 2 (PID2) register on page 247	RO	Peripheral ID2
0xFEC	12.3.2.1.27 Peripheral ID 3 (PID3) register on page 247	RO	Peripheral ID3
0xFF0	12.3.2.1.28 Component ID 0 (CID0) register on page 247	RO	Component ID0
0xFF4	12.3.2.1.29 Component ID 1 (CID1) register on page 248	RO	Component ID1
0xFF8	12.3.2.1.30 Component ID 2 (CID2) register on page 248	RO	Component ID2
0xFFC	12.3.2.1.31 Component ID 3 (CID3) register on page 248	RO	Component ID3

The Secure Enclave System Control registers only support 32-bit word aligned accesses. Access by other bit sizes or unaligned accesses, are treated as 32-bit word aligned accesses without generating an error reponse.

The Secure Enclave System Control registers have the following behavior:

- Any read to a write-only register is treated as RAZ.
- Any write to a read-only register is treated as WI.

In both cases no error is generated.

12.3.2.2.1 Secure Enclave Reset Syndrome (SE_RST_SYN) register

The following table gives a bit-level description of the Secure Enclave Reset Syndrome register for Secure Enclave internal sources.

Table 12-96: SE_RST_SYN register

Bits	Name	Description	Type	Reset
[31:3]	-	Reserved	RO	0x0000_0000
[2]	CA_ERR	Indicates the last reset of the Secure Enclave was caused by an error of the Crypto Accelerator or not: • 0b0: Last reset of the Secure Enclave was not caused by an error of the Crypto Accelerator. • 0b1: Last reset of the Secure Enclave was caused by an error of the Crypto Accelerator.	RO	0b0
[1]	WD_RESET	Indicates that the last reset cause of the Secure Enclave was caused by the Secure Enclave Watchdog or not: • 0b0: Last reset of the Secure Enclave was not caused by the Secure Enclave Watchdog. • 0b1: Last reset of the Secure Enclave was caused by the Secure Enclave Watchdog.	RO	0b0
[0]	SW_RESET	Indicates that the last reset cause of the Secure Enclave was caused by the Secure Enclave software reset request: • 0b0: Last reset of the Secure Enclave was not caused by the Secure Enclave software reset request. • 0b1: Last reset of the Secure Enclave was caused by the Secure Enclave software reset request.	RO	0b0


Note

This register must be used in conjunction with the SOC_RST_SYN register in the Secure Enclave Base System Control registers.

12.3.2.2.2 Secure Enclave Reset Mask (SE_RST_MSK) register

The following table gives a bit-level description of the Secure Enclave Reset Mask register for Secure Enclave internal sources.

Table 12-97: SE_RST_MSK register

Bits	Name	Description	Type	Reset
[31:1]	-	Reserved	RO	`0x0000_0000`
[2]	CA_ERR_MSK	Controls whether an error of the Crypto Accelerator generates an IPoR: • `0b0`: An error of the Crypto Accelerator generates an IPoR. • `0b1`: An error of the Crypto Accelerator does not generate an IPoR Software must only change the value of this field when there is no error of the Crypto Accelerator outstanding, otherwise it is **UNPREDICTABLE** whether an IPoR is generated or not.	RW	`0b0`
[1,0]		Reserved	RO	`0b0`

12.3.2.2.3 Secure Enclave Power Control (SE_PWR_CTRL) register

The following table gives a bit-level description of the SECENTOP Power Domain Control register.

Table 12-98: SE_PWR_CTRL register

Bits	Name	Description	Type	Reset
[31:1]	-	Reserved	RO	`0x0000_0000`
[0]	PWR_GATE_EN	SECENCTOP can enter OFF or MEM_RET: • `0b0`: SECENCTOP must remain in ON. • `0b1`: SECENCTOP can enter OFF or MEM_RET next time the Secure Enclave is idle.	RW	`0b0`

12.3.2.2.4 Secure Enclave General Purpose {0-3} (SE_GP{0-3}) register

The following table gives a bit-level description of the SE_GP{0-3} register.

Table 12-99: SE_GP{0-3} register

Bits	Name	Description	Type	Reset
[31:0]	GP	General-purpose data register	RW	0x0000_0000

12.3.2.2.5 Secure Enclave Clock Divider Control (SE_CLK_DIV) register

The following table gives a bit-level description of the SE_CLK_DIV register.

Table 12-100: SE_CLK_DIV register

Bits	Name	Description	Type	Reset
[31:21]	-	Reserved	RO	0x000
[20:16]	CLKDIV_CUR	Current value of integer divider applied to **SECENCCLK**: 0x00: Divide by 1 0x01: Divide by 2 ... 0x1F: Divide by 32	RO	UNKNOWN
[15:5]	-	Reserved	RO	0x000
[4:0]	CLK_DIVIDER	Controls the clock divider ratio: 0: 1:1 ratio between **SECENCDIVCLK** clock and **SECENCCLK** clock 1: 1:2 ratio between **SECENCDIVCLK** and **SECENCCLK** 2 - 31: Reserved	RW	0x01

12.3.2.2.6 Secure Enclave Build Configuration (SE_BLD_CFG) register

The following table gives a bit-level description of the SE_BLD_CFG register.

Table 12-101: SE_BLD_CFG register

Bits	Name	Description	Type	Reset
[31:16]	**RAM_SIZE**	Secure Enclave RAM size. The value of this field is the size of the Secure Enclave RAM in KB. A value of 0x00 is Reserved. For example, a value of 0x40 indicates a RAM of 64KB whilst a value of 0x80 indicates a RAM of 128KB.	RO	SEC_ENC_RAM_SIZE

Bits	Name	Description	Type	Reset
[15:0]	**ROM_SIZE**	Secure Enclave ROM size. The value of this field is the size of the Secure EnclaveROM in KB. A value of `0x00` is Reserved. For example, a value of `0x20` indicates a ROM of 32KB whilst a value of `0x40` indicates a ROM of 64KB.	RO	SEC_ENC_ROM_SIZE

12.3.2.2.7 Peripheral ID 4 (PID4) register

The following table gives a bit-level description of the Peripheral ID 4 register.

Table 12-102: PID4 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	`0x00_0000`
[7:4]	Size	Number of 4KB occupied by the System ID block. This field is deprecated.	RO	`0x0`
[3:0]	DES_2	JEP Continuation	RO	`0x4`

12.3.2.2.8 Peripheral ID 5 (PID5) register

The following table gives a bit-level description of the Peripheral ID 5 register.

Table 12-103: PID5 fegister

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	`0x0000_0000`

12.3.2.2.9 Peripheral ID 6 (PID6) register

The following table gives a bit-level description of the Peripheral ID 6 register.

Table 12-104: PID6 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	`0x0000_0000`

12.3.2.2.10 Peripheral ID 7 (PID7) register

The following table gives a bit-level description of the Peripheral ID 7 register.

Table 12-105: PID7 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	`0x0000_0000`

12.3.2.2.11 Peripheral ID 0 (PID0) register

The following table gives a bit-level description of the Peripheral ID 0 register.

Table 12-106: PID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PART_0	Bits [7:0] of part ID	RO	0x79

12.3.2.2.12 Peripheral ID 1 (PID1) register

The following table gives a bit-level description of the Peripheral ID 1 register.

Table 12-107: PID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	DES_0	Bits [3:0] of JEP 106 Identity	RO	0xB
[3:0]	PART_1	Bits [11:8] of part ID	RO	0x0

12.3.2.2.13 Peripheral ID 2 (PID2) register

The following table gives a bit-level description of the Peripheral ID 2 register.

Table 12-108: PID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVISION	Major revision of the System ID block.	RO	0x0
[3]	JEDEC	Indicates the use of JEDEC JEP106 identification scheme.	RO	0b1
[2:0]	DES_1	Bits [6:4] of JEP 106 Identity	RO	0b011

12.3.2.2.14 Peripheral ID 3 (PID3) register

The following table gives a bit-level description of the Peripheral ID 3 register.

Table 12-109: PID3 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVAND	Minor revision of the System ID block	RO	0x0
[3:0]	CMOD	Customer modification field	RO	0x0

12.3.2.2.15 Component ID 0 (CID0) register

The following table gives a bit-level description of the Component ID 0 register.

Table 12-110: CID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_0	Preamble 0	RO	0x0D

12.3.2.2.16 Component ID 1 (CID1) register

The following table gives a bit-level description of the Component ID 1 register.

Table 12-111: CID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	CLASS	Class of the component	RO	0xF
[3:0]	PRMBL_1	Preamble 1	RO	0x0

12.3.2.2.17 Component ID 2 (CID2) register

The following table gives a bit-level description of the Component ID 2 register.

Table 12-112: CID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_2	Preamble 2	RO	0x05

12.3.2.2.18 Component ID 3 (CID3) register

The following table gives a bit-level description of the Component ID 3 register.

Table 12-113: CID3 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_3	Preamble 3	RO	0xB1

12.3.3 Interrupt Router register summary

This section describes the register set of the Interrupt Router.

The Interrupt Router is in the AONTOP power domain. It is accessed at offset `0x1A50_0000` in the Host System memory map.

Table 12-114: Interrupt Router register summary

Offset	Short Name	Access	Name
`0x000`	INT_RTR_CTRL	RW	Interrupt Router Control 12.3.3.1 Interrupt Router Control (INT_RTR_CTRL) register on page 257
`0x004 – 0x00C`	-	RO	Reserved
`0x010`	LD_CTRL	RW	Lockdown Control 12.3.3.2 Lockdown Control (LD_CTRL) register on page 257
`0x014 – 0x0FC`	-	RO	Reserved
`0x100`	SHD_INT_INFO	RO	Shared Interrupt Information 12.3.3.3 Shared Interrupt Information (SHD_INT_INFO) register on page 258
`0x104`	SHD_INT_CFG	RW	Shared Interrupt Configuration 12.3.3.4 Shared Interrupt Configuration (SHD_INT_CFG) register on page 258
`0x108`	SHD_INT_LCTRL	RW	Shared Interrupt Lock Control 12.3.3.5 Shared Interrupt Lock Control (SHD_INT_LCTRL) register on page 259
`0x10C`	SHD_INT_SEL	RW	Shared Interrupt Select 12.3.3.6 Shared Interrupt Select (SHD_INT_SEL) register on page 259
`0x110 – 0xE8C`	-	RO	Reserved
`0xE90`	INT_RTR_TMP_ST	RW	Interrupt Router Tamper Status 12.3.3.7 Interrupt Router Tamper Status (INT_RTR_TMP_ST) register on page 259
`0xE94 – 0xF9C`	-	RO	Reserved
`0xFA0`	INT_RTR_CAP	RO	Interrupt Router Capability 12.3.3.8 Interrupt Router Capability (INT_RTR_CAP) register on page 260
`0xFA4 – 0xFAC`	-	RO	Reserved
`0xFB0`	INT_RTR_CFG	RO	Interrupt Router Configuration 12.3.3.9 Interrupt Router Configuration (INT_RTR_CFG) register on page 260
`0xFB4 – 0xFCC`	-	RO	Reserved
`0xFD0`	PID4	RO	Peripheral ID4 12.3.3.10 Peripheral ID 4 (PID4) register on page 261
`0xFD4`	PID5	RO	Peripheral ID5 12.3.3.11 Peripheral ID 5 (PID5) register on page 261
`0xFD8`	PID6	RO	Peripheral ID6 12.3.3.12 Peripheral ID 6 (PID6) register on page 261
`0xFDC`	PID7	RO	Peripheral ID7 12.3.3.13 Peripheral ID 7 (PID7) register on page 262
`0xFE0`	PID0	RO	Peripheral ID0 12.3.3.14 Peripheral ID 0 (PID0) register on page 262
`0xFE4`	PID1	RO	Peripheral ID1 12.3.3.15 Peripheral ID 1 (PID1) register on page 262
`0xFE8`	PID2	RO	Peripheral ID2 12.3.3.16 Peripheral ID 2 (PID2) register on page 262
`0xFEC`	PID3	RO	Peripheral ID3 12.3.3.17 Peripheral ID 3 (PID3) register on page 263

Offset	Short Name	Access	Name
0xFF0	CID0	RO	Component ID0 12.3.3.18 Component ID 0 (CID0) register on page 263
0xFF4	CID1	RO	Component ID1 12.3.3.19 Component ID 1 (CID1) register on page 263
0xFF8	CID2	RO	Component ID2 12.3.3.20 Component ID 2 (CID2) register on page 263
0xFFC	CID3	RO	Component ID3 12.3.3.21 Component ID 3 (CID3) register on page 264


Note

These registers only support 32-bit word aligned access. Any other attempts at access generate an error and are treated as RAZ/WI. For more information on access to the registers of the Interrupt Router see B. Interrupt Router on page 340.

12.3.3.1 Interrupt Router Control (INT_RTR_CTRL) register

The following table gives a bit-level description of the Interrupt Router Control register.

Table 12-115: INT_RTR_CTRL register

Bits	Name	Description	Type	Reset
[31:1]	-	Reserved	RO	0x0000_0000
[0]	ERR	Configures the response for Configuration Accesses which generate a Configuration Access Error: 0: No error 1: Error	RW	0b1

12.3.3.2 Lockdown Control (LD_CTRL) register

The following table gives a bit-level description of the Lockdown Control register.

Table 12-116: LD_CTRL register

Bits	Name	Description	Type	Reset
[31:3]	-	Reserved	RO	0x0000_0000
[2]	LDI_ST	Lockdown interface status. Indicates the current value of the Lockdown interface: 0: Lockdown interface is de-asserted 1: Lockdown interface is asserted	RO	UNKNOWN

Bits	Name	Description	Type	Reset
[1:0]	LOCK	Indicates the lock state of the Interrupt Router: 0b00: Open lockdown state. 0b01: Reserved and treated as 0b00 0b10: Partial lockdown state 0b11: Full lockdown state	RW	0b00

The LD_CTRL.LOCK field is not updateable, when the following are all true:

- Lockdown interface is asserted HIGH.
- Lockdown state of the Interrupt Router is Partial or Full.

12.3.3.3 Shared Interrupt Information (SHD_INT_INFO) register

The following table gives a bit-level description of the Shared Interrupt Information register.

Table 12-117: SHD_INT_INFO register

Bits	Name	Description	Type	Reset
[31:16]	-	Reserved	RO	0x0000
[15:0]	ICI_DST	Interrupt Controller Destination. Each bit indicates whether the interrupt selected by the SHD_INT_SEL.INT_SEL field, can be routed to ICI interface associated with the bit, starting with bit 0 for ICI0 to bit 3 for ICI3: 0: Shared interrupt cannot be routed to the ICI{x} 1: Shared interrupt can be routed to the ICI{x} The value of this field is defined by the SI{x}_ICI_DST configuration option.	RO	CFG_DEF

12.3.3.4 Shared Interrupt Configuration (SHD_INT_CFG) register

The following table gives a bit-level description of the Shared Interrupt Configuration Register.

Table 12-118: SHD_INT_CFG register

Bits	Name	Description	Type	Reset
[31:16]	-	Reserved	RO	0x0000

Bits	Name	Description	Type	Reset
[15:0]	ICI_EN	Interrupt Controller Enable. Each bit selects whether the interrupt selected by the SHD_INT_SEL.INT_SEL field, is routed to the ICI interface associated with the bit, starting with bit 0 for ICI0 to bit 3 for ICI3: 0: Shared interrupt not routed to the ICI{x} 1: Shared interrupt routed to the ICI{x} Bits, where the respective bit in the SHD)_INT_INFO.ICI_DST field is `0b0` are Reserved and treated as RAZ/WI. The default value of this field is defined by the SI{x}_DEF_ICI configuration option.	RW	CFG_DEF

12.3.3.5 Shared Interrupt Lock Control (SHD_INT_LCTRL) register

The following table gives a bit-level description of the Shared Interrupt Lockdown Control register.

Table 12-119: SHD_INT_LCTRL register

Bits	Name	Description	Type	Reset
[31]	LOCK	Control the lock status of the interrupt selected by the SHD_INT_SEL.INT_SEL field: `0b0`: Shared interrupt is not locked `0b1`: Shared interrupt is locked This field becomes read-only when this field is set to `0b1` and the LD_CTRL.LOCK field is set to `0b10` or `0b11`. For more information on the lock status, see B. Interrupt Router on page 340.	RW	`0b0`
[30:0]	-	Reserved	RO	`0x0000_0000`

12.3.3.6 Shared Interrupt Select (SHD_INT_SEL) register

The following table gives a bit-level description of the Shared Interrupt Select register.

Table 12-120: SHD_INT_SEL register

Bits	Name	Description	Type	Reset
[31:16]	-	Reserved	RO	`0x0000`
[15:0]	INT_SEL	Selects which interrupt the SHD_INT_INFO, SHD_INT_CFG and SHD_INT_LCTRL registers refer to. When INT_SEL is greater than NUM_SHD_INT, the fields in SHD_INT_INFO, SHD_INT_CFG and SHD_INT_LCTRL are Reserved and treated as RAZ/WI.	RW	`0x0000`

12.3.3.7 Interrupt Router Tamper Status (INT_RTR_TMP_ST) register

The following table gives a bit-level description of the Interrupt Router Tamper Status register.

Table 12-121: INT_RTR_TMP_ST register

Bits	Name	Description	Type	Reset
[31]	TMP_ST_VLD	Indicates whether the INT_RTR_TMP_ST register contains valid data or not: • 0: INT_RTR_TMP_ST does not contain valid data • 1: INT_RTR_TMP_ST contains valid data This field is written 1 to clear, writing a value of 0 has no effect.	RW	0b0
[30]	TMP_ST_OVERFLW	Indicates whether a tamper transaction occurred, while the INT_RTR_TMP_ST.TMP_ST_VLD was 1: • 0: No tamper transaction overflow occurred • 1: Tamper transaction overflow occurred This field is written 1 to clear, writing a value of 0 has no effect.	RW	0b0
[29:12]	-	Reserved	RO	0x0_0000
[11:2]	TMP_TRANS_ADDR	Address of the register accessed by the tamper transaction. When TMP_ST_VLD is 0 this field is not valid.	RO	UNKNOWN
[1:0]	-	Reserved	RO	0x00


Note

This register can only be accessed by Secure privileged access from the Security Monitor. In SSE-710 the Secure Enclave is the Security Monitor for the Interrupt Router. For more information, see B. Interrupt Router on page 340.

12.3.3.8 Interrupt Router Capability (INT_RTR_CAP) register

The following table gives a bit-level description of the Interrupt Router Capability register.

Table 12-122: INT_RTR_CAP register

Bits	Name	Description	Type	Reset
[31:4]	-	Reserved	RO	0x0000
[3:0]	LDE_LVL	Level of the Lockdown Extension implemented by the Interrupt Router. Read as 0x2 – LDE.2 is implemented. All other values are Reserved.	RO	0x2

12.3.3.9 Interrupt Router Configuration (INT_RTR_CFG) register

The following table gives a bit-level description of the Interrupt Router Configuration register.

Table 12-123: INT_RTR_CFG Register

Bits	Name	Description	Type	Reset
[31:20]	-	Reserved	RO	0x000
[19:16]	NUM_ICI	Number of Interrupt Controllers Interrupt interface (ICI)s supported by the Interrupt Router: 0x3: 4 ICIs supported	RO	0x3
[15:0]	NUM_SHD_INT	Number of shared interrupts supported by the Interrupt Router: 0x0: 1 Shared interrupt 0x1: 2 Shared interrupts ... 0x18B: 395 Shared interrupts	RO	CFG_DEF

12.3.3.10 Peripheral ID 4 (PID4) register

The following table gives a bit-level description of the Peripheral ID 4 register.

Table 12-124: PID4 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	Size	Number of 4KB blocks occupied by the System ID block. This field is deprecated.	RO	0x0
[3:0]	DES_2	JEP Continuation	RO	0x4

12.3.3.11 Peripheral ID 5 (PID5) register

The following table gives a bit-level description of the Peripheral ID 5 register.

Table 12-125: PID5 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.3.12 Peripheral ID 6 (PID6) register

The following table gives a bit-level description of the Peripheral ID 6 register.

Table 12-126: PID6 Register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.3.13 Peripheral ID 7 (PID7) register

The following table gives a bit-level description of the Peripheral ID 7 register.

Table 12-127: PID7 Register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.3.14 Peripheral ID 0 (PID0) register

The following table gives a bit-level description of the Peripheral ID 0 register.

Table 12-128: PID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PART_0	Bits [7:0] of part ID	RO	0x74

12.3.3.15 Peripheral ID 1 (PID1) register

The following table gives a bit-level description of the Peripheral ID 1 register.

Table 12-129: PID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	DES_0	Bits [3:0] of JEP 106 Identity	RO	0xB
[3:0]	PART_1	Bits [11:8] of part ID	RO	0x0

12.3.3.16 Peripheral ID 2 (PID2) register

The following table gives a bit-level description of the Peripheral ID 2 register.

Table 12-130: PID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVISION	Major revision of the System ID block	RO	0x0
[3]	JEDEC	Indicates the use of JEDEC JEP106 identification scheme	RO	0b1
[2:0]	DES_1	Bits [6:4] of JEP 106 Identity	RO	0b011

12.3.3.17 Peripheral ID 3 (PID3) register

The following table gives a bit-level description of the Peripheral ID 3 register.

Table 12-131: PID3 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVAND	Minor revision of the System ID block	RO	0x0
[3:0]	CMOD	Customer modification field	RO	0x0

12.3.3.18 Component ID 0 (CID0) register

The following table gives a bit-level description of the Component ID 0 register.

Table 12-132: CID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_0	Preamble 0	RO	0x0

12.3.3.19 Component ID 1 (CID1) register

The following table gives a bit-level description of the Component ID 1 register.

Table 12-133: CID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	CLASS	Class of the component	RO	0xF
[3:0]	PRMBL_1	Preamble 1	RO	0x0

12.3.3.20 Component ID 2 (CID2) register

The following table gives a bit-level description of the Component ID 2 register.

Table 12-134: CID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_2	Preamble 2	RO	0x05

12.3.3.21 Component ID 3 (CID3) register

The following table gives a bit-level description of the Component ID 3 register.

Table 12-135: CID3 Register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_3	Preamble 3	RO	0xB1

12.3.4 REFCLK Counter CNTControl register summary

This section summarizes the **REFCLK** Counter CNTControl registers.

REFCLK CNTControl is an implementation of a Memory Mapped Counter as defined by the *Arm® Architecture Reference Manual Armv8, for Armv8-A architecture profile*. The **REFCLK** Counter CNTControl and CNTRead frames are accessed at offset 0x1A20_0000 and 0x1A21_0000 respectively in the Host System memory map.

The CNTControl frame uses three additional undefined registers, for more information see the reference manuals stated above. The following table lists the additional three registers and their offset address in the CNTControl frame.

Table 12-136: REFCLK Counter register summary

Offset	Short name	Access	Name
0xC0	CNTSCR	RW	Counter Synchronization Control Register. 12.3.4.1 CNTControl Counter Synchronization Control (CNTSCR) register on page 264
0xC4	-	RO	Reserved
0xC8	CNTSVL	RO	Counter Synchronized Counter Lower Value Register 12.3.4.2 CNTControl Synchronized Counter Value (CNTSVL) register on page 265.
0xCC	CNTSVU	RO	Counter Synchronized Counter Upper Value Register. 12.3.4.3 CNTControl Synchronized Counter Value (CNTSVU) register on page 265

12.3.4.1 CNTControl Counter Synchronization Control (CNTSCR) register

The following table gives a bit-level description of the REFCLK CNTControl Sync Control register.

Table 12-137: CNTSCR register

Bits	Name	Description	Type	Reset
[31:1]	-	Reserved	RO	0x0000_0000
[0]	ENSYNC	Controls the way the Counter Control register EN bit operates: 0b0: The counter is enabled or disabled immediately 0b1: The enabling of the counter is delayed until just after the next rising edge at the **S32KCLK** Disabling the counter is not delayed.	RW	0b0

12.3.4.2 CNTControl Synchronized Counter Value (CNTSVL) register

The following table gives a bit-level description of the CNTControl Synchronized Counter Value Lower Word register.

Table 12-138: CNTSVL register

Bits	Name	Description	Type	Reset
[31:0]	CNTSVL	**S32KCLK**-sampled value of counter, lower word – CNTSV[31:0]	RO	0x0000_0000

12.3.4.3 CNTControl Synchronized Counter Value (CNTSVU) register

The following table gives a bit-level description of the Synchronized Counter Value Upper Word register.

Table 12-139: CNTSVU register

Bits	Name	Description	Type	Reset
[31:0]	CNTSVU	Reads back the value of the counter sampled on the rising of edge of **S32KCLK**. Returns the upper word of CNTSV[63:32].	RO	0x0000_0000

12.3.5 System ID register summary

This section summarizes the System ID registers.

The System ID block registers identify:

- SoC and its implementor
- Version and implementor of SSE-710

The registers are in the AONTOP domain and are accessed at offset 0x1A00_0000 in the Host System memory map.

Table 12-140: System ID register summary

Offset	Short name	Access	Name
0x000	BSYS_CFG0	RO	Base System Configuration 0 12.3.5.1 Base System Config 0 (BSYS_CFG0) register on page 266
0x004	BSYS_CFG1	RO	Base System Configuration 1 12.3.5.2 Base System Config 1 (BSYS_CFG1) register on page 267
0x008	BSYS_CFG2	RO	Base System Configuration 2 12.3.5.3 Base System Config 2 (BSYS_CFG2) register on page 268
0x00C	BSYS_CFG3	RO	Base System Configuration 3 12.3.5.4 Base System Config 3 (BSYS_CFG3) register on page 269
0x010 – 0x03C	-	RO	Reserved
0x040	SOC_ID	RO	SoC Identification 12.3.5.5 SoC Identification (SOC_ID) register on page 269
0x044 – 0xFC4	-	RO	Reserved
0xFC8	IIDR	RO	Implementer Identification Register 12.3.5.6 Implementer Identification Register (IIDR) on page 269
0xFCC	-	RO	Reserved
0xFD0	PID4	RO	Peripheral ID4 12.3.2.1.20 Peripheral ID 4 (PID4) register on page 245
0xFD4	PID5	RO	Peripheral ID5 12.3.2.1.21 Peripheral ID 5 (PID5) register on page 246
0xFD8	PID6	RO	Peripheral ID6 12.3.2.1.22 Peripheral ID 6 (PID6) register on page 246
0xFDC	PID7	RO	Peripheral ID7 12.3.2.1.23 Peripheral ID 7 (PID7) register on page 246
0xFE0	PID0	RO	Peripheral ID0 12.3.2.1.24 Peripheral ID 0 (PID0) register on page 246
0xFE4	PID1	RO	Peripheral ID1 12.3.2.1.25 Peripheral ID 1 (PID1) register on page 247
0xFE8	PID2	RO	Peripheral ID2 12.3.2.1.26 Peripheral ID 2 (PID2) register on page 247
0xFEC	PID3	RO	Peripheral ID3 12.3.2.1.27 Peripheral ID 3 (PID3) register on page 247
0xFF0	CID0	RO	Component ID0 12.3.2.1.28 Component ID 0 (CID0) register on page 247
0xFF4	CID1	RO	Component ID1 12.3.2.1.29 Component ID 1 (CID1) register on page 248
0xFF8	CID2	RO	Component ID2 12.3.2.1.30 Component ID 2 (CID2) register on page 248
0xFFC	CID3	RO	Component ID3 12.3.2.1.31 Component ID 3 (CID3) register on page 248

Note

The System ID registers support only 32-bit word-aligned access. Any access by other bit size or unaligned access is not supported. For any non 32-bit word aligned access, the following happens:

- Generates an error and is treated as RAZ/WI

The System ID registers have the following behavior:

- Any read to a write-only register is treated as RAZ
- Any write to a read-only register is treated as WI

In both cases no error is generated.

12.3.5.1 Base System Config 0 (BSYS_CFG0) register

The following table gives a bit-level description of the Base System Config 0 register.

Table 12-141: BSYS_CFG0 register

Bits	Name	Description	Type	Reset
[31:4]	-	Reserved	RO	0x000_0000
[3:0]	NUM_HOST_CPU	Number of Host processor cores used	RO	HOST_CPU_NUM_CORES

12.3.5.2 Base System Config 1 (BSYS_CFG1) register

The following table gives a bit-level description of the Base System Config 1 register.

Table 12-142: BSYS_CFG1 register

Bits	Name	Description	Type	Reset
[31:16]	-	Reserved	RO	0x0000
[15:12]	EXT_SYS3	External System 3 Configuration: 0x0: External System 3 not implemented 0x1: Reserved 0x2: External System 3 implemented without Arm® TrustZone® 0x3: External System 3 implemented with Arm® TrustZone® All other values are Reserved. For SSE-710 this field always reads as 0x0.	RO	0x0
[11:8]	EXT_SYS2	External System 2 Configuration: 0x0: External System 2 not implemented 0x1: Reserved 0x2: External System 2 implemented without Arm® TrustZone® 0x3: External System 2 implemented with Arm® TrustZone® All other values are Reserved. For SSE-710 this field always reads as 0x0.	RO	0x0

Bits	Name	Description	Type	Reset
[7:4]	EXT_SYS1	External System 1 Configuration: 0x0: External System 1 not implemented. 0x1: Reserved 0x2: External System 1 implemented without Arm TrustZone 0x3: External System 1 implemented with Arm® TrustZone® All other values are Reserved. For SSE-710 this field always reads as 0x3.	RO	0x3
[3:0]	EXT_SYS0	External System 0 Configuration: 0x0: External System 0 not implemented. 0x1: Reserved 0x2: External System 0 implemented without Arm® TrustZone® 0x3: External System 0 implemented with Arm® TrustZone® All other values are Reserved. For SSE-710 this field always reads as 0x3.	RO	0x3

12.3.5.3 Base System Config 2 (BSYS_CFG2) register

The following table gives a bit-level description of the Base System Config 2 register.

Table 12-143: BSYS_CFG2 register

Bits	Name	Description	Type	Reset
[31:25]	-	Reserved	RO	0x00
[24]	OCVM_EN	Indicates if the OCVM interface is supported: • 0b0: OCVM interface is not supported. • 0b1: OCVM interface is supported.	RO	0x1
[23:20]	NUM_EXP_MST	Number of Expansion Master interfaces: • 0x0: No Expansion Master interfaces. • 0x1: 1 Expansion Master interface • 0x2: 2 Expansion Master interfaces • 0x3: 3 Expansion Master interfaces • 0x4: 4 Expansion Master interfaces All other values are Reserved. For SSE-710 this field always reads as 0x2.	RO	0x2

Bits	Name	Description	Type	Reset
[19:16]	NUM_EXP_SLV	Number of Expansion Slave interfaces: • 0x0: No Expansion Slave interfaces • 0x1: 1 Expansion Slave interface • 0x2: 2 Expansion Slave interfaces • 0x3: 3 Expansion Slave interfaces • 0x4: 4 Expansion Slave interfaces All other values are Reserved. For SSE-710 this field always reads as 0x2.	RO	0x2
[15:0]	-	Reserved	RO	0x0000

12.3.5.4 Base System Config 3 (BSYS_CFG3) register

The following table gives a bit-level description of the Base System Config 3 register.

Table 12-144: BSYS_CFG3 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.5.5 SoC Identification (SOC_ID) register

The following table gives a bit-level description of the SoC Identification register.

Table 12-145: SOC_ID register

Bits	Name	Description	Type	Reset
[31:20]	PRODUCT_ID	User defined value identifying the SoC. This field is set to **SOCPRTID**.	RO	CFG_DEF
[19:16]	VARIANT	User defined value variant or major revision of the SoC. This field is set to **SOCVAR**.	RO	CFG_DEF
[15:12]	REVISION	Use defined value used to distinguish minor revisions of the SoC. This field is set to **SOCREV**.	RO	CFG_DEF
[11:0]	IMPLEMENTER	Contains the JEP106 code of the company that used the SoC: [11:8] JEP106 continuation code of implementer. Provided by bits 10:7 of **SOCIMPLID**. [7] Always 0. [6:0] JEP106 identity code of implementer. Provided by bits 6:0 of **SOCIMPLID**	RO	CFG_DEF

12.3.5.6 Implementer Identification Register (IIDR)

The following table gives a bit-level description of the Implementer Identification Register.

Table 12-146: IIDR register

Bits	Name	Description	Type	Reset
[31:20]	PRODUCT_ID	Product ID of SSE-710	RO	0x762
[19:16]	VARIANT	Variant or major revision of SSE-710	RO	0x0
[15:12]	REVISION	Minor revisions of SSE-710	RO	0x0
[11:0]	IMPLEMENTER	Contains the JEP106 code of the company that implemented SSE-710: [11:8] JEP106 continuation code of implementer [7] Always 0 [6:0] JEP106 identity code of implementer	RO	0x43B

12.3.5.7 Peripheral ID 4 (PID4) register

The following table gives a bit-level description of the Peripheral ID 4 register.

Table 12-147: PID4 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	Size	Number of 4KB occupied by the System ID block. This field is deprecated.	RO	0x0
[3:0]	DES_2	JEP Continuation	RO	0x4

12.3.5.8 Peripheral ID 5 (PID5) register

The following table gives a bit-level description of the Peripheral ID 5 register.

Table 12-148: PID5 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.5.9 Peripheral ID 6 (PID6) register

The following table gives a bit-level description of the Peripheral ID 6 register.

Table 12-149: PID6 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.5.10 Peripheral ID 7 (PID7) register

The following table gives a bit-level description of the Peripheral ID 7 register.

Table 12-150: PID7 Register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.5.11 Peripheral ID 0 (PID0) register

The following table gives a bit-level description of the Peripheral ID 0 register.

Table 12-151: PID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PART_0	Bits [7:0] of part ID	RO	0x62

12.3.5.12 Peripheral ID 1 (PID1) register

The following table gives a bit-level description of the Peripheral ID 1 register.

Table 12-152: PID1 Register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	DES_0	Bits [3:0] of JEP 106 Identity	RO	0xB
[3:0]	PART_1	Bits [11:8] of part ID	RO	0x7

12.3.5.13 Peripheral ID 2 (PID2) register

The following table gives a bit-level description of the Peripheral ID 2 register.

Table 12-153: PID2 Register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVISION	Major revision of the System ID block	RO	0x0
[3]	JEDEC	Indicates the use of JEDEC JEP106 identification scheme	RO	0b1
[2:0]	DES_1	Bits [6:4] of JEP 106 Identity	RO	0b011

12.3.5.14 Peripheral ID 3 (PID3) register

The following table gives a bit-level description of the Peripheral ID 3 register.

Table 12-154: PID3 Register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVAND	Minor revision of the System ID block	RO	0x0
[3:0]	CMOD	Customer modification field	RO	0x0

12.3.5.15 Component ID 0 (CID0) register

The following table gives a bit-level description of the Component ID 0 register.

Table 12-155: CID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_0	Preamble 0	RO	0x0D

12.3.5.16 Component ID 1 (CID1) register

The following table gives a bit-level description of the Component ID 1 register.

Table 12-156: CID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	CLASS	Class of the component	RO	0xF
[3:0]	PRMBL_1	Preamble 1	RO	0x0

12.3.5.17 Component ID 2 (CID2) register

The following table gives a bit-level description of the Component ID 2 register.

Table 12-157: CID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_2	Preamble 2	RO	0x05

12.3.5.18 Component ID 3 (CID3) register

The following table gives a bit-level description of the Component ID 3 register.

Table 12-158: CID3 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_3	Preamble 3	RO	0xB1

12.3.6 Boot register summary

This section summarizes the Boot Register registers.

The Boot Registers are accessed at offset 0x0000_0000 in the Host System memory map. The Boot Registers are written once only by the Secure Enclave. The Boot Register is in the AONTOP power domain.

Table 12-159: Boot register summary

Offset	Short name	Access	Name
0x000 – 0x00C	BIR{0-3}	RW	Boot Instruction Register {0-3} 12.3.6.1 Boot Instruction Register {0-3} (BIR{0-3}) on page 274
0x010 – 0x03C	BIR{4-15}	RO	Boot Instruction Register {4-15} 12.3.6.2 Boot Instruction Register {4-15} (BIR{4-15}) on page 274
0x040 – 0xFCC	-	RO	Reserved
0xFD0	PID4	RO	Peripheral ID4 12.3.6.3 Peripheral ID 4 (PID4) register on page 274
0xFD4	PID5	RO	Peripheral ID5 12.3.6.6 Peripheral ID 5 (PID5) register on page 275
0xFD8	PID6	RO	Peripheral ID6 12.3.6.4 Peripheral ID 6 (PID6) register on page 275
0xFDC	PID7	RO	Peripheral ID7 12.3.6.5 Peripheral ID 7 (PID7) register on page 275
0xFE0	PID0	RO	Peripheral ID0 12.3.6.7 Peripheral ID 0 (PID0) register on page 275
0xFE4	PID1	RO	Peripheral ID1 12.3.6.8 Peripheral ID 1 (PID1) register on page 276
0xFE8	PID2	RO	Peripheral ID2 12.3.6.9 Peripheral ID 2 (PID2) register on page 276
0xFEC	PID3	RO	Peripheral ID3 12.3.6.10 Peripheral ID 3 (PID3) register on page 276
0xFF0	CID0	RO	Component ID0 12.3.6.11 Component ID 0 (CID0) register on page 276
0xFF4	CID1	RO	Component ID112.3.6.12 Component ID 1 (CID1) register on page 277
0xFF8	CID2	RO	Component ID2 12.3.6.13 Component ID 2 (CID2) register on page 277
0xFFC	CID3	RO	Component ID3 12.3.6.14 Component ID 3 (CID3) register on page 277

Note

An error is generated if any of the following occur:

- Any attempt to write to the register from any master other than the Secure Enclave.
- Secure Enclave writes to the same register more than once.

- Any attempt to write to a read-only register, by any master.
- Any attempt to read a Reserved register.

The registers of the Boot Register support the following accesses:

- Read accesses support any size access, but must be aligned to the size of the access.
- Write accesses must be 32-bit word-aligned accesses.
- Any access not meeting these requirements generates an error and is treated as RAZ/WI.

12.3.6.1 Boot Instruction Register {0-3} (BIR{0-3})

The following table gives a bit-level description of the Boot Instruction Register {0-3}.

Table 12-160: BIR{0-3} register

Bits	Name	Description	Type	Reset
[31:0]	BOOT_INSTR	Write-once register. Sets the instructions for the Host processor.	RW	UNKNOWN

This is a write-once register by the Secure Enclave.

12.3.6.2 Boot Instruction Register {4-15} (BIR{4-15})

The following table gives a bit-level description of the Boot Instruction Register {4-15}.

Table 12-161: BIR{4-15} register

Bits	Name	Description	Type	Reset
[31:0]	BOOT_INSTR	Reads of this address return the value in the Boot Instruction Register {0-3} registers: BIR{4,8,12} return the value in BIR0 BIR{5,9,13} return the value in BIR1 BIR{6,10,14} return the value in BIR2 BIR{7,11,15} return the value in BIR3	RO	UNKNOWN

12.3.6.3 Peripheral ID 4 (PID4) register

The following table gives a bit-level description of the Peripheral ID 4 register.

Table 12-162: PID4 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	Size	Number of 4KB blocks occupied by the System ID block. This field is deprecated.	RO	0x0
[3:0]	DES_2	JEP Continuation	RO	0x4

12.3.6.4 Peripheral ID 6 (PID6) register

The following table gives a bit-level description of the Peripheral ID 6 register.

Table 12-163: PID6 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.6.5 Peripheral ID 7 (PID7) register

The following table gives a bit-level description of the Peripheral ID 7 register.

Table 12-164: PID7 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.6.6 Peripheral ID 5 (PID5) register

The following table gives a bit-level description of the Peripheral ID 5 register.

Table 12-165: PID5 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.6.7 Peripheral ID 0 (PID0) register

The following table gives a bit-level description of the Peripheral ID 0 register.

Table 12-166: PID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000

Bits	Name	Description	Type	Reset
[7:0]	PART_0	Bits [7:0] of part ID	RO	0x72

12.3.6.8 Peripheral ID 1 (PID1) register

The following table gives a bit-level description of the Peripheral ID 1 register.

Table 12-167: PID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	DES_0	Bits [3:0] of JEP 106 Identity	RO	0xB
[3:0]	PART_1	Bits [11:8] of part ID	RO	0x0

12.3.6.9 Peripheral ID 2 (PID2) register

The following table gives a bit-level description of the Peripheral ID 2 register.

Table 12-168: PID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVISION	Major revision of the System ID bloc	RO	0x0
[3]	JEDEC	Indicates the use of JEDEC JEP106 identification scheme	RO	0b1
[2:0]	DES_1	Bits [6:4] of JEP 106 Identity.	RO	0b011

12.3.6.10 Peripheral ID 3 (PID3) register

The following table gives a bit-level description of the Peripheral ID 3 register.

Table 12-169: PID3 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVAND	Minor revision of the System ID block	RO	0x0
[3:0]	CMOD	Customer modification field	RO	0x0

12.3.6.11 Component ID 0 (CID0) register

The following table gives a bit-level description of the Component ID 0 register.

Table 12-170: CID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000

Bits	Name	Description	Type	Reset
[7:0]	PRMBL_0	Preamble 0	RO	0x0D

12.3.6.12 Component ID 1 (CID1) register

The following table gives a bit-level description of the Component ID 1 register.

Table 12-171: CID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	CLASS	Class of the component	RO	0xF
[3:0]	PRMBL_1	Preamble 1	RO	0x0

12.3.6.13 Component ID 2 (CID2) register

The following table gives a bit-level description of the Component ID 2 register.

Table 12-172: CID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_2	Preamble 2	RO	0x05

12.3.6.14 Component ID 3 (CID3) register

The following table gives a bit-level description of the Component ID 3 register.

Table 12-173: CID3 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_3	Preamble 3	RO	0xB1

12.3.7 GPIO Control register summary

This section summarizes the GPIO Control registers.

Table 12-174: GPIO Control register summary

Offset	Access type	Name	Short name
0x000 – 0x07C	RW	General Purpose Output {0-31} Control 12.3.7.1 General Purpose Output {0-31} Control (GPO{0-31}_CTRL) register on page 279	GPO{0-31}_CTRL
0x080 – 0x0FC	RO	Reserved	-

Offset	Access type	Name	Short name
0x100 – 0x17C	RW	General Purpose Input {0-31} 12.3.7.2 General Purpose Input {0-31} Status (GPI{0-31}_ST) register on page 279	GPI{0-31}_ST
0x180 – 0xEFC	RO	Reserved	-
0xF00	RW	Integration Mode Control 12.3.7.3 Integration Mode Control register on page 279	ITCTRL
0xF0 - 0xF9C	RO	Reserved	-
0xFA0	RW	Claim Tag Set 12.3.7.4 Claim Tag Set register on page 280	CLAIMSET
0xFA4	RW	Claim Tag Clear 12.3.7.5 Claim Tag Clear register on page 280	CLAIMCLR
0xFA8	RO	Device Affinity 0 12.3.7.6 Device Affinity 0 register on page 280	DEVAFF0
0xFAC	RO	Device Affinity 1 12.3.7.7 Device Affinity 1 register on page 280	DEVAFF1
0xFB0	WO	Software Lock Access 12.3.7.8 Software Lock Access register on page 280	LAR
0xFB4	RO	Software Lock Status 12.3.7.9 Software Lock Status register on page 281	LSR
0xFB8	RO	Authentication Status 12.3.7.10 Authentication Status (AUTHSTATUS) register on page 281	AUTHSTATUS
0xFBC	RO	Device Architecture 12.3.7.11 Device Architecture (DEVARCH) register on page 283	DEVARCH
0xFC0	RO	Device Configuration 2 12.3.7.12 Device Configuration 2 (DEVID2) register on page 283	DEVID2
0xFC4	RO	Device Configuration 1 12.3.7.13 Device Configuration 1 (DEVID1) register on page 284	DEVID1
0xFC8	RO	Device Configuration 12.3.7.14 Device Configuration (DEVID) register on page 284	DEVID
0xFCC	RO	Device Type Identifier 12.3.7.15 Device Type Identifier (DEVTYPE) register on page 285	DEVTYPE
0xFD0	RO	Peripheral ID4 12.3.7.16 Peripheral ID 4 (PID4) register on page 285	PID4
0xFD4	RO	Peripheral ID5 12.3.7.17 Peripheral ID 5 (PID5) register on page 285	PID5
0xFD8	RO	Peripheral ID6 12.3.7.18 Peripheral ID 6 (PID6) register on page 286	PID6
0xFDC	RO	Peripheral ID7 12.3.7.19 Peripheral ID 7 (PID7) register on page 286	PID7
0xFE0	RO	Peripheral ID0 12.3.7.20 Peripheral ID 0 (PID0) register on page 286	PID0
0xFE4	RO	Peripheral ID1 12.3.7.21 Peripheral ID 1 (PID1) register on page 286	PID1
0xFE8	RO	Peripheral ID2 12.3.7.22 Peripheral ID 2 (PID2) register on page 286	PID2
0xFEC	RO	Peripheral ID3 12.3.7.23 Peripheral ID 3 (PID3) register on page 287	PID3
0xFF0	RO	Component ID0 12.3.7.24 Component ID 0 (CID0) register on page 287	CID0
0xFF4	RO	Component ID1 12.3.7.25 Component ID 1 (CID1) register on page 287	CID1
0xFF8	RO	Component ID2 12.3.7.26 Component ID 2 (CID2) register on page 287	CID2
0xFFC	RO	Component ID3 12.3.7.27 Component ID 3 (CID3) register on page 288	CID3

Note

The GPIO Control registers only support 32-bit word aligned accesses. Any access of other size or unaligned access is not supported. For any non 32-bit aligned word access, the resulting behavior is:

- Generates an error and is treated as RAZ/WI.

The GPIO Control registers have the following behavior:

- Any read to a write-only register is treated as RAZ.
- Any write to a read-only register is treated as WI.

In both cases, no error is generated.

12.3.7.1 General Purpose Output {0-31} Control (GPO{0-31}_CTRL) register

The following table gives a bit-level description of the General Purpose Output {0-31} Control register.

Table 12-175: GPO{0-31}_CTRL register

Bits	Name	Description	Type	Reset
[31:1]	-	Reserved	RO	0x0000_0000
[0]	GPO	Controls the value of the general-purpose output. When AUTHSTATUS.NSID is 0b10 it is **IMPLEMENTATION DEFINED** whether this field behaves as follows: Is treated as RAZ/WI with the respective GPO output driven to 0b0. Is treated a read-write, but with the respective GPO output driven to 0b0. For SSE-710 this field is Reserved and treated as RAZ/WI for GPO{x}_CTRL registers where x is greater than or equal to 1.	RW	0b0

12.3.7.2 General Purpose Input {0-31} Status (GPI{0-31}_ST) register

The following table gives a bit-level description of the General Purpose Input {0-31} Status register.

Table 12-176: GPI{0-31}_ST register

Bits	Name	Description	Type	Reset
[31:1]	-	Reserved	RO	0x0000_0000
[0]	GPI	Provides the status of the General Purpose input. When AUTHSTATUS.NSID is 0b10 the value in this field is always 0b0. For SSE-710 this field is Reserved and treated as RAZ/WI for all GPI{x}_ST registers.	RO	0b0

12.3.7.3 Integration Mode Control register

The following table gives a bit-level description of the Integration Mode Control register.

Table 12-177: Integration Mode Control register

Bits	Name	Description	Type	Reset
[31:1]	-	Reserved	RO	0x0000_0000

Bits	Name	Description	Type	Reset
[0]	IME	Integration Mode Enable: 0b0: Component must enter functional mode. 0b1: Component must enter integration mode. The GPIO does not support the integration mode and therefore this field ignores writes which set this field to 0b1.	RW	0b0

12.3.7.4 Claim Tag Set register

The following table gives a bit-level description of the Claim Tag Set register.

Table 12-178: Claim Tag Set register

Bits	Name	Description	Type	Reset
[31:0]	SET	The GPIO does not support and Claim Tags. Therefore, all bits are treated as RAZ/WI.	RW	0x0000_0000

12.3.7.5 Claim Tag Clear register

The following table gives a bit-level description of the Claim Tag Clear register.

Table 12-179: Claim Tag Clear register

Bits	Name	Description	Type	Reset
[31:0]	CLR	The GPIO does not support and Claim Tags. Therefore, all bits are treated as RAZ/WI.	RW	0x0000_0000

12.3.7.6 Device Affinity 0 register

The following table gives a bit-level description of the Device Affinity 0 register.

Table 12-180: Device Affinity 0 register

Bits	Name	Description	Type	Reset
[31:0]	DEVAFF0	The field reads as zero	RO	0x0000_0000

12.3.7.7 Device Affinity 1 register

The following table gives a bit-level description of the Device Affinity 1 register.

Table 12-181: Device Affinity 1 register

Bits	Name	Description	Type	Reset
[31:0]	DEVAFF1	This field reads as zero	RO	0x0000_0000

12.3.7.8 Software Lock Access register

The following table gives a bit-level description of the Software Lock Access register.

Table 12-182: Software Lock Access register

Bits	Name	Description	Type	Reset
[31:0]	KEY	Key value. The GPIO Control does not implement the software lock mechanism so writes to this field are treated as WI. Reads to this field return `0x0000_0000`.	WO	`0x0000_0000`

12.3.7.9 Software Lock Status register

The following table gives a bit-level description of the Software Lock Status register.

Table 12-183: Software Lock Status register

Bits	Name	Description	Type	Reset
[31:3]	-	Reserved	RO	`0x0000_0000`
[2]	nTT	This field always reads as 0 which indicates a 32-bit Software Lock Access register is implemented.	RO	`0b0`
[1]	SLK	Software lock status	RO	`0b0`
[0]	SLI	Software lock mechanism is implemented. This field always reads `0b0` as the GPIO Control does not implement the software lock mechanism.	RO	`0b0`

12.3.7.10 Authentication Status (AUTHSTATUS) register

The following table gives a bit-level description of the Authentication Status register.

Table 12-184: AUTHSTATUS register

Bits	Name	Description	Type	Reset
[31:12]	-	Reserved	RO	`0x00_0000`
[11:10]	HNID	Hypervisor non-invasive debug. `0b00`: Functionality not implemented or controlled elsewhere `0b01`: Reserved `0b10`: Functionality disabled `0b11`: Functionality enabled This field always reads as `0b00` for the GPIO.	RO	`0b00`

Bits	Name	Description	Type	Reset
[9:8]	HID	Hypervisor invasive debug: 0b00: Functionality not implemented or controlled elsewhere 0b01: Reserved 0b10: Functionality disabled 0b11: Functionality enabled This field always reads as 0b00 for the GPIO.	RO	0b00
[7:6]	SNID	Secure non-invasive debug: 0b00: Functionality not implemented or controlled elsewhere 0b01: Reserved 0b10: Functionality disabled 0b11: Functionality enabled This field always reads as 0b00 for the GPIO.	RO	0b00
[5:4]	SID	Secure invasive debug: 0b00: Functionality not implemented or controlled elsewhere 0b01: Reserved 0b10: Functionality disabled 0b11: Functionality enabled This field always reads as 0b00 for the GPIO.	RO	0b00
[3:2]	NSNID	Non-secure non-invasive debug: 0b00: Functionality not implemented or controlled elsewhere 0b01: Reserved 0b10: Functionality disabled 0b11: Functionality enabled This field always reads as 0b00 for the GPIO.	RO	0b00

Bits	Name	Description	Type	Reset
[1:0]	NSID	Non-secure invasive debug: `0b00`: Functionality not implemented or controlled elsewhere `0b01`: Reserved `0b10`: Functionality disabled `0b11`: Functionality enabled This field can take the following values: `0b10`: when **DBGEN** is 0 `0b11`: when **DBGEN** is 1 The reset value of this field depends on the value of **DBGEN** input.	RO	See description.

12.3.7.11 Device Architecture (DEVARCH) register

The following table gives a bit-level description of the Device Architecture register.

Table 12-185: DEVARCH register

Bits	Name	Description	Type	Reset
[31:21]	ARCHITECT	Defines the architect of the component: Bits[31:28] Indicates the JEP106 continuation code Bits[27:21] Indicates the JEP106 identification code See the Standard Manufacturers Identification Code for information about JEP106. For components where Arm is the architect, this 11-bit field returns `0x23B`.	RO	`0x000`
[20]	PRESENT	Indicates the presence of this register: `0b0`: DEVARCH is not present. Bits[31:0] must be RAZ. `0b1`: DEVARCH is present. This field always reads as `0b0` for the GPIO	RO	`0b0`
[19:16]	REVISION	Architecture revision. Returns the revision of the architecture that the ARCHID field specifies.	RO	`0x0`
[15:0]	ARCHID	Architecture ID. Returns a value that identifies the architecture of the component.	RO	`0x0000`

12.3.7.12 Device Configuration 2 (DEVID2) register

The following table gives a bit-level description of the Device Configuration 2 register.

Table 12-186: DEVID2 Register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.7.13 Device Configuration 1 (DEVID1) register

The following table gives a bit-level description of the Device Configuration 1 register.

Table 12-187: DEVID1 Register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.7.14 Device Configuration (DEVID) register

The following table gives a bit-level description of the Device Configuration register.

Table 12-188: DEVID register

Bits	Name	Description	Type	Reset
[31:14]	-	Reserved	RO	0_0000
[13:8]	NUM_GPI	Number of GPI{x}_ST registers which are implemented: 0x00: No GPI{x}_ST register are implemented 0x01: 1 GPI{x}_ST register is implemented, starting with GPI0_ST 0x02: 2 GPI{x}_ST registers are implemented, starting with GPI0_ST up to GPI1_ST ... 0x20: 32 GPI{x}_ST registers are implemented, starting with GPI0_ST up to GPI31_ST For SSE-710 this field always reads as 0x00.	RO	0x00
[7:6]	-	Reserved	RO	0b00

Bits	Name	Description	Type	Reset
[5:0]	NUM_GPO	Number of GPO{x}_CTRL registers which are implemented: 0x00: No GPO{x}_CTRL register are implemented 0x01: 1 GPO{x}_CTRL register is implemented, starting with GPO0_CTRL 0x02: 2 GPO{x}_CTRL registers are implemented, starting with GPO0_CTRL up to GPO1_CTRL ... 0x20: 32 GPO{x}_CTRL registers are implemented, starting with GPO0_CTRL up to GPO31_CTRL For SSE-710 this field always reads as 0x01.	RO	0x01

12.3.7.15 Device Type Identifier (DEVTYPE) register

The following table gives a bit-level description of the Device Type Identifier register.

Table 12-189: DEVTYPE register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	SUB	Sub type for the component device type: 0x0: Indicates Other, undefined	RO	0x0
[3:0]	MAJOR	Major type for the component device type: 0x0: Indicates Miscellaneous	RO	0x0

12.3.7.16 Peripheral ID 4 (PID4) register

The following table gives a bit-level description of the Peripheral ID 4 register.

Table 12-190: PID4 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	Size	Number of 4KB units occupied by the System ID block. This field is deprecated.	RO	0x0
[3:0]	DES_2	JEP Continuation	RO	0x4

12.3.7.17 Peripheral ID 5 (PID5) register

The following table gives a bit-level description of the Peripheral ID 5 register.

Table 12-191: PID5 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved.	RO	0x0000_0000

12.3.7.18 Peripheral ID 6 (PID6) register

The following table gives a bit-level description of the Peripheral ID 6 register.

Table 12-192: PID6 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.7.19 Peripheral ID 7 (PID7) register

The following table gives a bit-level description of the Peripheral ID 7 register.

Table 12-193: PID7 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

12.3.7.20 Peripheral ID 0 (PID0) register

The following table gives a bit-level description of the Peripheral ID 0 register.

Table 12-194: PID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PART_0	Bits [7:0] of part ID	RO	0xF0

12.3.7.21 Peripheral ID 1 (PID1) register

The following table gives a bit-level description of the Peripheral ID 1 register.

Table 12-195: PID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	DES_0	Bits [3:0] of JEP 106 Identity	RO	0xB
[3:0]	PART_1	Bits [11:8] of part ID	RO	0x9

12.3.7.22 Peripheral ID 2 (PID2) register

The following table gives a bit-level description of the Peripheral ID 2 register.

Table 12-196: PID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVISION	Major revision of the System ID block	RO	0x0
[3]	JEDEC	Indicates the use of JEDEC JEP106 identification scheme	RO	0b01
[2:0]	DES_1	Bits [6:4] of JEP 106 Identity	RO	0b011

12.3.7.23 Peripheral ID 3 (PID3) register

The following table gives a bit-level description of the Peripheral ID 3 register.

Table 12-197: PID3 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVAND	Minor revision of the System ID block	RO	0x0
[3:0]	CMOD	Customer modification field	RO	0x0

12.3.7.24 Component ID 0 (CID0) register

The following table gives a bit-level description of the Component ID 0 register.

Table 12-198: CID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_0	Preamble 0	RO	0x0D

12.3.7.25 Component ID 1 (CID1) register

The following table gives a bit-level description of the Component ID 1 register.

Table 12-199: CID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	CLASS	Class of the component	RO	0x9
[3:0]	PRMBL_1	Preamble 1	RO	0x0

12.3.7.26 Component ID 2 (CID2) register

The following table gives a bit-level description of the Component ID 2 register.

Table 12-200: CID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_2	Preamble 2	RO	0x05

12.3.7.27 Component ID 3 (CID3) register

The following table gives a bit-level description of the Component ID 3 register.

Table 12-201: CID3 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_3	Preamble 3	RO	0xB1

13. Boot

This chapter describes the boot requirements and flow of the SSE-710 subsystem.

13.1 Boot overview

The SSE-710 boot is designed to protect the subsystem.

When **PORESETn** is released on SSE-710, the Secure Enclave is the only system that initially boots.

Both the Host and External Systems are prevented from executing any instructions, but access to resources within the SSE-710 is allowed. For example, the Secure Enclave can access the counters in the Host System to initialize the time domains.

SSE-710 does not mandate a specific boot flow but certain steps are required to be completed during the boot process.

13.2 Boot requirements on page 289 defines the boot requirements for SSE-710. 13.3 Example boot flow on page 290 shows an example boot flow, for reference only.

13.2 Boot requirements

This section describes actions that must be performed during the boot process. Certain steps can be performed by more than one system.

The following table shows the actions required as part of the boot process. The order of the rows does not indicate any ordering in the steps required.

Table 13-1: Boot requirements

Step	Secure Enclave	Host System	External System
Configure the Secure Enclave, including: • Enable the SoC and Secure Enclave watchdogs • Configure the Secure Enclave Firewall, to allow access to the Host System	Y	N	N
Initialize Host System Firewall: program the Host System Firewall to gain access to the CVM, OCVM, and NVM			
Host System Firewall full programming: program the Host System Firewall for all masters and slaves	Y[a]	Y[a]	
Initialize S32K counter	Y[b]	Y[b]	

[a] Full programming of the Host System Firewall can be done by either the Host CPU or Secure Enclave, depending on the system requirements.
[b] Initializing the S32K and REFCLK counters is only required to be done by one system.

Step	Secure Enclave	Host System	External System
Initialize REFCLK counter	Y[bc]	Y[bc]	
Initialize NVM : configure the non-volatile memory, the Flash Controller	Y	Y	
Load Secure Enclave firmware into Secure Enclave RAM and authenticate		N	
Load Host System Secure firmware into volatile memory and authenticate			
Authenticate Host System Non-secure firmware and OS			
Authenticate External System firmware[d]			
Write to the Host System Boot Register			
Release the Host CPU: • Clear the CPUWAIT in the Secure Enclave Base System Control registers. • Wake the respective Host CPU core by performing an action which causes the Host CPU Core power domain to exit the OFF power mode.			
Release the External System: clear the CPUWAIT, for the External System, in the Host System Base System Control registers	Y[e]	Y[e]	

13.3 Example boot flow

This section contains an example boot flow for reference only.

1. The Secure Enclave is released from reset and starts executing from the Secure Enclave ROM.
2. This ROM code configures the Secure Enclave system, which includes programming the Secure Enclave Firewall to gain access to the Host System.
3. Secure Enclave performs initial programming of the Host System Firewall. This includes:
   a. Granting access to the volatile and non-volatile memory for both the Host CPU and Secure Enclave. At this stage Arm recommends that only Secure access is granted.
   b. Granting access to any peripherals which are required. For example, if the Secure Watchdog is to be enabled, grant access to the REFCLK counter and Secure Watchdog.


Note

For the Secure Watchdog to be used, REFCLK counter must be initialized beforehand.

4. The Secure Enclave configures access to the Non-volatile memory, for example the Flash Controller.

[c] The REFCLK counter needs to be re-initialized every time SSE-710 exits the BSYS.SLEEP1 or BSYS.OFF power states.
[d] If the External System firmware is stored in a separate non-volatile storage device to the one connected to the NVM interface, the Secure Enclave must be able to access the storage device to authenticate the image.
[e] The External System can be released by either Secure Enclave or Host CPU, but only once the External System firmware has been authenticated.

5. The Secure Enclave loads and authenticates its own firmware into the Secure Enclave RAM. Decryption can be applied at this point.


Note

To avoid time-of-check to time-of-use attacks, the Secure Enclave must load the firmware image into the Secure Enclave RAM before performing the authentication. This task can be processed in parallel to reduce the time taken.

6. If the authentication succeeds, the Secure Enclave starts to execute from the loaded image. If authentication fails, a backup image is used if available, or it enters a recovery mode.
7. The Secure Enclave loads the Host CPU Secure firmware into the on-chip volatile memory and authenticates it. No Non-secure firmware is required to be loaded at this point. Decryption can be applied at this stage. If authentication fails, the Secure Enclave can attempt to authenticate the backup image, if available, or enter a recovery mode.
8. The Secure Enclave writes the Boot Register with data values of instruction opcodes to cause the Host CPU to branch into the correct location in its Secure firmware stored in on-chip volatile memory.
9. The Secure Enclave sets HOST_SYS_RST_CTRL.CPUWAIT to `0b0` to allow the Host CPU processors to execute instructions.
10. The Secure Enclave writes `0b1` to HOST_CPU_WAKEUP.CORE0_WAKEUP to request Core 0 to start executing instructions from the Boot Register.


Note

In this example, Core 0 is the Core selected to perform the initial boot but any implemented Core or multiple Cores could be used.

11. The Host CPU writes `0b0` to the HOST_CPU_WAKEUP.CORE0_WAKEUP to remove the wakeup request for Core 0.
12. The Host CPU Secure firmware configures the system. For example, the clock and power settings for the SoC.
13. At this stage, the full programming of the Firewall must be performed. This can be done by either the Host CPU Secure firmware or the Secure Enclave.
14. The Host CPU Secure firmware requests authentication of the Non-secure firmware and Rich OS from the Secure Enclave. If authentication fails the Host CPU can request authentication of a backup image, if available, or enter a recovery mode.
15. The Host CPU boots the Non-secure firmware and Rich OS. It is **IMPLEMENTATION DEFINED**, whether the Non-secure firmware and Rich OS are:
    - Loaded into the On-chip volatile memory.
    - Loaded into the Off-chip volatile memory.
    - Executed in-place.
16. The Rich OS requests authentication of the External System firmware from the Secure Enclave. If required, any decryption or transferring to a local memory within the External System can be

performed. If authentication fails, the Rich OS can request authentication of a backup image if available, or enter a recovery mode.

17. The Rich OS sets the EXT_SYS{0-1}_RST_CTRL.CPUWAIT to `0b0` to allow the External System processors to start to execute instructions.


Note

If there is more than one External System, steps 16 and 17 must be performed for each External System. It is **IMPLEMENTATION DEFINED** whether the Rich OS performs this in series or in parallel.

14. Software sequences

This chapter describes the software sequences that control functions within the SSE-710 subsystem.

In this section, there are details of software sequences that control or perform certain functions in SSE-710 subsystem.

14.1 Power control

This section shows example software sequences for controlling the power modes of the domains within SSE-710.

Arm® recommends that these examples are used as starting points for development of SoC firmware.

Arm® recommends that all power control sequences are performed by Secure software, either on the Secure Enclave or the Host CPU.

In all software sequences for power control, software is responsible for ensuring that all clocks used within the domain, and by the power control logic, are configured to use a clock source which is available in the current power state of the SoC.

Software may be required to change the clock source before starting a power sequence.

For example, if the clock source for the SYSPLL is in the SYSTOP domain, whenever the SYSTOP domain enters the OFF or MEM_RET power-modes, the SYSPLL source is no longer available.

Software is required to switch any clock source that is currently using SYSPLL to another clock source before allowing SYSTOP to enter the OFF or MEM_RET power-mode. Software can then re-enable the use of the SYSPLL when the SYSTOP domain enters the FUNC_RET or ON power mode and the SYSPLL output has been locked.

For details of the SECENC power doman, see 14.8.1 SECENCTOP power domain on page 308.

For more information regarding power domains, see 6.4 Power domains on page 84.

14.1.1 CORE{0-3} and CLUSTOP power domains

This section describes the software sequences controlling the CORE{0-3} and CLUSTOP power domains.

The *Power Control System Architecture* describes the software sequences for the CORE{0-3} and CLUSTOP power domains, with the following exceptions or additional requirements:

- There is no *System Control Processor* (SCP), and the expected behaviour is the PPU policy and is left at its default value of dynamic OFF.
- Execution of a WFE instruction does not cause entry into the FULL_RET power mode for a core.
- Before the last Host CPU core enters the OFF-power mode software must:
  - Save the configuration of the Host GIC if the HOST_CPU_CLUS_PWR_REQ.PWR_REQ is set to `0b0`
  - If the BSYS_PWR_REQ.SYSTOP_PWR_REQ field is `0b000` or `0b001`, it must save the configuration of components in the SYSTOP domain
  - If the BSYS_PWR_REQ.REFCLK_REQ field is `0b0`, it must move any future **REFCLK** time events to the S32K time domain

The selection of the power mode, which the CLUSTOP power domain enters when it is quiescent and all CORE{0-3} domains are in OFF, is set by the HOST_CPU_CLUS_PWR_REQ register. Arm recommends that software sets this register when the last Host CPU core is entering the OFF power mode.

Other things to consider before the last Host CPU enters the OFF-power mode:

- Software should configure the wakeup conditions. For example, using the **REFCLK** timer 0 to generate an interrupt at a point in time:
  - If the software does not want the Host CPU to wake up when CLUSTOP is in OFF or MEM_RET power-mode, it must set the Host Base System Control BSYS_PWR_REQ.WAKEUP_EN bit to `0b0`.
- The response time of the system to wakeup events. This determines the minimum power states of the SSE-710 subsystem and the allowed power modes its domains can enter.

14.1.2 SYSTOP power domain

The SYSTOP power domain is shared between all systems within SSE-710 subsystem.


Note

Arm® assumes that the SYSTOP PPU power policy is set to dynamic OFF. Programming of another policy might lead to unexpected behavior.

Both the Secure Enclave and the Host System Base System Control registers contain the BSYS_PWR_REQ.SYSTOP_PWR_REQ field which enables software to request the minimal power mode of the SYSTOP domain.

The hardware automatically selects the power-mode which the SYSTOP domain enters based on hardware indicators and the values in the BSYS_PWR_REQ.SYSTOP fields. The hardware selects the minimum power mode which meets all the requirements. Software must be able to handle the case where SYSTOP does not enter the power-mode requested in the BSYS_PWR_REQ.SYSTOP field.

Software must do the following:

- Before the software attempts to access the volatile memory, it must set the BSYS_PWR_REQ.SYSTOP_PWR_REQ to a value other than `0b000`. Failure to do so can lead to loss of data when there are no outstanding transactions in the SYSTOP domain.
- Before the software sets the BSYS_PWR_REQ.SYSTOP_PWR_REQ field to a value where there could be loss of data, software makes sure that it has saved all required information. Software is then responsible for restoring the information when it sets the field to a value where the data cannot be loss.

Arm® recommends the following for controlling the SYSTOP power domain:

- Software does not change the default PPU_PWPR policy register for the SYSTOP PPU.
- Software uses the BSYS_PWR_REQ registers in the Host or Secure Enclave Base System Control blocks to control the power mode of the domain.
- During boot of the SoC, software sets the BSYS_PWR_REQ.SYSTOP_PWR_REQ to a value other than `0b000`. Software does not change this value unless it no longer requires the contents of the volatile memory.
- The BSYS_PWR_REQ register of the Host System Base System Control registers is only accessed by one agent, either secure software executing on the Host CPU or software executing on the Secure Enclave. Any request to make changes to the power modes of the External System are made via messages sent to either the Host CPU or Secure Enclave using the MHUs in the system. It is not required for the External System to send a message when it needs access to resources inside SYSTOP, but only when it requires that SYSTOP does not enter a specific power mode were lost of data could occur.
- Software considers the implications of powering down of the SYSTOP domain. For example, if **SYSPLL** is in the SYSTOP domain then any clocks which are generated from **SYSPLL** are switched to another clock source.
- Software programs the entry delay values of the SYSTOP PPU to a value which meets the requirements of the current operating conditions of the SoC. By default, SYSTOP PPU places the SYSTOP domain into the lowest power mode when all indications, both hardware and software, reveal that there is no requirement for SYSTOP to be in the ON power mode. This can impact the performance if the traffic pattern of accesses to the SYSTOP domain is regular, but has a large period in between. This can be the case where the External System is only issuing a single access at a time to the Host System.

14.1.3 DBGTOP power domain

The DBGTOP power domain is shared between all systems in SSE-710 subsystem.

This section is written with the expectation that the DBGTOP PPU power policy is set to dynamic OFF. Programming of another policy may lead to different behavior.

Both the Secure Enclave and Host System Base System Control registers contain the BSYS_PWR_REQ.DBG_PWR_REQ field, which allows software to request the minimal power mode of the DBGTOP power domain. The DBGTOP domain can also be controlled using the CDBGPWRUPREQ field of the DBGROM table. The expected usage is that software running on the SoC uses the Base System Control registers, while a JTAG or SWD debugger uses the CoreSight™ ROM tables.

Similarly to the SYSTOP domain, hardware autonomously transitions the DBGTOP domain using the values in the BSYS_PWR_REQ.DBGTOP_PWR_REQ fields, and hardware indicators, to select the lowest power mode which satisfies all the request.

The DBGTOP domain software or debug agent must do the following:

- Before attempting to access any component in the DBGTOP domain, it must either request the DBGTOP domain is powered using either the BSYS_PWR_REQ.DBGTOP_PWR_REQ of the Base System Control registers, or the CDBGPWRUPREQ field of the DBGROM table. Failure to do so leads to the transaction completing with an error, except for accesses to the CoreSight™ STM-500 Extend Stimulus Port, which are treated as RAZ/WI and do not generate an error.
- After setting the BSYS_PWR_REQ.DBGTOP_PWR_REQ field to `0b1`, software must wait for BSYS_PWR_ST.DBGTOP_PWR_ST field to read as `0b1`, indicating that the DBGTOP domain is ON, before accessing the components of the DBGTOP domain.

  Note

  The BSYS_PWR_REQ.DBG_PWR_REQ and BSYS_PWR_ST.DBGTOP form a handshake between software and the hardware.

- After setting the CDBGPWRUPREQ field of the DBGROM table to `0b1`, the debug agent must wait for the CDBGPWRUPACK field to read as `0b1`, indicating that the DBGTOP domain is ON, before accessing the components of the DBGTOP domain.

  Note

  The CDBGPWRUPREQ and CDBGPWRUPRACK fields form a handshake between the debug agent and the hardware.

  Arm® recommends in both cases that the debug agent only accesses the DBGTOP domain when both fields which form the handshake read as `0b1`.

- When either software or the debug agent no longer require the debug components, and needs to enter the DBGTOP domain into OFF, they must perform the following before:
  - Disable any trace sources they were using, (if any), and perform a flush of the trace infrastructure at the trace sink (TPIU, Host ETR or SoC ETR).
  - Disable any enabled triggers.

Arm® recommends the following for controlling the DBGTOP power domain:

- Software does not change the default PPU_PWPR policy register for the DBGTOP PPU.
- Software uses the BSYS_PWR_REQ registers in the Host or Secure Enclave Base System Control blocks to control the power mode of the domain.

- JTAG or SWD debug agent uses the CDBGPWRUPREQ and CDBGPWRUPACK fields of the DBGTOP ROM table to control the power mode of the domain.
- The BSYS_PWR_REQ register of the Host System Base System Control registers is only accessed by one agent, either secure software executing on the Host CPU or software executing on the Secure Enclave. Any request to make changes to the power modes of the External System are made through messages sent to either the Host CPU or Secure Enclave using the MHUs in the system.
- When the software or the debug agent has finished using the debug, it makes sure that no the configuration of the debug components exposes any information about the system.

14.2 Time domains

In SSE-710 subsystem there are two time domains: REFCLK and S32K.

The REFCLK is used by the Host CPU, REFCLK timers {0-3}, Secure and Non-secure watchdogs, and can also be used by other components added by the integrator.

In the BSYS.SLEEP1 power state, the **REFCLK** time domain no longer increments and it is the responsibility of the software to restore the REFCLK time domain from the S32K domain, which has continued to increment in the BSYS.SLEEP1 power state.


Note

On exit from the BSYS.OFF power state, none of the two time domains are valid and must be restored from another source.

The following equation describes the relationship between the S32K time domain and the REFCLK time domain. The current **REFCLK** time can be calculated from the combination of the following:

- Current S32K time
- Ratio between the frequency of **S32KCLK** and **REFCLK**
- Known time values of the S32K and the **REFCLK** counters at a single point

Figure 14-1: The relationship between S32KCLK and REFCLK

$$t_{REF} = \left(\left(\frac{f_{REF}}{f_{32K}} \right) (t_{32K} - T_{32K}) \right) + T_{REF}$$

Where:

- t_{REF}: current time value in the **REFCLK** time domain
- t_{32K}: current time value in the S32K time domain
- T_{REF}: time value in the **REFCLK** time domain or CoreSight timestamp time domain at a synchronization point in the past

- T_{32K}: time value in the S32K time domain at a synchronization point in the past
- f_{REF}: frequency of **REFCLK**, in Hz
- f_{32K}: frequency of **S32KCLK**, in Hz

This relationship enables **REFCLK** time to be linearly scaled and offset from S32K time. The scaling factor is constant. The offse, T_{REF} and T_{32K}, must be recalculated at each power down so that all applications and software updates are preserved, as the following figure shows.

Figure 14-2: Time domain relationship

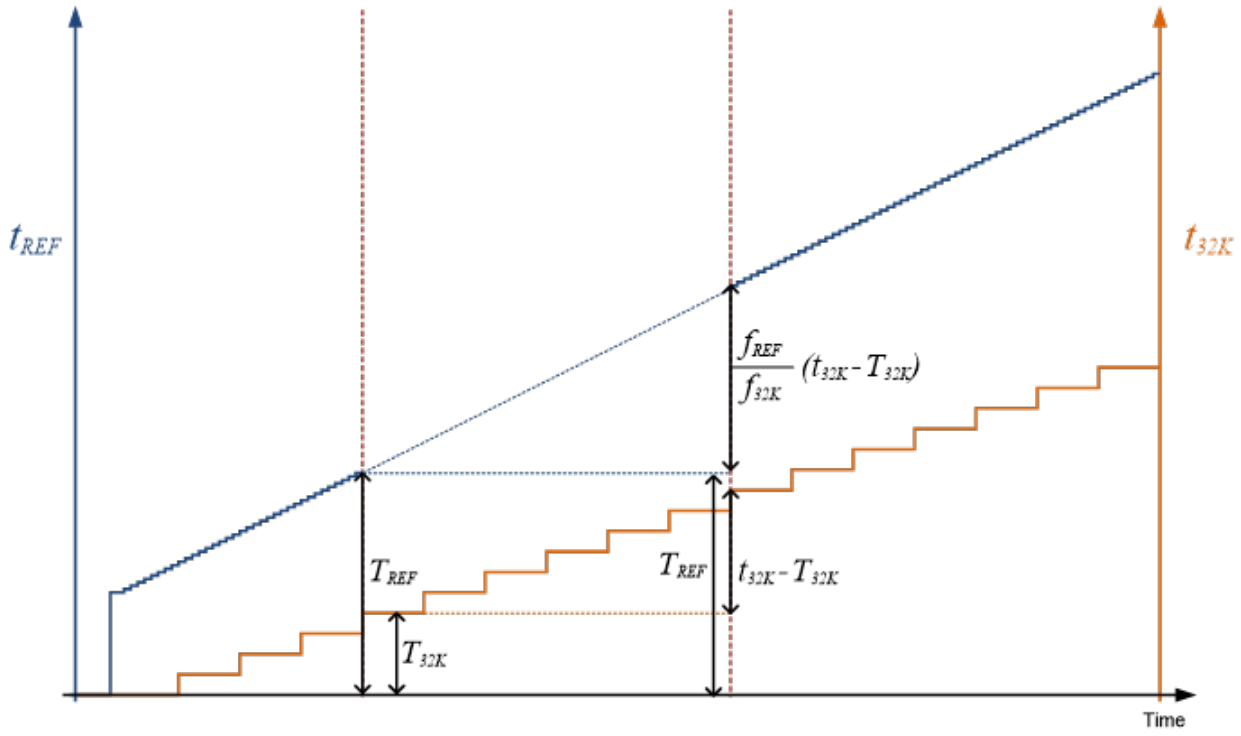

Restoring the REFCLK time domain

To restore the **REFCLK** time domain after it has been disabled, the software must calculate the new **REFCLK** domain time using the current S32K domain time and update the value of the **REFCLK** Counter.

For maximum accuracy this update can be timed to a **S32KCLK** clock edge. 12.3.4 REFCLK Counter CNTControl register summary on page 264 describes the registers that enable the **REFCLK** time value to be set against the rising edge of **S32KCLK**.

14.3 Debug agents

This section describes the SSE-710 debug agents.

14.3.1 Certificate injection

This section describes the procedure to inject a certificate to enable SSE-710 debug privileges.

The SSE-710 subsystem has a lifecycle state, which is used to define default values for CoreSight™ Authentication signals used within the SoC.

The lifecycle state of the SoC is managed by the Secure Enclave of SSE-710 subsystem. The lifecycle state can be one of the following:

- Chip Manufacture
- Device Manufacture
- Secure Enable
- Return Merchandise Authorization

For more information on the lifecycle states of the SoC, see 9.1.2 Lifecycle States (LCS) on page 148.

For example, in the Secure Enable lifecycle state all CoreSight™ Authentication signals are LOW indicating that all debug functionality is disabled.

To enable a JTAG/SWD debugger to access the debug infrastructure of the SSE-710 subsystem and wider SoC, the SSE-710 subsystem includes a CoreSight™ SDC-600. The CoreSight™ SDC-600 provides a standard method for a JTAG debugger to communicate with software running on the SoC.

The CoreSight™ SDC-600 is accessible by both JTAG/SWD and the software running on the SoC. Arm recommends that software does not use the CoreSight™ SDC-600 for certificate injection and instead communicates directly with the software controlling the CoreSight™ Authentication signals. For software outside the Secure Enclave, this communication requires the use of the MHUs to the Secure Enclave.

The CoreSight™ SDC-600 is split into two components EXT APBCOM and INT APBCOM. For more information on these components, see the *Arm® CoreSight™ SDC-600 Secure Debug Channel Technical Reference Manual*.

- The EXT APBCOM is used by the JTAG/SWD debug agent.
- The INT APBCOM is used by software in theSSE-710 subsystem. The INT APBCOM is part of the Host System memory map but the interrupt can be routed to either the Host CPU or the Secure Enclave, through the Interrupt Router.

It is **IMPLEMENTATION DEFINED**, whether software running on the Host CPU or the Secure Enclave manages the CoreSight™ SDC-600 INT APBCOM.

An example sequence for injecting a certificate using a JTAG/SWD debugger is as follows:

1. Debugger connects to the target SoC and reads the DP ROM table to discover the location of the EXT APBCOM.
2. Debugger uses the EXT APBCOM to establish a link with the software running on the target SoC, using the steps defined in *Advanced Communications Channel Architecture Specification*. The software running on the target SoC uses the INT APBCOM to receive messages from the debug agent.
3. When the link is established, the debugger transmits the certificate. Software receives the certificate and validates it and does one of the following:
   - If the certificate is valid, it updates the required Security Control Bits, which drive the associated CoreSight™ Authentication signals for the DAZ and DAACGs. It also provides an acknowledgment back to debugger indicating successful validation.
   - If the certificate is invalid, no update to the Security Control Bits occurs. It is **IMPLEMENTATION DEFINED** whether an error message is provided back to the debugger and what action the system and debugger takes when this occurs.
4. Debugger either:
   - If the requested debug privileges have been enabled, the debugger continues with the debug activity.
   - If the requested debug privileges were not enabled, the debugger retries or reports an error.

When the debugger has finished sending messages through CoreSight™ SDC-600, it must terminate the link using the steps defined in the *Advanced Communications Channel Architecture Specification*.

Arm recommends that before transitioning the device to a lifecycle state where debug functionality is disabled, the code used to perform certificate authentication has been tested.

SSE-710 only supports the following ways of injecting a certificate:

Runtime authentication

The certificate is inserted at any time and processed by the target system.

Boot-time authentication, where nSRST is used to enter boot mode

The certificate is inserted at boot time. The debug causes the SoC to perform a boot by causing an **nSRST** reset to occur. Triggering of **nSRST** can be done using either the **nSRST** or setting CSYSRSTREQ in the DP ROM to `0b1`.

14.3.2 Debug from Reset

Debug from Reset is the process of debugging specific CPUs from reset release in the SSE-710 subsystem.

The following CPUs can be debugged from reset release:

- Secure Enclave Cortex®-M0+
- Host CPU
- External System's CPUs

To perform Debug from Reset on the Secure Enclave Cortex®-M0+ the debug agent using the JTAG/SWD must follow the following steps:

1. Debug agent either asserts **nSRST** input or sets the CSYSRSTREQ bit in the DP ROM table and wait for the CSYSRSTACK bit to become `0b1`.
2. Debug agent configures the debug functionality required. It might be required that the debug agent perform power requests to be able to access certain debug registers. For example, the debug logic of the Host CPU core.
3. Once the debug agent has finished configuring the required debug functionality, it either de-asserts the **nSRST** input or clears the CSYSRSTREQ bit in the DP ROM table. At this point the boot process continues as normal.

It is possible for the CPU of the Host and External System to perform Debug from Reset. The method by which software running on SSE-710 enables Debug from Reset of the Host and External Systems is **IMPLEMENTATION DEFINED**.

Whether it is possible to perform Debug from Reset, depends on the settings of the DAZs in the SSE-710 subsystem. By default, in the Secure lifecycle the SECENCAUTH prevents debug access to the Secure Enclave Cortex®-M0+ and the debug agent is required to insert a certificate to request access. The processing of a certificate requires software to run on the SSE-710. It is not possible to insert the certificate and then cause a debug reset using either the **nSRST**, or DP ROM **CSYSRSTREQ/ACK** handshake as this resets the SCBs to the default value for the lifecycle state.

Arm® strongly recommends the following:

- The Secure Enclave ROM code includes the ability to handle certificate insertion using the CoreSight™ SDC-600 to allow for debugging of all non-ROM based boot code.
- The Secure Enclave Cortex®-M0+ checks for a possible certificate insertion as part of the boot process. This allows a debug agent to start inserting a certificate whilst preventing the Secure Enclave Cortex®-M0+ from executing software using either the **nSRST** input or the DP ROM CSYSRSTREQ/ACK handshake.
- Before transitioning to the Secure lifecycle state, the ability of the Secure Enclave Cortex®-M0+ to process a certificate is confirmed working, to allow for debug as early as possible in the boot process.

It might be required to de-assert **nSRST** input or clear CSYSRSTREQ field in the DP ROM earlier if the software is required to perform certain actions. For example, powering a specific domain.

If this is required, then an **IMPLEMENTATION DEFINED** software mechanism can be implemented to indicate that a debug agent is performing Debug from Reset on a specific target processor. If the debug agent is performing Debug from Reset on more than one processor in the SoC, and one of these processors is required to be released from reset to enable access, the debug agent must stage the configuration of the debug. At first, it programs the debug logic of the processor, which must be released from reset before de-asserting **nSRST** input or clearing CSYSRSTREQ field in the DP ROM. Once the other processor is accessible, it continues with the configuration of the debug.

A debugger can also use the **nSRST** input or DP ROM CSYSRSTREQ bit to reset the SoC to a known state without the need for Debug from Reset. In this case, the debug agent follows the steps above but does not perform any configuration at step2 of the process.

14.3.3 Using the External Debug Bus

The SSE-710 subsystem has an External Debug Bus enabling access to the Access Ports included in the SoC from the Host System, Secure Enclave or an External System.

The debug agent must not do the following:

- Use an Access Port to target a location on the External Debug Bus. For example, using the AXI-AP to access the Access Port of an External System. The debug agent should instead directly access the Access Port of the External System.
- Use the External Debug Bus to perform self-hosted debug of the system on which the debug agent is running. For example, to perform self-hosted debug of the Host system, the debug agent running on the Host CPU should access the debug components using the Host System Debug region of the Host System memory map.

Arm® expects the External Debug Bus to be used in one of the following scenarios:

- Debug agent running on the Host CPU to debug one or more of the External Systems
- Debug agent running on an External Systems to debug one or more of the other External Systems and/or the Host System

In both scenarios, the system on which the Debug agent is running can also be being debugged by the debug agent, however, the debug agent must not use the External Debug Bus to configure the debug components of its system or to perform any memory access to the system's memory map.

Recommendations

Arm® strongly recommends that the Secure Enclave is not used to run the debug agent. It is possible to have the debug agent run on the Secure Enclave. However, this is not the expected use-case because the Secure Enclaveprovides the Root of Trust to the system.

Arm® strongly recommends that the Secure Enclave is only debugged by a debug agent connecting through the Debug Port. It is possible for the Secure Enclave to be included in the list of systems being debugged in the scenarios above. However, this is not expected as the debug agent typically runs in an environment that is less trusted compared to the Secure Enclave. Allowing it to debug the Secure Enclave would compromise security.

14.4 Host and External System {0-1} reset request

For both the Host and External System {0-1} reset requests, there is a handshake between software and the Reset Controller using one of the following:

- HOST_SYS_RST_CTRL and HOST_SYS_RST_ST registers for the Host System reset request.
- EXT_SYS{0-1}_RST_CTRL and EXT_SYS{0-1}_RST_ST registers for the External System {0-1} reset request.

Software must follow this sequence:

1. Set RST_REQ to `0b1`.
2. Poll RST_ACK until the value is either `0b01` or `0b10`.
3. If the value is `0b10`, the requested reset has been accepted:
   a. Set RST_REQ to `0b0`.
   b. Wait for RST_ACK to become `0b00`.
4. If the value is `0b01`, the requested reset has been denied:
   a. Set RST_REQ to `0b0`.
   b. Wait for RST_ACK to become `0b00`.
   c. Software can then do one of the following:
      - Repeat the sequence again.
      - Request a large reset of the SoC:
        - For the External System {0-1} reset request, software can request a Host System or SoC reset
        - For the Host System reset request, software can request an SoC reset

In the sequence, above depending on whether software is attempting to reset the Host or an External System, the RST_REQ and RST_ACK fields are located in the:

- HOST_SYS_RST_CTRL and HOST_SYS_RST_ST respectively for the Host System reset request
- EXT_SYS{0-1}_RST_CTRL and EXT_SYS{0-1}_RST_ST respectively for the External System {0-1} reset request

If software sets RST_REQ back to `0b0` before RST_ACK becomes `0b01` or `0b10`, software cannot guarantee whether the reset request was applied or not.


Note

Arm recommends that software does not repeatedly request a reset of the system if it gets denied and requests a wider reset of the SoC instead.

To guarantee that the reset request is completed successfully, Arm recommends:

- To reset an External System the External System is in a quiescent state with no outstanding transactions from or to the External System.

- To reset the Host System there are no outstanding accesses from the Secure Enclave into the Host System address space.

14.5 Watchdog usage

SSE-710 subsystem implements four different watchdogs. Each one is used for a specific functionality:

Non-secure watchdog

The Non-secure watchdog is used by the Rich OS to guard against it becoming unresponsive. The Non-secure watchdog uses the REFCLK time domain and is not operational in the BSYS.SLEEP1 power state. On the first expiry of the watchdog an interrupt is generated to the Host CPU and is taken by the Rich OS. If the watchdog expires for a second time, a second interrupt is generated. The second interrupt is taken by the Host CPU Secure firmware which takes actions to handle the unresponsive Rich OS.

Secure watchdog

The Secure watchdog is used by the Host CPU Secure firmware to guard against it becoming unresponsive. The Secure watchdog uses the **REFCLK** time domain and is not operational in the BSYS.SLEEP1 power state. On the first expiry of the watchdog an interrupt is generated to the Host CPU and is taken by the Host CPU Secure firmware. If the watchdog expires for a second time, a second interrupt is generated and routed to the Secure Enclave. The Secure Enclave takes actions to handle the unresponsive Host CPU Secure firmware.

Secure Enclave watchdog

The Secure Enclave watchdog guards against the Secure Enclave becoming unresponsive. The Secure Enclave watchdog uses the **SECENCDIVCLK** and is not operational when the SECENCTOP power domain enters the OFF or MEM_RET power-modes. On first expiry of the watchdog, an interrupt is generated to the Secure Enclave Cortex®-M0+. If the watchdog expires for a second time a reset request is generated and causes a reset of the entire SoC.

SoC Watchdog

The SoC watchdog guards against the entire SoC becoming unresponsive. It is operational in all power states of SSE-710 and uses the **S32KCLK**. On first expiry of the watchdog an interrupt is generated to the Secure Enclave Cortex®-M0+. If the watchdog expires for a second time a reset request is generated and causes a reset of the entire SoC.

The SoC watchdog uses **S32KCLK** to decrement the watchdog, and is allowed to be enabled when the **S32KCLK** is not stable. This is different to the S32K time domain that can only be enabled once the **S32KCLK** is stable. However, Arm® recommends that software programs a value which is long enough so not to cause unexpected watchdog reset requests to be generated.


Note

For more information on the Secure Enclave watchdog and SoC watchdog, see 9.1.6.4 Watchdogs on page 155.

All watchdogs must be enabled before the watchdog generates any interrupt or reset request.

Arm® recommends the following:

- Software configures the debug infrastructure and watchdogs to halt when performing halted debug.
- Software disables the Non-secure or Secure watchdogs when entering the BSYS.SLEEP1 power state. This avoids a watchdog expiry occurring whilst returning to the BSYS.SLEEP0 or BSYS.RUN power states.
- Secure Enclave firmware disables the Secure Enclave watchdog before entering the SECENCTOP power domain into the OFF or MEM_RET power-modes.

14.6 Lifecycle

An SoC based on SSE-710 subsystem has a lifecycle state.

For details of lifecycle states, see 9.1.2 Lifecycle States (LCS) on page 148.

The Lifecycle state is managed by an **IMPLEMENTATION DEFINED** method and the method for advancing the lifecycle state is outside the scope of SSE-710 subsystem, however, Arm® makes the following recommendations:

- Before advancing the lifecycle to a state whereby debug functionality is disabled by default, there must be a method for a debug agent to request that the debug functionality be re-enabled, using the method described in 14.3.1 Certificate injection on page 299.
- Before advancing the lifecycle to a state where the Bypass interface of the Host System Firewall is driven to 0 by default, the Secure Enclave firmware has a method to enable the Bypass of the Firewall. To enable the Bypass on the Host System Firewall, the Secure Enclave firmware must update the SCB to drive bit [37] to `0b1`.

  
  Note

  Arm® recommends that this is only included as part of a debug certificate. For details of SCB, see 9.1.3 Security Control Bits (SCB) on page 150.

- Before advancing the lifecycle to a state whereby the Bypass interface of the Secure Enclave Firewall is driven to 0 by default, the Secure Enclave firmware is able to program the Firewall to an extent where Host CPU is able to boot.

14.7 Lock control

The SSE-710 provides several locks which software can use to prevent update to certain components and configuration registers.


Note

The value read in HOST_SYS_LCTRL_ST register can be a stale copy and software must handle this. The software sequences below take this into consideration.

Host CPU lock

The Host CPU lock is controlled using the HOST_CPU{0-3}_LOCK field in the HOST_SYS_LCTRL_{SET/CLR} registers. The status of the lock can be seen in the HOST_SYS_LCTRL_ST.HOST_CPU{0-3}_LOCK field.

The Host CPU lock is cleared automatically by hardware when:

- The Core{0-3} power domain enters one of the following power modes: OFF, OFF_EMU, or WARM_RST.
- The CLUSTOP power domain enters one of the following power modes: OFF, MEM_RET, or WARM_RST.

The software sequence to set the Host CPU lock is as follows:

1. Core enters the ON power-mode.
2. Software performs configuration of the registers which are affected by the **CP15SDISABLE** signal. For more information, see *Arm® Architecture Reference Manual Armv8, for Armv8-A architecture profile*. Software can perform other steps before this, but Arm® recommends this is kept to a minimum.
3. Software writes `0b1` to the HOST_SYS_LCTRL_SET.HOST_CPU{0-3}_LOCK field associated with the core.
4. Software checks the status of the lock by reading the HOST_SYS_LCTRL_ST.HOST_CPU{0-3}_LOCK field.


Note

Arm® strongly recommends the software performing this sequence is executing on the Host CPU core to be locked.

If the software performing the sequence is not executing on the Host CPU core to be locked, Arm® strongly recommends that the Host CPU core to be locked does not enter any of the following power modes: OFF, OFF_EMU, or WARM_RST once it has completed step 2 until it has performed step 4. This avoids race conditions between clearing and setting the lock.

The software sequence to clear the Host CPU lock is as follows:

1. Software writes `0b1` to HOST_SYS_LCTRL_CLR.HOST_CPU{0-3}_LOCK field associated with the core to be unlocked.
2. Software waits for the associated HOST_CPU{0-3}LOCK field in the HOST_SYS_LCTRL_ST register to become `0b0`, before performing any configuration of the locked registers in the Host CPU core.

Host GIC lock

The Host GIC lock is controlled using the HOST_GIC_LOCK field in the HOST_SYS_LCTRL_{SET/CLR} registers. The status of the lock can be seen in the HOST_SYS_LCTRL_ST.HOST_GIC_LOCK field.

The Host GIC lock is cleared automatically by hardware when the CLUSTOP power domain enters one of the following power modes: OFF, MEM_RET or WARM_RST.

The software sequence to set the Host GIC lock is as follows:

1. Configuration of the GIC configuration registers affected by the **CFGSDISABLE** signal. For more information, see *Arm® Generic Interrupt Controller Architecture Specification, architecture version 2.0*.
2. Software writes `0b1` to the HOST_SYS_LCTRL_SET.HOST_GIC_LOCK field.
3. Software checks the status of the lock by reading the HOST_SYS_LCTRL_ST.HOST_GIC_LOCK. If the field is `0b0` then software repeats the sequence.


Note

Arm® strongly recommends that software performing this sequence is excuting on one of the Host CPU cores.

If the software performing this sequence is not executing on one of the Host CPU cores, Arm® strongly recommends that the CLUSTOP domain does not enter any of the following power modes: OFF, MEM_RET or WARM_RST during the sequene. This avoids race conditions between clearing and setting the lock.

The software sequence to clear the HOST CPU lock is as follows:

1. Software writes `0b1` to HOST_SYS_LCTRL_CLR.HOST_GIC_LOCK field.
2. Software waits for the HOST_SYS_LCTRL_ST.HOST_GIC_LOCK field to become `0b0`, before performing any configuration of the locked registers in the Host GIC.

Other locks

The SSE-710 provides a Host Base System Control, Interrupt Router and Host System Firewall lock. The locks are controlled using the HOST_LOCK, INT_RTR_LOCK and HOST_FW_LOCK fields respectively, in the HOST_SYS_LCTRL_{SET/CLR} registers. The status of the lock can be seen in the same field in the HOST_SYS_LCTRL_ST register.


Note

For the Interrupt Router and Host System Firewall lock, the status of the lock can also be seen using registers in the component.

Lock Clear Disable

The default behavior of the locks in the SSE-710 allows software to remove the lock by writing `0b1` to the associated field in the HOST_SYS_LCTRL_CLR register. This behavior can be modified by setting the HOST_SYS_LCTRL_ST.LOCK_CLR_DIS field to `0b1`. This is done by software writing `0b1` to HOST_SYS_LCTRL_SET.LOCK_CLR_DIS field.

When the HOST_SYS_LCTRL_ST.LOCK_CLR_DIS field is `0b1` any writes to the HOST_SYS_LCTRL_CLR register are ignored and software cannot clear any locks.


Note

When the HOST_SYS_LCTRL_ST.LOCK_CLR_DIS is `0b1`, the HOST_SYS_LCTRL_ST.{HOST_CPU{0-3}_LOCK/HOST_GIC_LOCK} fields are still set to `0b0` when any of the conditions stated in sections Host CPU lock on page 306 and Host GIC lock on page 307 occur.

14.8 Secure Enclave software sequences

This section describes software sequences for the Secure Enclave.

14.8.1 SECENCTOP power domain

SSE-710 lets you control the SECENCTOP power domain.

SECENCTOP power domain is controlled using the following elements:

- PWR_GATE_EN field

  PWR_GATE_EN field in the Power Control register of the Secure Enclave System Control registers. For more information see 12.3.2.2 Secure Enclave System Control register summary on page 249.
- Power policy of the SECENC PPU
- Cortex®-M0+ in a DeepSleep state

Arm® strongly recommends that the SECENCTOP PPU policy is set to either dynamic OFF or dynamic MEM_RET, depending on whether software requires the retention of the contents of the Secure Enclave's SRAM. It is assumed that the PPU policy is either set to dynamic OFF or dynamic MEM_RET.


Note

Setting the SECENCTOP PPU policy to another policy may lead to different behavior.

To enter the SECENCTOP power domain into the MEM_RET or OFF power mode, software must:

1. Save any state which needs to be retained that otherwise is lost by entering either the MEM_RET or the OFF power state.
2. Set the PWR_GATE_EN bit `0b1`.
3. Set the SCR.SLEEPDEEP bit `0b1` and execute a WFI.

To select between entering MEM_RET or OFF, software can change the power policy in the PPU:

- To enter MEM_RET software must set the policy to dynamic MEM_RET.
- To enter OFF, software must set the policy to OFF.

Arm® strongly recommends that the SECENCTOP PPU is always programmed for dynamic transitions.

When the SECENCTOP power domain exits the MEM_RET or OFF power mode, software must set the PWR_GATE_EN bit `0b0` before performing any other actions.

14.8.2 Advancing lifecycle states

This section gives example sequences for advancing lifecycle states: one uses the SSE-710 **SOCLCC** signal, and the other uses the GPIO Control component.

14.8.2.1 Using the SOCLCC signal

The following example sequence uses the **SOCLCC** signal:

Procedure

1. Assert the **SOCLCC** signal of the SSE-710 subsystem HIGH.
2. De-assert **PORESETn** to the SSE-710 subsystem.
3. The Secure Enclave Cortex®-M0+ CPU performs an **IMPLEMENTATION DEFINED** software sequence to advance the lifecycle state. The software sequence depends on the Crypto Accelerator implemented in the Secure Enclave.
4. When the transition is complete, de-assert the **SOCLCC** signal LOW.

14.8.2.2 Using the Debug Port

The following example sequence uses the *Debug Port* (DP):

Procedure

1. Set the GPIO Control GPO0 output HIGH.
2. Cause a debug reset by either:
   - Asserting **nSRST** LOW.
   - Asserting **CSYSRSTREQ** of the DP ROM HIGH, and waiting for DP ROM **CSYSRSTACK** to become HIGH.
3. Release the SSE-710 subsystem from debug reset by:
   - De-asserting the **nSRST** HIGH.
   - De-asserting DP ROM **CSYSRSTREQ** LOW and waiting for the DP ROM **CSYSRSTACK** to become LOW.
4. The Secure Enclave Cortex®-M0+ CPU performs an **IMPLEMENTATION DEFINED** software sequence to advance the lifecycle state. This sequence depends upon the Crypto Accelerator implemented in the Secure Enclave.

Appendix A Message Handling Unit

This appendix describes the *Message Handling Unit v2.1* (MHUv2.1) component of the SSE-710 subsystem.

The MHUv2.1 is a memory mapped peripheral that provides a mechanism to facilitate interrupt based inter-processor message passing. The messages are passed in a single direction, from sender to receiver. The processors can be in different clock, reset and power domains. The MHUv2.1 can also handle unexpected resets. For example from a Watchdog.

Compliance

The MHUv2.1 complies with the following specification.

- *AMBA® APB Protocol Specification Version 2.0*
- *AMBA® Low Power Interface Specification, Arm® Q-Channel and P-Channel Interfaces*

A.1 Communication between systems in an SoC

SoCs can include multiple systems, each with a processing element and a different software stack. To perform the functions of the SoC, all these software entities must be able to communicate.

For example, a typical SoC for a cellular phone might contain the following systems, among others:

- Application processor: Executes a rich operating system, such as Linux.
- Communication processor: Manages communications with an external network through a modem or WiFi connection.
- System control processor: Handles system functions such as power and clock control.
- Secure Enclave: Provides security functionality to the rest of the SoC. To reduce the possibility of a malicious agent gaining access to secrets, a Secure Enclave is typically implemented as a separate system.

For the SoC to function as designed, the different software entities operating on each of these systems must communicate with one another. Various different communication methods can be used, for example:

- System calls, such as *Supervisor Calls* (SVCs) or *Secure Monitor Calls* (SMCs)
- Events
- Interrupts

Each type of communication has advantages and disadvantages, and is suited to different situations.

Only interrupt-based communication is described here. The use of system calls and events for communication is outside the scope of this section.

A.1.1 Interrupt-based communication

At the lowest level, interrupt communication depends on a sender passing an indication of an event to a receiver. This event indication is generated by a software operation, that is, an instruction or memory write.

Event indications are often passed to the receiver using a physical signal called an interrupt. Alongside the interrupt, additional information can be included either in-band or out-of-band. The interrupt and additional information together comprise a transfer, sent by the sender to the receiver. Several transfers then form a complete message, with the number of messages dependent on the message protocol that is used.

A.2 About the Message Handling Unit

The *Message Handling Unit* (MHU) is a hardware peripheral. It provides a mechanism to assert interrupt signals so that messages can be passed between different systems in an SoC.

The MHU is made up of two memory-mapped register frames. The first frame is used by the sender of the transfer, referred to as the Sender. The receiver of the transfer, denoted as the Receiver, uses the second frame. Each frame contains:

- Several channels that the Sender uses to pass transfers to the Receiver
- Control and identification registers

A.2.1 Channels

MHUs can be configured with multiple channels.

Each channel has the following registers:

- Channel Status (CH_ST): Shows the status of the channel. The way in which the status is displayed is determined by the transport protocol that is used. This register is included in both the Sender and Receiver channel windows.
- Channel Status Masked (CH_ST_MSK): Shows the status of the channel with the channel mask applied. This register is included in the Receiver channel window.
- Channel Set (CH_SET): Sets bits in the Channel Status and Channel Status Masked registers. This register is included in the Sender channel window.
- Channel Clear (CH_CLR): Clears bits in the Channel Status and Channel Status Masked registers. This register is included in the Receiver channel window.
- Channel Mask Status (CH_MSK_ST): Shows the status of the channel mask, which is used with the Channel Status register to generate the channel status mask. This register is included in the Receiver channel window.
- Channel Mask Set (CH_MSK_SET): Sets bits in the channel mask. This register is included in the Receiver channel window.
- Channel Mask Clear (CH_MSK_CLR): Clears bits in the channel mask. This register is included in the Receiver channel window.

- Channel Interrupt Status (CH_INT_ST): Shows whether a channel clear interrupt has occurred for the channel. This register is included in the Sender channel window.
- Channel Interrupt Clear (CH_INT_CLR): Clears the channel clear interrupt for the channel. This register is included in the Sender channel window.
- Channel Interrupt Enable (CH_INT_EN): Enables and disables the channel clear interrupt for the channel. This register is included in the Sender channel window.

An interrupt indicates that a transfer to the Receiver has occurred. A channel interrupt is generated when any bit in the CH_ST_MSK register is set to `0b1`. The interrupt is level-based and remains asserted until all bits in the register are set to `0b0`.

A channel can be subdivided into 32 flags, where each bit of the register controls a different flag or can be used as a transfer payload register. The way in which the channel is used depends on the transport protocol that is employed. For more information, see A.3 Transport protocols on page 334.

A.2.2 Transfers

MHU messages are composed of a series of transfers that are sent between the Sender and the Receiver.

In its most basic form, the Sender passes a transfer by writing `0b1` to one or more of the flag bits in the CH_SET register. This setting is reflected in:

- The CH_ST registers of both the Sender and the Receiver
- The CH_ST_MSK register of the Receiver, depending on the current setting of the CH_MSK_ST register. If any bit in the CH_ST_MSK register is set to `0b1`, then the interrupt for the channel is asserted.

The Receiver must clear the transfer by writing `0b1` to one or more flag bits in the CH_CLR register.

The way in which the Sender and the Receiver use the individual channels and flag bits depends on the transport protocol that is employed. For more information about the transport protocols that the MHU supports, see A.3 Transport protocols on page 334.

The number of transfers that are used to create a message depends on the:

- Size of the message that is sent
- Transport protocol that is used
- Number of channels that are implemented in the MHU
- Number of concurrent transfers between the Sender and the Receiver

A.2.3 Ready to Send protocol

Before sending a message, the Sender must ensure that the Receiver is ready receive the first transfer.

The Sender uses the ACCESS_REQUEST register to require that the Receiver enters a state in which the Receiver can accept a transfer. The Receiver indicates whether it is able to accept the transfer by setting the ACCESS_READY register.

When the **ACCESS_READY.ACC_RDY** signal is LOW, any attempt by the Sender to access the channel window registers is treated as RAZ/WI. When the RESP_CFG.NR_RDY field is set to `0b0`, an access attempt does not produce an error. But when the RESP_CFG.NR_RDY field is set to `0b1`, an error is generated.

Software requirements

Even when setting the ACCESS_REQUEST.ACC_REQ field to `0b1`, software is still responsible for guaranteeing that the Receiver has received the transfer using an **IMPLEMENTATION DEFINED** software protocol. The following steps describe an example message sequence:

1. The Sender configures the MHU to enable the *Not Ready to Ready* (NR2R) and *Ready to Not Ready* (R2NR) interrupts. The Sender enables the interrupts by setting the INT_EN.NR2R and INT_EN.R2NR fields to `0b1`.
2. The Sender sets the ACCESS_REQUEST.ACC_REQ field to `0b1`.
3. The Sender reads the status of the ACCESS_REQUEST.ACC_RDY field. If the field is set to `0b1`, then the Sender continues to the next step. Otherwise, the Sender waits for the NR2R interrupt to be triggered.
   An NR2R interrupt is generated when the Sender sets the ACCESS_REQUEST.ACC_REQ field to `0b1` and the ACCESS_READY.RDY field is already `0b1`. Software must clear this interrupt before continuing.
4. The Sender passes the transfer to the Receiver using one of the transport protocols that are defined in A.3 Transport protocols on page 334. While the Sender is performing the sequence of events for the transport protocol, the transfer is considered outstanding and is not guaranteed to reach the Receiver.
   While implementing the transport protocol, the Sender could receive an R2NR interrupt. Alternatively, if the RESP_CFG.NR_RDY field is set to `0b1`, the Sender might receive an error response to a channel window register access. If either of these situations occur, any transfers that were outstanding might be lost. So, the Sender must:
   a. Clear the R2NR interrupt and, if an error was generated, process the fault handler.
   b. Perform an **IMPLEMENTATION DEFINED** recovery sequence. This sequence can involve resending the transfer when the Receiver returns to a state in which it is able to accept transfers. Alternatively, other transfers can be sent to resynchronize the Sender and Receiver.
      When the Receiver enters a state in which it cannot accept a transfer, there is a possibility that the Receiver has been reset. Therefore, when the Receiver enters such a state, both the Sender and Receiver might need to resynchronize before they can continue to communicate.

5. If the Sender has more transfers to send, it repeats Step 4 until there are no more transfers to send. When the Sender has no more transfers to send, it sets the ACCESS_REQUEST.ACC_REQ field to `0b0`.
   The Sender can also enter a state in which it cannot complete the process of sending a transfer. For example, a watchdog might reset the Sender while it is writing to the CH_SET register. It is the responsibility of the Sender and the Receiver to clear and ignore any transfers that originated from before the Sender was reset.

   When the Sender enters a state in which it cannot complete the process of sending a transfer, there is a possibility that the Sender has been reset. Therefore, when the Sender enters such a state, both the Sender and Receiver might need to resynchronize before they can continue to communicate.

A.2.4 Interrupts

There are multiple interrupts in the MHU.

All interrupts in the MHU are level-based. Each interrupt has a register that indicates the interrupt status, a register for enabling the interrupt, and a register for clearing the interrupt.

The Sender frame has the following interrupts:

- NR2R (not ready to ready). This interrupt is asserted when the INT_EN.NR2R and INT_ST.NR2R fields are set to `0b1`. The INT_EN.NR2R field is set to `0b1` by the Sender to enable the NR2R interrupt. The INT_ST.NR2R field is set to `0b1` when:
  - The **ACC_RDY** signal goes HIGH and the **ACC_REQ** signal is `0b1`.
  - The **ACC_REQ** signal goes HIGH and the **ACC_RDY** signal is `0b1`.

  When software writes `0b1` to the INT_CLR.NR2R field in the Sender frame, the INT_ST.NR2R field is set to `0b0`.
- R2NR (ready to not ready). This interrupt is asserted when the INT_EN.R2NR and INT_ST.R2NR fields are set to `0b1`. The INT_EN.R2NR field is set to `0b1` by the Sender to enable the R2NR interrupt. The INT_ST.R2NR field is set to `0b1` when the **ACC_RDY** signal goes LOW. When software writes `0b1` to the INT_CLR.R2NR field in the Sender frame, the INT_ST.R2NR field is set to `0b0`.
- CHCOMB (combined). This interrupt is asserted when the INT_EN.CHCOMB field in the Sender frame is set to `0b1` and any of the following occur:
  - The CH_INT_EN.CH_CLR and CH_INT_ST.CH_CLR fields for any channel are both set to `0b1`. The CH_CLR bits relate to the channel clear interrupt.
  - The INT_EN.NR2R and INT_ST.NR2R fields are both set to `0b1`.
  - The INT_EN.R2NR and INT_ST.R2NR fields are both set to `0b1`.

The Receiver frame has a combined interrupt output that is set when both of the following occur:

- The INT_EN.CHCOMB field of the Receiver frame is set to `0b1`.
- At least one bit in the CH_ST_MSK register of any channel is set to `0b1`.

To clear the combined interrupt, all CH_ST_MSK bits for the channels must be cleared.

Channel clear interrupt

When the Receiver writes any value to the CH_CLR register of the channel, the channel clear interrupt is generated. As a result, it is not necessary for software to poll the status register to detect when the Receiver has finished processing the previous transfer. However, the Sender must still check the value of the CH_ST register, as defined by the transport protocol, when the Sender receives the clear interrupt.

The original message Sender is only required to receive the clear interrupt if the Ready to Send interface remains in the state where both signals are `0b1`. Therefore, software must not set the ACCESS_REQUEST.ACC_REQ bit to `0b0` until the clear interrupt is received. Generation of the clear interrupt is allowed when the Ready to Send interface is not in the state where both signals are `0b1`.

A.2.5 Programmers model

This section describes the MHU registers.

The MHU only supports 32-bit word-aligned read/write access. Any access attempt that does not meet this requirement is treated as an aligned 32-bit access.

A.2.5.1 Sender frame registers

The Sender frame contains registers for MHU configuration, interrupts, and identifiers.

Table A-1: Sender frame register summary

Offset	Name	Type	Description
`0x000-0xF7C`	-	-	See A.2.5.3.1 Sender channel window on page 318.
`0xF80`	MHU_CFG	RO	See A.2.5.4 MHU_CFG, Message Handling Unit Configuration register on page 324.
`0xF84`	RESP_CFG	RW	See A.2.5.5 RESP_CFG, Response Configuration register on page 324.
`0xF88`	ACCESS_REQUEST	RW	See A.2.5.6 ACCESS_REQUEST, Access Request register on page 325.
`0xF8C`	ACCESS_READY	RO	See A.2.5.7 ACCESS_READY, Access Ready register on page 325.
`0xF90`	INT_ST	RO	See A.2.5.8 INT_ST, Interrupt Status register on page 326.
`0xF94`	INT_CLR	WO	See A.2.5.9 INT_CLR, Interrupt Clear register on page 326.
`0xF98`	INT_EN	RW	See A.2.5.10 INT_EN, Interrupt Enable register on page 327.
`0xF9C`	-	RO	Reserved.
`0xFA0`	CHCOMB_INT_ST0	RO	Channel interrupt status for channels 0–31. See A.2.5.11 CHCOMB_INT_STn, Channel Combined Interrupt Status registers, n = 0–3 on page 327.
`0xFA4`	CHCOMB_INT_ST1	RO	Channel interrupt status for channels 32–63. See A.2.5.11 CHCOMB_INT_STn, Channel Combined Interrupt Status registers, n = 0–3 on page 327.
`0xFA8`	CHCOMB_INT_ST2	RO	Channel interrupt status for channels 64–95. See A.2.5.11 CHCOMB_INT_STn, Channel Combined Interrupt Status registers, n = 0–3 on page 327.
`0xFAC`	CHCOMB_INT_ST3	RO	Channel interrupt status for channels 96–123. See A.2.5.11 CHCOMB_INT_STn, Channel Combined Interrupt Status registers, n = 0–3 on page 327.

Offset	Name	Type	Description
0xFB0-0xFC4		RO	Reserved.
0xFC8	IIDR	RO	See A.2.5.12 IIDR, Implementer Identification Register on page 328.
0xFCC	AIDR	RO	See A.2.5.13 AIDR, Architecture Identification Register on page 329.
0xFD0	PID4	RO	See A.2.5.14 PID4, Peripheral ID 4 register on page 329.
0xFD4	PID5	RO	See A.2.5.15 PID5, Peripheral ID 5 register on page 330.
0xFD8	PID6	RO	See A.2.5.16 PID6, Peripheral ID 6 register on page 330.
0xFDC	PID7	RO	See A.2.5.17 PID7, Peripheral ID 7 register on page 330.
0xFE0	PID0	RO	See A.2.5.18 PID0, Peripheral ID 0 register on page 331.
0xFE4	PID1	RO	See A.2.5.19 PID1, Peripheral ID 1 register on page 331.
0xFE8	PID2	RO	See A.2.5.20 PID2, Peripheral ID 2 register on page 332.
0xFEC	PID3	RO	See A.2.5.21 PID3, Peripheral ID 3 register on page 332.
0xFF0	CID0	RO	See A.2.5.22 CID0, Component ID 0 register on page 332.
0xFF4	CID1	RO	See A.2.5.23 CID1, Component ID 1 register on page 333.
0xFF8	CID2	RO	See A.2.5.24 CID2, Component ID 2 register on page 333.
0xFFC	CID3	RO	See A.2.5.25 CID3, Component ID 3 register on page 334.

A.2.5.2 Receiver frame registers

The Receiver frame contains many of the same registers as the Sender frame, but without the RESP_CFG, ACCESS_REQUEST, and ACCESS_READY registers.

Table A-2: Receiver frame register summary

Offset	Name	Type	Description
0x000-0xF7C		-	See A.2.5.3.2 Receiver channel window on page 318.
0xF80	MHU_CFG	RO	See A.2.5.4 MHU_CFG, Message Handling Unit Configuration register on page 324.
0xF84-0xF8C		RO	Reserved
0xF90	INT_ST	RO	See A.2.5.8 INT_ST, Interrupt Status register on page 326.
0xF94	INT_CLR	RO	See A.2.5.9 INT_CLR, Interrupt Clear register on page 326.
0xF98	INT_EN	RW	See A.2.5.10 INT_EN, Interrupt Enable register on page 327.
0xF9C	-	RO	Reserved
0xFA0	CHCOMB_INT_ST0	RO	Channel interrupt status for channels 0–31. See A.2.5.11 CHCOMB_INT_STn, Channel Combined Interrupt Status registers, n = 0–3 on page 327.
0xFA4	CHCOMB_INT_ST1	RO	Channel interrupt status for channels 32–63. See A.2.5.11 CHCOMB_INT_STn, Channel Combined Interrupt Status registers, n = 0–3 on page 327.
0xFA8	CHCOMB_INT_ST2	RO	Channel interrupt status for channels 64–95. See A.2.5.11 CHCOMB_INT_STn, Channel Combined Interrupt Status registers, n = 0–3 on page 327.
0xFAC	CHCOMB_INT_ST3	RO	Channel interrupt status for channels 96–123. See A.2.5.11 CHCOMB_INT_STn, Channel Combined Interrupt Status registers, n = 0–3 on page 327.
0xFB0-0xFC4		RO	Reserved
0xFC8	IIDR	RO	See A.2.5.12 IIDR, Implementer Identification Register on page 328.
0xFCC	AIDR	RO	See A.2.5.13 AIDR, Architecture Identification Register on page 329.
0xFD0	PID4	RO	See A.2.5.14 PID4, Peripheral ID 4 register on page 329.

Offset	Name	Type	Description
0xFD4	PID5	RO	See A.2.5.15 PID5, Peripheral ID 5 register on page 330.
0xFD8	PID6	RO	See A.2.5.16 PID6, Peripheral ID 6 register on page 330.
0xFDC	PID7	RO	See A.2.5.17 PID7, Peripheral ID 7 register on page 330.
0xFE0	PID0	RO	See A.2.5.18 PID0, Peripheral ID 0 register on page 331.
0xFE4	PID1	RO	See A.2.5.19 PID1, Peripheral ID 1 register on page 331.
0xFE8	PID2	RO	See A.2.5.20 PID2, Peripheral ID 2 register on page 332.
0xFEC	PID3	RO	See A.2.5.21 PID3, Peripheral ID 3 register on page 332.
0xFF0	CID0	RO	See A.2.5.22 CID0, Component ID 0 register on page 332.
0xFF4	CID1	RO	See A.2.5.23 CID1, Component ID 1 register on page 333.
0xFF8	CID2	RO	See A.2.5.24 CID2, Component ID 2 register on page 333.
0xFFC	CID3	RO	See A.2.5.25 CID3, Component ID 3 register on page 334.

A.2.5.3 Channel windows

A channel window is a group of registers. The registers in the channel window vary between the Sender and Receiver frame views.

An MHU can contain between 1 and 124 channels. The number of channels that are implemented can be discovered from the MHU_CFG.NUM_CH field. Each channel occupies eight 32-bit words in both the Sender and Receiver register maps.

The address space that is allocated to channels that are not implemented is Reserved and treated as RAZ/WI.

A.2.5.3.1 Sender channel window

The Sender frame view of a channel window contains the register for setting the channel status and registers for controlling the channel interrupt.

Table A-3: Sender channel window register summary

Offset	Name	Short Name	Type	Description
0x00	Channel status	CH_ST	RO	See A.2.5.3.3 CH_ST, Channel Status register on page 319
0x04	Reserved	-	RO	-
0x08	Reserved	-	RO	-
0x0C	Channel Set	CH_SET	WO	See A.2.5.3.6 CH_SET, Channel Set register on page 320
0x10	Channel Interrupt Status	CH_INT_ST	RO	See A.2.5.3.10 CH_INT_ST, Channel Interrupt Status register on page 322
0x14	Channel Interrupt Clear	CH_INT_CLR	WO	See A.2.5.3.11 CH_INT_CLR, Channel Interrupt Clear register on page 323
0x18	Channel Interrupt Enable	CH_INT_EN	RW	See A.2.5.3.12 CH_INT_EN, Channel Interrupt Enable register on page 323
0x1C	Reserved	-	RO	-

A.2.5.3.2 Receiver channel window

The Receiver frame view of a channel window contains the register for clearing the channel status and registers for controlling the channel mask.

Table A-4: Receiver channel window register summary

Offset	Name	Short Name	Type	Description
0x00	Channel status	CH_ST	RO	See A.2.5.3.3 CH_ST, Channel Status register on page 319
0x04	Reserved	CH_ST_MSK	RO	See A.2.5.3.4 CH_ST_MSK, Channel Status Masked register on page 320
0x08	Reserved	CH_CLR	WO	See A.2.5.3.5 CH_CLR, Channel Clear register on page 320
0x0C	Channel Set	-	RO	-
0x10	Channel Interrupt Status	CH_MSK_ST	RO	See A.2.5.3.7 CH_MSK_ST, Channel Mask Status register on page 321
0x14	Channel Interrupt Clear	CH_MSK_SET	WO	See A.2.5.3.8 CH_MSK_SET, Channel Mask Set register on page 321
0x18	Channel Interrupt Enable	CH_MSK_CLR	WO	See A.2.5.3.9 CH_MSK_CLR, Channel Mask Clear register on page 322
0x1C	Reserved	-	RO	-

A.2.5.3.3 CH_ST, Channel Status register

The CH_ST register shows the state of the channel and is part of both the Receiver and Sender channel windows.

If the Receiver frame is reset, then the contents of the CH_ST register in the Sender frame are also reset. Software is responsible for handling any lost messages when the Receiver frame and CH_ST register are reset.

Configurations

This register is available in all configurations.

Bit descriptions

Table A-5: Channel Status bit descriptions

Bits	Name	Description	Type	Reset
[31:0]	FLAGn, n = 0–31	Display the status of channel flags. Each bit can be used as an individual flag or bits can be grouped. The way in which the register is used depends on the transport protocol that is employed. Bits in this register are set by writing 0b1 to the corresponding bits in the CH_SET register. Writing 0b1 to bits in the CH_CLR register clears the corresponding bits in the CH_ST register. If software: • Sets a bit that is already set, the bit remains set. • Clears a bit that is already cleared, the bit remains cleared. • Sets and clears a bit at the same time, the bit remains set. Arm® strongly recommends that software follows the transport protocols that are defined in A.3 Transport protocols on page 334.	RO	0x0000_0000

A.2.5.3.4 CH_ST_MSK, Channel Status Masked register

The CH_ST_MSK register shows the state of the channel with the channel mask applied, and is part of the Receiver channel window.

Configurations

This register is available in all configurations.

Table A-6: Channel Status Masked bit descriptions

Bits	Name	Description	Type	Reset
[31:0]	FLAG_MSKn, n = 0–31	Display the status of channel flags with the mask applied. When this register is nonzero, the interrupt for the Channel is asserted. The value in this register is equal to CH_ST and ~CH_MSK_ST, at the point at which the read occurs.	RO	`0x0000_0000`

A.2.5.3.5 CH_CLR, Channel Clear register

The CH_CLR register resets bits in the Channel Status and Channel Status Masked registers, and is part of the Receiver channel window.

Configurations

This register is available in all configurations.

Bit descriptions

Table A-7: Channel Clear bit descriptions

Bits	Name	Description	Type	Reset
[31:0]	FLAG_CLRn, n = 0–31	Clear the channel flags. • Writing `0b1` to bits in this register clears the corresponding bits in the CH_ST and CH_ST_MSK registers. • Writing `0b0` to bits in this register has no effect. Each bit always reads as `0b0`.	WO	`0x0000_0000`

A.2.5.3.6 CH_SET, Channel Set register

The CH_SET register writes bits in the Channel Status and Channel Status Mask registers, and is part of the Sender channel window.

Configurations

This register is available in all configurations.

Bit descriptions

Table A-8: Channel Set bit descriptions

Bits	Name	Description	Type	Reset
[31:0]	FLAG_SETn, n = 0–31	Set the channel flags. • Writing `0b1` to bits in this register sets the corresponding bits in the CH_ST register. • Writing `0b0` to bits in this register has no effect. Each bit always reads as `0b0`.	WO	`0x0000_0000`

A.2.5.3.7 CH_MSK_ST, Channel Mask Status register

The CH_MSK_ST register shows the state of the channel mask, and is part of the Receiver channel window. The channel mask is used with the Channel Status register to generate the channel status mask.

Configurations

This register is available in all configurations.

Table A-9: Channel Mask Status bit descriptions

Bits	Name	Description	Type	Reset
[31:0]	FLAG_MSKn, n = 0–31	Display the status of channel flag masks. • A channel mask bit that is set to `0b0` indicates that the corresponding flag bit is unmasked. When a bit is unmasked, the equivalent bits in the CH_ST and CH_ST_MSK registers have the same value. • A channel mask bit that is set to `0b1` indicates that the corresponding flag bit is masked. When a bit is masked, the equivalent bit in the CH_ST_MSK register always reads as `0b0`.	RO	`0x0000_0000`

A.2.5.3.8 CH_MSK_SET, Channel Mask Set register

The CH_MSK_SET register writes bits in the channel mask and is part of the Receiver channel window.

Configurations

This register is available in all configurations.

Bit descriptions

Table A-10: Channel Mask Set bit descriptions

Bits	Name	Description	Type	Reset
[31:0]	FLAG_MSK_SETn, n = 0–31	Set the channel flag masks. • Writing `0b1` to bits in this register sets the corresponding bits in the CH_MSK_ST register. • Writing `0b0` to bits in this register has no effect. Each bit always reads as `0b0`.	WO	`0x0000_0000`

A.2.5.3.9 CH_MSK_CLR, Channel Mask Clear register

The CH_MSK_CLR register resets bits in the channel mask and is part of the Receiver channel window.

Configurations

This register is available in all configurations.

Table A-11: Channel Mask Clear bit descriptions

Bits	Name	Description	Type	Reset
[31:0]	FLAG_MSK_CLRn, n = 0–31	Clear the channel flag masks. • Writing `0b1` to bits in this register clears the corresponding bits in the CH_MSK_ST register. • Writing `0b0` to bits in this register has no effect. Each bit always reads as `0b0`.	WO	`0x0000_0000`

A.2.5.3.10 CH_INT_ST, Channel Interrupt Status register

The CH_INT_ST register shows whether a channel clear interrupt has been generated for the channel.

Configurations

This register is available in all configurations.

Bit descriptions

Table A-12: Channel Interrupt Status bit descriptions

Bits	Name	Description	Type	Default
[31:1]	-	Reserved.	RO	`0x0000_0000`

Bits	Name	Description	Type	Default
[0]	CH_CLR	Displays the status of the channel clear interrupt. • `0b0` indicates that a channel clear interrupt has not occurred. • `0b1` indicates that a channel cleat interrupt is generated. This field is set to `0b1` when the Receiver writes to the CH_CLR register. There is no requirement for the field to be set to `0b1` if, at the point at which the write occurs, the Sender domain is: ◦ Not powered on ◦ Held in reset This field is set to `0b0` when the Sender writes `0b1` to the CH_INT_CLR.CH_CLR field. The process of setting this field to `0b1` takes priority over setting it to `0b0`. CH_INT_EN is not an interrupt mask register, but it enables and disables the interrupt generation itself.	RO	`0b0`

A.2.5.3.11 CH_INT_CLR, Channel Interrupt Clear register

The CH_INT_CLR register resets the channel clear interrupt for the channel.

Configurations

This register is available in all configurations.

Table A-13: Channel Interrupt Clear bit descriptions

Bits	Name	Description	Type	Reset
[31:1]	-	Reserved.	RO	`0x0000_0000`
[0]	CH_CLR	Clears the channel clear interrupt. • Writing `0b1` to this field clears the interrupt. • Writing `0b0` to this field has no effect. This field always reads as `0b0`.	WO	`0b0`

A.2.5.3.12 CH_INT_EN, Channel Interrupt Enable register

The CH_INT_EN register activates and deactivates generation of the channel clear interrupt for the channel.

Configurations

This register is available in all configurations.

Table A-14: Channel Interrupt Enable bit descriptions

Bit	Name	Description	Type	Reset
[31:1]	-	Reserved.	RO	0x0000_0000
[0]	CH_CLR	Enables and disables generation of the channel clear interrupt. • Writing 0b1 to this field enables generation of the interrupt. • Writing 0b0 to this field disables interrupt generation.	RW	0b0

A.2.5.4 MHU_CFG, Message Handling Unit Configuration register

The MHU_CFG register shows the number of channels that are implemented in the MHU.

Configurations

This register is available in all configurations.

Address offset

0xF80

Table A-15: Message Handling Unit Configuration bit descriptions

Bits	Name	Description	Type	Default
[31:7]	-	Reserved.	RO	0x000_0000
[6:0]	NUM_CH	Specifies the number of MHU channels that are implemented. The value of the field indicates the number of channels, up to a maximum of 124 (0x7C). The values 0x00, 0x7D, 0x7E, and 0x7F are reserved.	RO	CFG_DEF

A.2.5.5 RESP_CFG, Response Configuration register

The RESP_CFG register determines the response when the Sender attempts to access any channel window register and the Receiver is not ready to accept a transfer.

Configurations

This register is available in all configurations.

Address offset

0xF84

Table A-16: Response Configuration bit descriptions

Bits	Name	Description	Type	Default
[31:1]	-	Reserved.	RO	0x0000_0000

Bits	Name	Description	Type	Default
[0]	NR_RESP	Specifies the response that is generated when the Sender attempts to access any channel window register while the ACCESS_READY.ACC_RDY field is set to `0b0`. This setting indicates that the Receiver is not in a state in which it can accept a transfer. When the Receiver is not ready, channel window register access attempts by the Sender are treated as RAZ/WI. With the RESP_CFG.NR_RESP field set to `0b0`, an error is not generated. If this field is set to `0b1`, then an error is generated.	RW	`0b0`

A.2.5.6 ACCESS_REQUEST, Access Request register

The ACCESS_REQUEST register is used by the Sender to require that the Receiver enters a state in which the Receiver can accept a transfer.

Configurations

This register is available in all configurations.

Address offset

`0xF88`

Table A-17: Access Request bit descriptions

Bits	Name	Description	Type	Default
[31:1]	-	Reserved.	RO	`0x0000_0000`
[0]	ACC_REQ	Requests that the Receiver prepares to accept a transfer. • A setting of `0b0` for this field indicates that the Receiver does not need to prepare for a transfer. • If this field is set to `0b1`, then the Receiver is requested to prepare to accept a transfer.	RW	`0b0`

A.2.5.7 ACCESS_READY, Access Ready register

The ACCESS_READY register shows whether the Receiver is in a state in which it can accept a transfer from the Sender.

Configurations

This register is available in all configurations.

Address offset

`0xF8C`

Table A-18: Access Ready bit descriptions

Bits	Name	Description	Type	Default
[31:1]	-	Reserved.	RO	`0x0000_0000`

Bits	Name	Description	Type	Default
[0]	ACC_RDY	Specifies whether the Receiver is able to accept a transfer. • A setting of 0b0 for this field indicates that the Receiver is not able to accept a transfer. • If this field is set to 0b1, then the Receiver is able to accept a transfer.	RO	0b0

A.2.5.8 INT_ST, Interrupt Status register

The INT_ST register shows whether the ready to not ready, not ready to ready, and combined interrupts have been generated.

Configurations

This register is available in all configurations.

Attributes

Address offset

0xF90

Table A-19: Interrupt Status bit descriptions

Bits	Name	Description	Type	Default
[31:3]	-	Reserved.	RO	0x0000_0000
[2]	CHCOMB	Displays the status of the channel combined interrupt. • A setting of 0b0 for this field indicates that an interrupt has not occurred on any channel. • If this field is set to 0b1, then an interrupt has been generated on at least one channel. There is no corresponding bit in the INT_CLR register. To clear the combined interrupt, software must clear the underlying interrupt.	RO	0b0
[1]	R2NR	Displays the status of the ready to not ready interrupt. • A setting of 0b0 for this field indicates that a ready to not ready interrupt has not occurred. • If this field is set to 0b1, then a ready to not ready interrupt has been generated.	RO	0b0
[0]	NR2R	Displays the status of the not ready to ready interrupt. • A setting of 0b0 for this field indicates that a not ready to ready interrupt has not occurred. • If this field is set to 0b1, then a not ready to ready interrupt has been generated.	RO	0b0

A.2.5.9 INT_CLR, Interrupt Clear register

The INT_CLR register resets the ready to not ready and not ready to ready interrupts.

Configurations

This register is available in all configurations.

Address offset

0xF94

Table A-20: Interrupt Clear bit descriptions

Bits	Name	Description	Type	Default
[31:2]	-	Reserved.	RO	0x0000_0000
[1]	R2NR	Clears the ready to not ready interrupt. • Writing 0b1 to this field clears the ready to not ready interrupt. • Writing 0b0 to this field has no effect.	WO	0b0
[0]	NR2R	Clears the not ready to ready interrupt. • Writing 0b1 to this field clears the not ready to ready interrupt. • Writing 0b0 to this field has no effect.	WO	0b0

A.2.5.10 INT_EN, Interrupt Enable register

The INT_EN register activates and deactivates generation of the combined, ready to not ready, and not ready to ready interrupts.

Configurations

This register is available in all configurations.

Address offset

0xF98

Table A-21: Interrupt Enable bit descriptions

Bits	Name	Description	Type	Default
[31:3]	-	Reserved.	RO	0x0000_0000
[2]	CHCOMB	Enables and disables generation of the channel combined interrupt. • Writing 0b1 to this field enables generation of the channel combined interrupt. • Writing 0b0 to this field disables channel combined interrupt generation.	RW	0b1
[1]	R2NR	Enables and disables generation of the ready to not ready interrupt. • Writing 0b1 to this field enables generation of the ready to not ready interrupt. • Writing 0b0 to this field disables ready to not ready interrupt generation.	RW	0b0
[0]	NR2R	Enables and disables generation of the not ready to ready interrupt. • Writing 0b1 to this field enables generation of the not ready to ready interrupt. • Writing 0b0 to this field disables not ready to ready interrupt generation.	RW	0b0

A.2.5.11 CHCOMB_INT_STn, Channel Combined Interrupt Status registers, n = 0–3

The CHCOMB_INT_STn registers show whether channels have pending interrupts.

These registers are located at offsets of `0xFA0` to `0xFAC` in the Sender frame, starting with CHCOMB_INT_ST0 at `0xFA0` and continuing in ascending order. Each register contains the status for a different set of channels:

- CHCOMB_INT_ST0 holds the status for channels 0–31, starting with channel 0 at bit 0.
- CHCOMB_INT_ST1 holds the status for channels 32–63, starting with channel 32 at bit 0.
- CHCOMB_INT_ST2 holds the status for channels 64–95, starting with channel 64 at bit 0.
- CHCOMB_INT_ST3 holds the status for channels 96–123, starting with channel 96 at bit 0.

Configurations

These registers are available in all configurations.

Address offset

`0xFA0–0xFAC`

Table A-22: Channel Combined Interrupt Status bit descriptions

Bits	Name	Description	Type	Default
[31:0]	CH_INT_STn, n = 0–31	Display the status of channel interrupts. Each bit indicates whether there is a pending interrupt on the corresponding channel. Where a channel is not implemented, the corresponding bit is Reserved and treated as RAZ/WI. In the CHCOMB_INT_ST3 register, bits[31:28] are always Reserved and treated as RAZ/WI. The fields within the CH_INT_STn registers of the Sender frame are set to `0b1` when the associated CH_INT_ST.CH_CLR field is set to `0b1`.	RO	`0x0000_0000`

A.2.5.12 IIDR, Implementer Identification Register

The IIDR contains information about the MHU implementation.

Configurations

This register is available in all configurations.

Address offset

`0xFC8`

Table A-23: Implementer Identification Register bit descriptions

Bits	Name	Description	Type	Default
[31:20]	PRODUCT_ID	Specifies the MHU part identifier.	RO	`0x076`
[19:16]	VARIANT	Specifies the MHU major revision number.	RO	`0x0`

Bits	Name	Description	Type	Default
[15:12]	REVISION	Specifies the MHU minor revision number.	RO	0x0
[11:0]	IMPLEMENTER	Specifies the JEDEC JEP106 manufacturers identification code for Arm®. Bits[11:8] contain the JEP106 continuation code for the implementer, bit[7] must always be 0, and bits[6:0] give the JEP106 identity code for the implementer.	RO	0x43B

A.2.5.13 AIDR, Architecture Identification Register

The AIDR contains the MHU architecture version.

Configurations

This register is available in all configurations.

Address offset

0xFCC

Table A-24: Architecture Identification Register bit descriptions

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved.	RO	0x00_0000
[7:4]	ARCH_MAJOR_REV	Specifies the MHU major architecture revision number. • A value of 0x1 indicates that the MHU conforms to MHU architecture version 2. • The setting 0x0 is Reserved. When the ARCH_MAJOR_REV field is set to 0x0, the values in the ARCH_MINOR_REV field and IIDR register are RAZ. Software must determine in a platform-specific manner the MHU architecture version to which the component conforms.	RO	0x1
[3:0]	ARCH_MINOR_REV	Specifies the MHU minor architecture revision number. A value of 0x0 indicates that the architecture minor revision number is 0, while a setting of 0x1 specifies a minor revision number of 1. All other values are reserved.	RO	0x1

A.2.5.14 PID4, Peripheral ID 4 register

The PID4 register contains information about the JEDEC JEP106 configuration code for the peripheral.

Configurations

This register is available in all configurations.

Address offset

0xFD0

Bit descriptions

Table A-25: Peripheral ID 4 bit descriptions

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved.	RO	0x00_0000
[7:4]	Size	Specifies the number of 4KB blocks that the System ID block occupies. This field is deprecated.	RO	0x0
[3:0]	DES_2	Specifies the JEDEC JEP106 continuation code for the peripheral.	RO	0x4

A.2.5.15 PID5, Peripheral ID 5 register

The PID5 register is Reserved.

Configurations

This register is available in all configurations.

Address offset

0xFD4

Table A-26: Peripheral ID 5 bit descriptions

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

A.2.5.16 PID6, Peripheral ID 6 register

The PID6 register is Reserved.

Configurations

This register is available in all configurations.

Address offset

0xFD8

Table A-27: Peripheral ID 6 bit descriptions

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

A.2.5.17 PID7, Peripheral ID 7 register

The PID7 register is Reserved.

Configurations

This register is available in all configurations.

Address offset

`0xFDC`

Table A-28: Peripheral ID 7 bit descriptions

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	`0x0000_0000`

A.2.5.18 PID0, Peripheral ID 0 register

The PID0 register contains the first eight bits of the identifier for the peripheral.

Configurations

This register is available in all configurations.

Address offset

`0xFE0`

Table A-29: Peripheral ID 0 bit descriptions

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	`0x00_0000`
[7:0]	PART_0	Specifies bits[7:0] of the part identifier for the peripheral	RO	`0x76`

A.2.5.19 PID1, Peripheral ID 1 register

The PID1 register contains the first four bits of the JEDEC JEP106 identity code and the second eight bits of the identifier for the peripheral.

Configurations

This register is available in all configurations.

Address offset

`0xFE4`

Table A-30: Peripheral ID 1 bit descriptions

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	`0x00_0000`

Bits	Name	Description	Type	Reset
[7:4]	DES_0	Specifies bits[3:0] of the JEDEC JEP106 identity code for the peripheral	RO	0xB
[3:0]	PART_1	Specifies bits[11:8] of the part identifier for the peripheral	RO	0x0

A.2.5.20 PID2, Peripheral ID 2 register

The PID2 register specifies whether the MHU uses the JEDEC JEP106 identification scheme, and contains parts of the peripheral JEP106 identity code and System ID block version.

Configurations

This register is available in all configurations.

Address offset

0xFE8

Table A-31: Peripheral ID 2 bit descriptions

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVISION	Specifies the major revision of the System ID block	RO	0x0
[3]	JEDEC	Specifies whether the MHU uses the JEDEC JEP106 identification scheme	RO	0b1
[2:0]	DES_1	Specifies bits[6:4] of the JEDEC JEP106 identity code for the peripheral	RO	0b011

A.2.5.21 PID3, Peripheral ID 3 register

The PID3 register contains the RTL modification field and part of the System ID block version.

Configurations

This register is available in all configurations.

Address offset

0xFEC

Table A-32: Peripheral ID 3 bit descriptions

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVAND	Specifies the minor revision of the System ID block	RO	0x0
[3:0]	CMOD	Specifies the customer modification value, which can be provided if a customer updates the RTL for the MHU	RO	0x0

A.2.5.22 CID0, Component ID 0 register

The CID0 register contains segment 0 of the preamble to the component class identifier.

Configurations

This register is available in all configurations.

Address offset

`0xFF0`

Table A-33: Component ID 0 bit descriptions

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	`0x00_0000`
[7:0]	PRMBL_0	Specifies segment 0 of the preamble to the code that identifies the component class	RO	`0x0D`

A.2.5.23 CID1, Component ID 1 register

The CID1 register contains segment 1 of the preamble and the component class identifier.

Configurations

This register is available in all configurations.

Address offset

`0xFF4`

Table A-34: Component ID 1 bit descriptions

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	`0x00_0000`
[7:4]	CLASS	Specifies a code that identifies the component class	RO	`0xF`
[3:0]	PRMBL_1	Specifies segment 1 of the preamble to the component class identifier	RO	`0x0`

A.2.5.24 CID2, Component ID 2 register

The CID2 register contains segment 2 of the preamble to the component class identifier.

Configurations

This register is available in all configurations.

Address offset

`0xFF8`

Table A-35: Component ID 2 bit descriptions

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_2	Specifies segment 2 of the preamble to the code that identifies the component class	RO	0x05

A.2.5.25 CID3, Component ID 3 register

The CID3 register contains segment 3 of the preamble to the component class identifier.

Configurations

This register is available in all configurations.

Address offset

0xFFC

Table A-36: Component ID 3 bit descriptions

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_3	Specifies segment 3 of the preamble to the code that identifies the component class	RO	0xB1

A.2.6 Limitations

The MHU only guarantees that the Receiver gets messages under specific circumstances.

Receipt of MHU messages is only guaranteed if:

- The Ready to Send protocol is used
- Both the Sender and the Receiver complete the required transport protocol steps, as defined in A.3 Transport protocols on page 334
- The Sender and the Receiver are not reset during the process

A.3 Transport protocols

Several different protocols for transferring messages between different entities are defined for the MHU.


Note

This appendix does not define the message protocols, that is the format of the messages.

Transport protocols specify the methods that are used to transfer messages between two entities. Message protocols determine the format of the messages. MHU transport protocols can transfer any message protocol, even messages that comprise several transfers.

Three transport protocols are defined for the current version of the MHU (Doorbell, Single-wod transfer, and Multi-word transfer).You might be able to implement other transport protocols using the MHU. However, only these three transport protocols are guaranteed to be supported in future versions of the MHUv2.

A.3.1 Doorbell transport protocol

The doorbell transport protocol uses the MHU to generate an interrupt to the Receiver only.

Each flag bit of the CH_ST register is used to indicate when the Sender passes a transfer. After the Sender has passed a transfer, it can send further transfers using other flags within the same channel, or any flag within another channel. The Sender must not use the same flag bit until the Receiver has processed the transfer.

The message protocol determines whether there is data associated with the flag, transferred in out-of-band memory, or whether the flag bit indicates the full message. If out-of-band data is included, then the location of the data is pre-agreed between the two entities.

The following sequence diagram shows the events that are required to implement the doorbell transport protocol.

Figure A-1: Sequence of events for the doorbell transport protocol

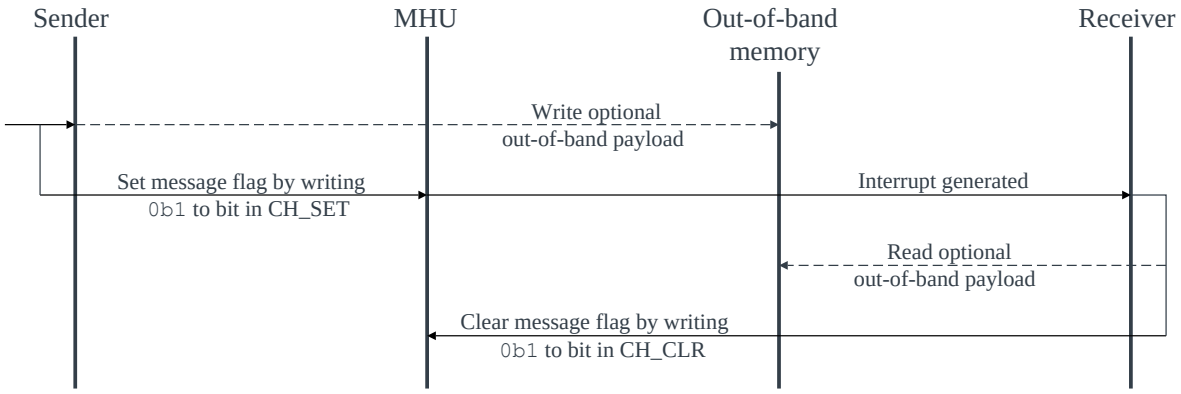

The following steps describe the doorbell transport protocol:

1. The Sender writes the optional payload for the transfer to the out-of-band memory. The memory location is pre-agreed using a software **IMPLEMENTATION DEFINED** method.
2. The Sender writes `0b1` to the correct bit in the CH_SET register.
3. The MHU generates the interrupt to the Receiver.
4. The Receiver gets the interrupt and, if necessary, reads the optional transfer payload in the out-of-band memory.

5. The Receiver clears the flag by writing `0b1` to the corresponding bit in the CH_CLR register.
6. The Sender waits for the bit in the CH_ST register that is set to `0b1` to be reset to `0b0`. This transition indicates that the Receiver has processed the transfer. The Sender can ensure that it is informed of the change by either:
   - Polling the CH_ST register and checking the bit that is set to `0b1` until the setting changes to `0b0`
   - Waiting to receive a clear interrupt for the channel before confirming that the bit that was set to `0b1` has been reset to `0b0`

The doorbell transport protocol has the following requirements:

- If out-of-band memory is used, then the location must be agreed between the Sender and the Receiver, and must be accessible to both.
- Any CH_ST register flag bits that are used to indicate a transfer must not be masked in the Receiver. For example, if bits[3:0] are used to indicate four different transfers, the bits[3:0] of the CH_MSK_ST register must be set to `0b0`.

A.3.2 Single-word transfer transport protocol

In the single-word transfer transport protocol, each transfer comprises one word that is transferred in-band using a single channel.

Sending a single message can involve multiple transfers and each transfer can be any value, except for zero. When the Sender has passed the transfer, the Receiver processes the transfer. The Sender must not use the same channel to send another transfer until the Receiver has processed the first transfer. While the Sender is waiting for the Receiver to process the transfer, the Sender can use different channels to send other transfers.

The following sequence diagram shows the events that are required to implement the single-word transfer transport protocol.

Figure A-2: Sequence of events for the single-word transfer transport protocol

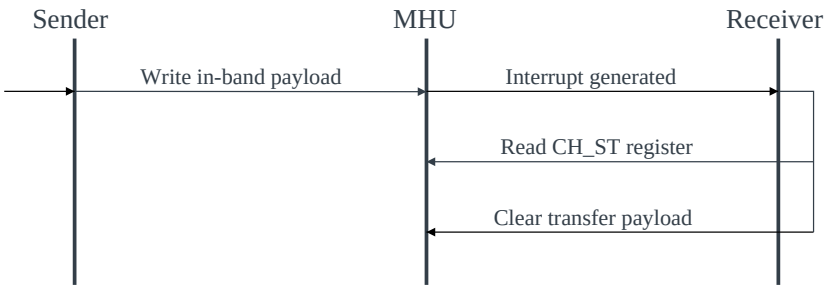

The following steps describe the single-word transfer transport protocol:

1. Using the CH_SET register, the Sender writes the in-band payload for the transfer, which can be any value except all zeros.

2. The MHU generates the interrupt to the Receiver.
3. The Receiver gets the interrupt and reads the in-band payload for the transfer in the CH_ST register.
4. The Receiver clears the in-band payload for the transfer by writing `0b1` to each bit of the CH_CLR register.
5. The Sender waits for the Receiver to finish processing the transfer, as indicated by all bits in the CH_ST register reading as `0b0`. The Sender can ensure that it is informed when processing is complete by either:
    - Polling the CH_ST register and checking whether it reads as all zeros.
    - Waiting to receive a clear interrupt for the channel before confirming that all bits in the CH_ST register read as `0b0`.

The single-word transfer transport protocol has the following requirements:

- All transfer payloads must have at least one bit set to `0b1`.
- The CH_MSK_ST register must have at least one bit set to `0b0`.
- For at least one bit that is set to `0b1` in the CH_ST register, the corresponding bit in the CH_MSK_ST register must be set to `0b0`.

Arm® strongly recommends that software either:

- Has all CH_MSK_ST register bits set to `0b0`.
- Uses a single bit to indicate a transfer to the Receiver and only that bit is required to be `0b0` in the CH_MSK_ST register. For example, if bit[31] is used to indicate a transfer, then only the CH_MSK_ST[31] field must be set to `0b0`.

A.3.3 Multi-word transfer transport protocol

In the multi-word transfer transport protocol, each transfer comprises two or more words.

The Sender writes every word of the transfer, but only the final word is used to indicate to the Receiver that a transfer is waiting. One or more channels can be used with the multi-word transfer transport protocol.

Not all the words must be transferred using MHU channels. Some words in a transfer can be in out-of-band memory. The location of the out-of-band memory can be either pre-agreed between the entities or referenced as part of the transfer payload.

The following sequence diagram shows the events that are required to implement the multi-word transfer transport protocol.

Figure A-3: Sequence of events for the multi-word transfer transport protocol

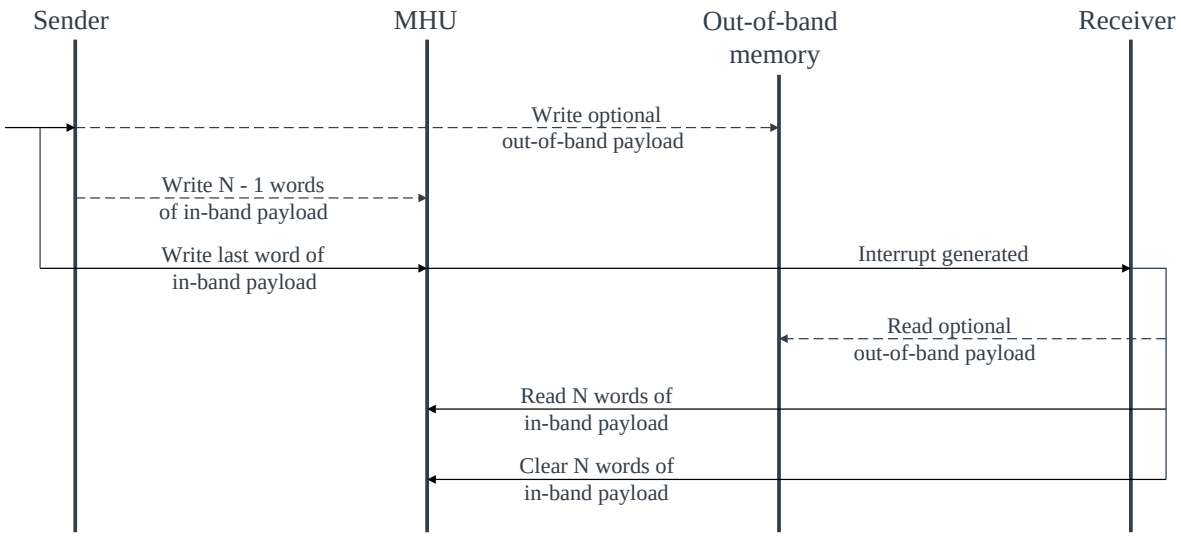

The following steps describe the multi-word transfer transport protocol:

1. The Sender writes the optional out-of-band payload for the transfer to the out-of-band memory. The memory location is either:
   - Pre-agreed using a software **IMPLEMENTATION DEFINED** method
   - Pointed from the payload that is transferred using the MHU channel
2. The Sender writes N - 1 words of the in-band transfer payload to the channels using the CH_SET register for each channel. The channel values can be zero.
3. The Sender writes the last word of the in-band transfer to the final channel. This value must be nonzero.
4. The MHU generates the interrupt to the Receiver.
5. The Receiver gets the interrupt and:
   a. Reads the optional transfer payload in the out-of-band memory, if necessary. The method by which the Receiver knows the location of the out-of-band memory is software **IMPLEMENTATION DEFINED**.
   b. Reads the in-band payload for the transfer in the CH_ST registers for the channels.
6. The Receiver clears the in-band payload for the transfer by writing `0b1` to each bit of the CH_CLR registers for the channels.
7. The Sender waits for the Receiver to finish processing the transfer, as indicated by all bits in the CH_ST registers for the channels reading as `0b0`. The Sender can ensure that it is informed when processing is complete by either:
   - Polling the CH_ST registers for the channels and checking whether they read as all zeros.
   - Waiting to receive a clear interrupt for each channel before confirming that all bits in the channel CH_ST registers read as `0b0`.

The following behaviors are software **IMPLEMENTATION DEFINED**:

- The order in which the Sender writes the out-of-band and N - 1 words of the in-band payload. However, the last word of the in-band payload must only be written when all in-band and out-of-band payload data is guaranteed to be visible to the Receiver.
- The order in which the Receiver reads the out-of-band and in-band payloads.
- The order in which the Receiver clears the in-band payload.

The multi-word transfer transport protocol has the following requirements:

- For channels that are not used to send the last written word of a transfer, all bits in the CH_MSK_ST registers must be set to `0b1`.
- For the channel that is used to send the last written word of a transfer, the CH_MSK_ST register must have at least one bit set to `0b0`.
- For at least one bit that is set to `0b1` in the CH_ST register, the corresponding bit in the CH_MSK_ST register of the last channel must be set to `0b0`.
- The Sender can write the words of the transfer in any order. However, the channel with at least one bit of the CH_MSK_ST register set to `0b0` must be used for the last word to be written. The last word written to a channel does not need to be the last word of the transfer. It can be the first word. That is, the Sender writes words 1, 2, and 3 of a four-word transfer first, then writes word 0.
- If out-of-band memory is used, then the location must be accessible to both the Sender and the Receiver.

Arm® recommends that the last channel that the Receiver clears is the channel with at least one bit of the CH_MSK_ST register set to `0b0`. This approach means that it is only necessary for the Sender to wait for that CH_ST register to read as all zeros. The Receiver can detect this state by either polling the register or waiting for the clear interrupt for the channel.

The number of words in a transfer might be greater than the number of channels that are allocated. In such cases, the transfer process is the same as for the single-word transfer transport protocol, except:

- The Sender must write all words of the out-of-band transfer payload before writing the last word of the transfer payload to the MHU channel.
- The Receiver must read the entire transfer payload, including words that are stored in out-of-band memory, before clearing the channels in the MHU. There is no requirement to clear the words in the out-of-band memory.

Appendix B Interrupt Router

This appendix describes the configuration and operation of the component that routes shared interrupts in SSE-710.

The Interrupt Router is a programmable router for shared interrupts that is located before the interrupt controllers in SSE-710. Software can select the interrupt controllers to which a specific shared interrupt is forwarded.

In the SSE-710 subsystem, the term interrupt controller can refer to:

- CoreLink™ GIC-400 in the Host System
- The interrupt controller in External System 0 or External System 1
- The interrupt collator in the Secure Enclave

In SSE-710, the Interrupt Router supports routing of *Shared Peripheral Interrupts* (SPIs) only. SPIs are signal-based interrupts that can be routed to any of the interrupt controllers in an SSE-710-based SoC.

B.1 About the Interrupt Router

The SSE-710 subsystem includes an Interrupt Router to allow for the routing of interrupts from shared peripherals to the interrupt controller of a specific system.

The following block diagram shows how interrupts are routed in SSE-710.

Figure B-1: Flow of interrupts through the Interrupt Router to the interrupt controllers

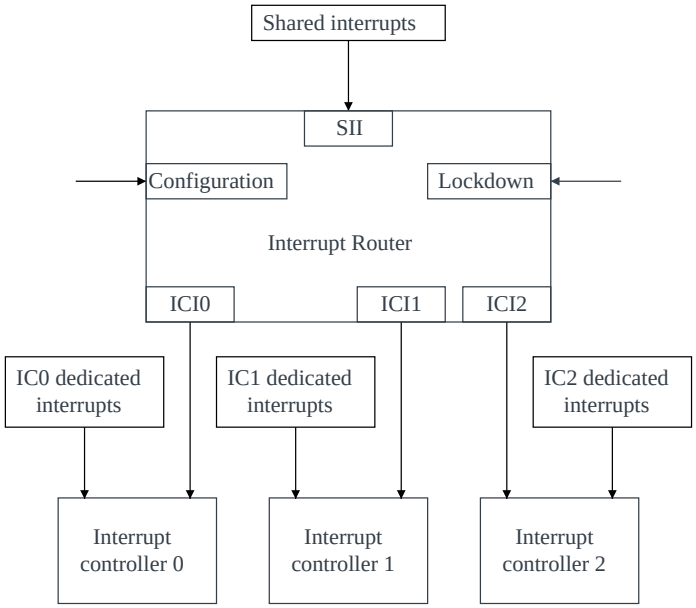

The Interrupt Router provides the following configuration options:

- Number of shared interrupts (NUM_SHD_INT) that the Interrupt Router supports. This value must be in the range 0–395.
- Interrupt Controller Interrupt Mask (SI{x}_ICI_DST) for each shared interrupt. The variable x is the shared interrupt number, which must be between 0 and (NUM_SHD_INT – 1).
- Interrupt Default Interrupt Controller Interface (SI{x}_DEF_ICI) for each shared interrupt. The variable x is the shared interrupt number, which must be between 0 and (NUM_SHD_INT – 1).

The Interrupt Router can be configured so that a shared interrupt can only be routed to a single *Interrupt Controller Interrupt* (ICI) interface.

A shared interrupt number is not the same as the final interrupt ID for the interrupt of the target interrupt controller. The final interrupt ID can be different depending on the interrupt controller to which the interrupt is routed. For example, an interrupt with the shared interrupt number 0 can have an interrupt ID of 32, 123, and 4 at different interrupt controllers within an SSE-710-based SoC.

B.2 Software configuration of the Interrupt Router

The Interrupt Router provides a window-based configuration scheme using the SHD_INT_{INFO,CFG,LCTRL} registers.

The SHD_INT_SEL register allows software to select the interrupt to which the SHD_INT_{INFO,CFG,LCTRL} registers refer. If the SHD_INT_SEL register selects an interrupt that is not implemented, the SHD_INT_{INFO,CFG,LCTRL} registers are Reserved and treated as RAZ/WI.

The SHD_INT_INFO.ICI_DST field indicates the ICI interfaces to which the interrupt can be routed. This field is set to the value of the SI{x}_ICI_DST configuration option for the shared interrupt that the SHD_INT_SEL.INT_SEL field selects.

The SHD_INT_CFG.ICI_EN field configures the ICI interfaces to which the interrupt is routed. The default value of this field depends on the SI{x}_DEF_ICI configuration option for the shared interrupt that the SHD_INT_SEL.INT_SEL field selects. If software changes the value in the SHD_INT_CFG.ICI_EN field when the interrupt source is active, it is **UNPREDICTABLE** whether the interrupt is routed to the correct destinations.

Arm strongly recommends that software disables the interrupt at source before making any changes to the SHD_INT_CFG.ICI_EN field.

The SHD_INT_LCTRL.LOCK field determines whether the configuration of the interrupt configuration is locked. For more information, see B.5 Lockdown Extension support on page 343.

B.3 Interrupt routing

Interrupts are routed to interrupt controllers according to a configurable routing mechanism.

Interrupts that the Interrupt Router receives are routed to the connected interrupt controllers based on:

- The subset of the allowed interrupt controllers to which a shared interrupt can be routed, as defined by the SI{x}_ICI_DST configuration option
- The software configuration for the shared interrupt

The following diagram shows how a shared interrupt is routed. The ICI to which the interrupt is routed is selected from the allowed ICIs using the SHD_INT_CFG.ICI_EN field.

Note

Shared interrupts can be routed to more than one interrupt controller at the same time.

Figure B-2: The routing mechanism for shared interrupts

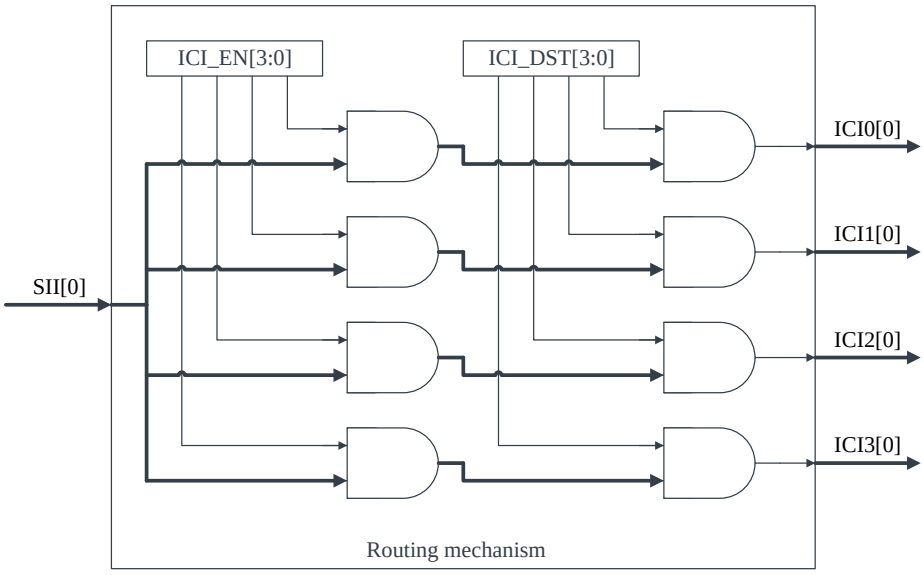

B.4 Configuration accesses

An access to the configuration registers of the Interrupt Router is called a configuration access.

A configuration access can cause a Configuration Access Error if:

- The configuration access attempts to access a Reserved location
- The configuration access is not a 32-bit word-aligned access
- The configuration access performs any of the actions that are described in B.5.3 Accesses to locked interrupt configurations on page 344

If there is a Configuration Access Error, the Interrupt Router must set the read data to all zeros and ignore the write data. Additionally, if the INT_RTR_CTRL.ERR field is set to `0b1`, the Interrupt Router must also generate an error response to the configuration access. If the INT_RTR_CTRL.ERR field is not set to `0b1`, the response to the configuration access can complete without an error.

A Configuration Access Error does not result in the generation of a tamper interrupt. A tamper interrupt is only generated if one of the conditions that are defined in B.5.3 Accesses to locked interrupt configurations on page 344 occur.

B.5 Lockdown Extension support

The Interrupt Router allows software to lock down shared interrupt configurations.

Software can lock down an individual shared interrupt configuration or all shared interrupt configurations. The options that are available depend on the *Lockdown Extension* (LDE) level that the implementation of the Interrupt Router supports.

SSE-710 supports *Lockdown Extension level 2* (LDE.2). Lockdown of interrupts is supported at the individual shared interrupt level and at the Interrupt Router level.

B.5.1 Lockdown states

The Interrupt Router has multiple lockdown states.

The following lockdown states are possible in the Interrupt Router.

Open

In this state, all registers can be updated. The only exception is that the SHD_INT_CFG register cannot be updated when the SHD_INT_LCTRL.LOCK field is set to `0b1`.

Partial

In this state, all registers can be updated except for:

- The SHD_INT_CFG register when the SHD_INT_LCTRL.LOCK field is set to `0b1`
- The SHD_INT_LCTRL.LOCK field, which is set to `0b1`

- The INT_RTR_CTRL register

Full

In this state, the following are read-only:

- The SHD_INT_CFG register
- The SHD_INT_LCTRL register
- The LD_CTRL.LOCK field if the Lockdown interface is asserted

The LD_CTRL.LOCK field is read-only if the Lockdown interface is asserted.

The Interrupt Router implements a Lockdown interface that is a single signal. When the Lockdown interface is asserted, it prevents the lockdown state of the Interrupt Router from being changed when the state is Partial or Full. The INT_RTR_LOCK bit of Host System Lock Control registers drives the Interrupt Router Lockdown interface. For more information, see 12.3.1.13 Host System Lock Control Status (HOST_SYS_LCTRL_ST) register on page 215.

B.5.2 Tamper Interrupt interface

The Interrupt Router reports tamper interrupts using a specific interface.

The Tamper Interrupt interface comprises a single level-sensitive interrupt to the Secure Enclave. The interface is asserted when either the INT_RTR_TMP_ST.TMP_ST_VLD bit or the INT_RTR_TMP_ST.TMP_SWT_OVERFLOW bit is set to `0b1`.

For more information about the tamper interrupt mapping of the Interrupt Router, see B.4 Configuration accesses on page 342.

B.5.3 Accesses to locked interrupt configurations

Configuration Access Errors can occur when locked configuration registers of the Interrupt Router are accessed.

Configuration Access Errors occur when:

- The configuration of an interrupt that is locked is updated. An interrupt configuration is locked when the SHD_INT_LCTRL.LOCK field is set to `0b1`.
- A locked interrupt is unlocked when the Interrupt Router lockdown state is set to Partial or Full. An interrupt configuration is unlocked by changing the SHD_INT_LCTRL.LOCK field from `0b0` to `0b1`.
- The Interrupt Router lockdown state is changed from either Partial or Full and the Lockdown interface is asserted, except for the value of SHD_INT_SEL register.
- The value of the INT_RTR_CTRL register is updated.
- The INT_RTR_TMP_ST register is accessed with a transaction that does not have the properties that are described in B.5.4 Access to tamper reports on page 345.

When one of the preceding events occurs, the Interrupt Router:

- Generates a Configuration Access Error.
- Asserts the Tamper Interrupt interface, unless it is already asserted, and either:
  - Generates a tamper report, updating the INT_RTR_TMP_ST register with the address of the configuration access. If the INT_RTR_TMP_ST.TMP_ST_VLD field is set to `0b0`, the Interrupt Router also sets this field to `0b1`.
  - Generates a tamper report overflow. If the INT_RTR_TMP_ST.TMP_STP_VLD field is set to `0b1`, the Interrupt Router also sets the INT_RTR_TMP_ST.TMP_ST_OVEFLW bit to `0b1`.

For more information about configuration accesses and Configuration Access Errors, see B.4 Configuration accesses on page 342.

B.5.4 Access to tamper reports

The INT_RTR_TMP_ST register holds information about any transactions that attempt to modify the configuration of an Interrupt Router that is in a lockdown state.

Access to the INT_RTR_TMP_ST register is only allowed for transactions with the following properties:

- Secure privileged transactions.
- Transactions from a Secure monitor agent.

  A Secure monitor agent is the most trusted master in the SoC. The tamper interrupt is routed to the interrupt controller that is associated with the Secure monitor.

If a transaction that is not either Secure privileged or from a Secure monitor agent attempts an access, the Interrupt Router:

- Generates a Configuration Access Error.
- Asserts the Tamper Interrupt interface, unless it is already asserted, and either:
  - Generates a tamper report, updating the INT_RTR_TMP_ST register with the address of the configuration access. If the INT_RTR_TMP_ST.TMP_ST_VLD field is set to `0b0`, the Interrupt Router also sets this field to `0b1`.
  - Generates a tamper report overflow. If the INT_RTR_TMP_ST.TMP_STP_VLD field is set to `0b1`, the Interrupt Router also sets the INT_RTR_TMP_ST.TMP_ST_OVEFLW bit to `0b1`.

Appendix C Firewall

This section gives a brief introduction to the firewall.

The firewall provides the monitoring and protection of the address space that is required between different entities within the SoC.


Note

For the firewall, an entity is either:

- A bus master
- A software application that executes on a bus master

Compartmentalization of the address space of a system between different entities provides protection between the entities. This compartmentalization also protects an entity from performing operations that are not allowed within its own compartment.

The firewall also provides address translation. For example, an external system needs to access the Host System memory map. If the external system has a dissimilar memory map, which could cause memory map address clashes between the systems. This might make it impossible for the external system to access locations in the Host System without address translation. Address translation provides an access window from the external system to the Host System, where an appropriate shared resource resides. The address from the external system is treated like a virtual address and is translated into a physical address in the Host System.

Compartmentalization is achieved through various methods, including traditional Memory Management Units (MMU). However, the memory footprint overhead in supporting page table descriptors is prohibitive in a memory constrained device. The SSE-710 Firewall provides some of the benefits of traditional memory management without the associated overhead of page table descriptors.

The monitoring functionality detects when accesses are attempted that cause errors to be generated in the system. The detection of an error can indicate that an entity is making illegal accesses. Detection and reporting of these errors, lets the software take corrective action or to limit the impact, if the errors are generated by a prohibited action.

C.1 Firewall usage

The firewall can be used as an address space protection and monitoring unit for memory and for peripherals in a system.

The following examples show various implementations of the firewall.

The firewall block represents all firewall Components and is placed between the master, memory and peripherals.

Figure C-1: Example of using firewall as protection and monitoring for a single Master

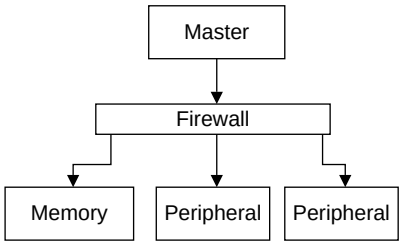

The firewall can also be used to monitor and protect the memory and peripherals from multiple masters as the following figure shows.

Figure C-2: Example of using firewall as protection and monitoring for multiple masters

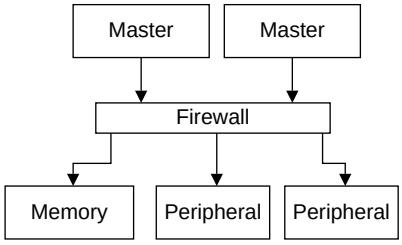

The firewall can also be used with firewall Components either side of an interconnect. This allows firewall Components after the interconnect (FC2-4 in the figure below) to define regions that are defined by firewall Components before the interconnect (FC1 in the figure below). By allowing refinement of regions, the granularity of how the address space is divided between masters in the system can be retained without increasing the number of regions that are required in each firewall Component.

The following figure shows an example system using firewall Components both sides of the interconnect: a CPU and two masters, A and B, attached to two peripherals, 0 and 1, and a small amount of memory. They are connected via an interconnect. This provides an error response to any transaction that targets a location that is not allocated to a downstream device.

Figure C-3: Example system with firewall Components both sides of the interconnect

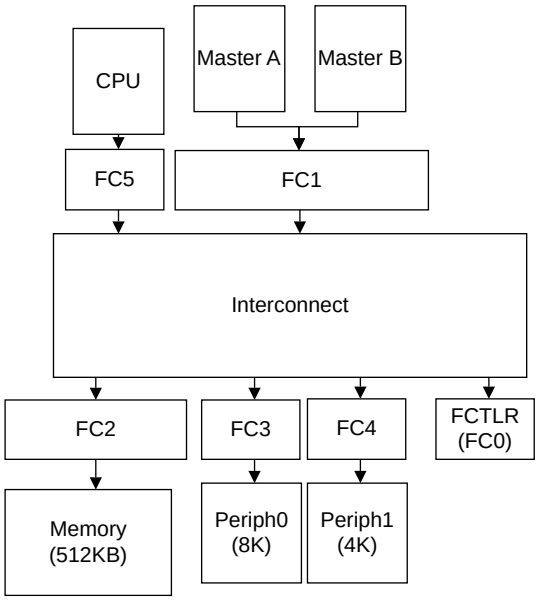

In this example, the following table shows the levels of extensions that the firewall components implement.

Table C-1: Example firewall Component levels of extensions

firewall Component	PE level	ME level	TE level
FC0	1	0	0
FC1	2	2	2
FC2	2	0	0
FC3	1	0	0
FC4	1	0	0
FC5	0	2	0

The firewall components can implement any level of extension, but for this example this has no effect on the behavior.

Masters A and B can be either masters, for example DMA or Display, or can be systems with their own address space. For this example, both masters are systems with their own independent address spaces. Each of these exposes one or more windows into the address space of the Host System, in this case the CPU address.

The following figure shows example memory maps for each address space.

Figure C-4: Memory maps for CPU, master A, and master B

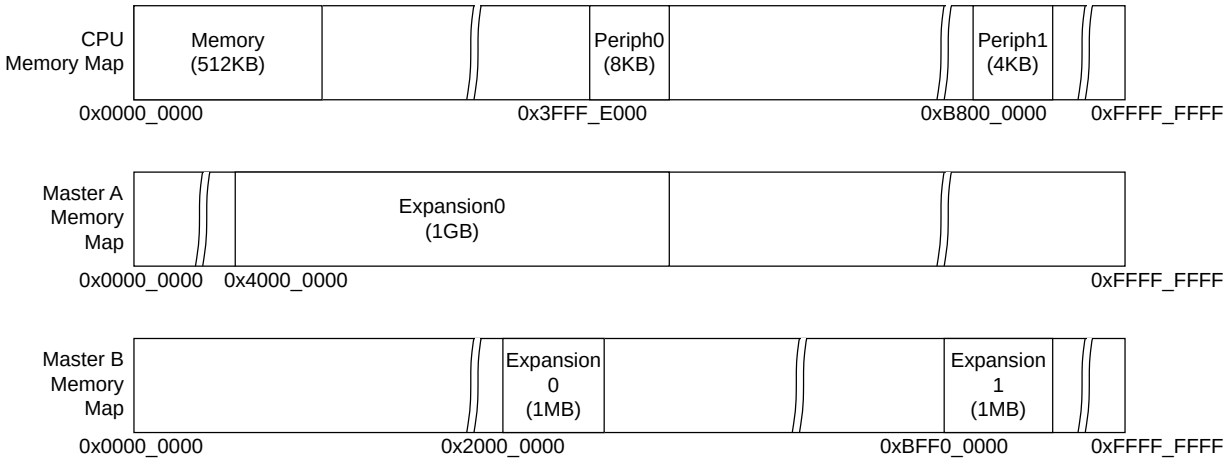

Transactions issued by master A and B are not altered as they enter the Host System.

There are five firewall Components in the system between the masters and peripherals, excluding the firewall Controller.

In this example software divides the resources of the Host System, as the following table shows.


Note

In this table, the Region Identifier column is not to be confused with the Region ID or Region Number, which is used in other sections of this chapter.

The “Access masters and privileges” column uses the following format for describing the allowed access types from each master: {Master_Name} – {Security Type} – {Privilege}. The Security type can be either Secure (S) or Non-secure (NS). The privilege is a combination of one or more Reads (R), Writes (W), or Executes (X). If a master is not present in the column, then it has no access.

Table C-2: Example memory regions requirements of example system

Region identifier	Description	Base address	Size of region	Access masters and privileges
0	CPU Secure code and data	0x0000_0000	8KB	CPU - S RWX
1	CPU Non-secure code and data	0x0000_4000	16KB	CPU - NS RWX
2	Master A to CPU mailbox	0x0002_0000	4KB	CPU - NS R Master A - NS RW
3	CPU to Master A mailbox	0x0002_1000	4KB	CPU - NS RW Master A - NS R
4	Master A Non-secure data	0x0002_2000	8KB	Master A - NS RWX
5	Master B to CPU mailbox	0x0003_0000	4KB	CPU - NS R Master B - NS RW

Region identifier	Description	Base address	Size of region	Access masters and privileges
6	CPU to Master B mailbox	0x0003_1000	4KB	CPU - NS RW Master B - NS R
7	Master B Non-secure data	0x0003_2000	8KB	Master B - NS RWX
8	Peripheral 0 allocated to Master A	0x7FFF_E000	8KB	Master A - NS RW
9	Peripheral 1 allocated to Master B	0xB800_0000	4KB	Master B - NS RW

.

To achieve this allocation of resources, the following regions must be defined in the firewall Components, as the following table shows.

Table C-3: Example Region definition for example system

Firewall Component	Region number	Region base address	Region size	Master ID	Security	Access types	Translation enable	Output base address
FC1	1	0x4000_0000	1GB	Master A	NS	RWX	Yes	0x0000_0000
	2	0x2000_0000	1MB	Master B	NS	RWX	Yes	0x0003_0000
	3	0xBFF0_0000	1MB	Master B	NS	RW	Yes	0xB800_0000
FC2	1	0x0000_0000	8KB	CPU	S	RWX	N/A	-
	2	0x0000_4000	16KB	CPU	NS	RWX	N/A	-
	3	0x0002_0000	4KB	CPU	NS	R	N/A	-
				Master A	NS	RW	N/A	-
	4	0x0002_1000	4KB	CPU	NS	RW	N/A	-
				Master A	NS	R	N/A	-
	5	0x0002_2000	8KB	Master A	NS	RW	N/A	-
	6	0x0003_0000	4KB	CPU	NS	R	N/A	-
				Master B	NS	RW	N/A	-
	7	0x0003_1000	4KB	CPU	NS	RW	N/A	-
				Master B	NS	R	N/A	-
	8	0x0003_2000	8KB	Master B	NS	RW	N/A	-
FC3	0	N/A	N/A	Master A	NS	RW	N/A	-
FC4	0	N/A	N/A	Master B	NS	RW	N/A	-

In this example, Firewall Components 0 and 5 are ignored:

- FC0 only defines regions to protect the firewall. For more information on firewall Component 0 regions, see C.15.1.1 Regions and RWE on page 450.
- FC5 is provided to monitor the bus transactions, as the CPU implements its own MMU.

Region 0 is not used for firewall Components 1 and 2, as this is the default region. Region 0 has different properties to other regions in the firewall. For firewall Components 3 and 4, region 0 is an ordinary region.

For more information on the default region, see C.8.2.3 Default regions on page 379.

C.2 Extensions

The firewall architecture has several extensions which provide additional features for the firewall.

An extension applies either at the individual Firewall Component, or the firewall level. Each extension defines at least two levels of support, starting at level 0. Level 0 indicates the extension is not implemented. An implementation of a firewall can select the level to which each extension is implemented.

The following table gives a summary of each extension.

Table C-4: Extension summary table

Extension	Level	Description
Protection Extension (PE)	0	No ability to compartmentalize the address space
	1	Ability to compartmentalize the address space, using predefined regions
	2	Ability to compartmentalize the address space, using software defined regions
Monitor Extension (ME)	0	No ability to detect errors
	1	Ability to detect errors, but with limited information about the transaction
	2	Ability to detect errors, with full information about the transactions
Region Size Extension (RSE)	0	Region's size must be of a power of 2, starting from the *Minimum Region Size* (MNRS)
	1	Region's size can be either an integer multiple of MNRS or a power of 2, starting from the MNRS
Translation Extension (TE)	0	Output transaction has the same address and transaction properties as the incoming transaction
	1	Output transaction has the same address as the incoming transaction but can have different transaction properties
	2	Output transaction can have different address and transaction properties to the incoming transaction
Lockdown Extension (LDE)	0	No ability to prevent update to configuration registers of the firewall
	1	Ability to prevent updates to configuration registers at the Firewall Component level
	2	Ability to prevent updates to configuration registers at the region and Firewall Component levels
Security Extension (SE)	0	firewall only support a single security world
	1	Firewall Components support two security worlds (Secure and Non-secure) for regions, but only a single security world for configuration
	-	-
Save and Restore Extension (SRE)	0	Firewall Component, when in the Disconnected state, lose all configuration and are inaccessible by software
	1	Firewall Component, when in the Disconnected state, are accessible by software. When a Firewall Component returns from the Disconnected state, its registers are restored automatically.

For more information on each extension see chapters C.8 Protection Extension on page 371 to C.14 Save and Restore Extension on page 445.

C.3 Bus protocol

This architecture lets you use any bus protocol on the interfaces to the Firewall Components.

The bus protocol must meet the following requirements:

- Bus protocol must be either a burst or beat based protocol:
  - For a burst-based protocol, the range of bytes accessed by the transaction must be provided at the start of the transaction. An example of a burst-based protocol is AXI.
  - For a beat-based protocol, the range of bytes accessed by the current beat must be provided. An example, of a beat-based protocol is AHB.
- Provides an identifier of the master or group of masters, which issued the transaction, when more than one master or groups of masters are connected to the same Firewall Component. For more information, see C.3.2 Stream ID on page 353.
- Identifies whether the transaction is a read or a write.
- Identifies whether the transaction is Secure or Non-secure, if SE.1 or greater is implemented.
- Able to indicate an error response to the master that issued the transaction.
- Except for the address of the transaction the other properties of the transaction are the same throughout the transaction.


Note

For a beat-based protocol the properties can change every beat.

This architecture allows a single firewall to implement different bus protocols or configurations of the same bus protocol, per each bus slave and master interface. It is **IMPLEMENTATION DEFINED** how the Firewall Components handle the conversion between the different bus protocols and configurations.


Note

Arm® recommends that if the bus protocol can indicate a decode error separate to a slave error, that the firewall uses:

- A decode error for transactions which are blocked by the firewall (this includes access to the firewall's owner configuration registers).
- A slave error for any transaction which attempts to access the configuration registers of the firewall and performs an operation which is not allowed.
- For transactions which the firewall allows through the response should be returned unaltered.

C.3.1 Transaction properties

This section describes the transactions properties associated with transactions processed by the Firewall.

The Firewall has the following properties that it associates with transactions it receives or issues on its interfaces:

- Address: The address range of the transaction.
- Type: Whether the access is read or write.
- Instruction: Whether the access is instruction or data access. Only applies to read transactions, all write transactions are considered data accesses.
- Privilege: The privilege level of the transaction.
- Memory Attribute: The memory type, cache allocation policy and whether it is transient or not.
- Security: The security of the transaction.

A Firewall Component never issues an outgoing transaction which has illegal properties, whether this is because of illegal incoming properties or bad software programming when translation is being applied. When outputting the transaction, the following rules are obeyed when the Firewall Component implements PE.1 or greater:

- If a transaction has a memory attribute of Normal with either the Inner or Outer cache level being Non-cacheable, the transaction is considered to be read-no-allocate, write-no-allocate and non-transient at that level, regardless of the programmed or incoming values. A non-cacheable access is considered read-no-allocate, write-no-allocate and non-transient. This is independent of programmed or incoming value.
- If a transaction has a memory attribute of Normal with a cacheable type that is read-no-allocate and write-no-allocate the memory has no transient attribute and is always considered non-transient independent of the programmed or incoming value.


Note

These requirements apply independently of:

- The level of TE implemented
- The value of the property's translation enable

For more information on memory attributes see C.10.1.2 Output transaction memory attribute property on page 432.

When PE.0 is implemented transactions are always considered to be issued without going through the Firewall Component.

C.3.2 Stream ID

The StreamID identifies the master or group of masters in the SoC.

The StreamID is sent alongside all transactions using the **AxMMUSID** signal.

The StreamID is passed through Firewall Component, unaltered, alongside the transaction.

C.3.3 Single master or group of masters

This section describes how Firewall supports a single master or a single group of masters.

When a Firewall Component is connected to a single master or a single group of masters, the transactions that are issued to the Firewall Component are not required to indicate the MasterID. A Firewall Component supports being configured for use with a single master, FC_CFG2.SINGLE_MST set to `0b1`.

When a Firewall Component is configured to support a single master, it does the following:

- Ignores the incoming MasterID of transactions
- The following fields are hardwired to the same value, set at design time:
  - RGN_MID
  - FE_MID
  - EDR_MID
  - FW_TMP_MID, for the Firewall Controller only
- When outputting the transaction, the MasterID is set to the same value in the fields listed above.

C.4 Firewall interfaces

This section describes firewall interfaces.

Table C-5: Firewall interfaces

Interface	Description
One or more Bus slave interfaces	Transactions to be processed, by the Firewall Component, arrive on this interface.
One or more Bus master interfaces	Transactions, which pass the checks performed by the Firewall Component, are output to the destination on this interface.
One Programming interface	Firewall configuration accesses are received and then processed on this interface.
One or more Power Control interfaces	Used to indicate the power mode of the domain the Firewall Component(s) resides in.
One or more Clock Control interfaces	Used to indicate the status of the clock domain the Firewall Component(s) resides in.
One Lockdown interface: used to provide configuration protection.	Only present if LDE.1 or greater is implemented.
One Interrupt interface	Used to indicate interrupts to the system.
One Tamper Interrupt interface	Used to indicate Tamper interrupts to the system. Only present if LDE.1 or greater is implemented.

Interface	Description
One or more Firewall Configuration interfaces	Used for communication between the Firewall Components.
One or more Protection Size interfaces	One is implemented per each Firewall Component which implements PE.2
One Bypass interface per each Firewall Component which implements PE.1 or greater	Used to bypass the checks performed by the Protection Extension.

C.4.1 Bus Slave and Bus Master interfaces

The number and type of bus protocols for the bus slave and bus master interfaces is **IMPLEMENTATION DEFINED**.


Note

The term *bus interfaces* refers to both the bus slave interface and bus master interface.

The number and type of bus protocols must meet the requirements described in C.3 Bus protocol on page 352.

The bus slave and master interfaces of the firewall, are associated with a Firewall Component within the firewall. Each Firewall Component has at least one bus slave and one bus master interface. The Firewall Component implements an interconnect between its bus slave and master interfaces.

A firewall can connect the bus master interface of one Firewall Component to the bus slave interface of another Firewall Component. This allows a firewall to stack Firewall Components. The topology of the interconnect implemented by the firewall and the Firewall Components, is **IMPLEMENTATION DEFINED**.

C.4.1.1 Bus Address Width

The requirements of Bus Address width depends on the PE implementation.

When a Firewall Component implements PE.0 there are no requirements on the address widths of the bus slave or master interfaces. This is because all transactions are considered to not pass through the Firewall Component even if, for integration reasons, the bus fabric routes both transaction requests and responses through the Firewall Component.

When a Firewall Component implements PE.1 or greater d there are requirements on the address width of the bus interfaces. C.8.1 Protection Size and bus address widths on page 372 describes these requirements.

C.4.1.2 StreamID

This architecture requires that a Firewall Component either passes the StreamID (MasterID) unmodified through the Firewall Component, or issues the StreamID (Fixed MasterID) when the transaction is issued.

All bus interfaces of a Firewall Component follow these rules for StreamID width:

- All bus master interfaces must have a MasterID field width greater than or equal to the maximum width of all bus slave interfaces of the Firewall Component.
- When a bus slave interface has a MasterID field width that is less than the maximum bus slave interface width, the MasterID field is zero extended.
- When a bus master interface has a MasterID field width that is greater than the maximum bus slave interface width, the MasterID field is zero extended.

Software can discover the MasterID width for the Firewall Component using the FC_CFG2.MST_ID_WIDTH field. This field is the maximum size of any bus interface, because the Firewall Component can support different bus interfaces with different-sized MasterID fields.

C.4.1.3 Bus Slave incoming transaction properties rules

The properties of a transaction received by a Firewall Component depend on the bus protocol that is used.

The bus protocol and the Firewall Component can support different properties. The table below shows the value used by the Firewall Component when it implements a property that is not supported by the bus slave interface.

Table C-6: Default bus property values, when supported by Firewall Component, but not by the bus protocol

Property	Firewall Component support	Value when not supported by bus protocol
Instruction	FC_CFG1.INST_SPT is 1	Data
Privilege	FC_CFG1.PRIV_SPT is 1	Unprivileged
Shareability	FC_CFG1.SH_SPT is 1	Non-shareable
Memory Attribute	FC_CFG1.MA_SPT is 1	Normal-iWB_RWA_nT-oWB_RWA_nT
Security	FC_CFG1.SEC_SPT is 1	Non-secure

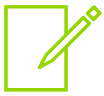
Note

These values are used in a future transaction translation. If the bus protocol supports a property that is not supported by the Firewall Component, then the Firewall Component ignores the value.

Arm® strongly recommends that the transaction properties supported by the Firewall Component are equal to or greater than the transaction properties supported by the bus protocol, which is used by the bus slave interfaces of the Firewall Component.

If the bus protocol for a transaction does not support one of these transaction properties, but you require a value not listed in the table above, then the Firewall Component must:

- Be configured with support for the property.
- Bus slave interface signal for the property is tied to the desired value

C.4.1.4 Bus Master output transactions properties rules

A Firewall Component must not issue an outgoing transaction that has illegal properties.This is true, either because of illegal incoming properties or because of bad software programming when applying translation.

When outputting the transaction, the following rules must always be obeyed when the Firewall Component implements PE.1 or greater:

- If a transaction has a memory attribute of Device or Normal Inner Non-Cacheable Outer Non-Cacheable the output shareability must be Outer Shareable.
- If a transaction has a memory attribute of Normal with either the Inner or Outer cache level being Non-cacheable, the transaction is considered to be read-no-allocate, write-no-allocate, and non-transient at that level regardless of the programmed or incoming values. A non-cacheable access is considered read-no-allocate, write-no-allocate, and non-transient. This is independent of programmed or incoming value.
- If a transaction has a memory attribute of Normal with a cacheable type that is read-no-allocate and write-no-allocate, the memory has no transient attribute and is always considered non-transient independent of the programmed or incoming value.

The above requirements apply independently of:

- The level of TE implemented.
- The value of the property's translation enable.

For more information on memory and shareability attributes see C.10.1.2 Output transaction memory attribute property on page 432 and C.10.1.2 Output transaction memory attribute property on page 432.

When PE.0 is implemented, transactions are always considered to be issued without going through the Firewall Component even, if for integration reasons, the bus fabric always routes both transaction requests and responses through the Firewall Component.

When the Firewall Components support a property which is not supported by the protocol for the Bus Master interface, the property is omitted when the transaction is issued on the interface.

When the bus protocol of the Bus Slave interface supports a property which the Firewall Component does not support, the Firewall Component ignores the property. Any output transactions have the property:

- Unmodified, if supported by the protocol of the Bus Master interface
- Omitted, if not supported by the protocol of the Bus Master interface

Independent of which option is selected, the above rules must always be followed for transaction output by Firewall Components implementing PE.1 or greater.

Arm® strongly recommends that the transaction properties supported by the Firewall Component are equal to or greater than those supported by the bus protocol used by the Bus Master interfaces of the Firewall Component.

C.4.2 Programming interface

The Programming Interface must be present on the Firewall Controller and is used to perform configuration accesses to the configuration of the firewall.

The Programming interface must adhere to the requirements described in section C.8.4 Transaction processing on page 385.

Arm® strongly recommends that the Firewall Controller and the bus protocol used for the Programming interface support the following transactions properties:

- Instruction
- Security
- Privilege

The Programming interface can be combined with one of bus slave interfaces of the Firewall Controller, provided the following conditions are met:

- All the requirements for both interfaces are met.
- The bus slave interface cannnot prevent a transaction for the Programming interface making progress.
- The Programming interface and the bus slave interface support the same bus protocol and level of properties.

C.4.3 Power Control interfaces

An implementation of the firewall is required to have a Power Control interface per hardware entity which makes up the Firewall.

Arm® strongly recommends that the Power Control interface is either a P or a Q-Channel interface and that it follows the guidelines in the *Arm® Power Control System Architecture Specification*.

The Power Control interface must be able to:

- Indicate the current power mode of the power domain where the Firewall Component resides.
- Minimally support the ability to indicate whether the Firewall Component is operational or not. The Firewall Component is considered operational when it is guaranteed not to lose any architectural state. It is considered non-operational when there is a possibility that a full or partial loss of architectural state could occur.
- Provide a mechanism for the Firewall Component to indicate its desired power mode.
- Provide a mechanism for the Firewall Component to accept or deny a change in power mode.

The power modes that are supported by the Firewall Component, and whether the Firewall Component is considered operational or non-operational in that power mode, is **IMPLEMENTATION DEFINED**.

Some power modes, which a Firewall Component can enter, do not cause the loss of any architectural state. However, the Firewall Component is unable to process transactions on the bus slave or Configuration interfaces. In these power modes the Firewall Component is still considered operational, as there is no state loss.

The Firewall Component must be able to exit these power modes and enter one where it is able to process new transactions on the bus slave or Configuration interfaces when they arrive. The method by which this is achieved is **IMPLEMENTATION DEFINED,** but must be achieved without any software interaction.


Note

A limited amount of software interaction might be required with the power controller. This depends on the power control of the system which the Firewall is integrated into, and is outside the scope of this specification. A Firewall Component can enter the Disconnected state and not lose any architectural state. Ths cause could be that power controller did not complete the transition to a power mode where state would be lost.

For example, an implementation can include support for ON, OFF and FULL_RET power modes. This allows for the Firewall Component to enter a retention state when there are no transactions to process on any of the bus slave or Firewall Configuration interfaces. In this case, ON and FULL_RET are considered operational mode,s as no architectural state is loss. Requirements are placed on the system integration to cause an exit from the FULL_RET power mode when either:

- A transaction arrives to be processed, on any Bus interface.
- An access arrives on the Firewall Configuration interface.


Note

Depending on the level of PE and ME implemented by the Firewall Component, processing the transaction refers to performing the protection checks, detection of errors or both.

Before a Firewall Component enters a power mode, which causes it to transition from being operational to non-operational or vice versa, the Firewall Component must perform the required Firewall Component state transition. For more information on Firewall Component states, see C.8.7.2 Firewall Component Status (PE_ST) on page 392.

If the SRE.1 is implemented, it is **IMPLEMENTATION DEFINED** whether the Firewall supports:

- Placing the shadow registers into retention when there are no accesses to the shadow registers.
- Firewall Controller is in a power mode, where it is considered non-operational, but retaining the contents of the shadow registers.

Arm® strongly advises that the Power Control interface, for the Firewall Controller, supports placing the shadow registers into retention when there are no accesses to the shadow registers.

In implementations which support placing the Firewall Controller in a non-operational mode, the FW_SR_CTRL.SR_PWR determines whether the shadow registers are required to retain their values or not. For more information, see C.8.7.2 Firewall Component Status (PE_ST) on page 392.

It is **IMPLEMENTATION DEFINED**, how the Firewall Controller knows the shadow register contain valid data after it leaves the Disconnected state.

C.4.4 Clock Control interfaces

An implementation of the firewall must have at least one Clock Control interface for each clock that is used by the Firewall.

It is **IMPLEMENTATION DEFINED** when multiple Firewall Components sharing the same clock, share a Clock Control interface. If the Clock Control interface is shared it is **IMPLEMENTATION DEFINED** how the Firewall Components are informed about the changes in the clock state.

The Clock Control interface must be capable of:

- Indicating whether the clock of the clock domain, is gated
- Provide a mechanism for the Firewall Component to indicate its desire for the clock to be gated or not
- Provide a mechanism for the Firewall Component to accept or deny the clock being gated

Arm® strongly recommends that the Clock Control interface is implemented as a Q-Channel interface and follows the guidelines in the *Arm® Power Control System Architecture Specification*.

The clock can only be gated if the Firewall Component accepts the clock gating request made on the Clock Control interface. When the clock is gated, the Firewall Component is not required to process any transactions on the Bus Slave or Configuration interfaces of the Firewall Component. The Firewall Component is required to request a clock when it has:

- Transactions outstanding on the Bus Slave or Configuration interfaces
- A request on the power control interface

Whether the clock is gated or not has no effect on anything other than the Firewall Component.


Note

This mechanism is in addition to any internal clock gating that may be implemented within the Firewall Component.

C.4.5 Lockdown interfaces

The lockdown interface is an optional interface and is only implemented if the firewall implements LDE.1 or greater. The Lockdown interface is only implemented on the Firewall Controller.

The lockdown interface prevents Firewall Components from changing their lockdown state, under certain conditions, when asserted high.


Note

For more information on the conditions where the Firewall Component's lockdown state can change see C.15.3 Lockdown Extension on page 454.

For more information about the interaction between the SE and LDE see C.12 Security Extension on page 440 and C.15.3 Lockdown Extension on page 454.

C.4.6 Interrupt interface

The Interrupt interface provides a method for the firewall to generate interrupts to the system. Interrupts that are generated by the firewall use a wired interrupt.

The Interrupt interface has the following requirements:

- Must be able to indicate when an interrupt has been generated.
- Indicates when the interrupt has been cleared.
- Remains asserted while any interrupt status field, it is associated with, is set to 1
- It must not be possible to receive the interrupt, without being able to observe the required registers have been updated. For more information see C.15.6 Interrupts on page 458 on the requirements for each interrupt type.

C.4.7 Tamper Interrupt interface

The Tamper Interrupt interface provides a method for the firewall to generate tamper interrupts to the system.

The Tamper Interrupt interface is an optional interface which is only implemented when LDE.1 or greater is implemented. This interface must be separate to the Interrupt interface else it be routed to another interrupt controller in the system. For example, the root of trust in the system. The Tamper Interrupt interface is similar to the Interrupt interface and must use wired interrupts.

The Tamper Interrupt interface must follow these requirements:

- Must be able to indicate when an interrupt has been generated.
- Indicates when the interrupt has been cleared.
- Remains asserted while either TR_VLD or TMP_ST_OVERFLW fields, in the FW_TMP_CTRL register, are set to 1.

- It must not be possible to receive the interrupt, without being able to observe that the required registers have been updated. For more information see C.15.6 Interrupts on page 458 on the requirements for each interrupt type.

C.4.8 Firewall Configuration interface

The Firewall Configuration interface is private to the firewall, however depending on how the firewall is implemented the interface may be external or internal.

The firewall, independent of whether the interface is external or internal, must meet the following requirements:

- Bi-directional communication between the Firewall Controller and the other Firewall Component
- Capable of performing the handshakes and request
- Capable of performing configuration accesses, received by the firewall on the Programming interface to the configuration registers of the Firewall Components. See section C.4.2 Programming interface on page 358 for more information on the Programming interface.
- Private to the firewall. It must only be accessible by Firewall Components, even if these components would not be able to access any firewall registers or observe any handshakes or requests sent over the interface.


Note

It is not required that a Firewall Component can communicate with another Firewall Component which is not the Firewall Controller.

C.4.9 Protection Size interface

Each Firewall Component, which implements PE.2 within a Firewall, must have a corresponding Protection Size interface.

The Protection Size Interface is an 8-bit signal, which allows setting the Protection Size of the Firewall Component. The Protection Size can be set, in power of twos, between the Minimum Region Size (MNRS) and the Maximum Region Size (MXRS). See C.8.1 Protection Size and bus address widths on page 372 for more information on the MNRS, MXRS and Protection Size. The Protection Size can also be set to 0 to force all transactions to be treated as failing the checks irrespective of the transaction and the regions defined in the Firewall Component. The encoding of the Protection Size interface is the same as the FC_CFG2.PROT_SIZE field. See C.6.2.3 Firewall Component Configuration Register 2 (FC_CFG2) on page 368 for more information of the encoding of FC_CFG2.PROT_SIZE.

The Protection Size interfaces are always present on the Firewall Controller. The signal is sampled when the Firewall Controller enters the Connected state and the shadow registers do not contain any valid values.


Note

When SRE.0 is implemented there are no shadow registers. Therefore, the Firewall Controller always samples the Protection Size interface on entry into the Connected state.

When this appendix refers to the Protection Size interface value, it refers to the sampled value.

As part of the Connect handshake, the sampled value of the Protection Size Interface must be communicated to the Firewall Component.

Once the Protection Size Interface is sampled the values of the following fields must be updated to reflect the sampled value:

- FC_CFG2.PROT_SIZE
- RGN_SIZE.SIZE for region 0 of a Firewall Component implementing PE.2

After the Protection Size Interface is sampled, any change in the value has no effect on the behavior of the Firewall Component until the interface is next sampled.

C.4.10 Bypass interface

There is a Bypass interface for each Firewall Component. It implements PE.1 or greater, and controls whether the Firewall Component is bypassed or not.

The value of the Bypass interface of the Bypass interface of the Firewall Component can be read by software using the PE_BPS.BYPASS_IF_ST field.

For more details, see C.8.6 Bypass interface on page 388

C.5 Overview

A Firewall Component can be configured to provide protection, monitoring or both depending on the level of PE and ME implemented by the Firewall Component.

This section describes the common features that all Firewall Components provide. Sections C.8 Protection Extension on page 371 to C.14 Save and Restore Extension on page 445 provide descriptions of each extension. Section C.15 Firewall Controller on page 448 provides a description of the differences between the Firewall Controller and a standard Firewall Component.

Each Firewall Component has 64KB allocated from the Firewall memory map for its own configuration registers, but only the first 4KB is used and the remaining 60KB is Reserved and generates a Configuration Access Error if accessed. In the following sections, the configuration registers of a Firewall Component are described, alongside the offset address within the allocated space.

For a full view of all the configuration registers of a Firewall Component or the memory map of the Firewall, see section C.17 Programmers model overview on page 478. Any unused address space within the 4KB is Reserved and generates a Configuration Access Error if accessed.

C.6 Common registers

All Firewall Components must implement certain common registers.

These registers control:

- The capabilities of the Firewall and individual Firewall Components.
- The configuration of the Firewall and individual Firewall Components.

C.6.1 Capability registers

A Firewall Component has the Firewall Component Capability registers, FC_CAP{0-3}.

Software can use these registers to find out the extension levels the Firewall and individual Firewall Component use.

Table C-7: Summary of Capability registers

Offset	Short Name	Access	Name
0xFA0	FC_CAP0	RO	Firewall Component Capability Register 0
0xFA4	FC_CAP1	RO	Firewall Component Capability Register 1
0xFA8	FC_CAP2	RO	Firewall Component Capability Register 2
0xFAC	FC_CAP3	RO	Firewall Component Capability Register 3

C.6.1.1 Firewall Component Capability Register 0 (FC_CAP0)

The following table gives a bit-level description of Firewall Component Capability Register 0.

Table C-8: Register 0 (FC_CAP0)

Bits	Name	Description	Type	Reset
[31:28]	-	Reserved	RO	0x0
[27:24]	SE_LVL	Level of the Security Extension used by the Firewall. 0x0: SE.0 is implemented 0x1: SE.1 is implemented All other values are Reserved. For Firewall Components, other than 0, this field is Reserved and treated as RAZ/WI.	RO	IMPL_DEF

Bits	Name	Description	Type	Reset
[23:20]	SRE_LVL	Level of the Save and Restore Extension used by the Firewall. 0x0: SRE.0 is implemented 0x1: SRE.1 is implemented All other values are Reserved. For Firewall Components, other than 0, this field is Reserved and treated as RAZ/WI.	RO	IMPL_DEF
[19:16]	LDE_LVL	Level of the Lockdown Extension implemented by the Firewall. 0x0: LDE.0 is implemented 0x1: LDE.1 is implemented 0x2: LDE.2 is implemented All other values are Reserved. For Firewall Components, other than 0, this field is Reserved and treated as RAZ/WI.	RO	IMPL_DEF
[15:12]	TE_LVL	Level of the Translation Extension implemented by Firewall Component 0. 0x0: TE.0 is implemented 0x1: TE.1 is implemented 0x2: TE.2 is implemented All other values are Reserved. This field must always be 0x0 when the FC_CAP0.PE_LVL is 0x0.	RO	IMPL_DEF
[11:8]	RSE_LVL	Level of the Region Size Extension implemented by Firewall Component. 0x0: RSE.0 is implemented 0x1: RSE.1 is implemented All other values are Reserved. This field must always be 0x0 when the FC_CAP0.PE_LVL is 0x0.	RO	IMPL_DEF
[7:4]	ME_LVL	Level of the Monitor Extension implemented by the Firewall Component. 0x0: ME.0 is implemented 0x1: ME.1 is implemented 0x2: ME.2 is implemented All other values are Reserved.	RO	IMPL_DEF

Bits	Name	Description	Type	Reset
[3:0]	PE_LVL	Level of the Protection Extension implemented by the Firewall Component. 0x0: PE.0 is implemented 0x1: PE.1 is implemented 0x2: PE.2 is implemented All other values are Reserved.	RO	IMPL_DEF

The values of FC_CAP0.{PE_LVL,ME_LVL,RSE_LVL,TE_LVL} can be different for each Firewall Component within the Firewall.

C.6.1.2 Firewall Component Capability Register 1(FC_CAP1)

The following table gives a bit-level description of Firewall Component Capability Register 1.

Table C-9: Register 1(FC_CAP1)

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

C.6.1.3 Firewall Component Capability Register 2 (FC_CAP2)

The following table gives a bit-level description of the Firewall Component Capability Register 2.

Table C-10: Register 2 (FC_CAP2)

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

C.6.1.4 Firewall Component Capability Register 3 (FC_CAP3)

The following table gives a bit-level description of Firewall Component Capability Register 3.

Table C-11: Register 3 (FC_CAP3)

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

C.6.2 Configuration registers

A Firewall Component has the Firewall Component Configuration registers, FC_CFG{0-3}.

Software can use these registers to find out how the Firewall Component has been configured, for example, the address width of the Bus interfaces.

Table C-12: Configuration registers FC_CFG{0-3}

Offset	Short Name	Access	Name
0xFB0	FC_CFG0	RO	Firewall Component Configuration Register 0
0xFB4	FC_CFG1	RO	Firewall Component Configuration Register 1
0xFB8	FC_CFG2	RO	Firewall Component Configuration Register 2
0xFBC	FC_CFG3	RO	Firewall Component Configuration Register 3

C.6.2.1 Firewall Component Configuration Register 0 (FC_CFG0)

The following table gives a bit-level description of Firewall Component Configuration Register 0.

Table C-13: Register 0 (FC_CFG0)

Bits	Name	Description	Type	Reset
[31:5]	-	Reserved	RO	0x000_0000
[4:0]	FC_ID	Firewall Component ID	RO	IMPL_DEF

C.6.2.2 Firewall Component Configuration Register 1 (FC_CFG1)

The following table gives a bit-level description of Firewall Component Configuration Register 1.

Table C-14: Register 1 (FC_CFG1)

Bits	Name	Description	Type	Reset
[31:21	-	Reserved	RO	0x000
[20]	SEC_SPT	Firewall Component support for checking the security of the incoming transaction and setting the security of the outgoing transaction. 0b0: Not supported 0b1: Supported This field is 0b0 when SE.0 is implemented and 1 when SE.1 is implemented.	RO	IMPL_DEF
[19]	MA_SPT	Firewall Component support for setting the memory type of the outgoing transaction. 0b0: Not supported 0b1: Supported This field is 0b0 when TE.0 is implemented.	RO	IMPL_DEF
[18]	SH_SPT	Firewall Component support for setting the shareability of the outgoing transaction. 0b0: Not supported 0b1: Supported This field is 0b0 when TE.0 is implemented.	RO	IMPL_DEF

Bits	Name	Description	Type	Reset
[17]	INST_SPT	Firewall Component support for checking whether the incoming transaction is an instruction or data access and setting the instruction property of the outgoing transaction. 0b0: Not supported 0b1: Supported	RO	IMPL_DEF
[16]	PRIV_SPT	Firewall Component support for checking the privileged level of the incoming transaction and setting the privileged level of the outgoing transaction. 0b0: Not supported 0b1: Supported	RO	IMPL_DEF
[15:14]	-	Reserved	RO	0x00
[13:12]	NUM_MPE	Number of MPEs implemented, per region, in the Firewall Component. 0b00: 1 MPE per region 0b01: 2 MPEs per region 0b10: 3 MPEs per region 0b11: 4 MPEs per region This field is Reserved and treated as RAZ/WI when PE.0 is implemented.	RO	IMPL_DEF
[11]	-	Reserved	RO	0x0
[10:8]	MNRS	Minimum Region Size. 0x1: 64B 0x2: 128B 0x3: 256B 0x4: 512B 0x5: 1KB 0x6: 2KB 0x7: 4KB This field is Reserved and treated as RAZ/WI when PE.0 is implemented.	RO	IMPL_DEF
[7:0]	NUM_RGN	Number of regions implemented in the Firewall Component. The number of regions implemented is NUM_RGN+1. This field is Reserved and treated as RAZ/WI when PE.0 is implemented.	RO	IMPL_DEF

C.6.2.3 Firewall Component Configuration Register 2 (FC_CFG2)

The following table gives a bit-level description of Firewall Component Configuration Register 2.

Table C-15: Register 2 (FC_CFG2)

Bits	Name	Description	Type	Reset
[31]	SINGLE_MST	Whether the Firewall Component supports a single MasterID. 0b0: Firewall Component supports more than one MasterID. 0b1: Firewall Component supports only one MasterID.	RO	IMPL_DEF
[30:21]	-	Reserved	RO	0x0_0000
[20:13]	PROT_SIZE	Protection Size The value of this field indicates the range of addresses which the Firewall Component protects. 0x00: 0B 0x05: 32B 0x06: 64B ... 0x0C: 4K 0x0D: 8KB ... 0x40: 16EB The value of this field depends on the level of Protection Extension implemented by the Firewall Component. PE.0: 0x00 PE.1: Same value as FC_CFG2.MXRS. PE.2: Sampled value of the Protection Size interface for the Firewall Component.	RO	IMPL_DEF
[12:8]	MST_ID_WIDTH	Maximum MasterID field width of the Bus interfaces of the Firewall Component. 0x00: 1 bit of MasterID 0x01: 2 bits of MasterID ... 0x1F: 32 bits of MasterID When a Firewall Component supports more than one Bus Slave or Master interface this field is set to the largest width used.	RO	IMPL_DEF
[7]	-	Reserved	RO	0x0

Bits	Name	Description	Type	Reset
[6:0]	MXRS	The Maximum Region Size the Firewall Component Supports. 0x05: 32B 0x06: 64B ... 0x0C: 4KB 0x0D: 8KB ... 0x40: 16EB All other values are Reserved. This field is Reserved and treated as RAZ/WI invalid for a Firewall Component which implements PE.0.	RO	IMPL_DEF

Arm recommends that when FC_CFG2.SINGLE_MST reads as 1, that FC_CFG1.NUM_MPE reads as 00, as defining more than one MPE for a region would lead to a programming error as both entries would have the same MasterID value.

C.6.2.4 Firewall Component Configuration Register 3 (FC_CFG3)

The following table gives a bit-level description of Firewall Component Configuration Register 3.

Table C-16: Register 3 (FC_CFG3)

Bits	Name	Description	Type	Reset
[31:6]	-	Reserved	RO	0x000_0000
[5:0]	NUM_FC	Number of Firewall Components implemented. The number of Firewall Components which are implemented in the Firewall is NUM_FC+1. For Firewall Components, other than 0, this field is Reserved and treated as RAZ/WI.	RO	IMPL_DEF

C.7 Firewall Component states

Each Firewall Component can be in one of several states.

The following states are supported:

Disconnected

Firewall Component is unable to process transaction and configuration accesses are dependent on the level of SRE implemented:

SRE level	Description
SRE.0	Access to configuration registers of a Firewall component, in the Disconnected state, generate a Configuration Access Error response.
SRE.1	Access to configuration register of a Firewall Component, in the Disconnected state complete without generating a Configuration Access Error response. Unless the access would cause a Configuration Access Error if the Firewall Component was in the Connected state. Depending on whether the register and fields accessed by the access, are saved and restored, the access may either: • For read access return 0s or the value of the field • For write access cause an update or be treated as WI For more information on the behavior of the different registers when SRE.1 is implemented see section C.14.3 Register Behavior when SRE.1 Implemented on page 446.

Connecting

Firewall Component is transitioning from Disconnected to Connected. For configuration accesses it is **IMPLEMENTATION DEFINED**, which one of the following the implementation implements:

- Option 1: Accesses behave as if the Firewall Component, is still in the Disconnected state. Depending on the level of the SRE implemented:

  Option 1a:

  SRE.0: Generate a Configuration Access Error.

  Option 1b

  SRE.1: Allow the access. If the access updates a register, which has not been restored by the Firewall Controller, it must not be possible for the restore process to complete and the old value be used.

- Option 2: Accesses is stalled, at the Programming interface of the Firewall Controller, whilst the Connect handshake is completed.


Note

Arm® recommends that an implementation selects:

- Option 1: when SRE.0 is implemented.
- Option 2 :when SRE.1 is implemented.
- Transactions are not stalled indefinitely to avoid deadlocking the system.

Connected

Firewall Component is able to process transactions and register accesses complete as normal.

Disconnecting

Firewall Component is transitioning from Connected to Disconnected. Configuration accesses behavior as described for the Connecting state.

A Firewall Component transitions between these states based on the power mode of the domain which the Firewall Component resides in.

C.8 Protection Extension

The Protection Extension enables the Firewall Component to compartmentalize the address space of the Bus Master interfaces of the masters connected to the Bus Slave interfaces of the Firewall Component


Note

In this section, any references to a Firewall Component implicitly implies that the Firewall Component implements PE.1 or greater.

Compartmentalization is achieved using regions. A Firewall Component implements between 1 and 256 regions. Each region defines the access permissions to an address range based on the master which issued the transaction. Using the Region Window Entry (RWE) registers, software can program a region. For more information on regions and Region Window Entry see sections C.8.2 Regions on page 374 and 12.3.2 Secure Enclave Registers on page 236 respectively.

Alongside the regions, a Firewall Component which implements PE.1 or greater, also implements at least one fault entry. Each fault entry provides details about a transaction which has faulted. Software accesses the individual fault entries using the Fault Window Entry (FWE). For more information on fault entries and Fault Window Entry, see sections C.8.3 Fault entries on page 383 and 12.3.2 Secure Enclave Registers on page 236 respectively.

When in the Connected state, a Firewall Component processes transactions as described in section C.8.4 Transaction processing on page 385. When processing, the Firewall Component performs checks on the transactions. Transactions which pass proceed to their destination are referred to as an allowed transactions. For transactions which fail, also referred to as faulting transactions, the Firewall Component blocks the transaction from continuing to its destination. This is referred to as terminating the transaction.

C.8.1 Protection Size and bus address widths

This section gives the rules of Protection Size and bus address width.

A Firewall Component, which supports PE.1 or greater, has the following properties:

- Minimum Region Size, MNRS: smallest region size which can be defined
- Maximum Region Size, MXRS: largest region size which can be defined
- Protection Size, PROT_SIZE: the range of address bits which the Firewall Component uses to match the incoming transaction against a region when performing address translation. This value represents the bits which can be modified.
- Bus Slave interface address width, BSAW, per Bus Slave interface implemented
- Bus Master interface address width, BMAW, per Bus Master interface implemented

The following rules apply:

- MXRS must always be >= MNRS

- MNRS must be one of the following: 32B, 64B, 128B, 256B, 512B, 1KB, 2KB, 4KB
- MXRS must be between 32B and 16EB in powers of 2
- Protection Size must be between MNRS and MXRS inclusive, or 0:
  - When PE.1 is implemented Protection Size is always equal to MXRS.
  - When PE.2 is implemented Protection Size is set by the sampled value of the Protection Size interface:
    - If the Protection Size is set to a value less than MNRS, except for 0, it is treated as if it was set to the MNRS value.
    - If the Protection Size is set to a value greater than MXRS, it is treated as if it was set to the MXRS value.
    - When the Protection Size is set to 0, all transactions are treated as failing the checks when the Firewall Component is not bypassed.
- Max(BSAW) must never be greater than Min(BMAW).
- A Firewall Component treats all incoming transactions as having an address width of 64 bits:
  - When a Bus Slave interface address width is less than 64 bits the address is zero extended to 64 bits.
  - When a Bus Master interface address width is less than 64 bits the address is truncated to match the width of the Bus Master interface.
- When a transaction is received on a Bus Slave interface there are the following requirements on the address bits of the incoming transactions:
  - Address bits log2(MXRS) to Max(BSAW)-1, if BSAW is greater, are ignored.


Note

This can lead to aliasing if Firewall Component's Max(BSAW) is greater than log2(MXRS). It is the responsibility of the integrator of the Firewall to avoid alias if it can occur. It is only considered aliasing if the bits of the address greater than the log2(MXRS)-1 can take more than a single value. If, due to an upstream component, the address bits of all transactions above this are always the same it is not considered aliasing.

  - Address bits PROT_SIZE to log2(MXRS)-1 must be 0, otherwise the transaction is automatically considered to fail the protection checks. Depending on the value of the PE_ST.FLT_CFG the transaction is either:
    - Terminated without a fault entry
    - Terminated with a fault entry. If a fault entry is generated, then a transaction fault is generated.


Note

When the Firewall Component is bypassed this requirement no longer applies.

- When a transaction is to be issued on a Bus Master interface address bits log2(MXRS) to Max(BMAW)-1, if BMAW is greater, are passed through unaltered.

C.8.2 Regions

A region has properties which software can program using the *Region Window Entry* (RWE).

A Firewall Component has between 1 and 256 regions. The number of regions implemented in a Firewall Component is set in FC_CFG1.NUM_RGN. For more information on programming a region, see section 12.3.2 Secure Enclave Registers on page 236.

C.8.2.1 Region properties

A region defines an area of memory which certain masters can perform certain memory operations on.

A region can have the following properties:

- Base address
- Upper address
- Size
- 1-4 *Master Permission Entries* (MPE), each containing:
  - MasterID
  - Enables for combinations of Secure and Non-secure, privileged and unprivileged, read, write and execute accesses.

Not all region properties are programmable by software. Some of the properties are dependent on:

- The level of extensions implemented
- The design time configuration of Firewall Component
- How software has programmed the Firewall Component

C.8.2.1 Region properties on page 374 to C.8.2.1.3 Master Permission Entries (MPEs) on page 377 describes each property in detail.

C.8.2.1.1 Base and upper address

This section describes the base and upper address of a region.

The base address property of a region sets the lower address boundary that the region matches on incoming transactions. The base address is considered inclusive to the region, i.e. the address of the incoming transaction must be greater than or equal to the base address.

The upper address property of a region sets the upper address boundary that the region will match incoming transactions on. The upper address is considered inclusive to the region, i.e. the address of the incoming transaction must be less than or equal to the upper address.

The base and upper address properties can be either:

- Fixed: the value of the property is defined at design time.
- Directly programmable: the value of the property can be programmed by software by writing to the field.
- Indirectly programmable: the value of the property can be programmed by software setting the size property. This only applies to the upper address, which becomes equal to the base address plus the size minus 1.

The table below shows the behavior of the property against the various levels of extension, configuration and programming options.

Table C-18: Behavior of a region's base and upper address

Property	Level of RSE	Level of PE	RGN_SIZE.MULnPO2	Behavior
Base	X	1	X	Fixed
	X	2	X	Directly programmable
Upper	X	1	X	Fixed
	0	2	X	Indirectly programmable, via the size property
	1	2	0	Indirectly programmable, via the size property
	1	2	1	Directly programmable

In all cases:

- Both the base and upper address must be an integer multiple of the MNRS. An attempt to program a value which is not an integer multiple of MNRS is round down to the nearest integer multiple of MNRS to the value.
- The base address must be within the Protection Size of the Firewall Component. If the base address is set to a value greater than the Protection Size, no transactions can match against the region.

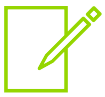
Note

Arm® recommends that the upper address is never programmed to be less than the base address. If a region is programmed like this, then no transaction matches against the region as it always fails the check. The base and upper address may be set to the same value to request a region which is equal to the MNRS in size starting at the base address.

The bits of the base and upper address used when matching an incoming transaction to the region depend on the configuration of the region:

- When either RSE.0 is or RSE.1 is implemented and MULnPO2 is 0, the Firewall Component treats the base address as being aligned to the size of the region.
- When RSE.1 is implemented and MULnPO2 is 1, the Firewall Component treats the base and upper addresses as being aligned to MNRS of the Firewall Component.


Note

Arm® strongly recommends that software always sets the base and upper address values to match the behavior of the Firewall Component.

For more information on how the base and upper addresses are used in region matching, see section C.8.2.4 Region matching on page 380.

C.8.2.1.2 Size

Each region has a size in bytes associated with it. The size can be set either using software or at design time

The size property of a region can be either fixed or programmable.

Table C-19: Behavior and legal values of a region's size property

Level of PE	Legal values for size	Behavior
1	MNRS to MXRS inclusive and 0	Fixed
2	MNRS to MXRS inclusive and 0	Programmable

The size of a region can be defined as 0 to prevent any transaction matching against the region.

If a regions size is:

- Less than MNRS: it is treated as if the size was set to 0
- Greater than MXRS: it is treated as if the size was set to 0

For more information on MNRS and MXRS, see section C.8.1 Protection Size and bus address widths on page 372.

For example, in a Firewall Component with:

- MNRS of 4KB
- MXRS of 1MB

The table below shows whether certain values of region size are legal or not.

Table C-20: Region Size legality example

Region size	MULnPO2	Legal	Notes
32B	X	No	Size is below MNRS
4KB	X	Yes	Size is between the MNRS and MXRS and a power of 2
8KB	X	Yes	Size is between the MNRS and MXRS and a power of 2
12KB	0	No	Region is defined as a power of 2
	1	Yes	Region is defined with a base and upper address. Both must be an integer multiple of MNRS.
1GB	X	No	Size is above MXRS

C.8.2.1.3 Master Permission Entries (MPEs)

There are up to four MPEs, MPE{0-3} for each region that allow up to four different permissions to be defined.

Each MPE defines the access permission for a specific master or, in the case of MPE0, for any master to that region. The MPE is only used if:

- The MasterID of the transaction and the MPE match.
- The MPE is configured to match any MasterID, only for MPE0.

The number of MPEs implemented, per region by a Firewall Component, can be discovered using the FC_CFG2.NUM_MPE field.

If a transaction matches more than one region and MPE pair, the Firewall Component generates a programming fault

If more than four masters must have access to a single area of memory, then software should define more regions with the same base address, size or upper address and use the MPE entries of the additional regions. It is valid to define the same area of memory in multiple regions so long as the same master is not defined in more than one MPE for the same area of memory.


Note

Arm® recommends that if software programs MPE0 to match against any MasterID that:

- No other MPEs are enabled for that region.
- No other regions, for the same address range are defined.

Otherwise a programming fault is generated.

Each MPE is made up of:

- MasterID, to be matched against the incoming transaction
- *Master Permissions List* (MPL)

The ranges of permissions provided, depends on the levels of extensions implemented and the Firewall Component support for transaction properties. The table below shows the permissions.

Table C-21: Master permissions supported by Firewall Component

Permission	Description
NSUR	Non-secure unprivileged read
NSUW	Non-secure unprivileged write
NSUX	Non-secure unprivileged execute
NSPR	Non-secure privileged read
NSPW	Non-secure privileged write
NSPX	Non-secure privileged execute

Permission	Description
SUR	Secure unprivileged read
SUW	Secure unprivileged write
SUX	Secure unprivileged execute
SPR	Secure privileged read
SPW	Secure privileged write
SPX	Secure privileged execute

Each permission type allows the defined master to perform access of that type when set to 1. For example, a non-secure privileged read transaction from Master 0 when NSPR is set to 1, passes the checks, but a Secure unprivileged write from Master 0 when SUW is set to 0 fails the checks.

Both the permissions and the MasterID are ignored unless the enable bit for the MPE is set to 1.

C.8.2.2 Enables

This section summarizes region enables.

C.8.2.2.1 Region enable

The region enable is always implemented. It is used to control whether incoming transactions can match against this region.

When the region enable is set to 1, then any write to modify the following region properties generates a Configuration Access Error:

- Base address
- Size
- Upper address
- Translation address
- Translation properties

C.8.2.2.2 MPE enable

There is a Master Permission Enable per MPE which is implemented for the region.

This enable controls whether a transaction can match against the MPE for the region. When the MPE enable is set to 1, any write to modify the associated MPE is ignored.

C.8.2.2.3 Region lock

The region lock is implemented when the LDE.2 is implemented.

When the region lock is set to 1 any attempt to change any property of the region generates:

- A Configuration Access Error

- A Tamper interrupt and Tamper report, if there is no valid Tamper report present. Otherwise, a Tamper Overflow interrupt is generated.

The behavior of an attempt to update the region lock depends on the value of the region lock and the lockdown state of the Firewall Component, when it is set to 1, depends on the lockdown state of the Firewall Component:

- Open: Write is allowed
- Partial and region is locked generates:
  - A Configuration Access Error
  - A Tamper interrupt and Tamper report, if there is no valid Tamper report present. Otherwise, a Tamper Overflow interrupt is generated.
- Full: Generates:
  - A Configuration Access Error
  - A Tamper interrupt and Tamper report, if there is no valid Tamper report present. Otherwise, a Tamper Overflow interrupt is generated.

For more information on LDE see section C.15.3 Lockdown Extension on page 454.

C.8.2.3 Default regions

When PE.2 is implemented, the Firewall Component implements a default region (region 0), which behaves differently to the other regions within the Firewall Component.

The default region has the following differences to the other regions:

- Base address is fixed and set to a value of all 0s.
- Size is fixed and set equal to the FC_CFG2.PROT_SIZE field.
- RGN_TCFG2.ADDR_TRANS_EN is 0 and treated as RAZ/WI.
- RGN_SIZE.MULnPO2 is 0 and treated as RAZ/WI.


Note

When PE.2 is implemented and FC_CFG1.NUM_RGN is 0, then only the default region is implemented.

C.8.2.3.1 Default region matching

A transaction only matches against the default region when the default region is enabled, and the transaction does not match any other region.

If a transaction matches a region, but fails to find any MPE with a MasterID matching the transaction, or MPE0 is not configured to match any MasterID, then the transaction can be checked against the default region.

A transaction does not match against the default region in the following case:

If a transaction fails the permission checks, despite matching a region, and finding either:

- An MPE which has the MasterID matching the transaction
- MPE0 of the region is configured to match any MasterID

If a transaction matches more than one MPE in the default region, then a Programming Error is generated for the transaction.

C.8.2.4 Region matching

Region matching is matching a transaction against the enabled regions of the Firewall Component.

The Firewall Component performs checks that are referred to as protection logic checks, or checks. This section describes the process for performing these checks. These checks are only performed if the Firewall Component's protection logic is enabled and not bypassed. For more information on the behavior of a Firewall Component and bypass, see section C.8.6 Bypass interface on page 388.

When the protection logic is disabled and not bypassed, it treats all transactions as failing the checks and enters the transactions into the Fault state. How the Firewall Component handles the transaction in the Fault state, depends on the values of the PE_ST.{FLT_CFG,ERR,RAZ} fields.

When the protection logic is enabled, the protection logic matches a transaction against the enabled regions of the Firewall Component. The Firewall Component checks the address and properties of the transaction, against the enabled regions of the Firewall Component. A region is enabled only when RGN_ST.EN bit is 1.


Note

Before performing the region matching, the Firewall Component checks that the transaction is entirely within the Protection Size of the Firewall Component. For more information on Protection Size, see section C.8.1 Protection Size and bus address widths on page 372.

A transaction only matches if all the locations which it accesses, are contained within the region. It is permissible for a transaction to match in multiple regions at this stage. When performing the region matching, the Firewall Component uses the following:

- Transaction lower address (TLA) and upper address (TUA): Depending on the bus protocol, the protection logic needs to calculate the transaction lower and upper address from other properties of the bus protocol. For example, in AXI the lower and upper address is calculated by using the AxADDR, AxSIZE, AxLEN and AxBURST transaction properties.
- Region base address (RBA) and upper address (RUA): Depending on the configuration of the region, the protection logic needs to calculate the RUA using the RBA, size of the region, and whether it is configured as a power of 2 region.
- Address mask (MSK): This is calculated as follows:

- When MULnPO2 is 0, the value is `0xFFFF_FFFF_FFFF_FFFF` left shifted by the size of the region. If the size is set to `0x00` then the region is not checked.
- When MULnPO2 is 1 the value is `0xFFFF_FFFF_FFFF_FFFF` left shifted by MNRS.


Note

In both cases, the left shift causes 0s to be added in the least significant bit position.

The protection logic uses this expression when performing the region match:

Region Match = ((RBA & MSK) <= (TLA & MSK)) && ((TUA & MSK) <= (RUA & MSK))


Note

The Firewall always zero extends address to 64-bits.

After matching the transaction with at least one region, the enabled MPEs of the regions that match are checked for an entry which has:

- The same MasterID as the master which issued the transaction
- The MPE has been programmed to match any MasterID, only for MPE0 of each region


Note

If the Firewall Component FC_CFG2.SINGLE_MST is 1 then the MasterID is ignored.

If a single MPE is found the transaction is checked against the MPL. Only enabled MPEs, of the regions which matched, are checked. An MPE is enabled when the associated RGN_ST.MPE_EN bit is 1. If there is more than one MPE found to match the transaction, a Programming Error is generated to software.


Note

If an MPE is programmed to match all MasterIDs, this is treated as if there was an MPE entry for each MasterID. This means that if an MPE is programmed to match all MasterIDs, and there is another MPE with the same MasterID as the transaction, a Programming Error is generated to software.

Software can program limits for the transactions allowed for a master using a combination of any of the supported permissions. For example, a region can read, write or execute permissions to Secure privileged, but only give Secure unprivileged read access and no access to both Non-secure privileged and unprivileged respectively.

The following figure shows an example, where a Firewall Component has 4 regions defined, regions 1 to 4.

Figure C-5: Example of transaction matching a region

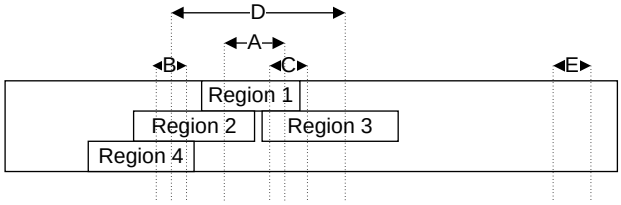

Transactions A to E are issued to the Firewall Component. The arrows indicate the range of address space which the transaction is attempting to access. The results of the region matching are as follows:

- Transaction A matches with region 1 and uses the MPEs associated with region 1 to perform the MPL check.
- Transaction B matches with regions 2 and 4. It uses both regions for the next step. However, only one of the regions can define an MPE with the MasterID associated with the transaction. If both regions define an MPE entry with the same MasterID, that results in a Programming Error being generated and transaction being treated as if it failed the checks.
- Transaction C matches with only region 3. It does not match with region 1 as the upper address of the transaction is greater than the upper address of the region.
- Transaction D does not match with any region, even though the address range it accesses is covered by region 1 to 4. This is because its lower address is below the base address for regions 1 and 3 and its upper address is above the upper address for regions 2 and 4.
- Transaction E matches no regions.

Note

The Firewall Component has an additional region - region 0, which is not shown in this diagram. The behavior of region 0 varies depending on the level of the Protection Extension implemented by the Firewall Component. See section C.8.2.3 Default regions on page 379 for more information.

Note

Even though it is possible that two regions are defined to create a single contiguous region, a transaction which accesses across the boundary would not be contained exclusively within a region. In this case the transaction does not match either of the regions.

Arm® recommends to avoid transactions crossing region boundaries by either:

- Setting MNRS of a Firewall Component to be greater than or equal than the largest size transaction which a master can issue.
- Software configuring the regions within the Firewall and the master, so that no transaction matches in two regions at once.

C.8.3 Fault entries

A Firewall Component must have at least one fault entry. Each fault entry provides details about the transactions which have entered the Faulted state.

A fault entry can be:

Valid

Contains information about a faulting transaction which has not yet been acknowledged by software.

Invalid

Contains **UNKNOWN** value.

All fault entries start in the invalid state then transition into valid state when a faulting transaction's details are loaded. Once software acknowledges the fault entry, it transitions into invalid state. After this, the fault entry transitions between valid and invalid states on transaction fault and software acknowledgement respectively.

Fault entries are not accessed directly by software but are instead accessed using the FWE. Only valid fault entries are accessible in the FWE.

For information about how software can use the fault entries see C.16.1 Fault usage model on page 473.

Fault types

A fault entry can contain either a Transaction fault or a Program fault:

- Transaction fault: details about a transaction which has failed the checks.
- Program fault: details about a transaction which has matched more than one region and MPE pairing.

The Firewall Component, if configured by the PE_ST.FLT_CFG, generates a fault entry for both types of faults. The only difference between the two fault types is the interrupt which is generated alongside the fault entry, all other rules for fault entry properties, generation, acknowledgement and overflow.

- For a transaction fault an Access Error interrupt is generated.
- For a programming fault a Programming Error interrupt is generated.

Fault entry properties

A fault entry provides the following information:

- Address of the transaction
- Transaction properties:
  - MasterID
  - Privilege level, if FC_CFG1.PRIV_SPT is 1

- Data or instruction, if FC_CFG1.INST_SPT is 1
- Read or write
- Security, if SE.1 is implemented

- Fault type

Fault entry generation

Depending on the value of PE_ST.FLT_CFG at the time the transaction is checked, an implementation must generate a fault entry.

See sections C.8.3.1 PE_ST.FLT_CFG is 0b10 on page 385 to C.8.3.2 PE_ST.FLT_CFG is 0b11 on page 385.

Depending on the bus protocol used for the Bus Slave interface and the implementation of the Firewall Component, it may be possible for a Firewall Component to process more than one transaction in a cycle. This can occur when:

- Read and write transactions are issued using different signals
- Within a single Bus Slave interface, transactions are allocated to a queue. Whereas transactions within different queues have no ordering requirements. Meanwhile transactions within the same queue are required to be processed in the order received.

When the protection logic processes more than one transaction in the same cycle, the order in which the fault entries are generated, for the parallel processed transactions, is **IMPLEMENTATION DEFINED**. It is also **IMPLEMENTATION DEFINED**, whether more than one fault entry per queue is generated at the same time. The protection logic must attempt to generate the fault entry for subsequent transactions in the queue when a transaction reaches the front of the queue.


Note

For two or more transactions, to be processed, at the same time there cannot be any dependency between the transactions.

Fault entry acknowledgement

A fault entry remains valid until software acknowledges the fault entry.

Fault entry overflow

A Firewall Component can generate more faults than it has fault entries. This leads to a Fault Entry Overflow and a Fault Entry Overflow interrupt is generated.

Fault entry and power

Fault entries are lost when the Firewall Component enters the Disconnected state.

Software can use the PE_CTRL.FE_PWR to set the PE_ST.FE_PWR field and prevent Firewall Components from entering the Disconnected state when the FWE contains a valid fault entry.

Fault window entry behavior

The FWE behaves as a FIFO with fault entries automatically added as they are generated, provided an invalid fault entry exists.

When software acknowledges a fault entry, it is removed from the FIFO and the FWE then points to the next valid fault entry, if one exists. If no more valid fault entries exist, the FWE:

- Indicates that there are no more valid fault entries, by the FE_CTRL.FE_VLD bit being 0
- FE_TAL, FE_TA and FE_TP all read as `0x0`
- Writes to the FE_CTRL.ACK field are ignored

The FE_CTRL register implements a field, FE_CTRL.LAST_FE, which indicates if the current fault entry is the last one in the FWE. When the FWE points to the last entry this field reads as 1, otherwise it reads as 0.

C.8.3.1 PE_ST.FLT_CFG is `0b10`

When a transaction faults and the PE_ST.FLT_CFG is `0b10`, the transaction enters the Faulted state.

The transaction is completed by the Firewall Component. Software has requested that a fault entry be generated. If when the protection logic attempts to generate the fault entry:

- All the fault entries are currently valid the protection logic does the following:
  - Transaction enters the Faulted state without generating a fault entry
  - Generates a Fault Entry Overflow interrupt
- At least one fault entry is invalid the protection logic does the following:
  - Transaction enters the Faulted state
  - Generates the fault entry. The fault entry must be added to the FWE before the terminated response is provided to the issuing master.
  - Generates an Access or Programming Error interrupt

Arm® recommends that the Firewall Components attempt to generate a fault entry and Access or Programming Error interrupt, for transactions which enter the Faulted state, as soon as possible after the checks have been performed. This removes the need to hold the transaction information after the transaction has entered the Retired state.

C.8.3.2 PE_ST.FLT_CFG is `0b11`

When a transaction faults and the PE_ST.FLT_CFG is `0b11` the transaction enters the Faulted state.

The transaction is completed by the Firewall Component. No fault entry or Access or Programming Error interrupt are generated.

C.8.4 Transaction processing

Each transaction is processed by a sequence of operations as they it arrives on the Bus Slave interface.

The default behavior is that a transaction enters the Faulted state unless it passes all the checks.

The figure below shows the sequence of operations to process transactions as they arrive on the Bus Slave interface, of a Firewall Component which supports PE.1 or greater.

Figure C-6: Protection logic transaction processing flow

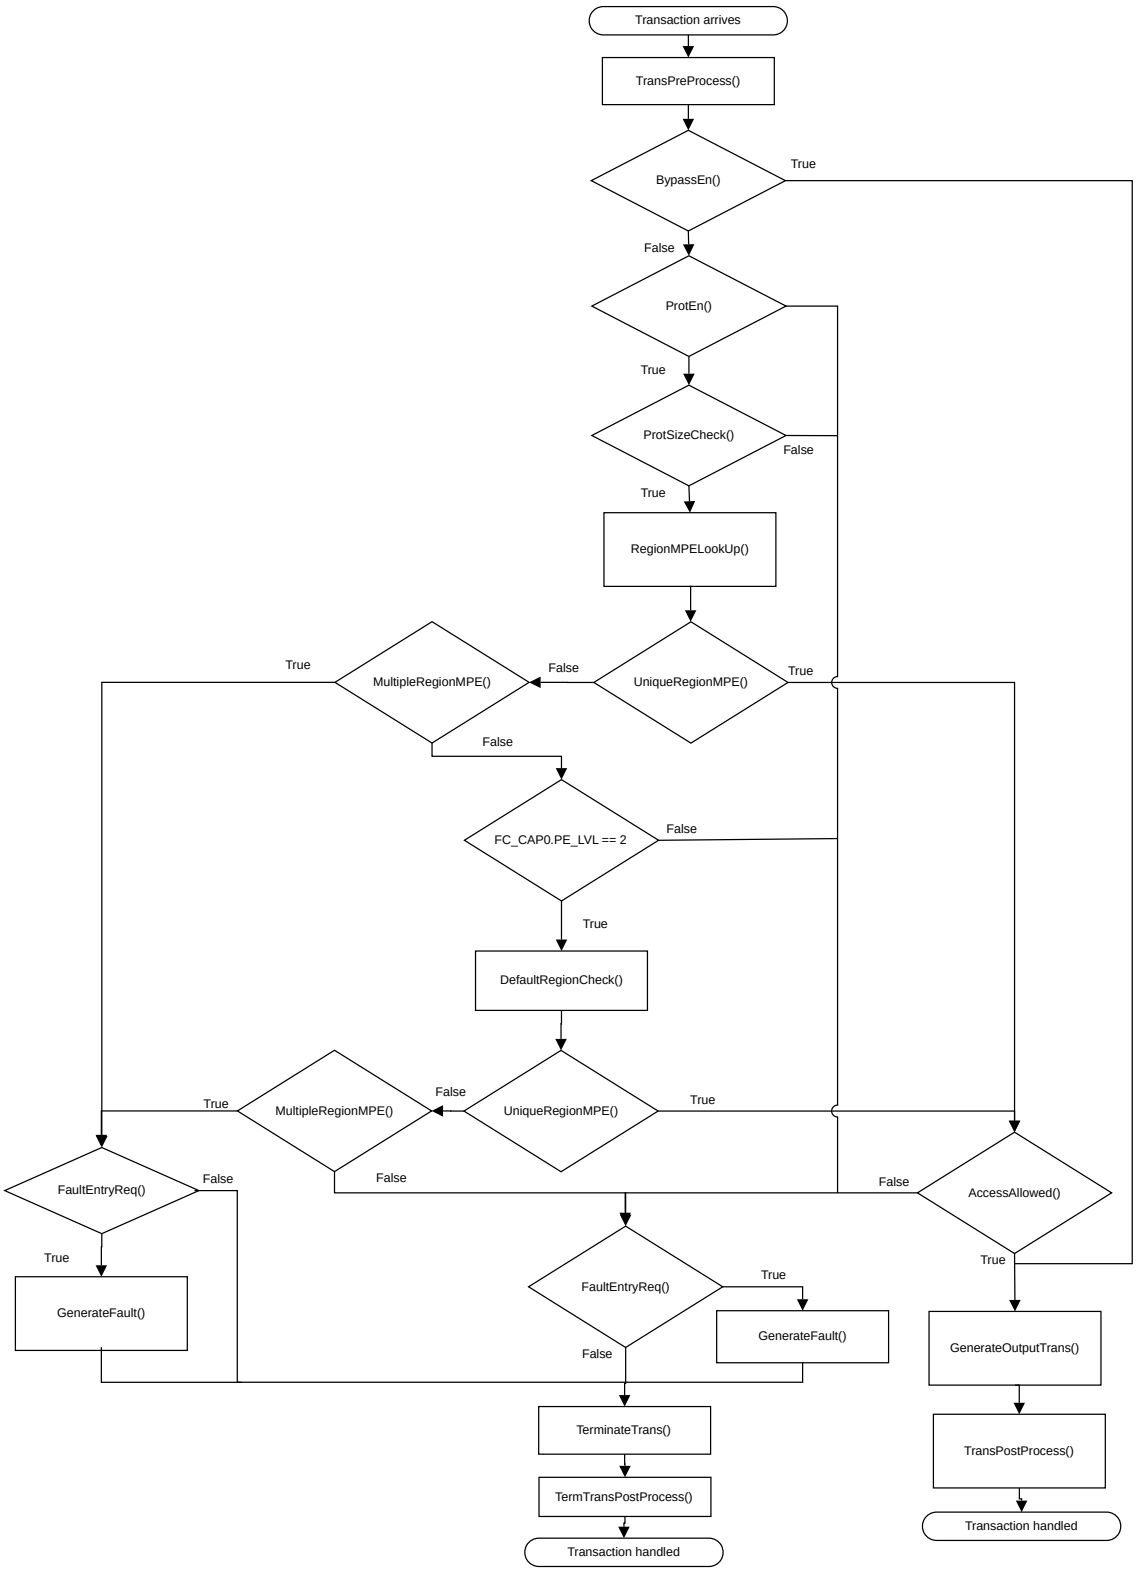

The following table describes the functions in the above figure.

Table C-22: Transaction processing functions

Processing function	Description
TransPreProcess()	Convert the incoming bus protocol into the format used by the protection logic. This function can block waiting for a new transaction.
TransPostProcess()	Convert the transaction format used by the protection logic, into the format used by the Bus Master interface.
TermTransPostProcess()	Convert the transaction format used by the Firewall Component, into the format used by the Bus Slave interface.
ProtEn()	This function checks whether the protection logic is enabled or not. It returns true if the logic is enabled otherwise it returns false.
BypassEn()	This function checks whether the Firewall Component protection logic is bypass or not. It returns true if the protection logic is bypassed otherwise it returns false.
ProtSizeCheck()	This function checks that the transaction is within the Protection Size defined for the Firewall Component's protection logic. It is applied when the transaction arrives, before any region lookup. It returns true when the transaction address is within the Protection Size, otherwise it returns false.
RegionMPELookUp()	This function performs the following operations: • Checks the transaction against all enabled regions, except for the default regions for Firewall Components which implemented PE.2, looking for a region which the transaction matches. • Checks the transaction against the MPEs, of the regions the transaction matches against, which have a MasterID the same as the transaction, or match any master. This function returns a count of the number of region and MPE pairs it matches against. When the transaction matches against a unique region MPE pair, the function returns the region and MPE it matches.
UniqueRegionMPE()	This function checks the results of the RegionMPELookUp() function. It returns true if a unique region and MPE pair was found. Otherwise it returns false.
MultipleRegionMPE()	This function checks the results of the RegionMPELookUp() function. It returns true if multiple region and MPE pairs are found. Otherwise it returns false.
DefaultRegionCheck()	The function is only implemented for Firewall Components, which implement PE.2. It performs the same operation as RegionMPELookUp() but only on the default region.
AccessAllowed()	This function takes the MPE found by either the RegionMPELookUp() or the DefaultRegionCheck() function and applies the permission check on the transaction. It returns true if the transaction is allowed, otherwise it returns false.
GenerateOutputTrans()	This function takes the incoming transaction and the translation information, from the region associated with the MPE used for the AccessAllowed() function, to generate the output transaction.
FaultEntryReq()	This function returns true if the value of PE_ST.FLT_CFG indicates a fault entry is required.
GenerateFault()	This function attempts to generate the fault entry.
TerminateTrans()	This function generates the correct response to a transaction which the Firewall Component has terminated.

C.8.5 Protection size interface

When a Firewall Component implements PE.2, a Protection Size interface is implemented on the Firewall Controller, allowing for the modification of the Protection Size of the Firewall Component.

The value of the Protection Size interface configures the Protection Size of the Firewall Component's protection logic, see section C.8.1 Protection Size and bus address widths on page 372.

C.8.6 Bypass interface

The Firewall Component has a Bypass interface.

The value of the Bypass interface and the PE_ST.BYPASS_MSK field control whether the Firewall Component's protection logic is bypassed or not. Software can determine whether the Firewall Component's protection logic is bypassed or not using the PE_BPS.BYPASS_ST field. The table below shows how the values of the Bypass interface and PE_ST.BYPASS_MSK affect the value of PE_BPS.BYPASS_ST.

Table C-23: Firewall Component's protection logic behavior for bypass

Bypass interface	PE_ST.BYPASS_MSK	PE_BPS.BYPASS_ST
`0b0`	`0b0`	`0b0`
`0b1`	`0b0`	`0b1`
X	`0b1`	`0b0`

The value of PE_BPS.BYPASS_ST affects the behavior of all Firewall Components in the Firewall, as follows:

- `0b0`: The Firewall Component is not bypassed. Transactions are treated as passing or failing the protection logic checks depending on the results of the region matching.
- `0b1`: The Firewall Component is bypassed. Transactions are treated as passing the protection logic checks without performing a region matching. The transaction is then issued on a Bus Master interface of the Firewall Component, without any translation, if supported being applied.

C.8.7 Registers

This section summarizes the registers which are defined by the Protection Extension.

The Protection Extension defines three types of registers:

- Protection Control and Status
- Region Window Entry (RWE)
- Fault Window Entry (FWE)

The Firewall Component's protection logic can be either enabled or disabled by software. When the protection logic is enabled it processes transactions.

Alongside enabling and disabled the protection logic of the Firewall Component, the behavior of the Firewall Component when a transaction fails the checks can also be configured.

The Region Window Entry (RWE) is the interface which software uses to program the regions within the Firewall Component. The RWE occupies 21 consecutive 32-bit words. Software uses the RWE_CTRL register to select the region, within the Firewall Component the RWE refers to. For more information on the RWE_CTRL register, see section C.8.7.4 RWE Control (RWE_CTRL) on page 394.

Software accesses the fault entries using the Fault Window Entry (FWE). The FWE occupies five consecutive 32-bit words. The FE_CTRL register allows software to acknowledge the fault entry referenced by the FWE.

Summary of Protection Extension registers

Table C-24: Control and status

Offset	Short Name	Access	Name
0x100	PE_CTRL	RW	Protection Extension Control
0x104	PE_ST	RO	Protection Extension Status
0x108	PE_BS	RW	Protection Extension Bypass

Table C-25: Region Window Entry (RWE)

Offset	Short Name	Access	Name
0x10C	RWE_CTRL	RW	Region Window Entry Control
0x110	RGN_CTRL0	RW	Region Control 0
0x114	RGN_CTRL1	RW	Region Control 1
0x118	RGN_LCTRL	RW	Region Lock Control
0x11C	RGN_ST	RO	Region Status
0x120	RGN_CFG0	RW	Region Config 0
0x124	RGN_CFG1	RW	Region Config 1
0x128	RGN_SIZE	RW	Region Size
0x12C	-	RO	Reserved
0x130	RGN_TCFG0	RW	Region Translation Config 0
0x134	RGN_TCFG1	RW	Region Translation Config 1
0x138	RGN_TCFG2	RW	Region Translation Config 2
0x13C	-	RO	Reserved
0x140	RGN_MID0	RW	Region MasterID 0
0x144	RGN_MPL0	RW	Region Master Permission List 0
0x148	RGN_MID1	RW	Region MasterID 1
0x14C	RGN_MPL1	RW	Region Master Permission List 1
0x150	RGN_MID0	RW	Region MasterID 2
0x154	RGN_MPL2	RW	Region Master Permission List 2
0x158	RGN_MID0	RW	Region MasterID 3
0x15C	RGN_MPL3	RW	Region Master Permission List 3

Table C-26: Fault Window Entry (FWE)

Offset	Short Name	Access	Name
0x180	FE_TAL	RO	Fault Entry Transaction Address Lower
0x184	FE_TAU	RO	Fault Entry Transaction Address Upper
0x188	FE_TP	RO	Fault Entry Transaction Properties
0x18C	FE_MID	RO	Fault Entry MasterID
0x190	FE_CTRL	RW	Fault Entry Control

Some of the registers of the RWE are only implemented depending on the configuration and extensions which are implemented. For example, the RGN_TCFG{0-2} registers are only implemented when:

- RGN_TCFG{0-1} are implemented when RSE.1 or TE.2 are implemented.
- RGN_TCFG2 is implemented when TE.1 or greater is implemented.

When a register is not implemented in the RWE it is Reserved and generates a Configuration Access Error when accessed.

C.8.7.1 Protection Extension Control (PE_CTRL)

The following table gives a bit-level description of the Protection Extension Control (PE_CTRL) register.

The Protection Extension Control register allows software to enable or disable the Firewall Component's protection logic and select the behavior of the protection logic when a fault occurs.

Table C-27: PE_CTRL register

Bits	Name	Description	Type	Reset
[31]	EN	Request the Firewall Component's protection logic enables or disables. `0b0`: Request the Firewall Component's protection logic becomes disabled. `0b1`: Request the Firewall Component's protection logic becomes enabled. The reset value of this field is `0b0` for all Firewall Component's other than the Firewall Controller which resets to `0b1`. For more information on the Firewall Controller see section C.15 Firewall Controller on page 448.	RW	See description
[30:6]	-	Reserved	RO	`0x000_0000`
[5]	BYPASS_MSK	Request Firewall Component behavior to Bypass interface. `0b0`: Firewall Component uses the value of the Bypass interface to calculate whether it is bypassed or not. `0b1`: Firewall Component ignores the value of the Bypass interface and treats the value as `0b0` to calculate whether it is bypassed or not.	RW	`0b0`
[4]	FE_PWR	Request Fault Entry power behavior `0b0`: Fault Entry does not prevent entry into a Disconnected state. `0b1`: Fault Entry does prevent entry into a Disconnected state.	RW	`0b0`

Bits	Name	Description	Type	Reset
[3:2]	FLT_CFG	Requested Fault Configuration Configures the behavior of the Firewall Component when a transaction enters the Faulted state. 0b00: Reserved and treated as 10 0b01: Reserved and treated as 10 0b10: Terminate transaction, generate a fault entry and Access or Programming Error interrupt. 0b11: Terminate transaction, but no fault entry or Access or Programming Error interrupt are generated.	RW	0b10
[1]	RAZ	Requested behavior for read data returned for read transactions terminated by the Firewall Component. 0b0: Read data is based on the StreamID 0b1: Read data all 0s	RW	0b0
[0]	ERR	Requested behavior for responses for transactions terminated by the Firewall Component. 0b0: No error 0b1: Error	RW	0b1

Arm recommends the following:

- No outstanding transactions are being processed by the Firewall Component when the values of the PE_CTRL.{FLT_CFG,RAZ,ERR} are changing (for Firewall Components other than the Firewall Controller). The method which software uses to guarantee this, is outside the scope of Appendix C .
- There are no outstanding transactions, except for configuration accesses to change the value of the PE_CTRL.{FLT_CFG,RAZ,ERR}, when Firewall Controller fields are changed.

C.8.7.2 Firewall Component Status (PE_ST)

The following table gives a bit-level description of the Firewall Component Status (PE_ST) register.

Table C-28: PE_ST register

Bits	Name	Description	Type	Reset
[31]	EN	Status of the Firewall Component's protection logic. 0b0: Firewall Component's protection logic is disabled. 0b1: Firewall Component's protection logic is enabled. When SRE.1 is implemented and the Firewall Component has entered the Disconnected state this field matches the value in PE_CTRL.EN. As if the request to change from enabled to disabled or disabled to enabled has completed. The reset value matches the reset value of PE_CTRL.EN.	RO	See description
[30:6]	-	Reserved	RO	0x0000_000

Bits	Name	Description	Type	Reset
[5]	BYPASS_MSK	Firewall Component behavior to Bypass interface. 0b0: Firewall Component uses the value of the Bypass interface to calculate whether it is bypassed or not. 0b1: Firewall Component ignores the value of the Bypass interface and treats the value as 0b0 to calculate whether it is bypassed or not.	RO	0
[4]	FE_PWR	Fault Entry power behavior. 0b0: Fault Entry does not prevent entry into a Disconnected state. 0b1: Fault Entry does prevent entry into a Disconnected state.	RO	0b0
[3:2]	FLT_CFG	Fault Configuration Behavior of the Firewall Component when a transaction enters the Faulted state. 0b00: Reserved 0b01: Reserved 0b10: Terminate transaction, generate a fault entry and Access or Programming Error interrupt. 0b11: Terminate transaction, but no fault entry or Access or Programming Error interrupt are generated.	RO	0b10
[1]	RAZ	Value returned for read accesses when the Firewall Component terminates the transaction. 0b0: Read data is based on the StreamID 0b1: Read data is all 0s	RO	0b0
[0]	ERR	Response for a transaction terminated by the Firewall Component. 0b0: No error 0b1: Error	RO	0b1

The value of PE_ST.EN defines whether the protection logic is enabled or disabled. When the protection logic is disabled, transactions do not match against the regions of the Firewall Component and all enter the Fault state.

The values of PE_ST.{FLT_CFG,RAZ,ERR} define the fault behavior of the Firewall Component's protection logic. The reset values of the PE_ST.{FLT_CFG,RAZ,ERR} fields match the reset values of the associated field in the PE_CTRL register.


Note

See section C.8.9 Protection logic terminated transaction response on page 415 for more information on the transaction response generated for transactions terminated by the Firewall Component's protection logic.

C.8.7.3 Protection Extension Bypass (PE_BPS)

The following table gives a bit-level description of the Protection Extension Bypass (PE_BPS) register.

Table C-29: PE_BPS register

Bits	Name	Description	Type	Reset
[31]	BYPASS_VLD	Indicates whether the values in the BYPASS_ST and BYPASS_IF_ST is valid or not. `0b0`: Values are not valid `0b1`: Values are valid The behavior of this field depends the level of SRE implemented by the Firewall. SRE.0: This field always reads as `0b1` SRE.1: This field reads as `0b1` only when the Firewall Component is in the Connected state. Otherwise this field reads as `0b0`.	RO	See description
[30:2]	-	Reserved	RO	`0x0000_0000`
[1]	BYPASS_ST	Indicates if the Firewall Component's protection logic is bypassed or not. `0b0`: Firewall Component's protection logic is not bypassed. `0b1`: Firewall Component's protection logic is bypassed.	RO	0
[0]	BYPASS_IF_ST	Bypass interface status The reset value of this field depends on the value of the Bypass interface of the Firewall Component.	RO	See description

C.8.7.4 RWE Control (RWE_CTRL)

The following table gives a bit-level description of the RWE Control (RWE_CTRL) register.

Table C-30: RWE Control register description

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	`0x00_0000`
[7:0]	RGN_INDX	Region Index Selects the region to which the RWE refers to. The width of this field is dependent on the log2(FC_CFG1.NUM_RGN+1) rounded up to the nearest whole number. Any unused bits are Reserved and treated as RAZ/WI.	RW	`0x00`

Configuration Access Errors are generated for any access to any register of the RWE, when the RWE_CTRL.RGN_INDEX refers to a region which is not implemented by the Firewall Component. For example, a Firewall Component with 6 regions has an RGN_INDX field width of 3. This allows software to program the values 0 to 7, however only regions 0 to 5 exists within the Firewall Component.

C.8.7.5 Region Control 0 (RGN_CTRL0)

The following table gives a bit-level description of the Region Control 0 (RGN_CTRL0) register.

Table C-31: Region Control 0 register

Bits	Name	Description	Type	Reset
[31:1]	-	Reserved	RO	0x0000_0000
[0]	EN	Region enable 0b0: Request to disable the region 0b1: Request to enable the region	RW	0

C.8.7.6 Region Control 1 (RGN_CTRL1)

The following table gives a bit-level description of Region Control 1 (RGN_CTRL1) register.

Table C-32: RGN_CTRL1 register

Bits	Name	Description	Type	Reset
[31:5]	-	Reserved	RO	0x000_0000
[4]	MPE3_EN	Master enable for MPE3 0b0: Request to disable Master permission entry 3. 0b1: Request to enable Master permission entry 3. Only implemented when FC_CFG2.NUM_MPE is 3. Otherwise this field is Reserved and treated as RAZ/WI.	RW	0
[3]	MPE2_EN	Master enable for MPE2 0b0: Request to disable Master permission entry 2. 0b0: Request to enable Master permission entry 2. Only implemented when FC_CFG2.NUM_MPE is 2 or greater. Otherwise this field is Reserved and treated as RAZ/WI.	RW	0
[2]	MPE1_EN	Master enable for MPE1 0b0: Request to disable Master permission entry 1. 0b1: Request to enable Master permission entry 1. Only implemented when FC_CFG2.NUM_MPE is 1 or greater. Otherwise this field is Reserved and treated as RAZ/WI.	RW	0
[1]	MPE0_EN	Master enable for MPE0 0b0: Request to disable Master permission entry. 0 0b1: Request to enable Master permission entry 0.	RW	0
[0]	-	Reserved	RO	0

C.8.7.7 Region Lock Control (RGN_LCTRL)

The following table gives a bit-level description of the Region Lock Control (RGN_LCTRL) register.

Table C-33: RGN_LCTRL register

Bits	Name	Description	Type	Reset
[31:1]	-	Reserved	RO	0x0000_0000
[0]	LOCK	Control the lock status of the region. 0b0: Region is unlocked 0b1: Region is locked When this field is 1 a configuration access which attempts to update any of the following fields: • RGN_CTRL{0-1} • RGN_CFG{0-1} • RGN_SIZE • RGN_TCFG{0-2} • RGN_MID{0-3} • RGN_MPL{0-3} Generates: • A Configuration Access Error • A Tamper interrupt and Tamper report, if there is no valid Tamper report present. Otherwise, a Tamper Overflow interrupt is generated. For more information see C.15.3 Lockdown Extension on page 454. This field is Reserved and treated as RAZ/WI when LDE.0 or LDE.1 is implemented. When any of the following occur: • Firewall Components enter the Full lockdown state • Firewall Components enter the Partial lockdown state and this field is set to 1 This field becomes RO and any attempt to update this field generates: • A Configuration Access Error • A Tamper interrupt and Tamper report, if there is no valid Tamper report present. Otherwise, a Tamper Overflow interrupt is generated. For more information see section C.15.3 Lockdown Extension on page 454. For more information on the Firewall Component lockdown state see C.13.1 Firewall Component lockdown on page 441.	RW	0x0

C.8.7.8 Region Status (RGN_ST)

The following table gives a bit-level description of the Region Status (RGN_ST) register.

Table C-34: RGN_ST register

Bits	Name	Description	Type	Reset
[31:5]	-	Reserved	RO	0x000_0000
[4]	MPE3_EN	Master enable for MPE3: • 0b0: Master permission entry 3 is disabled. • 0b1: Master permission entry 3 is enabled. Only when a master permission entry is enabled is an incoming transaction be allowed to match against the entry. When this field is 1, RGN_MID3 and RGN_MPL3 are read-only and attempts to update the registers will generate a Configuration Access Error.	RO	0
[3]	MPE2_EN	Master enable for MPE2: • 0b0: Master permission entry 2 is disabled. • 0b1: Master permission entry 2 is enabled. Only when a master permission entry is enabled will an incoming transaction be allowed to match against the entry. When this field is 1, RGN_MID2 and RGN_MPL2 are read-only and attempts to update the registers will generate a Configuration Access Error.	RO	0
[2]	MPE1_EN	Master enable for MPE1: • 0b0: Master permission entry 1 is disabled. • 0b1: Master permission entry 1 is enabled. Only when a master permission entry is enabled will an incoming transaction be allowed to match against the entry. When this field is 1, RGN_MID1 and RGN_MPL1 are read-only and attempts to update the registers will generate a Configuration Access Error.	RO	0
[1]	MPE0_EN	Master enable for MPE0: • 0b0: Master permission entry 0 is disabled. • 0b1: Master permission entry 0 is enabled. Only when a master permission entry is enabled will an incoming transaction be allowed to match against the entry. When this field is 1, RGN_MID0 and RGN_MPL0 are read-only and attempts to update the registers will generate a Configuration Access Error.	RO	0

Bits	Name	Description	Type	Reset
[0]	EN	Region enable: • 0b0: Region disable • 0b1: Region enable When set to 1 the following registers are read-only: • RGN_CFG0 • RGN_CFG1 • RGN_TCFG0 • RGN_TCFG1 • RGN_TCFG2 • RGN_SIZE Any attempt to write to these registers, when the region is enabled, generates a Configuration Access Error.	RO	0

C.8.7.9 Region Config {0,1} (RGN_CFG{0,1})

The following table gives a bit-level description of the Region Config {0,1} register.

The base address of a region is programmed using the RGN_CFG{0,1} registers, with the lower 32-bits of the address being in RGN_CFG0 and the upper 32-bits in RGN_CFG1.

Table C-35: RGN_CFG{0,1} register

Bits	Name	Description	Type	Reset
[63:5]	BASE_ADDR	The base address of the region. The width of this field depends on the MXRS and MNRS properties for the Firewall Component. Number of bits implemented is log2(MXRS)-1 to log2(MNRS), starting at bit log2(MNRS). Any unimplemented bits are Reserved and treated as RAZ/WI. If MXRS and MNRS are equal, then all bits in this register are Reserved and treated as RAZ/WI.	See below	See below
[4:0]	-	Reserved	RO	0x00

Depending on the Firewall Component type, the access type and reset value of BASE_ADDR is:

- Read-write with an UNKNOWN reset value for a Firewall Component, implementing PE.2
- Read-only with a reset value equal to the base address of the region, bits log2(MXRS)-1 to log2(MNRS) only, for Firewall Component, implementing PE.1


Note

If either RGN_CFG{0,1} registers have all bits Reserved and treated as RAZ/WI, then the register is Reserved.

C.8.7.10 Region Size (RGN_SIZE)

The following table gives a bit-level description of the Region Size (RGN_SIZE) register.

Table C-36: RGN_SIZE register

Bits	Name	Description	Type	Reset
[31:9]	-	Reserved	RO	0x00_0000
[8]	MULnPO2	Selects whether the region is defined with a base address and size or a base and upper address. 0b0: Region defined by a base address and size, which must be a power of 2. The value of RGN_SIZE.SIZE is used. 0b1: Region defined by a base and upper address, which must both be an integer multiple of MNRS. The value of RGN_SIZE.SIZE is ignored. This field is Reserved and treated as RAZ/WI when RSE.0 is implemented.	See Additional information on page 399	See Additional information on page 399
[7:0]	SIZE	The size of the region. 0x00: 0B 0x05: 32B 0x06: 64B ... 0x0C: 4KB 0x0D: 8KB ... 0x40: 16KB The legal values this field can be set to depend on the MNRS and MXRS of the Firewall Component. If software attempts to set this field to a value which is: • Greater than 0x40 • Less than MNRS of the Firewall. Component • Greater than MXRS of the Firewall Component The field is set to 0x00 instead. When this field reads as 0x00 does not match against any transactions. The value of the RGN_SIZE field is only used if RGN_SIZE.MULnPO2 is 0. Otherwise the field is ignored.	See Additional information on page 399	See Additional information on page 399

Additional information

Depending on the level of Protection Extension implemented:

- MULnPO2 field is:
  - When PE.2 is implemented:
    - Read-write with an UNKNOWN reset value for all regions, except for the default region
    - Read-only with a reset value of 0 for the default region
  - When PE.1 is implemented:
    - Read-only with a reset value set at design time
- SIZE field is:
  - When PE.2 is implemented:
    - Read-write with an UNKNOWN reset value for all regions, except for the default region
    - Read-only with a reset value which matches the value in the FC_CFG2.PROT_SIZE field, for the default region
  - When PE.1 is implemented and MULnPO2 is set to 0:
    - Read-only with a reset value set at design time
  - When PE.1 is implemented and MULnPO2 is set to 1:
    - Read-only with an UNKNOWN reset value

C.8.7.11 Region Translation Config {0,1} (RGN_TCFG{0,1})

The following table gives a bit-level description of the Region Translation Config {0,1} register.

The output address or the upper address of a region is programmed using the RGN_TCFG{0,1} registers, with the lower 32-bits of the address being in RGN_TCFG0, and the upper 32-bits in RGN_TCFG1.

These registers are implemented when either RSE.1 or TE.2 is implemented, otherwise they are Reserved and generate a Configuration Access Error when accessed.

Table C-37: RGN_TCFG{0,1} register

Bits	Name	Description	Type	Reset
[63:5]	OUTPUT_ADDR/ UPPER_ADDR	The output address or the upper address range of the region. The width of this field depends on the MXRS and MNRS properties for the Firewall Component. Number of bits implemented is log2(MXRS)-1 to log2(MNRS), starting at bit log2(MNRS). Any unimplemented bits are Reserved and treated as RAZ/WI.	See below	See below
[4:0]	-	Reserved	RO	0x00

Depending on the Firewall Component configuration, the OUTPUT_ADDR/UPPER_ADDR field behaves as follows:

- When PE.2 is implemented:
  - Read-write field with an UNKNOWN reset value for all regions, except for the default region.

- Read-only field with an UNKNOWN reset value for the default region.

- When PE.1 and TE.0 or TE.1 are implemented and RGN_SIZE.MULnPO2 is 0:
  - Read-only field with an UNKNOWN reset value.
- When PE.1 and TE.2 are implemented and RGN_SIZE.MULnPO2 is 0:
  - Read-write field with an UNKNOWN reset value.
- When PE.1 is implemented and RGN_SIZE.MULnPO2 is 1:
  - Read-only field with a reset value set at design time.


Note

If either RGN_TCFG{0,1} registers have all bits Reserved and treated as RAZ/WI, then the register is Reserved.

C.8.7.12 Region Translation Config 2 (RGN_TCFG2)

The following table gives a bit-level description of the Region Translation Config 2 (RGN_TCFG2) register.

This register is implemented when TE.1 or greater is implemented, otherwise it is Reserved and generates a Configuration Access Error when accessed.

Table C-38: RGN_TCFG2 register

Bits	Name	Description	Type	Reset
[31:18]	-	Reserved	RO	`0x0000`
[17]	ADDR_TRANS_EN	Address Translation enable: `0b0`: Address translation is disabled. The output transaction has the same address as the incoming transaction. `0b1`: Address translation is enabled. The output transaction has an address is set based on the formula in C.10.1.1 Output transaction translation address on page 431 This field is Reserved and treated as RAZ/WI when: • TE.1 or lower is implemented • PE.2 is implemented for the Default Region • Regions 0-2 for Firewall Controller	RW	`0b0`

Bits	Name	Description	Type	Reset
[16]	MA_TRANS_EN	Memory Attribute Translation enable 0b0: Memory attribute translation is disabled. The output transaction has the same memory attribute as the incoming transaction. 0b1: Memory attribute translation is enabled. The output transaction's memory attribute is as defined in RGN_TCG2.MA. This field is Reserved and treated as RAZ/WI when either: • FC_CFG1.MA_SPT is 0b0 • TE.0 is implemented	RW	0b0
[15:14]	INST	Output transaction instruction or data: 0b00: Use incoming transaction value. 0b01: Reserved and treated as 00. 0b10: Data 0b11: Instruction This field only affects read transactions. All write transactions are considered as data and outgoing write transactions are outputted as data access. This field is Reserved and treated as RAZ/WI when either: • FC_CFG1.INST_SPT is 0b0 • TE.0 is implemented	RW	0b00
[13:12]	PRIV	Output transaction privileged level 0b00: Use incoming privileged level. 0b01:Reserved and treated as 00. 0b10: Unprivileged 0b11: Privileged This field is Reserved and treated as RAZ/WI when either: • FC_CFG1.PRIV_SPT is 0b0 • TE.0 is implemented	RW	0b00
[11:4]	MA	Output transaction memory attribute Defines the memory type, cache allocation policy and whether it is transient or not for the output transaction. For the list of values see section C.10.1.2 Output transaction memory attribute property on page 432. This field is Reserved and treated as RAZ/WI when either: • FC_CFG1.MA_SPT is 0b0 • TE.0 is implemented	RW	0x00

Bits	Name	Description	Type	Reset
[3:2]	SH	Output transaction shareability. 0b00: Non-shareable. 0b01: Use incoming shareability. 0b10: Outer shareable. 0b11: Inner shareable. This field is Reserved and treated as WI and as reads as 0b01, when either: FC_CFG1.SH_SPT is 0. or TE.0 is implemented.	RW	0b01
[1:0]	NS	Output transaction security 0b00: Output transaction is marked with the same security as incoming transaction. 0b01: Reserved and treated as 0b00. 0b10: Output transaction is marked as Secure. 0b11: Output transaction is marked as Non-secure. This only applies if the incoming transaction was Secure. This field is Reserved and treated as RAZ/WI when any of the following are true: • FC_CFG1.SEC_SPT is 0b0 • TE.0 is implemented	RW	0b00

When the Firewall Component implements PE.1 and the region has MULnPO2 set to 0b1, all fields in this register are read-only with a reset value as defined in the table above.

C.8.7.13 Region MasterID {0-3} (RGN_MID{0-3})

The following table gives a bit-level description of the Region MasterID {0-3} (RGN_MID{0-3}) register.

The number of Region MasterID implemented, per region, depends on the value of the FC_CFG1.NUM_MPE. When a Region MasterID is not implemented it is Reserved and generates a Configuration Access Error when accessed.

Table C-39: RGN_MID{0-3} register

Bits	Name	Description	Type	Reset
[31:0]	MST_ID	MasterID. The value of the MasterID part of the StreamID which the transaction must have to match this MPE. The width of this field depends on the value of FC_CFG2.MST_ID_WIDTH. Any unused bits are Reserved and treated as RAZ/WI. This field is read-only when FC_CFG2.SINGLE_MST is `0b1`.	RW	The reset value of this field depends on the value of FC_CFG2.SINGLE_MST: `0b0`: UNKNOWN `0b1`: MST_ID is a fixed value defined at design time

C.8.7.14 Region Master Permission List {0-3} (RGN_MPL{0-3})

The following table gives a bit-level description of the Region Master Permission List {0-3} (RGN_MPL{0-3}) register.

The number of Region Master Permission Lists implemented per region, depends on the value of the FC_CFG1.NUM_MPE. When a Region Master Permission List is not implemented, it is Reserved and generates a Configuration Access Error when accessed.

Table C-40: RGN_MPL{0-3} register

Bits	Name	Description	Type	Reset
[31:13]	-	Reserved	RO	`0x0_0000`
[12]	ANY_MST	Selects whether entry is used for all transactions, irrespective of MasterID. `0b0`: MPE only used if MasterID of transaction and MPE match `0b1`: MPE used irrespective of transaction MasterID This field is read-only when FC_CFG2.SINGLE_MST is 1. This field is Reserved and treated as RAZ/WI for RGN_MPL{1-3}.	RW	The reset value of this field depends on the value of FC_CFG2.SINGLE_MST: `0b0`: UNKNOWN `0b1`: ANY_MST is 0
[11]	SPX	Secure privilege execute enable `0b0`: Secure privileged instruction fetches are not allowed. `0b1`: Secure privileged instruction fetches are allowed. This field is Reserved and treated as RAZ/WI when any of the following are true: • FC_CFG1.SEC_SPT is 0 • FC_CFG1.PRIV_SPT is 0 • FC_CFG1.INST_SPT is 0	RW	UNKNOWN

Bits	Name	Description	Type	Reset
[10]	SPW	Secure privilege write enable `0b0`: Secure privileged data write operations are not allowed. `0b1`: Secure privileged data write operations are allowed. This field is Reserved and treated as RAZ/WI when any of the following are true: • FC_CFG1.SEC_SPT is 0 • FC_CFG1.PRIV_SPT is 0	RW	UNKNOWN
[9]	SPR	Secure privilege read enable `0b0`: Secure privileged data read operations are not allowed. `0b1`: Secure privileged data read operations are allowed. This field is Reserved and treated as RAZ/WI when any of the following are true: • FC_CFG1.SEC_SPT is 0 • FC_CFG1.PRIV_SPT is 0	RW	UNKNOWN
[8]	SUX	Secure unprivileged execute enable `0b0`: Secure unprivileged instruction fetches are not allowed. `0b1`: Secure unprivileged instruction fetches are allowed. This field is Reserved and treated as RAZ/WI when any of the following are true: • FC_CFG1.SEC_SPT is 0 • FC_CFG1.INST_SPT is 0	RW	UNKNOWN
[7]	SUW	Secure unprivileged write enable `0b0`: Secure unprivileged data write operations are not allowed. `0b1`: Secure unprivileged data write operations are allowed. This field is Reserved and treated as RAZ/WI when FC_CFG1.SEC_SPT is 0.	RW	UNKNOWN

Bits	Name	Description	Type	Reset
[6]	SUR	Secure unprivileged read enable `0b0`: Secure unprivileged data read operations are not allowed. `0b1`: Secure unprivileged data write operations are allowed. This field is Reserved and treated as RAZ/WI when FC_CFG1.SEC_SPT is 0.	RW	UNKNOWN
[5]	NSPX	Non-secure privilege execute enable `0b0`: Non-secure privileged instruction fetches are not allowed. `0b1`: Non-secure privileged instruction fetches are allowed. This field is Reserved and treated as RAZ/WI when any of the following are true: • FC_CFG1.PRIV_SPT is 0 • FC_CFG1.INST_SPT is 0	RW	UNKNOWN
[4]	NSPW	Non-secure privilege write enable `0b0`: Non-secure privileged data write operations are not allowed. `0b1`: Non-secure privileged data write operations are allowed. This field is Reserved and treated as RAZ/WI when FC_CFG1.PRIV_SPT is 0.	RW	UNKNOWN
[3]	NSPR	Non-secure privilege read enable `0b0`: Non-secure privileged data read operations are not allowed. `0b1`: Non-secure privileged data read operations are allowed. This field is Reserved and treated as RAZ/WI when FC_CFG1.PRIV_SPT is 0.	RW	UNKNOWN
[2]	NSUX	Non-secure unprivileged execute enable `0b0`: Non-secure unprivileged instruction fetches are not allowed. `0b1`: Non-secure unprivileged instruction fetches are allowed. This field is Reserved and treated as RAZ/WI when FC_CFG1.INST_SPT is 0.	RW	UNKNOWN

Bits	Name	Description	Type	Reset
[1]	NSUW	Non-secure unprivileged write enable `0b0`: Non-secure unprivileged data write operations are not allowed. `0b1`: Non-secure unprivileged data write operations are allowed	RW	UNKNOWN
[0]	NSUR	Non-secure unprivileged read enable `0b0`: Non-secure unprivileged data read operations are not allowed. `0b1`: Non-secure unprivileged data read operations are allowed.	RW	UNKNOWN

C.8.7.15 Fault Entry Transaction Address Lower (FE_TAL)

The following table gives a bit-level description of the Fault Entry Transaction Address Lower (FE_TAL) register.

Table C-41: FE_TAL register

Bits	Name	Description	Type	Reset
[31:0]	FAULT_ADDR_LWR	Fault transaction address lower. This field is RAZ when FE_CTRL.FE_VLD is 0.	RO	UNKNOWN

C.8.7.16 Fault Entry Transaction Address Upper (FE_TAU)

The following table gives a bit-level description of the Fault Entry Transaction Address Upper register.

Table C-42: FE_TAU register

Bits	Name	Description	Type	Reset
[31:0]	FAULT_ADDR_UPR	Fault transaction address upper. This field is RAZ when FE_CTRL.FE_VLD is 0.	RO	UNKNOWN

C.8.7.17 Fault Entry Transaction Properties (FE_TP)

The following table gives a bit-level description of the Fault Entry Transaction Properties register.

Table C-43: FE_TP register

Bits	Name	Description	Type	Reset
[31:22]	-	Reserved	RO	`0x0000`

Bits	Name	Description	Type	Reset
[21]	W	Indicates whether the transaction was a read or write: 0b0: Read 0b1: Write	RO	UNKNOWN
[20:19]	-	Reserved	RO	0x00
[18]	INST	Indicates whether the transaction was an instruction or data access: 0b0: Data 0b1: Instruction	RO	UNKNOWN
[17]	PRIV	Indicates the privileged level of the transaction: 0b0Unprivileged 0b1: Privileged	RO	UNKNOWN
[16]	NS	Indicates the security level of the transaction: 0b0: Secure 0b1: Non-secure When SE.0 is implemented this field is Reserved and is treated as RAO/WI.	RO	UNKNOWN
[15:0]	-	Reserved	RO	0x0000

When FE_CTRL.FE_VLD is 0, this register reads as zero.

C.8.7.18 Fault Entry MasterID (FE_MID)

The following table gives a bit-level description of the Fault Entry MasterID (FE_MID) register.

Table C-44: FE_MID register

Bits	Name	Description	Type	Reset
[31:0]	MST_ID	The reset value of this field depends on the value of FC_CFG2.SINGLE_MST: Indicates the MasterID of the master which issued the transaction. The width of this field depends on the value of FC_CFG2.MST_ID_WIDTH. Any unused bits are Reserved and treated as RAZ/WI.	RO	0b0: **UNKNOWN** 0b1: is a fixed value defined at design time.

When FE_CTRL.FE_VLD is 0b0 this register reads as zero.

C.8.7.19 Fault Entry Control (FE_CTRL)

The following table gives a bit-level description of the Fault Entry Control (FE_CTRL) register.

Table C-45: FE_CTRL register

Bits	Name	Description	Type	Reset
[31]	LAST_FE	Indicates if this fault entry is the last valid entry: `0b0`: Entry is the not last valid fault entry. `0b1`: Entry is the last valid fault entry.	RO	`0x0`
[30]	FE_VLD	Indicates whether the FWE is pointing to a valid fault entry: `0b0`: FWE is pointing to an invalid fault entry. `0b1`: FWE is pointing to a valid fault entry. When this field is `0b0` the following fields, in the FWE, read as `0x0`: • FE_TAL • FE_TAU • FE_TP • FE_MID • FE_CTRL.FLT_TYPE	RO	`0x0`
[29:4]	-	Reserved	RO	`0x000_0000`
[3]	FLT_TYPE	Indicates the fault type: `0b0`: Transaction fault `0b1`: Programming fault	RO	`0x0`
[2:1]	-	Reserved	RO	`0x00`
[0]	ACK	Acknowledge the transaction. This field always reads as `0b0`. Writes to this field behave as follows: `0b0`: Ignored `0b1`: Fault transaction is acknowledged Writes to this register are ignored if FE_CTRL.FE_VLD is `0b0`.	WO	`0x0`

When software acknowledges a fault entry, the FWE points to the next fault entry.

C.8.8 Changing Configuration Settings of Protection Logic

This section covers the requirements for changing the configuration settings of the protection logic and individual regions.

C.8.8.1 Enabling and disabling the Firewall Component's protection logic

The Firewall Component's protection logic can be enabled and disabled using the PE_CTRL.EN field.

The status of the Firewall Component's protection logic is indicated by the PE_ST.EN field. When PE_CTRL.EN and PE_ST.EN have different values. It indicates that the Firewall Component is either performing an enablement or disablement of the protection logic:

Enablement

PE_CTRL.EN is `0b1` and PE_ST.EN is `0b0`

Disablement

PE_CTRL.EN is `0b0` and PE_ST.EN is `0b1`

The value of PE_ST.EN only updates to the value in the PE_CTRL.EN field, when the enablement or disablement process completes:

- The enablement process completes when the write to the PE_CTRL.EN completes.
- The disablement process completes when any new transaction detects the protection logic has been disabled and is treated as failing the checks.

If software attempts to set the PE_CTRL.EN field back to its previous value while an enablement or disablement is in-progress, the Firewall treats the update as a Configuration Access Error and does not update the PE_CTRL register.

Arm® strongly recommends:

- When software has set the value of PE_CTRL.EN, it waits for the value to be reflected in PE_ST.EN before changing the value again.
- That before enabling the protection logic, software waits for any previous configuration of the Firewall Component to complete. If software enables the Firewall Component protection logic while a previous configuration of the Firewall Component's protection logic is ongoing, it is **UNPREDICTABLE** whether processed transactions use the new or old configuration values.
- Software waits for the disablement process to complete before attempting any other configuration of the Firewall Component's protection logic. If software attempts to update the configuration of the Firewall Component's protection logic, before the disablement process completes, it is **UNPREDICTABLE** whether transactions which are processed, see the new or old configuration values.

C.8.8.2 Changing fault entry power behavior

You can change the fault entry power behavior.

The Firewall Component behavior can be configured for when the following conditions occur:

- Firewall Component has at least one valid fault entry
- The Power Control interface of the Firewall Component, makes a request to enter a power mode where the Firewall Component would be considered non-operational.

Using the PE_CTRL.FE_PWR field, software can configure whether the Firewall Component accepts or denies the power request under these conditions. Upon writing this field, the PE_ST.FE_PWR field is updated automatically to match the value written to the PE_CTRL.FE_PWR field. The Firewall Component uses the value in the PE_ST.FE_PWR to select whether it accepts or denies the power request.

C.8.8.3 Changing fault transaction behavior

The fault behavior of the Firewall Component can be configured by software, using the PE_CTRL.FLT_CFG field.

The current fault behavior of the Firewall Component, is indicated in the PE_ST.FLT_CFG field.

When the value of the PE_CTRL.FLT_CFG is different to the PE_ST.FLT_CFG field, the Firewall Component is changing its fault behavior. If software attempts to change the value of the PE_CTRL.FLT_CFG field, whilst the Firewall Component is changing its fault behavior, Firewall treats the update as a Configuration Access Error and does not update the PE_CTRL register.

The value of the PE_ST.FLT_CFG field updates to the value in the PE_CTRL.FLT_CFG field, when the Firewall Component treats any new transaction faults as configured in the PE_CTRL.FLT_CFG field. Any transaction which has previously been faulted, continues to behave as defined by the value of PE_ST.FLT_CFG at the point when it failed the checks. It is **UNPREDICTABLE** whether a transaction which fails the protection logic checks, while the values of PE_CTRL.FLT_CFG and PE_ST.FLT_CFG are different, uses the new or old value of PE_ST.FLT_CFG.

Arm® strongly recommends:

- When software changes the value of either PE_CTRL.FLT_CFG, it waits for the value of PE_ST.FLT_CFG to reflect this change, before changing the value again.
- When software changes the values of PE_CTRL.FLT_CFG there are no outstanding transactions to the Firewall Component. The software method used to achieve this is outside the scope of this document. If there are outstanding transactions, it is **UNPREDICTABLE** whether faulted transactions use the new or old value of the PE_ST.FLT_CFG field.


Note

The PE_CTRL.FLT_CFG field has Reserved values which are treated as other values. The Firewall Component can update the value of PE_ST.FLT_CFG with no delay, if the Reserved value written is treated as the current value of the PE_ST.FLT_CFG. For example, `0b01` is treated as `0b10` if the current value of PE_ST.FLT_CFG is `0b10` and software writes the PE_CTRL.FLT_CFG to `0b01`, the value of PE_ST.FLT_CFG is updated to `0b10` without delay.

C.8.8.4 Changing terminated transaction response type

When a transaction is terminated by the Firewall Component, its configuration specifies whether it generates an error response or not. Software uses the PE_CTRL.ERR field to control this behavior.

The PE_ST.ERR field indicates the current response type to terminated transactions generated by the Firewall Component. When PE_CTRL.ERR and PE_ST.ERR are different, the Firewall Component is currently changing its terminated transaction response behavior.

The value of PE_ST.ERR only updates to the value of PE_CTRL.ERR, when any new transactions terminated by the Firewall Component are treated as per the value configured in the PE_CTRL.ERR field. Any transaction which has been terminated continues to behave as defined by the value of PE_ST.ERR at the point it failed the checks While the values of PE_CTRL.ERR and PE_ST.ERR are different, it is **UNPREDICTABLE** whether a transaction which enters the Faulted state uses the new or old value of PE_ST.ERR.

If software attempts to change the value of PE_CTRL.ERR back to its previous value, when it differs from PE_ST.ERR, the firewall treats the update as a Configuration Access Error and does not update the PE_CTRL register.

Arm® strongly recommends:

- When software changes the value of the PE_CTRL.ERR field, it waits for the value of PE_ST.ERR to reflect this change before changing the value again.
- When software changes the values of PE_CTRL.ERR there are no outstanding transactions to the Firewall Component. The method software uses to achieve this is system-dependent. If there are outstanding transactions it is **UNPREDICTABLE** whether transactions terminated by the Firewall Component, use the new or old value of the PE_ST.ERR field.

C.8.8.5 Changing terminated transaction read data response

When a read transaction is terminated by the Firewall Component, it is configurable whether the read data is set to all 0s or a value dependent on the StreamID of the transaction.

Software uses the PE_CTRL.RAZ field to control this behavior. The PE_ST.RAZ field indicates the current value which the Firewall Component sets the read data to for terminated transactions. When PE_CTRL.RAZ and PE_ST.RAZ are different the Firewall Component is changing its terminated read response behavior.

The value of PE_ST.RAZ only updates to the value of PE_CTRL.RAZ, when any new read transactions terminated by the Firewall Component are treated as per the value configured in the PE_CTRL.RAZ field. Any transaction which has previously been terminated continues to behave as defined by the value of PE_ST.RAZ at the point it failed the checks. While the values of PE_CTRL.RAZ and PE_ST.RAZ are different, it is **UNPREDICTABLE** whether a read transaction which enters the Faulted state, uses the new or old value of PE_ST.RAZ.

If software attempts to change the value of PE_CTRL.RAZ back to its previous value, when it differs from PE_ST.RAZ, the Firewall treats the update as a Configuration Access Error and does not update the PE_CTRL register.

Arm® strongly recommends:

- When software changes the value of PE_CTRL.RAZ, it waits for the value of PE_ST.RAZ to reflect this change, before changing the value again.
- When software changes the values of PE_CTRL.RAZ there are no outstanding transactions to the Firewall Component. The method software uses to achieve this is system-dependent. If there are outstanding transactions it is **UNPREDICTABLE** whether read transactions terminated by the Firewall Component use the new or old value of PE_ST.RAZ field.

C.8.8.6 Enabling and disabling regions

The regions of the Firewall Component can be enabled and disabled using the RGN_CTRL0.EN field.

The status of the region is indicated by the RGN_ST.EN field. When RGN_CTRL0.EN and RGN_ST.EN have different values, it indicates that the Firewall Component is either performing an enablement or disablement of the region:

Enablement

RGN_CTRL0.EN is 1 and RGN_ST.EN is 0

Disablement

RGN_CTRL0.EN is 0 and RGN_ST.EN is 1

The value of RGN_ST.EN only updates, to the value in the RGN_CTRL0.EN field when the enablement or disablement process completes.

- The enablement process completes when the write to RGN_CTRL0.EN completes.
- The disablement process completes when no new transactions lookup against the region.

It is **UNPREDICTABLE** whether a transaction matches against the region, while the values of RGN_CTRL0.EN and RGN_ST.EN are different.

If software attempts to set the RGN_CTRL0.EN field, back to its previous value while an enablement or disablement is in-progress, the Firewall treats the update as a Configuration Access Error and does not update the RGN_CTRL0 register.

Arm® strongly recommends:

- When software sets the value of RGN_CTRL0.EN, it waits for the value to be reflected in RGN_ST.EN before changing the value again.
- Before enabling the region, software waits for any previous configuration of the region to complete. If software enables the region while a previous configuration of the region is pending it is **UNPREDICTABLE** whether transactions, which are processed by the Firewall Component, see the new or old configuration values.
- Software waits for the disablement process to complete before attempting any other configuration of the region.
- Software waits for the enablement process to complete before issuing or allowing to be issued, any transactions which must match against the region.

- Software makes sure there are no outstanding and does not allow any transactions to be issued, which can match against the region being disabled.

C.8.8.7 Enabling and disabling MPEs

The MPEs of a region can be enabled and disabled using the RGN_CTRL1.MPE_EN{0-3} fields.

The status of the MPE is indicated by the RGN_ST.MPE_EN{0-3} fields. When RGN_CTRL1.MPE_EN{0-3} and RGN_ST.MPE_EN{0-3} have different values, it indicates that the Firewall Component is either performing an enablement or disablement of the MPE:

Enablement

RGN_CTRL1.MPE_EN{0-3} is 1 and RGN_ST.MPE_EN{0-3} is 0

Disablement

RGN_CTRL1.MPE_EN{0-3} is 0 and RGN_ST.MPE_EN{0-3} is 1

The value of RGN_ST.MPE_EN{0-3} only updates, to the value in the RGN_CTRL1.MPE_EN{0-3} field, when the enablement or disablement process completes. The enablement process completes when the write to RGN_CTRL1.MPE_EN{0-3} completes. The disablement process completes when no new transactions lookup against the MPE. It is **UNPREDICTABLE** whether a transaction matches against the MPE, while the values of RGN_CTRL1.MPE_EN{0-3} and RGN_ST.MPE_EN{0-3} are different.

If software attempts to set the RGN_CTRL1.MPE_EN{0-3} field back to its previous value, while an enablement or disablement is in-progress, the Firewall treats the update as a Configuration Access Error and does not update any field in the RGN_CTRL1 register.

Arm® strongly recommends:

- When software sets the value of RGN_CTRL1.MPE_EN{0-3}, it waits for the value to be reflected in RGN_ST.MPE_EN{0-3} before changing the value again.
- Software waits for the disablement process to complete before attempting to update the RGN_MID{0-3} and RGN_MPL{0-3} registers associated with the MPE.
- Software waits for the enablement process to complete before issuing or allowing to be issued, any transaction which match against the MPE.
- Software makes sure there are no outstanding transactions, and does not allow any transactions to be issued, which can match against the MPE being disabled.

C.8.8.8 Changing bypass mask

The Firewall Component is configured to use the value of its Bypass interface or not, for calculating whether it has been bypassed or not.

Software configures this using the PE_CTRL.BYPASS_MSK field. On writing this field, the PE_ST.BYPASS_MSK field is updated automatically to match the value written to the PE_CTRL.BYPASS_MSK field. The Firewall Component uses the value in the PE_ST.BYPASS_MSK to determine the use of the Bypass interface.

C.8.8.9 Changing bypass

The Firewall Component's protection logic is bypassed using the Bypass interface and the PE_ST.BYPASS_MSK.

The PE_BPS.BYPASS_ST field shows if the Firewall Component's protection logic is bypassed. The value of the field can change if ther is:

- A change in the value of PE_ST.BYPASS_MSK field
- A change in the value of the Firewall Component's Bypass interface, shown in the PE_BPS.BYPASS_IF_ST field

The value of the PE_BPS.BYPASS_ST only updates due to one of the above causes, when it is guaranteed that any new transactions processed by the Firewall Component's protection logic are treated as per the updated value of PE_BPS.BYPASS_ST.

Arm® strongly recommends:

- Only changes PE_CTRL.BYPASS_MSK or the Bypass interface value at once.
- When software changes the value of PE_CTRL.BYPASS_MSK from 0 to 1, it waits for PE_BPS.BYPASS_ST to become 0 before changing the value of PE_CTRL.BYPASS_MSK again.
- When changing either the value of PE_CTRL.BYPASS_MSK or the Bypass interface, there are no outstanding transactions to the Firewall Component, as it is **UNPREDICTABLE** whether the transactions see the new or old value of the PE_BPS.BYPASS_ST field.
- Software only sets PE_CTRL.BYPASS_MSK to 1, when there are no requirements to perform SoC debug without software co-operation.

C.8.9 Protection logic terminated transaction response

Firewall Component's protection logic can generate multiple types of response when it terminates a transaction which failed the checks and the protection logic.

Table C-46: Protection logic termination response

Transaction Type	PE_ST.ERR	PE_ST.RAZ	Response
Read	0	0	A non-error read response is generated with the read data set to a value specific to the StreamID
	0	1	A non-error read response is generated with the read data set to all 0s
	1	0	An error read response is generated with the read data set to a value specific to the StreamID
	1	1	An error read response is generated with the read data set to all 0s
Write	0	X	A non-error write response is generated
	1	X	An error write response is generated

The StreamID specific read data is **IMPLEMENTATION DEFINED**.

It is **IMPLEMENTATION DEFINED** whether a Firewall Component indicates whether the transaction has been terminated by the Firewall Component, or by the slave.

C.9 Monitor Extension (ME)

When a Firewall Component implements ME.1 or greater, it includes monitor logic to detect errors generated by transactions issued by Masters connected to the Firewall Component.


Note

In this section any references to a Firewall Component, it is implicit that the Firewall Component implements ME.1 or greater.

When a Firewall Component's monitor logic is enabled, it checks the transaction responses for errors and generate an error detection report with an Error Detection interrupt. For read transactions the read data can be either:

- Forwarded unmodified
- Replaced with a value specified by the StreamID of the master which issued the transaction

When a Firewall Component's monitor logic is disabled, the transaction responses are passed through to the issuing master without altering the read data or generating an error detection report and interrupt.

The monitor logic, can ignore transaction responses, using an **IMPLEMENTATION DEFINED** method, where the transaction response has:

- Been generated by another Firewall Component's protection logic terminating the transaction
- Already been logged by another Firewall Component's monitor logic

For more information see section C.9.3 Error response ignore on page 420.

A Firewall Component must implement at least 1 error detection report. Software accesses the individual error detection reports using the Error Detection Window (EDW). For more information on error detection report and Error Detection Window, see C.9.5 Registers on page 423.

C.9.1 Error detection

The monitor logic checks the responses on the Bus Master interfaces of the Firewall Component.

If the response indicates an error, the Firewall Component either:

- Logs the transaction and forwards the response to the issuing master. When the response is forwarded, for a read transaction the read data can be either set to a specific value or left unmodified.
- Forwards the response, without logging or modifying the read data.

C.9.1.1 Bus Protocol Error Response

Depending on the bus protocol used, the responses received may be either:

- Single cycle, for example, AXI for a write transaction
- Multi-cycle, for example, AXI for a read transaction


Note

When referring to multi-cycle responses, this does not include any cycles where the response is being stalled, for example in AXI the RREADY is low.

Some bus protocols support treating responses to transaction as either single or multi-cycle responses. For example, AHB can be considered as either. Each beat of the request has an associated response beat, except where the master terminates the burst early due to an error in a previous beat of the burst. The monitoring logic can treat these types of protocol as either single or multi-cycle, depending on the following rules:

- If the Firewall Component implements PE.1, or greater, if the protection checks are:
  - Applied on the entire transaction, the monitoring logic treats the response as multi-cycle
  - Applied on each beat of the transaction, the monitoring logic treats the response as a single cycle

In either case the information reported in the error detection report, for the address and transaction properties must be the same as the information reported in a fault entry, if the transaction failed the checks.

If the Firewall Component implements PE.0, it is **IMPLEMENTATION DEFINED**, whether it is treated as single or multi-cycle.

For multi-cycle responses the monitoring logic, implements a simple OR function for detection of errors. If any cycle of a response generates an error, then the whole of the response is considered to be an error and is reported. The address, reported in the error detection report, when ME.2 is implemented, depends on the bus protocol type:

- Burst-based protocols: Either:
  - Value of the beat which generated the error
  - Value provided at the start of the burst
- Beat-based protocols: Value provided of the beat, which generated the error.


Note

Arm® recommends for burst-based the value provided is the value at the start of the transaction.

C.9.2 Error detection report

An error detection report contains the information about the transaction which generated the error response.

An error detection report can be:

- Valid: contains information about a transaction which caused an error response, which has not been acknowledge by software
- Invalid: contains an **UNKNOWN** value

All error detection reports start in the invalid state, then transition into valid state when an error detection report details are loaded. Once software acknowledges the error detection report, it transitions into invalid state. After this the error detection report transitions between valid and invalid states on error response detection and software acknowledgement respectively.

Error detection reports are not accessed directly by software, but are instead accessed using the EDW. Only valid error detection reports are accessible in the EDW.

For information about how software can use error detection reports see section C.16.2 Error detection report usage on page 473.

C.9.2.1 Error detection report properties

An error detection report contains the following information:

- MasterID
- Privilege level
- Data or instruction
- Security
- Read or write
- Address of transaction, when ME.2 is implemented

The information contained in the error detection report is of the incoming transaction and is not affected by any translation applied by the Firewall Component detecting the error.

C.9.2.2 Error detection report generation

When the monitor logic detects an error response, if::

- All error detection reports are valid then the monitor logic:
  - Generates an Error Detection Overflow interrupt
  - Does not generate the error detection report
  - Issues the response to the master which issued the transaction

- At least one error detection report is invalid then the monitor logic:
  - Generates the error detection report
  - Generates the Error Detection Interrupt
  - Issues the response to the master which issued the transaction


Note

For more information on how the monitor logic issues the response to the master see C.9.5.1 Monitor Extension Control (ME_CTRL) on page 424

C.9.2.3 Error Detection Report Acknowledgement

Overview of error detection report acknowledgement.

An error detection report remains valid until software acknowledges the error detection report.

C.9.2.4 Error detection report overflow

It is possible for a Firewall Component to receive more error responses, than it has error detection reports.

When this occurs the monitor logic generates an Error Detection Overflow interrupt.

C.9.2.5 Error detection report and power

Error detection reports are lost when the Firewall Component enters the Disconnected state.

Software can use the ME_CTRL.EDR_PWR to set the ME_ST.EDR_PWR field and prevent Firewall Components from entering the Disconnected state when the EDW contains a valid error detection report.

C.9.2.6 Error detection window behavior

The EDW behaves as a FIFO with error detection reports automatically added when they are generated, provided an invalid error detection report exists.

When software acknowledges an error detection report, it is removed from the FIFO and the EDW then points to the next valid error detection report if one exists. If no more valid error detection report exists, the EDW:

- Indicates that there are no more valid error detection reports, by setting EDR_CTRL.EDR_VLD bit to 0
- EDR_TAL, EDR_TAU, EDR_TP and EDR_MID all read as 0x0
- Writes to the EDR_CTRL.ACK field are ignored

The EDR_CTRL register implements a field, EDR_CTRL.LAST_EDR, which indicates if the current error detection report is the last one in the EDW. When the EDW points to the last report this field reads as 1, otherwise it reads as 0.

C.9.3 Error response ignore

This section describes when an error response could be ignored.

The monitor logic can use an **IMPLEMENTATION DEFINED** method to ignore error responses, which have either:

- Been generated by another Firewall Component's protection logic, as part of a transaction being terminated
- Already been logged by another Firewall Component's monitor logic

This reduces the number of locations where either a terminated transaction or error response is logged, when Firewall Components are connected in parallel, as the following figure shows.

Example C-1: System with Firewall Components both sides of the interconnect

This example assumes the following:

- A Firewall Component, implementing PE.1 or greater, is configured to generate fault entries and error detection reports.
- A Firewall Component, implementing ME.1 or greater, is configured with its monitor logic enabled.
- The Firewall Controller is configured to generate an error for any configuration access which generates a Configuration Access Error.

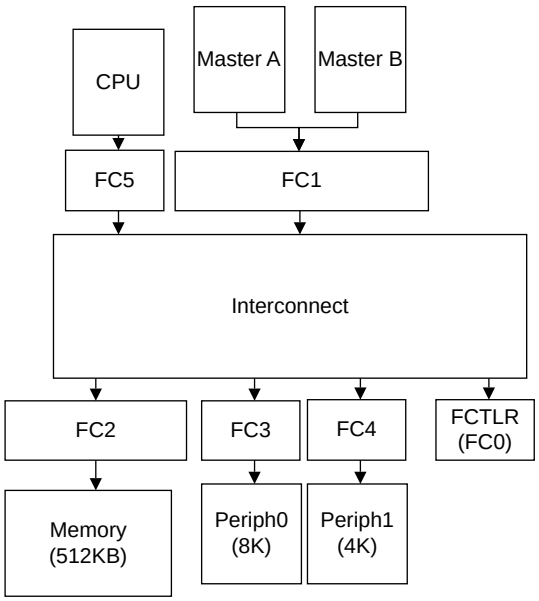

Only FC1 and FC5 have ME.2 implemented. All other FCs have ME.0 implemented. When any of the following happens, the transaction is logged in the following location:

- For a transaction terminated by FC1, a fault entry is generated in the FC1.
- For a transaction terminated by FC{0,2-4}, a fault entry is generated in the Firewall Component which terminated the transaction. The response also passes through either FC1 or FC5, depending on the master which issued the transaction. The Firewall component that terminated the transaction uses its **IMPLEMENTATION DEFINED** method to indicate that a fault has occurred. The error response should be ignored by other Firewall Components which the responses passes through and not generate an error detection report.
- For an error generated by either the memory or Periph 0 or 1, an error detection report is generated in the Firewall Component, FC1 or FC5, depending on the master which issued the transaction.
- For an error generated by FC0, an error detection report is generated in the Firewall Component, FC1 or FC5, depending on the master which issued the transaction.

C.9.4 Response processing

This section describes response processing sequence by a Firewall component.

Figure C-8: Monitor logic transaction response processing flow

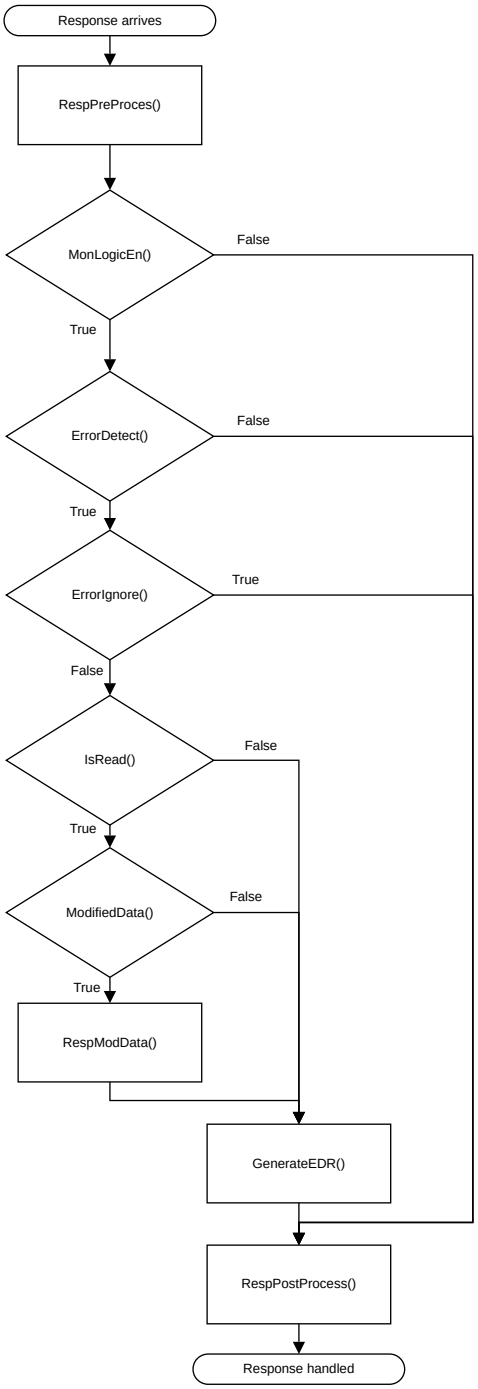

The following table gives a brief description of each function.

Table C-47: Response processing functions

Function	Description
RespPreProcess()	**IMPLEMENTATION DEFINED** function to convert the incoming bus protocol into the format used by the monitor logic. This function can block waiting for a new transaction.
RespPostProcess()	**IMPLEMENTATION DEFINED** function to covert the transaction format used by the monitor logic, into the format used by the Bus Slave interface.
MonLogicEn()	This function returns true if the monitor logic is enabled, otherwise it returns false.
ErrorDetect()	This function returns true if the transaction response indicates an error, otherwise it returns false.
ErrorIgnore()	This function returns true if the transaction response indicates that it has either: • Been generated by another Firewall Component's protection logic, as part of a transaction being terminated • Already been detected by another Firewall Component's monitor logic Otherwise, it returns false.
IsRead()	This function returns true if the transactions which generated the response was a read transaction, otherwise it returns false.
ModifiedData()	This function returns true if the monitor logic is configured to modify the read data value, based on the StreamID of the transaction. Otherwise it returns false.
RespModData()	This function replaces the read data, with a value associated with the StreamID of the transaction.
GenerateEDR()	This function attempts to generate the error detection report and Error Detection interrupt, if there is an invalid error detection report in the EDW of the Firewall Component. Otherwise, it generates an Error Detection Overflow interrupt.

C.9.5 Registers

This section summarizes the registers which are defined by the Monitor Extension.

The Monitor Extension defines two types of registers:

- Monitor Control and Status
- Error Detection Window (EDW)

The Firewall Component's monitor logic can be either enabled or disabled by software. When the monitor logic is disabled, the Firewall Component, allows responses to pass through the Firewall Component unaltered. When the monitor logic is enabled it monitors the response to transactions.

A Firewall Component which implements ME.1 or greater, has the registers listed in the table below. These registers allow software to configure the monitor logic and access any error detection reports which have been generated. When ME.0 is implemented all these registers are Reserved and generate a Configuration Access Error when accessed.

Summary of Monitor Extension registers

Table C-48: Control and Status registers

Offset	Short name	Access	name
0x200	ME_CTRL	RW	Monitor Extension Control
0x204	ME_ST	RO	Monitor Extension Status

Table C-49: Error Detection Window (EDW) registers

Offset	Short name	Access	name
0x260	EDR_TAL	RO	Error Detection Report Transaction Address Lower
0x264	EDR_TAU	RO	Error Detection Report Transaction Address Upper
0x268	EDR_TP	RO	Error Detection Report Transaction Properties
0x26C	EDR_MID>	RO	Error Detection Report MasterID
0x270	EDR_CTRL	RW	Error Detection Report Control

The EDR_TAL and EDR_TAU registers are only implemented when ME.2 is implemented, otherwise the registers are Reserved.

C.9.5.1 Monitor Extension Control (ME_CTRL)

The following table gives a bit-level description of the Monitor Extension Control (ME_CTRL) register.

The Monitor Extension Control register allows software to enable the Firewall Component's monitor logic to monitor transactions responses and select the behavior of the monitor logic when an error is detected.

Table C-50: ME_CTRL register

Bits	Name	Description	Type	Reset
[31]	EN	Request the Firewall Component's monitor logic enables or disables. 0b0: Request to disable the Firewall Component's monitor logic 0b1: Request to enable the Firewall Component's monitor logic	RW	0x0
[30:5]	-	Reserved	RO	0x000_0000
[4]	EDR_PWR	Request error detection reports power behavior. 0b0: Error detection report does not prevent entry into a Disconnected state 0b1: Error detection report does prevent entry into a Disconnected state	RW	0x0
[3:2]	-	Reserved	RO	0x00
[1]	RDUM	Request behavior for read data returned for a read transaction which has caused an error. 0b0: Read data is based on the StreamID 0b1: Read data is left unmodified	RW	0x0
[0]	-	Reserved	RO	0x0

Arm recommends that when the value of the ME_CTRL.RDUM is changed there are no outstanding transactions. The method which software uses to guarantee this is outside the scope of Appendix C.

C.9.5.2 Monitor Extension Status (ME_ST)

The following table gives a bit-level description of the Monitor Extension Status (ME_ST) register.

Table C-51: ME_ST register

Bits	Name	Description	Type	Reset
[31]	EN	Status of the Firewall Component's monitor logic enables or disables. `0b0`: Firewall Component's monitor logic is disabled. `0b1`: Firewall Component's monitor logic is enabled. When SRE.1 is implemented and the Firewall Component has entered the Disconnected state this field matches the value in ME_CTRL.EN. As if the request to change from enabled to disable or disabled to enabled has completed.	RO	`0x0`
[30:5]	-	Reserved	RO	`0x000_0000`
[4]	EDR_PWR	Error detection reports power behavior. `0b0`: Error detection report does not prevent entry into a Disconnected state. `0b1`: Error detection report does prevent entry into a Disconnected state.	RO	0
[3:2]	-	Reserved	RO	`0x00`
[1]	RDUM	Request behavior for read data returned for a read transaction which has caused an error. `0b0`: Read data is based on the StreamID `0b1`: Read data is left unmodified	RO	0
[0]	-	Reserved	RO	0

See section C.9.7 Monitor Logic Response Forwarding on page 430 for more information on how transaction responses are forwarded to the issuing master when the Firewall Component's monitor logic detects an error.

C.9.5.3 Error Detection Report Transaction Address Lower (EDR_TAL)

The following table gives a bit-level description of the Error Detection Report Transaction Address Lower (EDR_TAL) register.

This register is only implemented when ME.2 is implemented, otherwise it is Reserved and generates a Configuration Access Error when accessed.

Table C-52: EDR_TAL register

Bits	Name	Description	Type	Reset
[31:0]	ERROR_ADDR_LWR	Error transaction address lower. This field is RAZ when EDR_CTRL.EDR_VLD is 0.	RO	UNKNOWN

C.9.5.4 Error Detection Report Transaction Address Upper (EDR_TAU)

The following table gives a bit-level description of the Error Detection Report Transaction Address Upper (EDR_TAU) register.

This register is only implemented when ME.2 is implemented, otherwise it is Reserved and generates a Configuration Access Error when accessed.

Table C-53: EDR_TAU register

Bits	Name	Description	Type	Reset
[31:0]	ERROR_ADDR_UPR	Error transaction address upper. This field is RAZ when EDR_CTRL.EDR_VLD is 0.	RO	UNKNOWN

C.9.5.5 Error Detection Transaction Properties (EDR_TP)

The following table gives a bit-level description of the Error Detection Transaction Properties (EDR_TP) register.

Table C-54: EDR_TP register

Bits	Name	Description	Type	Reset
[31:22]	-	Reserved	RO	0x0000
[21]	W	Indicates whether the transaction, which caused the error, was a read or write. 0b0: Read 0b1: Write	RO	UNKNOWN
[20:19]	-	Reserved	RO	0x00
[18]	INST	Indicates whether the transaction, which caused the error, was an instruction or data access. 0b0: Data 0b1: Instruction	RO	UNKNOWN
[17]	PRIV	Indicates the privileged level of the transaction, which caused the error. 0b0: Unprivileged. 0b1: Privileged.	RO	UNKNOWN
[16]	NS	Indicates the security level of the transaction. 0b0: Secure 0b1: Non-secure	RO	UNKNOWN
[15:0]	-	Reserved	RO	0x00


Note

When EDR_CTRL.EDR_VLD is 0 this register reads as zero.

C.9.5.6 Error Detection Report MasterID (EDR_MID)

The following table gives a bit-level description of the Error Detection Report MasterID (EDR_MID) register.

Table C-55: EDR_MID register

Bits	Name	Description	Type	Reset
[31:0]	MST_ID	Indicates the MasterID of the transaction which caused the error. The width of this field depends on the value of FC_CFG2.MST_ID_WIDTH. Any unused bits are Reserved and treated as RAZ/WI.	RO	The reset value of this field depends on the value of FC_CFG2. SINGLE_MST: `0b0`: **UNKNOWN** `0b1`: Fixed valued defined at design time. When EDR_CTRL.EDR_VLD is 0 this register reads as zero.

C.9.5.7 Error Detection Report Control (EDR_CTRL)

The following table gives a bit-level description of the Error Detection Report Control (EDR_CTRL) register.

Table C-56: EDR_CTRL register

Bits	Name	Description	Type	Reset
[31]	LAST_EDR	Indicates if this error detection entry is the last valid entry. `0b0`: Report is not the last valid error detection report. `0b1`: Report is the last valid error detection report.	RO	`0b0`
[30]	EDR_VLD	Indicates whether the EDW is pointing to a valid error detection report. `0b0`: EDW is pointing to an invalid error detection report. `0b1`: EDW is pointing to a valid error detection report. When this field is `0b0` the values in the following registers read as 0: • EDR_TAL • EDR_TAU • EDR_TP • EDR_MID	RO	`0b0`
[29:1]	-	Reserved	RO	`0x0000_0000`

Bits	Name	Description	Type	Reset
0	ACK	Acknowledge the error transaction. This field always reads as `0b0`. Writes to this field behave as follows: `0b0`: Ignored `0b1`: Error detection report is acknowledged. Writes to this register are ignored if EDR_CTRL.EDR_VLD is `0b0`.	WO	`0b0`

Software uses the EDR_CTRL to:

- Know if there is a valid error detection report in the EDW
- Acknowledge the current error detection report what EDW is currently pointing to
- Know if the current error detection report is the last one in the EDW

C.9.6 Changing configuration settings of monitor logic

This section describes the requirements for changing the configuration settings of the monitor logic.

C.9.6.1 Enabling and disabling monitor logic

The Firewall Component's monitor logic can be enabled and disabled using the ME_CTRL.EN field.

The status of the Firewall Component's monitor logic is indicated by the ME_ST.EN field. When ME_CTRL.EN and ME_ST.EN have different values, it indicates that the Firewall Component is either performing an enablement or disablement of the monitor logic:

- Enablement: ME_CTRL.EN is 1 and ME_ST.EN is 0
- Disablement: when ME_CTRL.EN is 0 and ME_ST.EN is 1

The value of ME_ST.EN only updates, to the value in the ME_CTRL.EN field, when the enablement or disablement process completes. The enablement process completes when the write to ME_CTRL.EN completes. The disablement process completes when any new transactions which complete with an error, do not generate an error detection report. It is **UNPREDICTABLE** whether a transaction which completes with an error, whilst the values of ME_CTRL.EN and ME_ST.EN are different, uses the old or new value of ME_ST.EN. This includes for multicycle responses, where the error response is received before disablement process starts, but the last cycle of the response is received after the disablement process started.

If software attempts to set the ME_CTRL.EN field, back to its previous value, whilst an enablement or disablement is in-progress, the Firewall treats the update as a Configuration Access Error and does not update the ME_CTRL register.

Arm strongly recommends:

- That when software has set the value of ME_CTRL.EN, it waits for the value to be reflected in ME_ST.EN before changing the value again.
- That before enabling the monitor logic, software waits for any previous configuration of the Firewall Component monitor logic to complete. If software enables the Firewall Component monitor logic, whilst a previous configuration of the Firewall Component monitor logic is pending, it is **UNPREDICTABLE** whether transactions which are processed by the Firewall Component use the new or old configuration values.
- Software waits for the disablement process to complete before attempting any other configuration of the Firewall Component's monitor logic. If software attempts to update the configuration of the Firewall Component's monitor logic, before the disablement process completes it is **UNPREDICTABLE** whether the processed transaction responses use the new or old configuration values.

C.9.6.2 Changing error detection report power behavior

The Firewall Component can be configured as to its behavior when the following conditions occur:

- Firewall Component has at least one valid error detection report
- Power Control interface, of the Firewall Component, makes a request to enter a power mode where the Firewall Component would be considered non-operational

Software can configure whether the Firewall Component accepts or denies the power request, under these conditions, using the ME_CTRL.EDR_PWR field. On writing this field the ME_ST.EDR_PWR field is updated automatically to match the value, written to the ME_CTRL.EDR_PWR field. The Firewall Component uses the value in the ME_ST.EDR_PWR to select whether it accepts or denies the power request.

C.9.6.3 Changing error transaction read data response

When the Firewall Component detects a response to a read transaction, which indicates an error has occurred, the Firewall Component can be configured whether to set the read data to all 0s or a value dependent on the StreamID of the transaction.

Software uses the ME_CTRL.RDUM field to configure this behavior. The ME_ST.RDUM field indicates how the Firewall Component treats any response to a read transaction, which indicates an error has occurred. When ME_CTRL.RDUM and ME_ST.RDUM are different, the Firewall Component is changing its error transaction response behavior.

The value of ME_ST.RDUM only updates, to the value of ME_CTRL.RDUM, when any new response to a transaction is treated as the value configured in the ME_CTRL.RDUM field. Any response which has previously been detected as an error behaves as defined by the value of ME_ST.RDUM at the point it was detected. It is **UNPREDICTABLE** whether an error read response, whilst the values of ME_CTRL.RDUM and ME_ST.RDUM are different, uses the old or new value of ME_ST.RDUM.

If software attempts to change the value of ME_CTRL.RDUM, back to its previous value, when it differs from ME_ST.RDUM, the Firewall treats the update as a Configuration Access Error and does not update the ME_CTRL register.

Arm® strongly recommends:

- When software changes the value of either ME_CTRL.RDUM, it waits for the value of ME_ST.RDUM to reflect this change, before changing it again.
- When software changes the values of ME_CTRL.RDUM there are no outstanding transactions to the Firewall Component. The software method used to achieve this is system-dependent. If there are outstanding transactions it is **UNPREDICTABLE** whether a read transaction response, marked with an error, uses the new or old value of the ME_ST.RDUM field.

C.9.7 Monitor Logic Response Forwarding

The following table summarizes response forwarding in monitor logic.

Table C-57: Summary of response

Transaction type	ME_ST.EN	ME_ST.RDUM	Transaction response	Response
Read	0	X	X	Response is forwarded unmodified
	1	0	No error	Response is forwarded unmodified
			Error	Response is forwarded with the read data set to a value specific to the StreamID
	1	1	X	Response is forwarded unmodified
Write	X	X	X	Response is forwarded unmodified

The StreamID specific read data is set by an **IMPLEMENTATION DEFINED** method.

It is **IMPLEMENTATION DEFINED** whether a Firewall Component indicates it has logged the error or not.

This reduces the number of places that the transaction logs. For more information see section C.9.3 Error response ignore on page 420.

C.10 Translation Extension

The Translation Extension allows the Firewall Component to modify properties of the outgoing transaction.

The Translation Extension level 1 or greater can only be implemented when PE.1 or greater is implemented.

Translation is only applied after a transaction has passed the checks and the region, used to perform the checks, was configured as a power of 2 in size (RGN_SIZE.MULnPO2 set to 0). The translation settings for the region, used to perform the checks, are used to generate the output transaction.

C.10.1 Region Properties

The Translation Extension adds the following properties to a region:

- Output transaction translation address
- Output transaction memory attribute property
- Output transaction security property
- Output transaction instruction property
- Output transaction privilege level property

Some of the properties are dependent on:

- The level of extensions implemented
- The design time configuration of Firewall Component
- How software has programmed the Firewall Component

In C.10.1.1 Output transaction translation address on page 431 to C.10.1.5 Output transaction privilege level property on page 435 each property is covered in detail.

C.10.1.1 Output transaction translation address

A region can define an output transaction translation address, which the Firewall Component uses to modify the address of the output transaction.

The following properties are used to calculate the output transaction's address (OTA):

- Translation address (TA)
- Region size (RGN_SIZE)
- Incoming transaction's address (ITA), zero extended to 64 bits
- Firewall Component's Protection Size (PROT_SIZE)

The following formula shows how the OTA is calculated:

OTA<63:0> = ITA<63:PROT_SIZE>:TA<PROT_SIZE-1:RGN_SIZE>:ITA<RGN_SIZE-1:0>

The address translation is only applied when the address translation enable is set to 1, otherwise the output transaction's address is the same as the incoming transaction's address.

The translation address is only programmable when TE.2 is implemented and the RGN_SIZE.MULnPO2 is 0. Otherwise the output transaction's address is the same as the incoming transaction's address.

The table below shows when address translation is applied.

Table C-58: Behavior of address translation

Level of TE	RGN_SIZE.MULnPO2	RGN_TCFG2.ADDR_TRANS_EN	Translation applied
0	0	X	No
1		X	No
2		0	No
		1	Yes
X	1	X	No

C.10.1.2 Output transaction memory attribute property

This section describes how the output transaction memory attribute properties are determined.

A region defines an output transaction memory attribute property which is programmable. The table below shows the outgoing transaction memory attribute value depending on the configuration of the Firewall Component and the programming of the region.

Table C-59: Behavior of memory attribute property translation

Level of TE	FC_CFG1.MA_SPT	RGN_TCFG2.MA_TRANS_EN	RGN_SIZE.MULnPO2	Outgoing transaction memory attribute
0	X	X	0	If the bus protocol supports memory attribute, it remains unchanged from the incoming transaction
>=1	0	X		Bus protocol does not support memory attribute
	1	0		Outgoing transaction has the same memory attribute as the incoming transaction
		1		Outgoing transaction has the memory attributed defined by RGN_TCFG2.MA
X	X	X	1	Region is not a power of 2 so translation is not applied. Incoming memory attributes are used unchanged


Note

The outgoing transaction may have different values for the memory attribute property than defined in the table above.

The memory property is defined as an 8-bit value which defines the inner and outer memory type, cache allocation policy, and whether it is transient or not.

The table below contains the list all legal values and the memory type.

Table C-60: Memory types

7:4	3:0	Device or Normal	Outer Cacheability	Inner Cacheability
0b0000	0b0000	Device-nGnRnE	N/A	N/A
	0b0100	Device-nGnRE	N/A	N/A
	0b1000	Device-nGRE	N/A	N/A
	0b1100	Device-GRE	N/A	N/A
0b00RW, RW not 0b00	0b00RW, RW not 0b00	Normal	Write-Through Transient	Write-Through Transient
	0b0100			Non-cacheable
	0b01RW, RW not 0b00			Write-Back Transient
	0b10RW			Write-Through Non-transient
	0b11RW			Write-Back Non-transient
0b0100	0b00RW, RW not 0b00		Non-cacheable	Write-Through Transient
	0b0100			Non-cacheable
	0b01RW, RW not 0b00			Write-Back Transient
	0b10RW			Write-Through Non-transient
	0b11RW			Write-Back Non-transient
0b01RW, RW not 0b00	0b00RW, RW not 0b00		Write-Back Transient	Write-Through Transient
	0b0100			Non-cacheable
	0b01RW, RW not 0b00			Write-Back Transient
	0b10RW			Write-Through Non-transient
	0b11RW			Write-Back Non-transient
0b10RW	0b00RW, RW not 0b00		Write-Through Non-transient	Write-Through Transient
	0b0100			Non-cacheable
	0b01RW, RW not 0b00			Write-Back Transient
	0b10RW			Write-Through Non-transient
	0b11RW			Write-Back Non-transient
0b11RW	0b00RW, RW not 0b00		Write-Back Non-transient	Write-Through Transient
	0b0100			Non-cacheable
	0b01RW, RW not 0b00			Write-Back Transient
	0b10RW			Write-Through Non-transient
	0b11RW			Write-Back Non-transient

R indicates the read allocation policy. Bit 5 indicates the outer read allocation, whilst bit 1 indicates the inner.

W indicates the write allocation policy. Bit 4 indicates the outer write allocation, whilst bit 0 indicates the inner.

R or W:

- 0 – No allocate
- 1 – Allocate

Any values not listed in the table above are Reserved and behave as follows:

- When bits 7:4 == 0000, bits 1:0 are ignored and treated as 00. For example, a value of `0x0F` is treated as Device-GRE.
- When bits 7:4 != 0000, the memory is treated as Normal with the inner cache policy set to Write-Through Transient Read and Write allocate. The outer cache policy is set by bits 7:4. For example, a value of `0x70` is treated as Normal memory outer Write-Back Transient Read and Write allocate, inner Write-Through Transient Read and Write allocate.


Note

Arm strongly recommends software never sets the memory type to a Reserved value.

C.10.1.3 Output transaction security property

This section describes how the output transaction security property is determined.

A region defines an output transaction security property which is programmable. The table below shows the security of the outgoing transaction depending on the configuration of the Firewall Component and the programming of the region.

Table C-61: Behavior and legality of security property translation

Level of TE	Level of SE	FC_CFG0.SEC_SPT	Incoming transaction security	RGN_TCFG2.NS	RGN_SIZE.MULnPO2	Outgoing transaction security
0	X	X	X	X	0b0	If the bus protocol supports a security property it remains unchanged.
>=1	0	0b0	X	X		
	0	0b1	X	X		Illegal Firewall Component configuration.
	1	0b0	X	X		
		0b1	X	0b00		If the bus protocol supports a security property it remains unchanged.
		0b1	S	0b10		S
				0b11		NS
			NS	X		NS
X	X	X	X	X	0b1	Region is not a power of 2 so translation is not applied. Incoming security is used unchanged.

C.10.1.4 Output transaction instruction property

This section describes how the output transaction instruction property is determined.

A region defines an output transaction instruction property (whether the access is an instruction or data access) which is programmable. The table below shows the instruction property of

the outgoing transaction depending on the configuration of the Firewall Component and the programming of the region.

Table C-62: Behavior of instruction property translation

Level of TE	FC_CFG0.INST_SPT	RGN_TCFG2.INST	RGN_SIZE.MULnPO2	Outgoing transaction instruction property
0	X	X	0	If the bus protocol supports instruction property, it remains unchanged from the incoming transaction
>=1	0	X		Firewall Component does not support the property
	1	0b00		Incoming instruction property is used unchanged
		0b10		Data
		0b11		Instruction
X	X	X	1	Region is not a power of 2 so translation is not applied. Incoming instruction property is used unchanged.


Note

The translation of the instruction property only applies to read transactions. All write transactions are considered to be data accesses.

C.10.1.5 Output transaction privilege level property

This section describes how the output transaction privilege level property is determined.

A region defines an output transaction privilege level property (whether the access is from a privileged or unprivileged master), which is programmable. The table below shows the privilege of the outgoing transaction depending on the configuration of the Firewall Component and the programming of the region.

Table C-63: Behavior of privilege level property translation

Level of TE	FC_CFG0.PRIV_SPT	RGN_TCFG2.PRIV	RGN_SIZE.MULnPO2	Outgoing transaction privilege property
0	X	X	0b0	If the bus protocol supports privilege property, it remains unchanged from the incoming transaction
>=1	0b0	X	-	Firewall Component does not support the property
	0b1	0b00		Incoming privilege property is used unchanged
		0b10		Unprivileged
		0b11		Privileged
X	X	X	0b1	Region is not a power of 2 so translation is not applied. Incoming privilege property is used unchanged.

C.10.1.5.1 Translation enables

When TE.1 is implemented a region implements the following translation enables:

- Memory Attribute
- Instruction
- Privilege
- Shareability
- Security

When TE.2 is implemented a region also implements the address translation enable.

Translation of the properties are only applied when RGN_SIZE.MULnPO2 is 0. Otherwise the outgoing transaction has the same values as the incoming transaction, irrespective of the value of the translation enable.

C.10.2 Property translation

When the Firewall Component supports TE.1 or greater, it can modify the following properties of the outgoing transaction:ion property translation.

- Shareability
- Memory attribute
- Security (for a Secure transaction only)
- Instruction
- Privilege

Software can select, for each property, to use either the incoming value, if the value is not supplied on the incoming Bus Slave interface the default value is used instead, or a software defined value.

It is possible for outgoing transaction to have illegal combination of properties, due to a combination of software programming and incoming values.

As it is possible for software to perform translation on any power of 2 size granule, it is possible for two different masters to access the memory with different attributes. This can cause visibility and coherency issues. Software is responsible for making sure that the memory attribute properties used by the different masters are correct and if any cache maintenance operations are required to make sure that updates made by one master are visible to another.

C.10.3 Address translation

When the Firewall Component supports TE.2, it can modify part of the address of the transaction.


Note

The Translation Extension of the Firewall is not a replacement for a *Memory Management Unit* (MMU).

The Firewall considers the incoming address of the transaction as the virtual address of the transaction, and the output transaction as the physical or intermediate physical address, just like an MMU. The difference between the Firewall and an MMU is the range of address that can be translated across, and the number of translations that can be configured.

An MMU allows any incoming address to map to any outgoing address, and in some cases allows a larger outgoing address range than an incoming address range. The Firewall defines the range of output addresses a transaction can target based on the Protection Size of the Firewall Component. It is only possible to translate transactions within the region of the address space which a region can be defined in.

The range of output addresses is limited by the Protection Size of the Firewall Component. Only the bits that are defined by the Protection Size of the Firewall Component are translated. For more information on how address translation is performed see C.10.1.1 Output transaction translation address on page 431.

C.10.3.1 Address translation and protection extension level

Address translation depends upon the protection extension level#

For a Firewall Component which supports PE.1 the incoming address range for each region is fixed at design time. Therefore, software is only able to define the output address.

For a Firewall Component which supports PE.2 software can define the input and output address for a region.

The following figure shows the difference between a Firewall Component, implementing TE.2 and PE.1 in contrast to implementing TE.2 and PE.2.

Figure C-9: Difference between a Firewall Component, implementing PE.1 or PE.2, for address translation

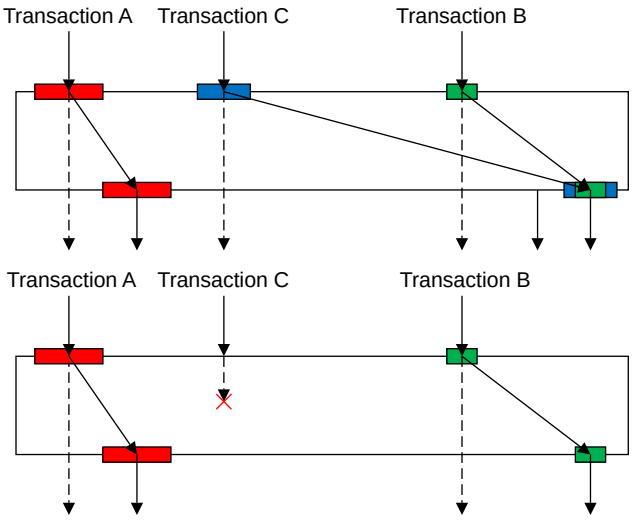

Both Firewall Components have two regions defined. The red region allows Transaction A to be remapped in both cases from the input to output address as the figure shows. The green region allows Transaction B to be remapped in both cases from the input to the output address. When Transaction C arrives, it is required to define a new region. For the Firewall Component, implementing PE.2, software defines the new region, blue, with the input and output address. In this example, Transaction C is mapped to the same location as Transaction B. But for a Firewall Component, implementing PE.1, software cannot define new regions and instead is only able to remap regions which have been predefined. This means that Transaction C always causes a fault, and therefore the transaction must be terminated.

In the following figure, the Protection Size of the Firewall Components is defined to match the address range of the incoming and output bus fabric. However, it is allowed for the Protection Size to be smaller. This can be either at design time or changed using the Protection Size interface. For example, the following figure shows the Protection Size of the Firewall Component is cut in half. This causes the following:

- Transaction B always fails the protection logic checks.
- Transaction C translates to a different address.

In the following figure, the dotted lines show the original Protection Size and translation. The solid lines show the translation when the Protection Size has been halved. For Transaction A the translation is the same in both cases.

Figure C-10: Example of a PE.2 Firewall Component where the Protection Size has been halved

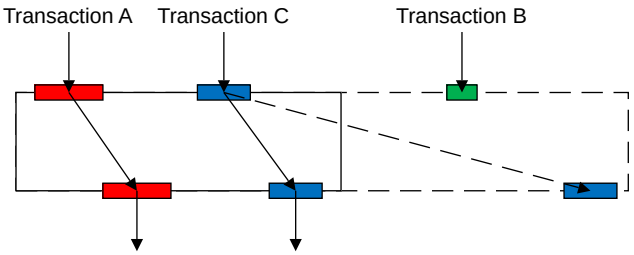

C.10.4 Registers

The Translation Extension adds the RGN_TCFG{0-2} registers.

C.8.7.11 Region Translation Config {0,1} (RGN_TCFG{0,1}) on page 400 to C.8.7.12 Region Translation Config 2 (RGN_TCFG2) on page 401 describe these registers. When TE.1 is implemented RGN_TCFG2 is the only register implemented, with RGN_TCFG{0-1} Reserved.

When TE.2 is implemented RGN_TCFG{0-2} are all implemented.

Note

RGN_TCFG{0-1} are also implemented if RSE.1 is implemented. See C.11 Region Size Extension on page 439 for more information on the RSE.

C.11 Region Size Extension

The Region Size Extension level 1 can only be implemented when PE.1 or greater is also implemented by the Firewall Component.

The Region Size Extension allows regions to have a size which is not equal to a power of 2. This allows regions to be defined to better match the area of memory it is protecting. For example, if there was a buffer in memory 12KB in size. When RSE.0 is implemented the region is defined as 16KB, this means that the memory between 12KB and 16KB is not able to be used for another device. However, when RSE.1 is implemented the region can be defined to be 12KB and no memory is wasted.

Note

Translation is only applied if the RGN_SIZE.MULnPO2 is 0.

C.11.1 Registers

The Region Size Extension adds the following registers and fields:

- RGN_TCFG{0-1} registers, if not already implemented by the Translation Extension
- RGN_SIZE.MULnPO2

For more information on the RGN_TCFG{0-1} register see C.8.7.11 Region Translation Config {0,1} (RGN_TCFG{0,1}) on page 400 and C.8.7.13 Region MasterID {0-3} (RGN_MID{0-3}) on page 403. For more information on the RGN_SIZE.MULnPO2, see C.8.7.10 Region Size (RGN_SIZE) on page 399.

C.12 Security Extension

The Security Extension is implemented at the Firewall level, with all Firewall Components implementing the same features.

A Firewall, which implements SE.0 does not take into consideration the security property of transactions it is processing or monitoring.

Arm® strongly recommends that SE.0 is only implemented in a system that does not support Arm TrustZone.

A Firewall, which implements SE.1 has Firewall Components which:

- Performs checks on the security of the transaction when PE.1, or greater, is implemented.
- Reports the security of a transaction, which fails the checks when PE.1, or greater, is implemented.
- Reports the security of a transaction, which caused an error response when ME.1, or greater, is implemented.
- Translating the security of a transaction, from Secure to Non-secure, when TE.1, or greater, is implemented.

C.13 Lockdown Extension

The Lockdown Extension is implemented at the firewall level, with all Firewall Components within the same firewall implementing the same features. The Lockdown Extension prevents registers of the Firewall Component from being updated. Depending on the level of the extension implemented, different granularities of lockdown can be applied.

All Firewall Components have a lockdown state, as defined in C.13.1 Firewall Component lockdown on page 441. Depending on the level of LDE implemented defines the lockdown states the Firewall Component can enter:

LDE.0
Always in Open lockdown state

LDE.1
Open or Full

LDE.2
Open, Partial or Full

Alongside the Firewall Component lockdown state, when LDE.2 is implemented, each region has a lockdown state: Unlocked or Locked.

When a Firewall Component is in the Partial or Full lockdown state or a region is in the Locked state, a configuration access which attempts to update certain registers behave as follows:

- Configuration Access Error is generated.
- Tamper interrupt and Tamper report is generated, if there is not a valid Tamper report present. Otherwise a Tamper Overflow interrupt is generated.

C.13.1 Firewall Component lockdown

This section describes which registers and fields can be updated in the different lockdown states.

Open

All registers can be updated, except for RWE registers, when the RWE_CTRL points to a region which is in the Locked state.

For more information on region lockdown see C.13.2 Region lockdown on page 442.

Partial

All registers, except for the following, are read-only:

- RWE_CTRL
- FE_CTRL
- EDR_CTRL
- RWE registers, when the RWE_CTRL points to a region which is in the Unlocked state. For more information on region lockdown see C.13.2 Region lockdown on page 442.
- The LD_CTRL register is read-only, when in the Partial lockdown state, and the LD_CTRL.LDI_ST of the Firewall Controller reads as 1.

Any attempt to:

- Update a register which is read-only due to the lockdown state of the Firewall Component
- Update the value of the LD_CTRL register when the Firewall Component is in Partial lockdown state and the LD_CTRL.LDI_ST of the Firewall Controller reads 1

Generates a:

- Configuration Access Error

- Tamper interrupt and Tamper report, if there is no valid Tamper report present. Otherwise, a Tamper Overflow interrupt is generated.


Note

This does not include registers which are Reserved due to the configuration of the Firewall Component or registers which are always read-only.

Full

All registers, except for the following, are read-only:

- RWE_CTRL
- FE_CTRL
- EDR_CTRL

The LD_CTRL register is read-only, when in the Full lockdown state, and the LD_CTRL.LDI_ST of the Firewall Controller reads 1.

Any attempt to:

- Update a register which is read-only due to the lockdown state of the Firewall Component
- Update the value of the LD_CTRL register when the Firewall Component is in Partial lockdown state and the LD_CTRL.LDI_ST of the Firewall Controller reads 1

Generates a:

- Configuration Access Error
- Tamper interrupt and Tamper report, if there is no valid Tamper report present. Otherwise, a Tamper Overflow interrupt is generated.


Note

This does not include registers which are Reserved due to the configuration of the Firewall Component or registers which are always read-only.

C.13.2 Region lockdown

Regions can be fully or partially locked down.

When a region is locked, RGN_LCTRL.LOCK set to 1, and the following registers are read-only:

- RGN_CTRL{0-1}
- RGN_CFG{0-1}
- RGN_SIZE
- RGN_TCFG{0-2}

- RGN_MID{0-3}
- RGN_MPL{0-3}

The RGN_LCTRL.LOCK field can be prevented from being updated by entering the Firewall Component into the Partial or Full lockdown state. When in the Partial lockdown state unlocked regions can still be programmed and locked, but locked regions cannot be programmed or have the RGN_LCTRL.LOCK field updated.

When in the Full lockdown state, no region can be programmed and the RGN_LCTRL.LOCK field cannot be updated, irrespective of the lockdown state of the region.

Any attempt to:

- Update any of the above registers, which are implemented and not Reserved, when the RGN_LCTRL.LOCK is 1

  Note
  This also applies if the RGN_CTRL{0-1}, RGN_MID{0-3} or RGN_MPL{0-3} are read-only due to the configuration of the Firewall Component.

- Update the RGN_LCTRL.LOCK field when RGN_LCTRL.LOCK is 1 and the Firewall Component is in the Partial lockdown state
- Update the RGN_LCTRL.LOCK field when the Firewall Component is in the Full lockdown state

Generates a:

- Configuration Access Error
- Tamper interrupt and Tamper report, if there is no valid Tamper report present. Otherwise, a Tamper Overflow interrupt is generated.

Note
It is possible to lock a region when the region is not enabled.

For information on how software locks and unlocked regions see C.13.5 Changing lockdown state of Firewall Component and regions on page 445.

C.13.3 Lockdown interface

When LDE.1, or greater, is implemented the Firewall Controller implements the Lockdown interface.

C.4.5 Lockdown interfaces on page 360 defines the Lockdown interface.

C.13.4 Registers

This section describes the registers defined for Lockdown Extension.

Table C-64: Summary of Lockdown registers

Offset	Short Name	Access	Name
0x010	LD_CTRL	RW	Lockdown Control

C.13.4.1 Lockdown Control Register (LD_CTRL)

The following table gives a bit-level description of the Lockdown Control (LD_CTRL) register.

The Lockdown Control register, LD_CTRL, allows software to configure the lockdown state of the Firewall Component, see C.13.1 Firewall Component lockdown on page 441 for more information on lockdown states of Firewall Components. The register is only implemented when LDE.1 or greater is supported by the Firewall, otherwise the register is Reserved and generates a Configuration Access Error when accessed.

Table C-65: LD_CTRL register

Bits	Name	Description	Type	Reset
[31:3]	-	Reserved	RO	0x0000_0000
[2]	LDI_ST	Lockdown interface status. Indicates the current value of the Lockdown interface. 0b0: Lockdown interface is de-asserted. 0b1: Lockdown interface is asserted. This field reads as 0b0 for Firewall Components other than the Firewall Controller. The value of this register is dependent on the value of the Lockdown interface of the Firewall.	RO	See description
[1:0]	LOCK	Indicates the lock state of the Firewall Component. 0b00: Open lockdown state 0b01: Reserved and treated as 0b00 0b10: Partial lockdown state. When LDE.1 is implemented this value is Reserved and treated as 0b00. 0b11: Full lockdown state	RW	0x00

The LD_CTRL.LOCK field is not updateable, when the following are all true:

- LD_CTRL.LDI_ST, of the Firewall Controller, is 0b1
- Lockdown state of the Firewall Component is Partial or Full

Any attempt to update the LD_CTRL.LOCK field, under these conditions generates:

- A Configuration Access Error
- A Tamper interrupt and Tamper report, if there is no valid Tamper report present. Otherwise, a Tamper Overflow interrupt is generated.

C.13.5 Changing lockdown state of Firewall Component and regions

This section describes changing the lockdown state of Firewall Component and regions.

Changing lockdown state of Firewall Component

When software wants to change the lockdown state of the Firewall Component it writes to the LD_CTRL.LOCK field.

Arm® recommends that software checks the value of the LD_CTRL.LDI_ST field, in the Firewall Controller, before attempting to update the LD_CTRL.LOCK field. Software must use a system-dependent manner to change the value of Lockdown interface.

Locking and unlocking regions

When LDE.2 is implemented regions of the Firewall Component can be locked using the RGN_LCTRL.LOCK field. Arm recommends software only locks a region after all other region programming is complete.

C.14 Save and Restore Extension

This section describes the details of Save and Restore Extension.

C.14.1 Shadow Registers

When the firewall implements the SRE.1 extension it contains shadow registers. Shadow registers are used to preserve the state of specific registers when the associated Firewall Component enters the Disconnected state.

The registers which are preserved are:

- PE_CTRL
- RWE_CTRL
- ME_CTRL
- LD_CTRL.LOCK
- FW_CTRL
- FC{0-31}_INT_MSK
- For each region independent of whether it is enabled or not:
  - RGN_CTRL{0-1}

- RGN_LCTRL
- RGN_CFG{0-1}
- RGN_SIZE
- RGN_TCFG{0-2}
- RGN_MPL{0-3}
- RGN_MID{0-3}

- Sampled value of all Protection Size interfaces, which are associated with a Firewall Component which implements PE.2

C.14.2 Register Behavior when SRE.0 Implemented

This section describes register behavior when you implement SRE.0.

Register behavior when SRE.0 is implemented is as follows:

- When the component is in the Disconnected state, all accesses to its configuration registers return either:
  - A Configuration Access Error for a Firewall Component other than the Firewall Controller.
  - Enter the Pending state for the Firewall Controller.
- When the component is in the Connecting or Disconnecting state, the access is stalled until the component enters the Connected state.
- When the component is in the Connected state all accesses must complete as normal.

C.14.3 Register Behavior when SRE.1 Implemented

This section describes register behavior when you implement SRE.1.

Register behavior when SRE.1 is implemented as follows:

- Accesses to the registers of the Firewall Controller, which is in the Disconnected state, all enter the Pending state.


Note

When the firewall is integrated into a system it is possible for a deadlock to occur if an access is made to the firewall whilst the Firewall Controller is in the Disconnected state.

- Accesses to the registers of Firewall Components other than the Firewall Controller, which is in the Disconnected state, behave as follows:
  - Read accesses to the following registers or field return the reset value:
    - FE_CTRL.{LAST_FE,FE_VLD}
    - EDR_CTRL.{LAST_EDR,EDR_VLD}
  - Read accesses to the following registers or field return 0:

        - PE_BPS
        - FE_TAL
        - FE_TAU
        - FE_TP
        - FE_MID
        - FE_CTRLACK
        - EDR_TAL
        - EDR_TAU
        - EDR_TP
        - EDR_MID
        - EDR_CTRL.ACK
    - Reads accesses to the following register or fields return a value dependent on another register:
        - PE_ST returns the value in the PE_CTRL register.
        - ME_ST returns the value in the ME_CTRL register.
        - RGN_ST.EN returns the value in RGN_CTRL0.EN field.
        - RGN_ST.MPE_EN{0-3} returns the value in RGN_CTRL1.MPE_EN{0-3} field.
    - Reads accesses to any other registers contained in the shadow registers, return the current value of the register in the shadow register.
    - Read accesses to any FC_CFG2.PROT_SIZE returns the value sampled on the Protection.
    - Read accesses to any registers not contained in the shadow registers, return their reset values.
    - Write accesses to the following registers or fields are ignored, but do not generate a Configuration Access Error:
        - FE_CTRL.ACK.
        - EDR_CTRL.ACK
    - Write accesses to the other registers or fields update the shadow register. The Firewall Controller is responsible for making sure that the most up-to-date value is restored when the Firewall Component enters the Connected state and before it starts processing transactions.
- When the component is in the Connecting state, this only applies to Firewall Components.
    - Reads and write accesses are stalled until the Firewall Component enters the Connected state.
- When the component is in the Connected state read and write accesses are as normal.
- When the component is in the Disconnecting state, this only applies to Firewall Components.
    - Read and write accesses are stalled until the Firewall Component enters the Disconnected state.

The above description is only for registers which are implemented by the Firewall Component, independent of the state of the Firewall Component. If a location which is Reserved due to level of extensions implemented by the Firewall Component, the access always generates a Configuration Access Error.

If a field within a register is Reserved, it always returns its Reserved value when the register is accessed independent of the state of the Firewall Component.

C.14.4 Shadow Register Initialization

The shadow registers might need to be initialized before use.

During the initialization period software can only access the following registers:

- Accesses to the FW_SR_CTRL register.
- Accesses to the FC_CAP{0-3} or FC_CFG{0-3} registers of all Firewall Components.
- Accesses to the Peripheral and Component Identification registers of the Firewall Controller.


Note

FC_CAP{0-3} and FC_CFG{0-3} are read-only registers. A write access to any of these registers remains unaffected by the shadow register initialization. For accesses to all other registers the firewall generates a Configuration Access Error.


Note

Accesses to address locations within the firewall which are Reserved, are unaffected by the shadow register initialization.

Whilst the shadow registers are being initialized, the FW_SR_CTRL.SR_RDY field reads as 0. FW_SR_CTRL.SR_RDY becomes 1 when the initialization is complete.

The FW_SR_CTRL.SR_RDY field remains 1 unless the following occurs:

- Firewall Controller enters the Disconnected state and the shadow registers were not required to retain their values. For example, the Firewall Controller enters a non-operational mode when FW_SR_CTRL.SR_PWR is 0.

Arm® recommends that software polls on the FW_SR_CTRL.SR_RDY field to become 1 before attempting to program the firewall.

C.15 Firewall Controller

The Firewall Controller (FCLTR) is a modified Firewall Component.


Note

The terms Firewall Controller and Firewall Component 0 refer to the same thing in this document.

The FCLTR provides the following features to the firewall:

- Access control to the Firewall Components, within the firewall
- Save and Restore functionality to the Firewall Components, if SRE.1 is implemented
- Lockdown interface, if LDE.1 or greater is implemented
- Interrupt interface
- Tamper Interrupt interface, if LDE.1 or greater is implemented
- Protection Size interfaces, for any Firewall Component which implements PE.2

The main task of the Firewall Controller is to implement Firewall Component 0. Firewall Component 0 provides access control to all Firewall Components in the firewall. It is possible to provide access control to regions which are not part of the firewall. However, care must be taken, in system integration and usage of the firewall, not to affect the protection Firewall Component 0 provides to the firewall.

The Firewall Controller, like other Firewall Components, is allocated 64KB from the Firewall memory map. The Firewall Controller occupies the first 64KB of the memory allocated to the Firewall.

Sections C.15 Firewall Controller on page 448 to C.15.8 Changing Configurations Settings of Firewall on page 471 describe the difference between the Firewall Controller and any other Firewall Component.

C.15.1 Protection Extension

The Firewall Controller always supports PE.1.

The Firewall Controller is implemented as described in C.8 Protection Extension on page 371, with the following differences:

- Reset value of the PE_CTRL.EN and PE_ST.EN is 1.
- It always supports at least three regions.
- Region 0 is enabled by default.

Arm® strongly recommends that software never sets the PE_CTRL.EN to 0 when the PE_ST.BYPASS_MSK is 1 or PE_BPS.BYPASS_ST is 0. If PE_CTRL.EN is set to 0 under either of those conditions, the Firewall Controller becomes disabled and all transaction are considered to

fail the protection logic checks. Therefore, software is unable to perform any other configuration accesses to the Firewall, once the Firewall Controller is disabled. The only way to reverse setting PE_CTRL.EN to 0 once PE_ST.EN is 0, is to enter the Firewall Controller into a Disconnected state before returning it back to the Connected state. The method by which software achieves this is outside the scope of C. Firewall on page 346.

C.15.1.1 Regions and RWE

The following table describes the regions defined for the Firewall Controller.

Table C-66: Region summary for Firewall Component 0

Region number	Description	Components
0	The region is enabled by default and all fields are read-only. The region has a pre-populated MPE which allows a single StreamID to access the Firewall. The agent allocated with this StreamID is considered the Configuration Master.	Firewall Controller
1	Standard region, except for the following fields, of the RWE, are always RAZ/WI: • Secure Privileged Execute (RGN_MPL{0-3}.SPX) • Secure Unprivileged Execute (RGN_MPL{0-3}.SUX) • Non-secure Privileged Execute (RGN_MPL{0-3}.NSPX) • Non-secure Unprivileged Execute (RGN_MPL{0-3}NSUX) • Address Translation Enable (RGN_TCFG2.ADDR_TRANS_EN)	Firewall Controller
2	Standard region, except for the following fields, of the RWE, are always RAZ/WI: • Secure Privileged Execute (RGN_MPL{0-3}.SPX) • Secure Unprivileged Execute (RGN_MPL{0-3}.SUX) • Non-secure Privileged Execute (RGN_MPL{0-3}.NSPX) • Non-secure Unprivileged Execute (RGN_MPL{0-3}NSUX) • Address Translation Enable (RGN_TCFG2.ADDR_TRANS_EN)	All other Firewall Components

The following table shows the default values of fields, and whether the values can be updated for Regions 0-2 for the Firewall Controller. The tables do not show the RGN_ST register. C.8.7.8 Region Status (RGN_ST) on page 396 defines the register. The reset values matching the reset value of the RGN_CTRL{0,1} register.

Table C-67: Region 0 values for Firewall Component 0

Register name	Field name	Value	Update
RGN_CTRL0	EN	0b1	N
RGN_CTRL1	MPE0_EN	0b1	N
	MPE1_EN	0b0	N
	MPE2_EN	0b0	N
	MPE3_EN	0b0	N
RGN_LCTRL	LOCK	0b00	N
RGN_CFG{0,1}	BASE_ADDR	0x0	N
RGN_SIZE	SIZE	**UNKNOWN**	N
	MULnPO2	0b1	N

Register name	Field name	Value	Update
RGN_TCFG{0,1}	OUTPUT_ADDR/UPPER_ADDR	0x1_0000 – 2^(FC_CFG1.MNRS+5)	N
RGN_TCFG2	MA_TRANS_EN	0b0	N
	ADDR_TRANS_EN	0b0	N
	INST	0b00	N
	PRIV	0b00	N
	NS	0b00	N
	SH	0b01	N
	MA	0x00	N
RGN_MID0	MST_ID	Configuration MasterID	N
RGN_MPL0	NSUR	0b0	N
	NSUW		N
	NSUX	0b0	N
	NSPR	0b0	N
	NSPW		N
	NSPX	0b0	N
	SUR	0b0	N
	SUW	**UNKNOWN**	N
	SUX	0b0	N
	SPR	0b1	N
	SPW		N
	SPX	0b0	N
	ANY_MST	0b0	N
RGN_MID{1-3}	MST_ID	**UNKNOWN**	N
RGN_MPL{1-3}	NSUR	**UNKNOWN**	N
	NSUW	**UNKNOWN**	N
	NSUX	0b0	N
	NSPR	**UNKNOWN**	N
	NSPW	**UNKNOWN**	N
	NSPX	0b0	N
	SUR	**UNKNOWN**	N
	SUW	**UNKNOWN**	N
	SUX	0b0	N
	SPR	**UNKNOWN**	N
	SPW	**UNKNOWN**	N
	SPX	0b0	N
	ANY_MST	**UNKNOWN**	N


Note

RGN_MID{1-3} and RGN_MPL{1-3} are shown here for completeness. Depending on the value of FC_CFG1.NUM_MPE, they may not be implemented and may be Reserved.

Table C-68: Region 1 values for Firewall Component 0

Register name	Field name	Value	Update
RGN_CTRL0	EN	0b0	Y
RGN_CTRL1	MPE0_EN	0b0	Y
	MPE1_EN	0b0	Y
	MPE2_EN	0b0	Y
	MPE3_EN	0b0	Y
RGN_LCTRL	LOCK	0b0	Y
RGN_CFG{0,1}	BASE_ADDR	0x0	N
RGN_SIZE	SIZE	**UNKNOWN**	-
MULnPO2	0b1	N	-
RGN_TCFG0	OUTPUT_ADDR/UPPER_ADDR	0x1_0000 – 2^(FC_CFG1.MNRS+5)	N
RGN_TCFG2	MA_TRANS_EN	0b0	N
	ADDR_TRANS_EN	0b0	N
	INST	0b00	N
	PRIV	0b00	N
	NS	0b00	N
	SH	0b01	N
	MA	0x00	N
RGN_MID{0-3}	MST_ID	**UNKNOWN**	Y
RGN_MPL{0-3}	NSUR	**UNKNOWN**	Y
	NSUW	**UNKNOWN**	Y
	NSUX	0b0	N
	NSPR	**UNKNOWN**	Y
	NSPW	**UNKNOWN**	Y
	NSPX	0b0	N
	SUR	**UNKNOWN**	Y
	SUW	**UNKNOWN**	Y
	SUX	0b0	N
	SPR	**UNKNOWN**	Y
	SPW	**UNKNOWN**	Y
	SPX	0b0	N
	ANY_MST	**UNKNOWN**	Y


Note

Only bits log2(MXRS)-1 to log2(MNRS) of the base and upper address are present in the BASE_ADDR, and OUTPUT_ADDR/UPPER_ADDR, fields. All other bits are RAZ.

RGN_MPL{1-3} are shown here for completeness. Depending on the value of FC_CFG1.NUM_MPE, they may not be implemented and may be Reserved.

Table C-69: Region 2 values for Firewall Component 0

Register name	Field name	Value	Update
RGN_CTRL0	EN	0b0	Y
RGN_CTRL1	MPE0_EN	0b0	Y
	MPE1_EN	0b0	Y
	MPE2_EN	0b0	Y
	MPE3_EN	0b0	Y
RGN_LCTRL	LOCK	0b0	Y
RGN_CFG{0,1}	BASE_ADDR	0x1_0000	N
RGN_SIZE	SIZE	**UNKNOWN**	N
	MULnPO2	0b1	N
RGN_TCFG0	OUTPUT_ADDR/UPPER_ADDR	0x1_0000 + (NUM_FC * 0x1_0000) – 2^(FC_CFG1.MNRS+5)	N
RGN_TCFG2	MA_TRANS_EN	0b0	N
	ADDR_TRANS_EN	0b0	N
	INST	0b00	N
	PRIV	0b00	N
	NS	0b00	N
	SH	0b01	N
	MA	0x00	N
RGN_MID{0-3}	MST_ID	**UNKNOWN**	Y
RGN_MPL{0-3}	NSUR	**UNKNOWN**	Y
	NSUW	**UNKNOWN**	Y
	NSUX	0b0	N
	NSPR	**UNKNOWN**	Y
	NSPW	**UNKNOWN**	Y
	NSPX	0b0	N
	SUR	**UNKNOWN**	Y
	SUW	**UNKNOWN**	Y
	SUX	0b0	N
	SPR	**UNKNOWN**	Y
	SPW	**UNKNOWN**	Y
	SPX	0b0	N
	ANY_MST	**UNKNOWN**	Y


Note

Only bits log2(MXRS)-1 to log2(MNRS) of the base and upper address are present in the BASE_ADDR, and OUTPUT_ADDR/UPPER_ADDR, fields. All other bits are RAZ.

RGN_MPL{1-3} are shown here for completeness. Depending on the value of FC_CFG1.NUM_MPE, they may not be implemented and may be Reserved.

Fault entries and FWE

C.8.3 Fault entries on page 383 and 12.3.2 Secure Enclave Registers on page 236 define the fault entries and FWE.

C.15.2 Security Extension

The Security Extension affects the Firewall Controller in the same way as the Firewall Component.

See C.12 Security Extension on page 440 for more information.

C.15.3 Lockdown Extension

The Firewall Controller has the same behavior for the Lockdown Extension as the Firewall component.

For more details, see C.13 Lockdown Extension on page 440.

The registers added to the Firewall Controller behave as follows:

- The FC{0-31}_INT_ST registers can be updated in any lockdown state.
- The FC{0-31}_INT_MSK registers are not updateable when in the Partial or Full lockdown state.
- The FW_TMP_CTRL can be updated in any lockdown state.
- The FW_CTRL register is not updated when in the Partial or Full lockdown state.

C.15.3.3 Registers on page 456 defines the registers added to the Firewall Controller.

C.15.3.1 Tamper Interrupt Interface

The Firewall Controller, when LDE.1 or greater is implemented, implements a Tamper Interrupt interface.

The Tamper Interrupt interface indicates when a Tamper or Tamper Overflow interrupt has occurred as indicated by the FW_TMP_CTRL.{TR_VLD,TR_OVERFLW} fields being 1. Software clears the Tamper Interrupt interface by writing 1 to FW_TMP_CTRL.ACK bit.

C.15.3.2 Tamper report

When LDE.1 or greater is implemented, when a configuration access which meets the following requirements:

- Has passed the protection logic checks of the Firewall Controller.
- Is a supported Configuration Access as defined by the implementation.
- Occurs when FW_SR_CTRL.SR_RDY is 0b1.

Attempts to do one of the following:

- Update an implemented register of a Firewall Component which is in the Partial or Full lockdown state, except for:
  - RWE_CTRL
  - FE_CTRL
  - EDR_CTRL
  - FC{0-31}_INT_ST
  - FW_TMP_CTRL
  - LD_CTRL, when Lockdown interface is not asserted.

  For more information see C.13.1 Firewall Component lockdown on page 441.
- Update an implemented register of a region which is in the locked state. For more information see C.13.2 Region lockdown on page 442.


Note

This includes if the register is a read-only register.

- Update the value of the RGN_LCTRL.LOCK field when it is set to 1 and the Firewall Component is in the Partial lockdown state.
- Update the value of the RGN_LCTRL.LOCK field when the Firewall Component is in the Full lockdown state.
- Change the lockdown state of a Firewall Component, when it is Partial or Full, when LD_CTRL.LDI_ST is 1 in the Firewall Controller.
- Access the Tamper Report registers with an access without the correct properties or from a master with a MasterID other than in region 0.

When one of these events occurs the Firewall Controller:

- Generates a Configuration Access Error to configuration access.
- Generates a Tamper interrupt and Tamper report, if there is not a valid Tamper report present. Otherwise a Tamper Overflow interrupt is generated.

A Tamper report is valid when the FW_TMP_CTRL.TR_VLD is set to 1. Software is required to acknowledge the Tamper report, only when it has collected the information contained in the FW_TMP_{TA,TP,MID,CTRL} registers. Once software has acknowledged the Tamper report, the values in the FW_TMP-{TA,TP,MID,CTRL} registers become 0. Except if another Tamper report is generated in the same cycle.

C.15.3.3 Registers

The table below summaries the registers added to the Firewall Control for the Tamper Report.

Table C-70: Tamper Report register summary

Offset	Short name	Access	Name
0xE90	FW_TMP_TA	RO	Firewall Tamper Transaction Address
0xE94	-	RO	Reserved
0xE98	FW_TMP_TP	RO	Firewall Tamper Transaction Properties
0xE9C	FW_TMP_MID	RO	Firewall Tamper MasterID
0xEA0	FW_TMP_CTRL	RW	Firewall Tamper Control

These registers are only accessible by transactions with the following properties:

- MasterID as defined in region 0
- When SE.1 or greater is implemented, with a security property set to secure
- Privileged

Any access to these registers which does not meet these requirements:

- Generates a Configuration Access Error
- Generates a tamper report:
  - If there is not already a valid tamper report
  - If there is a valid tamper report, which software is acknowledging this cycle
- Generates a tamper report overflow, if there is already a valid tamper report, which software is not acknowledging this cycle.

These registers are only implemented when LDE.1 or greater is implemented, otherwise they are Reserved and generate a Configuration Access Error if accessed.

C.15.3.3.1 Firewall Tamper Transaction Address (FW_TMP_TA)

The following table gives a bit-level description of the Firewall Tamper Transaction Address (FW_TMP_TA) register.

Table C-71: Firewall Tamper Transaction Address register

Bits	Name	Description	Type	Reset
[31:21]	-	Reserved	RO	0x00
[20:2]	TMP_TRANS_ADDR	Address of the accessed register which caused the tamper report to be generated	RO	UNKNOWN

Bits	Name	Description	Type	Reset
[1:0]	-	Reserved	RO	00

C.15.3.3.2 Firewall Tamper Transaction Properties (FW_TMP_TP)

The following table gives a bit-level description of the Firewall Tamper Transaction Properties (FW_TMP_TP) register.

Table C-72: Firewall Tamper Transaction Properties register

Bits	Name	Description	Type	Reset
[31:20]	-	Reserved	RO	0x000
[19:18]	-	Reserved	RO	0x00
[17]	PRIV	Indicates the privileged level of the transaction which generated the tamper report. 0b0: Unprivileged 0b1: Privileged	RO	**UNKNOWN**
[16]	NS	Indicates the security level of the transaction which generated the tamper report. 0b0: Secure 0b1: Non-secure	RO	**UNKNOWN**
[15:0]	-	Reserved	RO	0x0000

C.15.3.3.3 Firewall Tamper MasterID (FW_TMP_MID)

The following table gives a bit-level description of the Firewall Tamper MasterID (FW_TMP_MID) register.

Table C-73: FW_TMP_MID register

Bits	Name	Description	Type	Reset
[31:0]	MST_ID	Indicates the MasterID of the transaction which caused the tamper report.The width of this field depends on the value of FC_CFG2. MST_ID_WIDTH. Any unused bits are Reserved and treated as RAZ/WI.	RO	The reset value of this field depends on the value of FC_CFG2. SINGLE_MST: 0b0: **UNKNOWN** 0b1: Fixed valued defined at design time.

C.15.3.3.4 Firewall Tamper Control (FW_TMP_CTRL)

The following table gives a bit-level description of the Firewall Tamper Control (FW_TMP_CTRL).

Table C-74: FW_TMP_CTRL register

Bits	Name	Description	Type	Reset
[31]	TR_VLD	Indicates whether there is a valid tamper report or not. `0b0`: No valid tamper report `0b1`: Valid tamper report When this field is 0 the values in the following registers read as 0: • FW_TMP_TA • FW_TMP_TP • FW_TMP_MID	RO	0
[30]	TR_OVERFLW	Indicates whether a tamper report overflow has occurred `0b0`: No tamper report overflow has occurred `0b1`: Tamper report overflow has occurred	RO	0
[29:1]	-	Reserved	RO	`0x0000_0000`
[0]	ACK	Acknowledge the Tamper report This field always reads as 0. Writes to this field behave as follows: `0b0`: Ignored `0b1`: Tamper report acknowledged Writes to this register are ignored if FW_TMP_CTRL.TR_VLD is 0.	WO	0

C.15.4 Translation Extension

A Firewall Controller can support any level of the Translation Extension. However, no translation can be enabled for Regions 0 to 2. For these regions, the RGN_TCFG2 register is read-only.

Translation Extension support is the same as for a Firewall Component. See C.10 Translation Extension on page 430.

C.15.5 Region Size Extension

A Firewall Controller always supports RSE.1.

For more information on RSE see C.11 Region Size Extension on page 439.

C.15.6 Interrupts

The Firewall Controller receives Interrupt requests from the Firewall Components, on the Firewall Configuration interfaces.

When the Firewall Controller receives an Interrupt requests, and the interrupt is not masked, it must:

- Generate an interrupt, on the Interrupt interface
- Update the interrupt status registers

Changing the value of the interrupt mask does not affect any interrupts which have already been asserted.

All interrupts in the firewall are considered level sensitive, and require software to acknowledge the interrupt by writing a value of 1 to the corresponding interrupt status field. If the Firewall Controller receives a request from a Firewall Component to generate an interrupt at the same time as software writes 1 to the status bit:

- The interrupt must remain asserted
- The interrupt status field is set to 1

This includes the case where software acknowledges the interrupt in the same cycle as the request is received.

C.15.6.1 Generation

Interrupts are generated by the Firewall Components and reported to the Firewall Controller.

A Firewall Component can generate the following types of interrupts:

- When PE.1 or greater is implemented:
  - Access Error
  - Programming Error
  - Fault Entry Overflow
- When ME.1 or greater is implemented:
  - Error Detection
  - Error Detection Overflow
- Tamper, when LDE.1 or greater is implemented.
- Tamper Overflow, when LDE.1 or greater is implemented.

C.15.6.1.1 Access Error

The Access Error interrupt is generated when:

- A transaction enters the Faulted state and the Firewall Component is configured to generate a fault entry

- If there is at least one invalid fault entry in the Firewall Component

If all fault entries are valid, then a Fault Entry Overflow interrupt is generated instead, see C.15.6.1.3 Fault Entry Overflow on page 460.

C.15.6.1.2 Programming Error

The Programming Error interrupt is generated when a transaction matches more than one region and MPE pair.

The Programming Error is also generated:

- When the Firewall Component is configured to generate a fault entry
- If there is at least one invalid fault entry in the Firewall Component

If all fault entries are valid, then a Fault Entry Overflow interrupt is generated instead, see C.15.6.1.2 Programming Error on page 460.

C.15.6.1.3 Fault Entry Overflow

The Fault Entry Overflow interrupt is generated when the Firewall Component attempts to generate a fault entry, due to:

- The Firewall Component having the conditions to generate an Access Error interrupt, except for all fault entries being valid.
- The Firewall Component having the conditions to generate a Programming Error interrupt, except for all fault entries being valid.

C.15.6.1.4 Error Detection

The Error Detection interrupt is generated when the Firewall Component detects an error response to a transaction, the monitoring logic is enabled, and there is at least one invalid error detection report.

If all error detection reports are valid, then an Error Detection Overflow interrupt is generated instead, see C.15.6.1.4 Error Detection on page 460.

C.15.6.1.5 Error Detection Overflow

The Error Detection Overflow interrupt is generated when the Firewall Component detects an error response to a transaction, the monitoring logic is enabled, and all error detection reports are valid.

C.15.6.1.6 Tamper

The Tamper interruption is generated under certain conditions.

The Tamper interrupt is generated when a configuration access attempts to do any of the following and FW_TMP_CTRL.TR_VLD bit is 0:

- Update any register of a Firewall Component when LDE.1 or greater is implemented and the Firewall Component is in the Partial or Full lockdown state, except for the following registers:
  - RWE_CTRL
  - FE_CTRL
  - EDR_CTRL
  - FC{0-31}_INT_ST
  - FW_TMP_CTRL
  - LD_CTRL, when Lockdown interface is not asserted
- Update a field of a region which is in the Locked state, when LDE.2 and PE.1 or greater, are implemented.
- Update the value of the RGN_LCTRL.LOCK field when it is set to 1 and the Firewall Component is in the Partial lockdown state.
- Update the value of the RGN_LCTRL.LOCK field when the Firewall Component is in the Full lockdown state.
- Change the lockdown state of a Firewall Component, when it is Partial or Full, when LD_CTRL.LDI_ST is 1 in the Firewall Controller.
- Access the Tamper Report registers with an access without the correct properties or from a master with a MasterID other than in region 0.

If FW_TMP_CTRL.TR_VLD bit is 1, then a Tamper Overflow interrupt is generated instead, see section C.15.6.1.6 Tamper on page 460.

The Tamper interrupt is only supported when LDE.1 or greater is implemented.

Tamper Overflow

The Tamper Overflow interrupt is generated when a configuration access attempts to perform an event listed in C.15.6.1.6 Tamper on page 460 and the FW_TMP_CTRL.TR_VLD field is 1.

C.15.7 Registers

This section describes the extra registers defined for Firewall Controller compared with a Firewall Component.

A Firewall Controller includes all the registers present in a Firewall Component. In this section the additional registers are covered.

The Firewall Controller implements registers for:

- Firewall Control and Status of the whole Firewall
- Interrupts
- Identification

The Control and Status registers allow software to configure:

- Behavior of the Firewall Controller, when a Configuration Access Error is generated by itself or another Firewall Component.
- When SRE.1 is implemented, whether the shadow registers are required to retain their values when the Firewall Controller enters the Disconnected state.

The Interrupt registers allow software to configure and receive interrupts generated by the Firewall Components in the Firewall.

The Tamper register, logs information about tamper transactions.

The identification registers provide the following information:

- Architecture version the implementation of the Firewall is compliant to
- Details about the implementor of the Firewall
- Identification registers

The table below summarizes the registers added to the Firewall Controller over a standard Firewall Component.


Note

This excludes the Tamper Report registers which were defined in C.15.3.3 Registers on page 456.

Summary of the Firewall Controller registers

Table C-75: Control and Status registers

Offset	Short name	Access	Name
0x000	FW_CTRL	RW	Firewall Control
0x004	FW_ST	RO	Firewall Status
0x00C	FW_SR_CTRL	RW	Firewall Shadow Register Control

Table C-76: Interrupts registers

Offset	Short name	Access	Name
0xD00 – 0xD7C	FC{0-31}_INT_ST	RW	Firewall Component 0-31 Interrupt Status
0xD90	FC_INT_ST	RO	Firewall Interrupt Status
0xE00 – 0xE7C	FC{0-31}_INT_MSK	RW	Firewall Component 0 -31 Interrupt Mask

Table C-77: Identification registers

Offset	Short name	Access	Name
0xFC8	IIDR	RO	Implementation Identification Register
0xFCC	AIDR	RO	Architecture Identification Register
0xFD0	PID4	RO	Peripheral ID4
0xFD4	PID5	RO	Peripheral ID5
0xFD8	PID6	RO	Peripheral ID6

Offset	Short name	Access	Name
0xFDC	PID7	RO	Peripheral ID7
0xFE0	PID0	RO	Peripheral ID0
0xFE4	PID1	RO	Peripheral ID1
0xFE8	PID2	RO	Peripheral ID2
0xFEC	PID3	RO	Peripheral ID3
0xFF0	CID0	RO	Component ID0
0xFF4	CID1	RO	Component ID1
0xFF8	CID2	RO	Component ID2
0xFFC	CID3	RO	Component ID3

C.15.7.1 Firewall Control (FW_CTRL) register

The following table gives a bit-level description of the Firewall Control (FW_CTRL) register.

Table C-78: FW_CTRL register

Bits	Name	Description	Type	Reset
[31:2]	-	Reserved	RO	0x0000_0000
[1]	RAZ	Configures the value returned for read accesses which generate a Configuration Access Error. 0b0: Read data is based on the StreamID 0b1: Read data is all 0s	RW	0
[0]	ERR	Configures the response for transactions which generate a Configuration Access Error. 0b0: No error 0b1: Error	RW	1

C.15.7.2 Firewall Status (FW_ST) register

The following table gives a bit-level description of the Firewall Status (FW_ST) register.

Table C-79: FW_ST register

Bits	Name	Description	Type	Reset
[31:2]	-	Reserved	RO	0x0000_0000
[1]	RAZ	Indicates the value returned for read accesses which generate a Configuration Access Error. 0b0: Read data is based on the StreamID 0b1: Read data is all 0s	RO	0
[0]	ERR	Indicates the response for a transaction which generate a Configuration Access Error. 0b0: No error 0b1: Error	RO	1


Note

When the Firewall Controller implements ME.1, or greater, and FW_ST.ERR is 1, any access which generates a Configuration Access Error, will also generate an error detection report and associated Error Detection interrupt, if the monitor logic is enabled.

C.15.7.3 Firewall Shadow Register Control (FW_SR_CTRL) register

The following table gives a bit-level description of the Firewall Shadow Register Control (FW_SR_CTRL) register.

This register is only implemented when SRE.1 is implemented. Otherwise, it is Reserved and generates a Configuration Access Error when accessed.

Table C-80: FW_SR_CTRL register

Bits	Name	Description	Type	Reset
[31]	SR_RDY	Indicates when the shadow registers are ready. `0b0`: Shadow registers are not ready. `0b1`: Shadow registers are ready. For more information on this field see C.14.4 Shadow Register Initialization on page 448.This field is Reserved and treated as RAZ/WI when SRE.0 is implemented.	RO	0
[30:1]	-	Reserved	RO	`0x0000_0000`
[0]	SR_PWR	Shadow register power request when Firewall Controller enters the Disconnected state: `0b0`: No request for the shadow registers to retain their values the next time the Firewall Controller enters the Disconnected state. `0b1`: Request for the shadow registers to retain their values the next time the Firewall Controller enters the Disconnected state. This field is Reserved and treated as RAZ/WI when SRE.0 is implemented.	RW	0

C.15.7.4 Firewall Component {0-31} Interrupt Status (FC{0-31}_INT_ST) register

The following table gives a bit-level description of the Firewall Component {0-31} Interrupt Status (FC{0-31}_INT_ST) register.

There is an interrupt status register per Firewall Component implemented in the Firewall. The fields within this register indicate the status of the interrupts generated by that Firewall Component. The register is only implemented when the associated Firewall Component is implemented, otherwise this register is Reserved and generates a Configuration Access Error when accessed.


Note

Software can find out how many Firewall Components are implemented in the Firewall by reading the FC_CFG3.NUM_FC field of the Firewall Controller.

Table C-81: FC{0-31}_INT_ST register

Bits	Name	Description	Type	Reset
[31:5]	-	Reserved	RO	0x000_0000
[4]	ED_OVRFLW_ST	Indicates whether an error detection report overflow has occurred in the Firewall Component. 0b0: No error detection report overflow has occurred 0b1: Error detection report overflow has occurred This field is write 1 to clear, writing 0 to this field has no effect.	RW	0
[3]	ED_ST	Indicates whether the Firewall Component has detected an error response and generated an error detection report. 0b0: No error response has been detected 0b1: Error response has been detected This field is write 1 to clear, writing 0 to this field has no effect.	RW	0
[2]	FLT_OVRFLW_ST	Indicates whether a fault entry overflow has occurred in the Firewall Component. 0b0: No fault entry overflow has occurred 0b1: Fault entry overflow has occurred This field is write 1 to clear, writing 0 to this field has no effect.	RW	0
[1]	PROG_ERR_ST	Indicates whether the Firewall Component has detected a Programming Error. 0b0: No Programming Error has occurred. 0b1: Programming Error has occurred. This field is write 1 to clear, writing 0 to this field has no effect.	RW	0
[0]	ACC_ERR_ST	Indicates whether the Firewall Component has detected an Access Error. 0b0: No Access Error has occurred. 0b1: An Access Error has occurred. This field is write 1 to clear, writing 0 to this field has no effect.	RW	0

For all the fields in the FC{0-31}_INT_ST, if software attempts to clear the interrupt at the same time as the Firewall Controller receives a new interrupt request, the field remains set.

C.15.7.5 Firewall Interrupt Status (FW_INT_ST) register

The following table gives a bit-level description of the Firewall Interrupt Status (FW_INT_ST) register.

The fields within this register indicate the status of the interrupts generated by the connected Firewall Component. When any bit, in the associated FC{0-31}_INT_ST register is asserted, the corresponding bit in the FW_INT_ST register is asserted.

Table C-82: FW_INT_ST register

Bits	Name	Description	Type	Reset
[31:0]	FC_INT_ST{x}	Interrupt status for the associated Firewall Component. There is a bit per Firewall Component which indicates if any bit in the associated FC{0-31}_INT_ST register is 1. Any bits associated with a Firewall Component which is not implemented are Reserved and treated as RAZ/WI: • `0b0`: No interrupt is asserted by Firewall Component. • `0b1`: Interrupt is asserted by Firewall Component. Bit 0 is for Firewall Component 0, while bit 31 is for Firewall Component 31.	RO	`0x0000_0000`


Note

Software can find out how many Firewall Components are implemented in the Firewall by reading the FC_CFG3.NUM_FC field of the Firewall Controller.

C.15.7.6 Firewall Component {0-31} Interrupt Mask (FC{0-31}_INT_MSK) register

The following table gives a bit-level description of the Firewall Component {0-31} Interrupt Mask (FC{0-31}_INT_MSK) register.

There is an interrupt mask register per Firewall Component implemented in the Firewall. The fields within this register, control whether an interrupt is generated (for that type of interrupt) for the respective Firewall Component. The register is only implemented when the associated Firewall Component is implemented, otherwise this register is Reserved and generates a Configuration Access Error when accessed.


Note

Software can find out how many Firewall Components are implemented in the Firewall by reading the FC_CFG3.NUM_FC field of the Firewall Controller.

Table C-83: FC{0-31}_INT_MSK register

Bits	Name	Description	Type	Reset
[31:5]	-	Reserved	RO	`0x000_000`
[4]	ED_OVRFLW_MSK	Selects whether Error Detection Overflow interrupts are masked. 0 – Error Detection Overflow interrupts are reported to the system. 1 – Error Detection Overflow interrupt is not reported to the system. This field is Reserved and treated as RAZ/WI when ME.0 is implemented.	RW	0

Bits	Name	Description	Type	Reset
[3]	ED_MSK	Selects whether Error Detection interrupts are masked. 0 – Error Detection interrupts are reported to the system. 1 – Error Detection interrupt is not report to the system. This field is Reserved and treated as RAZ/WI when ME.0 is implemented.	RW	0
[2]	FLT_OVRFLW_MSK	Selects whether Fault Event Overflow interrupts are masked. 0 – Fault Event Overflow interrupts are reported to the system. 1 – Fault Event Overflow interrupts are not reported to the system. This field is Reserved and treated as RAZ/WI when PE.0 is implemented.	RW	0
[1]	PROG_ERR_MSK	Selects whether Programming Error interrupts are masked. 0 – Programming Error interrupts are reported to the system. 1 – Programming Error interrupts are not reported to the system. This field is Reserved and treated as RAZ/WI when PE.0 is implemented.	RW	0
[0]	ACC_ERR_MSK	Selects whether Access Error interrupts are masked. 0 – Access Error interrupts are reported to the system. 1 – Access Error interrupts are not reported to the system. This field is Reserved and treated as RAZ/WI when PE.0 is implemented.	RW	0

When an interrupt is masked the associated bit in the FC{0-31}_INT_ST register is also not set. However, if the associated bit in the FC{0-31}_INT_ST is already 1, setting the mask bit to 1 does not change the value of the bit in the FC{0-31}_INT_ST register.

C.15.7.7 Implementation Identification (IIDR) register

The following table gives a bit-level description of the Implementation Identification (IIDR) register.

Table C-84: IIDR register

Bits	Name	Description	Type	Reset
[31:20]	PRODUCT_ID	Firewall part ID	RO	`0x075`
[19:16]	VARIANT	Firewall varient	RO	`0x0`
[15:12]	REVISION	Firewall revision	RO	`0x1`
[11:0]	IMPLEMENTER	Contains the JEP106 code of Arm. • [11:8] JEP106 continuation code of implementer • [7] Always 0 • [6:0] JEP106 identity code of implementer	RO	`0xx43B`

C.15.7.8 Architecture Identification (AIDR) register

The following table gives a bit-level description of the Architecture Identification (AIDR) register.

Table C-85: AIDR register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	ARCH_MAJOR_REV	Firewall Architecture Major Revision 0x0 – Major Revision 0 All other values are reserved.	RO	0x0
[3:0]	ARCH_MINOR_REV	Firewall Architecture Minor Revision 0x0 – Minor Revision 0 All other values are reserved.	RO	0x0

C.15.7.9 Peripheral ID 4 (PID4) register

The following table gives a bit-level description of the Peripheral ID 4 (PID4) register.

Table C-86: PID4 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	Size	Number of 4KB occupied by the System ID block. This field is deprecated.	RO	0x0
[3:0]	DES_2	JEP Continuation. For Arm this field is 0x4.	RO	0x4

C.15.7.10 Peripheral ID 5 (PID5) register

The following table gives a bit-level description of the Peripheral ID 5 (PID5) register.

Table C-87: PID5 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

C.15.7.11 Peripheral ID 6 (PID6) register

The following table gives a bit-level description of the Peripheral ID 6 (PID6) register.

Table C-88: PID6 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

C.15.7.12 Peripheral ID 7 (PID7) register

The following table gives a bit-level description of the Peripheral ID 7 (PID7) register.

Table C-89: PID7 register

Bits	Name	Description	Type	Reset
[31:0]	-	Reserved	RO	0x0000_0000

C.15.7.13 Peripheral ID 0 (PID0) register

The following table gives a bit-level description of the Peripheral ID 0 (PID0) register.

Table C-90: PID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PART_0	Bits [7:0] of part ID	RO	0x75

C.15.7.14 Peripheral ID 1 (PID1) register

The following table gives a bit-level description of the Peripheral ID 1 (PID1) register.

Table C-91: PID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	DES_0	Bits [3:0] of JEP 106 Identity	RO	0xB
[3:0]	PART_1	Bits [11:8] of part ID	RO	0x0

C.15.7.15 Peripheral ID 2 (PID2) register

The following table gives a bit-level description of the Peripheral ID 2 register.

Table C-92: PID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVISION	Major revision of the System ID block.	RO	0x0
[3]	JEDEC	Indicates the use of JEDEC JEP106 identification scheme.	RO	0b1
[2:0]	DES_1	Bits [6:4] of JEP 106 Identity. For Arm this value is 0b011.	RO	0b011

C.15.7.16 Peripheral ID 3 (PID3) register

The following table gives a bit-level description of the Peripheral ID 3 (PID3) register.

Table C-93: PID3 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	REVAND	Minor revision of the System ID block	RO	0x1
[3:0]	CMOD	Customer modification field	RO	0x0

C.15.7.17 Component ID 0 (CID0) register

The following table gives a bit-level description of the Component ID 0 (CID0) register.

Table C-94: CID0 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_0	Preamble 0	RO	0x0D

C.15.7.18 Component ID 1 (CID1) register

The following table gives a bit-level description of the Component ID 1 (CID1) register.

Table C-95: CID1 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:4]	CLASS	Class of the component.	RO	0xF
[3:0]	PRMBL_1	Preamble 1	RO	0x0

C.15.7.19 Component ID 2 (CID2) register

The following table gives a bit-level description of the Component ID 2 (CID2) register.

Table C-96: CID2 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_2	Preamble 2	RO	0x05

C.15.7.20 Component ID 3 (CID3) register

The following table gives a bit-level description of the Component ID 3 (CID3) register.

Table C-97: CID3 register

Bits	Name	Description	Type	Reset
[31:8]	-	Reserved	RO	0x00_0000
[7:0]	PRMBL_3	Preamble 3	RO	0xB1

C.15.8 Changing Configurations Settings of Firewall

This section describes how to change certain configuration settings of the Firewall

C.15.8.1 Changing Configuration Access Error Response Type

This section describes the requirements of changing the Configuration of the Access Errors response type.

The Firewall Components generates Configuration Access Errors when illegal transactions are generated. When a Configuration Access Error is generated by a transaction, the Firewall Controller can be configured to generate an error response or not, when responding to the transaction. Software uses the FW_CTRL.ERR field to configure this behavior. The FW_ST.ERR field indicates the current response type the Firewall Controller uses to respond to Configuration Access Errors generated by the Firewall Components. When FW_CTRL.ERR and FW_ST.ERR are different, the firewall is changing its Configuration Access Error response behavior.

The value of FW_ST.ERR only updates to the value of FW_CTRL.ERR, when any new Configuration Access Errors use the new value configured in the FW_CTRL.ERR field. Any transaction which has already generated a Configuration Access, behaves as defined by the value of FW_ST.ERR at the moment that the Firewall Controller is notified about the Configuration Access Error.

If software attempts to change the value of FW_CTRL.ERR back to its previous value, when the value differs from FW_ST.ERR, the firewall treats the update as a Configuration Access Error and does not update the FW_CTRL register.

Arm® strongly recommends:

- After software changes the value of FW_CTRL.ERR, the software waits for the value of FW_ST.ERR to reflect this change, before changing the value again.
- When software changes the values of FW_CTRL.ERR there are no outstanding transactions to the Firewall Controller, other than the configuration access performing the update. The method software uses to achieve this is system-dependent. If there are other outstanding transactions it is **UNPREDICTABLE** whether any Configuration Access Errors use the new or old value of the FW_ST.ERR field.

C.15.8.2 Changing Configuration Access Error Read Data Response

This section describes the requirements of changing Configuration Access Errors response type.

When a read transaction, to the registers of the Firewall, generates a Configuration Access Error, it is configurable whether the read data is set to all 0s or a value dependent on the StreamID of the transaction. Software uses the FW_CTRL.RAZ field to configure this behavior. The FW_ST.RAZ field indicates what the read data for a read transaction, which generates a Configuration Access Error, is set to. When FW_CTRL.RAZ and FW_ST.RAZ are different, the Firewall Controller is changing its Configuration Access Error read response behavior.

The value of FW_ST.RAZ only updates, to the value of FW_CTRL.RAZ, when any new read transactions which become a Configuration Access Error, is treated as the value configured in the FW_CTRL.RAZ field. Any transaction which has already generated a Configuration Access, behaves as defined by the value of FW_ST.RAZ at the point the Firewall Controller is notified about the Configuration Access Error.

If software attempts to change the value of FW_CTRL.RAZ, back to its previous value, when it differs from FW_ST.RAZ, the firewall treats the update as a Configuration Access Error and does not update the FW_CTRL register.

Arm® strongly recommends:

- Once software changes the value of FW_CTRL.RAZ, the software waits for the value of PE_ST.RAZ to reflect this change, before changing the value again.
- When software changes the values of PE_CTRL.RAZ there are no outstanding transactions to the Firewall Controller, other than the configuration access performing the update. The method software uses to achieve this is system-dependent. If there are other outstanding transactions it is **UNPREDICTABLE** whether any Configuration Access Errors use the new or old value of the FW_ST.RAZ field.

C.15.8.3 Changing Shadow Register Power Behavior

This section describes the requirements of changing shadow register power behavior.

When SRE.1 is implemented, the FW_SR_CTRL.SR_PWR field enables software to configure if the Firewall Controller accepts entering a power mode where the shadow registers may lose their contents.

C.16 Software usage

This section covers the software usage model for the firewall.


Note

Arm® strongly advises that if the firewall programming code is updateable, then the software performs checks to validate the code before executing the code. Failure to validate can reduce the effectiveness of the Firewall.

C.16.1 Fault usage model

When a transaction is terminated by the Firewall Component (PE_ST.FLT_CFG is `0b10`), a fault entry may be generated alongside an Access or Programming Error interrupt.

Software should read the fault entry and log the details of the transaction which faulted. Once software has logged the details, software is expected to acknowledge the fault entry. When the fault entry is acknowledged, the fault entry becomes invalid and is available to be used for another faulting transaction.

C.16.2 Error detection report usage

Error detection reports are intended to allow software to discover when a master in the SoC, has performed an operation which a slave in the SoC has generated an error.

An example, of a slave in a SoC which generates errors to transactions is a Default Slave. A Default Slave is a peripheral, which transactions are routed to when the address of the transaction does not map to any other slave in the SoC. Typically, a Default Slave generates an error to complete the transaction and avoid a bus lockup. For a read transaction this includes generating the read data which is typically all 0s. For many processors, a value of all 0s is treated as a NOP operation if the data is executed. This means that it is possible for malicious software agents to use slaves, like the Default Slave, to perform a NOP slide attack and gain access to areas or privileges that they would not normally be allowed.

Using the error detection reports generated by the firewall, a software agent can detect when this is occurring and investigate and prevent the NOP slide attack from gaining the access or privilege.

It is expected, that when software receives an Error Detection interrupt, it reads the contents of the error detection report and either:

- Logs the information for analysis.
- Takes an action, such as disabling the master.
- Performs both.


Note

Software can be implemented to log the information by default. However, after a certain number of error transactions from the same master in a certain amount of time, it disables the master and investigates further.

Independent of what action software takes, Arm® strongly recommends that software acknowledges the error detection report in a timely manner to reduce the chance of getting an Error Detection Overflow interrupt.

C.16.3 Region programming

You must follow a series of steps to program a region.

When programming a region, software must follow these steps:

1. Select the correct region using the RWE_CTRL register.
2. Program the region base address using the RGN_CFG{0,1} registers.
3. Program the region size or upper address using the RGN_SIZE or RGN_TCFG{0,1} registers.

For regions in a Firewall Component implementing PE.1, setting the base address and size are not required, as the region base address, size and upper address are pre-defined.

4. Program and enable the translation properties, if required, using the RGN_TCFG2 register. If the region is required to perform address translation software must use the RGN_TCFG{0,1} to define the address translation.
5. Program the required Master Permission Entries using the RGN_MID{0-3} and RGN_MPL{0-3} registers.
6. Enable the required Master Permission Entries using the RGN_CTRL1.MPE_EN{0-3} bits.
7. Wait for the associated RGN_ST.MPE_EN{0-3} bits to become 1.
8. Enable the region, by setting the RGN_CTRL0.EN to 1.
9. Wait for RGN_ST.EN to become 1.

If software is required to update a region or MPEs settings, a break-before-make method is used.

The break-before-make method means that software sets certain enables to 0 before updating the values. This prevents a transaction being checked against the regions or MPEs previous values.

Depending on the settings of the region which are required to be updated, the sequence to follow for the break before-make-method is different.

C.16.3.1 Update Master Permission Entry

This section descibes the required software sequence of updating a *Master Permission Entry* MPE.

When software is only updating an MPE the sequence is as follows:

1. Select the correct region using the RWE_CTRL register.
2. Disable the MPE, by setting the associated RGN_CTRL1.MPE_EN{0-3} bit to 0. Software may need to unlock the region first.
3. Wait for the associated RGN_ST.MPE_EN{0-3} bit to become 0.

4. Update the values in the RGN_MID{0-3} and RGN_MPL{0-3} register.
5. Enable the MPE, by setting the associated RGN_CTRL1.MPE_EN{0-3} bit to 1.
6. Wait for the associated RGN_ST.MPE_EN{0-3} bit to become 1.

Arm® recommends that before software starts the above sequence. It requests that the master or masters associated with the MasterID in the RGN_MID{0-3} to be updated, stop issuing new transactions and wait for any outstanding transactions to complete. This removes the possibility for a transaction being checked against an incorrect or missing MPE programming.

C.16.3.2 Update region address range or translation

You can use software to update any region setting, other than an MPE.

When software is required to update any region setting, other than an MPE, the sequence is as follows:

1. Select the correct region using the RWE_CTRL register.
2. Disable the region, by setting the RGN_CTRL0.EN bit to 0. Software may need to unlock the region first.
3. Wait for the RGN_ST.EN bit to become 0.
4. Update the required settings in the RGN_SIZE, RGN_CFG{0,1} and RGN_TCFG{0-2} registers, as required.
5. Enable the region, by setting the RGN_CTRL0.EN bit to 1. At this point software can also lock, or re-lock the region.
6. Wait for the RGN_ST.EN bit to become 1.

Arm® recommends that before software starts the above sequence, it requests that all masters which can generate accesses that are checked against the region, stop issuing new transactions and wait for any outstanding transaction to complete. This removes the possibility for a transaction being checked against the incorrect or missing region programming.

C.16.3.3 Master and region restrictions

This section describes software considerations when multiple masters are implemented and multiple regions are defined.

Depending on the configuration and integration of the Firewall in a system, software may be required to:

- Program the regions of the Firewall so that no transaction issued by a master ever matches in more than one region.
- Prevent a master issuing a legal transaction which is not wholly contained within a single region of Firewall Components it passes through.

For example, there are three masters in a system A, B and C. Master A communicates with master B and C using separate 1KB regions which are contiguous in memory. Master A can generate any

size transactions as long as it does not cross a 2KB boundary. Masters B and C can only generate any size transaction as long as it does not cross 1KB boundary. Software must either:

- Define 2 regions:
  - 1 region allowing master A and B access to a 1KB region.
  - 1 region allowing master A and C access to a 1KB region.

  Also prevent master A from ever issuing a transaction which crosses the 1KB boundary.
- Define 3 regions:
  - 1 region allowing master A access to a 2KB region.
  - 1 region allowing master B access to the lower 1KB of that region.
  - 1 region allowing master C access to the upper 1KB of master A region.

C.16.4 Bypass

A Firewall Component which supports PE.1 or greater also supports a mechanism to bypass the protection logic checks, for the purpose of SoC debug.


Note

Arm® strongly recommends that software uses the bypasses feature carefully. The software disables all protection provided by the firewall, and can lead to exposing information to agents who would not normally have access.

C.16.4.1 Bypass mask

The PE_ST.BYPASS_MSK field is intended to mask the Bypass interface value of the Firewall Component.

Software can use this field to:

- Prevent the Firewall Component from becoming bypassed.

  This is typically the case in a production device where software sets the field to 1 during initial configuration of the Firewall Component. Arm® strongly recommends that software only sets PE_ST.BYPASS_MSK during initial configuration with a condition. The condition is that the software can guarantee that there is never be a requirement to perform SoC debug without assistant software that is able to re-configure the firewall.
- During debug, return a Firewall Component which is bypassed, to not being bypassed.

  This is typically the case during SoC debug, where the debug agent wants to re-enable functionality of the firewall. This could be to enable translation of transactions or to test the configuration of the firewall programming.

C.16.4.2 PE_BPS usage

The PE_BPS register provides an indication of whether the Firewall Component is bypassed or not and the current value of the Bypass interface.

The value read from this register must be considered as a snapshot in time, software cannot rely on the value as the only indication of whether a Firewall Component is bypassed or not. The value of PE_BPS must be used in the context of the wider system.

If the firewall supports SRE.1 then the value read when a Firewall Component is not in the Connected state could indicate the Firewall Component is not bypassed. However, after the Firewall Component returns to the Connected state, it could become bypassed. Software must ignore the value of PE_BPS register if PE_BPS.BYPASS_VLD is 0.

Arm® strongly recommends the following:

- When bypassing a Firewall Component the following sequence is followed:
  - Set PE_CTRL.BYPASS_MSK to 0, if not already done.
  - Update the value of the Bypass interface. The method by which this is achieved is system-dependent and outside the scope of this document.
  - Check that PE_BPS.BYPASS_ST reads as 1.
- When returning a Firewall Component to not-be-bypassed, one of the following sequences is performed:
  - Updated the value of the Bypass interface and wait for PE_BPS.BYPASS_ST to read as 0.
  - Set PE_CTRL.BYPASS_MSK to 1 and wait for PE_BPS.BYPASS_ST to read as 0.

  Which method is used depends on the reason why the Firewall Component is returning to not-to-be bypassed and how the firewall is integrated into the system.

C.16.5 Unknown values

Some fields in the Firewall have **UNKNOWN** reset values.

Software must never rely on the **UNKNOWN** reset value and must always set fields to a known value before performing other activities. Failure to do so can lead to **UNPREDICTABLE** results. Software must obey the following rules:

- Before enabling an MPE, it must:
  - Set the RGN_MPL fields with an **UNKNOWN** reset value to a known value.
  - Set either:
    - RGN_MID{0-3} to a known value.
    - RGN_MPL.ANY_MST to 1.
- Before enabling a region, it must:
  - Set the RGN_CFG{0,1} to a known value.
  - Set the RGN_TCFG{0,1} to a known value, if:

- RGN_TCFG2.ADDR_TRANS_EN is set to `0b1`.
- RGN_SIZE.MULnPO2 is set to `0b1`.


Note

When the Firewall implements SRE.1, fields which are part of the shadow registers can have a different value between the copy of the field in the shadow registers and the Firewall Component until the register is set to a value by software.

Read-only fields which have an **UNKNOWN** are not used by the hardware and software must not rely on the value in these fields.

C.17 Programmers model overview

This section summarizes the programmers model.

The programmers model for the Firewall is dependent upon the number and type of Firewall Components which are part of the Firewall. The Firewall occupies 2MB of address space in total. Each Firewall Component is allocated 64KB page, but only uses the first 4KB, with the remaining 60KB being Reserved. The Firewall Components are arranged by the Firewall Component ID, in ascending order. With the Firewall Controller allocated to the first 64KB page. The Firewall always occupies 2MB with any unused 64KB pages being Reserved.

Any memory marked as Reserved generates a Configuration Access Error if accessed. This includes the following conditions:

- Locations which are part of the Firewall memory map but are not allocated to an implemented Firewall Component.
- Locations which are always Reserved, as defined in the tables in sections C.17.2 Firewall Controller register summary on page 479 and C.17.3 Firewall Components register summary on page 482 for the Firewall Controller and Firewall Component respectively.
- Locations which are Reserved due to the level of extension implemented in the Firewall Component. For example, the RWE registers when PE.0 is implemented.
- Locations which are Reserved due to the configuration of the Firewall Component. For example, accessing a register of the RWE when the RWE_CTRL.RGN_INDEX refers to a region which is not implemented.

The above does not include the case where a field within a register is Reserved, due to the configuration of a level of extensions implemented in the Firewall Component. In this case the Reserved field returns its Reserved value.

The table below shows the memory layout of a Firewall.

Table C-98: Summary of Firewall memory map

Offset	Component
0x00_0000	Firewall Controller

Offset	Component
0x01_0000	Firewall Component 1/Reserved
0x02_0000	Firewall Component 2/Reserved
...	Firewall Component N/Reserved
0x1F_0000	Firewall Component 31/Reserved

C.17.1 Configuration access

This section describes supported configuration access and configuration access responses.

C.17.1.1 Supported configuration access

The firewall must support aligned 32-bit word accesses to the registers of the firewall.

Any other accesses generate a Configuration Access Error. Accesses to the firewall registers should be treated as Device-nRnGnE or Device-nRnGE, as defined by *Arm® Architecture Reference Manual Armv8, for Armv8-A architecture profile*, to guarantee that accesses occur in program order and do not cause any unexpected side effects.

C.17.1.2 Configuration access responses

The section describes the Configuration access responses.

The following table shows the response generated by the firewall for configuration access which generate a Configuration Access Error.

Table C-99: Firewall Controller behavior for configuration accesses which generate a Configuration Access Error

Transaction type	FW_ST.ERR	FW_ST.RAZ	Response
Read	0b0	0b0	A non-error read response is generated with the read data set to a value specific to the StreamID
	0b0	0b1	A non-error read response is generated with the read data set to all 0s
	0b1	0b0	An error read response is generated with the read data set to a value specific to the StreamID
	0b1	0b1	An error read response is generated with the read data set to all 0s
Write	0b0	X	A non-error write response is generated
	0b1	X	An error write response is generated

The StreamID specific read data is set by an **IMPLEMENTATION DEFINED** method.

C.17.2 Firewall Controller register summary

This section summaries the Firewall Controller registers.

Table C-100: Firewall Controller registers

Offset	Short name	Lockable	Save and restore
0x000	FW_CTRL	Y	Y
0x004	FW_ST	N	N
0x008	Reserved	-	-
0x00C	FW_SR_CTRL	Y	Y
0x010	LD_CTRL	a	Y
0x014 – 0x0FC	Reserved	-	-
0x100	PE_CTRL	Y	Y
0x104	PE_ST	N	N
0x108	PE_BPS	N	N
0x10C	RWE_CTRL	N	Y
0x110	RGN_CTRL0	b	c
0x114	RGN_CTRL1	b	c
0x118	RGN_LCTRL	b	c
0x11C	RGN_ST	N	N [d]
0x120	RGN_CFG0	b	c
0x124	RGN_CFG1	b	c
0x128	RGN_SIZE	b	c
0x12C	Reserved	-	-
0x130	RGN_SIZE	b	c
0x134	RGN_TCFG1	b	c
0x138	RGN_TCFG2	b	c
0x13C	Reserved	-	-
0x140	RGN_MID0	b	c
0x144	RGN_MPL0	b	c
0x148	RGN_MID1	b	c
0x14C	RGN_MPL1	b	c
0x150	RGN_MID2	b	c
0x154	RGN_MPL2	b	c
0x158	RGN_MID3	b	c
0x15C	RGN_MPL3	b	c
0x160-0x17C	Reserved	-	-
0x180	FE_TAL	N	N
0x184	FE_TAU	N	N
0x188	FE_TP	N	N
0x18C	FE_MID	N	N
0x190	FE_CTRL	N	N

Offset	Short name	Lockable	Save and restore
0x194-0x1FC	Reserved		
0x200	ME_CTRL	Y	Y
0x204	ME_ST	N	N
0x208 – 0x25C	Reserved	-	-
0x260	EDR_TAL	N	N
0x264	EDR_TAU	N	N
0x268	EDR_TP	N	N
0x26C	EDR_MID	N	N
0x270	EDR_CTRL	N	N
0x274 – 0xCFC	Reserved	-	-
0xD00 - 0xD7C	FC{0-31}_INT_ST	N	N
0xD80 – 0xD8C	Reserved	-	-
0xD90	FW_INT_ST	N	N
0xD94 – 0xDFC	Reserved	-	-
0xE00-0xE7C	FC{0-31}_INT_MSK	Y	Y
0xE80-0xE8C	Reserved	-	-
0xE90	FW_TMP_TA	N	N
0xE94	Reserved	-	-
0xE98	FW_TMP_TP	N	N
0xE9C	FW_TMP_MID	N	N
0xEA0	FW_TMP_CTRL	N	N
0xEA4-0xF9C	Reserved		
0xFA0-0xFAC	FC_CAP{0-3}	N	N
0xFB0-0xFBC	FC_CFG{0-3}	N	N d
0xFC0	Reserved		
0xFC4	Reserved		
0xFC8	IIDR	N	N
0xFCC	AIDR	N	N
0xFD0	PID4	N	N
0xFD4	PID5	N	N
0xFD8	PID6	N	N
0xFDC	PID7	N	N
0xFE0	PID0	N	N
0xFE4	PID1	N	N
0xFE8	PID2	N	N
0xFEC	PID3	N	N
0xFF0	CID0	N	N
0xFF4	CID1	N	N
0xFF8	CID2	N	N
0xFFC	CID3	N	N

a

The LD_CTRL register is only lockable when the Lockdown interface is asserted and the LD_CTRL.LOCK field reads as `0b10` or `0b11`.

b

A region is locked under the following conditions:

1. When LDE.1 is implemented and LD_CTRL.LOCK is `0b11`
2. When LDE.2 is implemented and LD_CTRL.LOCK is `0b11`
3. When LDE.2 is implemented and RGN_ST.LOCK, for the region, is `0b1`

c

Only fields which are programmable are required to be saved and restored.

d

The value of FC_CFG2.PROT_SIZE is not required to be saved and restored, however the sampled values of the associated interface are required to.

C.17.3 Firewall Components register summary

This section summarizes Firewall Components registers

Summary of registers in a Firewall Component

Offset	Short name	Lockable	Save and restore
`0x000 – 0x00C`	Reserved	-	-
`0x010`	LD_CTRL	a	Y
`0x014 – 0x0FC`	Reserved	-	-
`0x100`	PE_CTRL	Y	Y
`0x104`	PE_ST	N	N
`0x108`	PE_BPS	N	N
`0x10C`	RWE_CTRL	N	Y
`0x110`	RGN_CTRL0	b	c
`0x114`	RGN_CTRL1	b	c
`0x118`	RGN_LCTRL	b	c
`0x11C`	RGN_ST	N	N
`0x120`	RGN_CFG0	b	c
`0x124`	RGN_CFG1	b	c
`0x128`	RGN_SIZE	b	c
`0x12C`	Reserved	-	-
`0x130`	RGN_TCFG0	b	c
`0x134`	RGN_TCFG1	b	c
`0x138`	RGN_TCFG2	b	c
`0x13C`	Reserved	-	-
`0x140`	RGN_MID0	b	c

Offset	Short name	Lockable	Save and restore
0x144	RGN_MPL0	b	c
0x148	RGN_MID1	b	c
0x14C	RGN_MPL1	b	c
0x150	RGN_MID2	b	c
0x154	RGN_MPL2	b	c
0x158	RGN_MID3	b	c
0x15C	RGN_MPL3	b	c
0x160 –0x17C	Reserved	-	-
0x180	FE_TAL	N	N
0x184	FE_TAU	N	N
0x188	FE_TP	N	N
0x18C	FE_MID	N	N
0x190	FE_CTRL	N	N
0x194 – 0x1FC	Reserved		
0x200	ME_CTRL	Y	Y
0x204	ME_ST	N	N
0x208– 0x25C	Reserved	-	-
0x260	EDR_TAL	N	N
0x264	EDR_TAU	N	N
0x268	EDR_TP	N	N
0x26C	EDR_MID	N	N
0x270	EDR_CTRL	N	N
0x274 – 0xF9C	Reserved	-	-
0xFA0 – 0xFAC	FC_CAP{0-3}	N	N
0xFB0 – 0xFBC	FC_CFG{0-3}	N	Nd
0xFC0 – 0xFFC	Reserved	-	-

a

The LD_CTRL register is only lockable when the Lockdown interface is asserted and the LD_CTRL.LOCK field reads as `0b10` or `0b11`.

b

A region is locked under the following conditions:

- When LDE.1 is implemented and LD_CTRL.LOCK is `0b11`
- When LDE.2 is implemented and LD_CTRL.LOCK is `0b11`
- When LDE.2 is implemented and RGN_ST.LOCK, for the region, is `0b1`

c

Only fields which are programmable are required to be saved and restored.

d

The value of FC_CFG2.PROT_SIZE is not required to be saved and restored, however the sampled values of the associated interface are required to.

Appendix D Revisions

This appendix describes changes between released issues of this book.

D.1 Revisions

This appendix describes changes between released issues of this book.

Table D-1: Issue 0000-01

Change	Location	Affects
First release for EAC.	-	-

Table D-2: Differences between 0000-01 and 0000-02

Change	Location	Affects
First release for REL.	-	No technical changes