ARM® Generic Interrupt Controller Architecture Specification

GIC architecture version 3.0 and version 4.0

ARM®

ARM IHI 0069B (ID121715)

ARM Generic Interrupt Controller Architecture Specification

GIC architecture version 3.0 and version 4.0

Release Information

The following changes have been made to this document.

Change History

Date	Issue	Confidentiality	Change
June 2015	A	Non-confidential	First release of GICv3 and GICv4 issue A
December 2015	B	Non-confidential	First release of GICv3 and GICv4 issue B

Some of the information in this specification was previously published in *ARM® Generic Interrupt Controller, Architecture version 2.0, Architecture Specification*.

Proprietary Notice

LES-PRE-20327

In this document, where the term ARM is used to refer to the company it means “ARM or any of its subsidiaries as appropriate”.

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.

Product Status

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

Web Address

`http://www.arm.com`

Contents
ARM Generic Interrupt Controller Architecture Specification GIC architecture version 3.0 and version 4.0

Preface

This preface introduces the *ARM® Generic Interrupt ControllerArchitecture Specification*. It contains the following sections:

About this specification

This specification describes the *ARM Generic Interrupt Controller* (GIC) architecture. It defines version 3.0 (GICv3) and version 4.0 (GICv4) of the GIC architecture.

Throughout this document, references to *the GIC* or *a GIC* refer to a device that implements this GIC architecture. Unless the context makes it clear that a reference is to an IMPLEMENTATION DEFINED feature of the device, these references describe the requirements of this specification.

Intended audience

This specification is written for users who want to design, implement, or program the GIC in a range of ARM-compliant implementations from simple uniprocessor implementations to complex multiprocessor systems. It does not assume familiarity with previous version of the GIC.

The specification assumes that users have some experience of ARM products, and are familiar with the terminology that describes the ARMv8 architecture. See the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for more information.

Using this specification

This specification is organized into the following chapters:

Chapter 1 ***Introduction***

Read this for an overview of the GIC, and information about the terminology used in this document.

Chapter 2 ***GIC Partitioning***

Read this for an overview of the GIC partitioning and information about the GIC logical components.

Chapter 3 ***Distribution and Routing of Interrupts***

Read this for information about how the GIC uses affinity routing to distribute interrupts.

Chapter 4 ***Physical Interrupt Handling and Prioritization***

Read this for information about how the GIC handles physical interrupts.

Chapter 5 ***Virtual Interrupt Handling and Prioritization***

Read this for information about how the GIC handles virtual interrupts.

Chapter 6 ***Locality-specific Peripheral Interrupts and the ITS***

Read this for a description of *Locality-specific Peripheral Interrupts* (LPIs) and use of the *Interrupt Translation Service* (ITS).

Chapter 7 ***Power Management***

Read this for information about GIC power management.

Chapter 8 ***Programmers' Model***

Read this for a description of the GIC register interfaces, and all GIC registers.

Chapter 9 ***System Error Reporting***

Read this for information about GIC support for error reporting.

Chapter 10 ***Legacy Operation and Asymmetric Configurations***

Read this for information about GIC support for legacy operation and asymmetric configurations.

Appendix A ***GIC Stream Protocol interface***

Read this for a description of the AXI4-Stream protocol standard message-based interface that the GIC Stream Protocol interface uses.

Appendix B ***Pseudocode Definition***

Read this for a definition of the pseudocode that is used in this specification.

Appendix C ***Revisions***

Read this for a description of the technical changes between released issues of this specification.

Glossary Read this for definitions of some of the terms used in this specification.

Conventions

The following sections describe conventions that this book can use:

- *Typographic conventions*.
- *Signals*.
- *Numbers*.
- *Pseudocode descriptions*.

Typographic conventions

The typographical conventions are:

italic Introduces special terminology, and denotes citations.

bold Denotes signal names, and is used for terms in descriptive lists, where appropriate.

`monospace` Used for assembler syntax descriptions, pseudocode, and source code examples.

Also used in the main text for instruction mnemonics and for references to other items appearing in assembler syntax descriptions, pseudocode, and source code examples.

SMALL CAPITALS

Used for a few terms that have specific technical meanings, and are included in the *Glossary*.

Colored text Indicates a link. This can be:

- A URL, for example `http://infocenter.arm.com`.
- A cross-reference, that includes the page number of the referenced information if it is not on the current page, for example, *About the Generic Interrupt Controller (GIC)* on page 1-16.
- A link, to a chapter or appendix, or to a glossary entry, or to the section of the document that defines the colored term, for example, *Banked register* or GICC_CTLR.

Signals

In general this specification does not define processor signals, but it does include some signal examples and recommendations. 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 denotes an active-LOW signal.

Numbers

Numbers are normally written in decimal. Binary numbers are preceded by `0b`, and hexadecimal numbers by `0x`. In both cases, the prefix and the associated value are written in a `monospace` font, for example `0xFFFF0000`.

Pseudocode descriptions

This specification uses a form of pseudocode to provide precise descriptions of the specified functionality. This pseudocode is written in a `monospace` font, and follows the conventions described in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* and the *ARM® Architecture Reference Manual, ARMv7-A and ARMv7-R edition*.

Additional reading

This section lists relevant publications from ARM and third parties.

See the Infocenter, http://infocenter.arm.com for access to ARM documentation.

ARM publications

- *AMBA® 4 AXI4-Stream Protocol Specification* (ARM IHI 0051).
- *ARM® Architecture Reference Manual, ARMv7-A and ARMv7-R edition* (ARM DDI 0406).
- *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* (ARM DDI 0487).
- *ARM® Generic Interrupt Controller, Architecture version 2.0, Architecture Specification* (ARM IHI 0048).
- *ARM® CoreSight™ Architecture Specification v2.0* (ARM IHI 0029).
- *ARM® Debug Interface Architecture Specification ADIv5.0 to ADIv5.2* (ARM IHI 0031).
- *ARM® Server Base System Architecture (SBSA)* (ARM-DEN-0029).
- *ARM® System Memory Management Unit Architecture Specification, SMMU architecture version 2.0* (ARM IHI 0062).
- *GICv3 Software Overview* (DAI 0492)

Other publications

The following books are referred to in this manual, or provide more information:

- JEDEC Solid State Technology Association, *Standard Manufacture's Identification Code*, JEP106.

Feedback

ARM welcomes feedback on its documentation.

Feedback on this manual

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

- The title.
- The number, ARM IHI 0069B.
- The page numbers to which your comments apply.
- A concise explanation of your comments.

ARM also welcomes general suggestions for additions and improvements.

Chapter 1
Introduction

This chapter provides an introduction to the GIC architecture. It provides an overview of the GIC architecture, and of the features that are new to the architecture. It also provides definitions of the terminology that is used throughout this document. It contains the following sections:

1.1 About the Generic Interrupt Controller (GIC)

The GICv3 architecture is designed to operate with ARMv8-A and ARMv8-R compliant *processing elements*, PEs.

The *Generic Interrupt Controller* (GIC) architecture defines:

- The architectural requirements for handling all interrupt sources for any PE connected to a GIC.
- A common interrupt controller programming interface applicable to uniprocessor or multiprocessor systems.

The GIC is an architected resource that supports and controls interrupts. It provides:

- Registers for managing interrupt sources, interrupt behavior, and the routing of interrupts to one or more PEs.
- Support for:
  - The ARMv8 architecture.
  - Locality-specific Peripheral Interrupts (LPIs).
  - Private Peripheral Interrupts (PPIs).
  - Software Generated Interrupts (SGIs).
  - Shared Peripheral Interrupts (SPIs).
  - Interrupt masking and prioritization.
  - Uniprocessor and multiprocessor systems.
  - Wakeup events in power management environments.

For each PE, the GIC architecture describes how IRQ and FIQ interrupts can be generated from different types of interrupts within the system. The ARMv8-A Exception model then describes how the PE handles these IRQ and FIQ interrupts.

Interrupt handling also depends on other aspects of the ARMv8 architecture, such as the Security state, and, for Non-secure interrupts, support for virtualization. The ARM architecture provides two Security states, each with an associated physical memory address space:

- Secure state.
- Non-secure state.

The GIC architecture supports the routing and handling of interrupts that are associated with both Security states. See *Interrupt grouping and security on page 4-59* for more information.

The GIC architecture supports the ARMv8-A model for handling virtual interrupts that are associated with a *virtual machine*, VM. ARMv8-A supports virtualization in Non-secure state only. A virtualized system has:

- A hypervisor that must include a component executing at EL2, which is responsible for switching between VMs.
- Several VMs executing at Non-secure EL1.
- Applications executing at Non-secure EL0 on a VM.

For more information about the ARMv8 architecture, see *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*. For more information about VMs, see *About GIC support for virtualization on page 5-78*.

This specification defines version 3.0 (GICv3) and version 4.0 (GICv4) of the GIC architecture. Version 2.0 (GICv2) is only described in terms of the GICv3 optional support for legacy operation, see *GICv3 with legacy operation on page 1-26*. For detailed information about the GICv2 architecture, see the *ARM® Generic Interrupt Controller, Architecture version 2.0, Architecture Specification*.

Note

Because GICv4 is an extension of GICv3, all references to GICv3 in this manual apply equally to GICv4, unless explicitly indicated otherwise.

1.1.1 Changes to the GIC architecture from GICv2

GIC scalability

The GICv2 architecture only supports a maximum of eight PEs, and so has features that do not scale to a large system. GICv3 addresses this by changing the mechanism by which interrupts are routed, called *affinity routing*, and by introducing a new component to the interrupt distribution, called a *Redistributor*. See Chapter 2 *GIC Partitioning* for more information.

Affinity routing for a Security state is enabled by setting GICD_CTLR.ARE_S or GICD_CTLR.ARE_NS to 1.

Interrupt grouping

Interrupt grouping is the mechanism that is used by GICv3 to align interrupt handling with the ARMv8 Exception model:

- Group 0 physical interrupts are expected to be handled at the highest implemented Exception level.
- Secure Group 1 physical interrupts are expected to be handled at Secure EL1.
- Non-secure Group 1 physical interrupts are excepted to be handled at Non-secure EL2 in systems using virtualization, or at Non-secure EL1 in systems not using virtualization.

These interrupt groups can be mapped onto the ARMv8 FIQ and IRQ signals as described in *Interrupt grouping on page 4-58*, using configuration bits from the ARMv8 architecture and configuration bits within the GICv3 architecture.

In GICv3, interrupt grouping supports:

- Configuring each interrupt as Group 0, Secure Group 1, or Non-secure Group 1.
- Signaling Group 0 physical interrupts to the target PE using the FIQ exception request.
- Signaling Group 1 physical interrupts to the target PE in a manner that allows them to be handled using the IRQ handler in their own Security state. The exact handling of Group 1 interrupts depends on the current Exception level and Security state, as described in Chapter 4 *Physical Interrupt Handling and Prioritization*.
- A unified scheme for handling the priority of Group 0 and Group 1 interrupts.

Interrupt Translation Service (ITS)

The Interrupt Translation Service, ITS, provides functionality that allows software to control how interrupts that are forwarded to the ITS are translated into:

- Physical interrupts, in GICv3 and GICv4.
- Virtual interrupts, in GICv4 only.

The ITS also allows software to determine the target Redistributor for a translated interrupt. Software can control the ITS through a command interface and associated table-based structures in memory. The outputs of the Interrupt Translation Service (ITS) are always LPIs, which are a form of message-based interrupt. See *The ITS on page 6-99*.

Locality-specific Peripheral Interrupts (LPIs)

LPIs are a new class of interrupt that significantly extends the interrupt ID space that the GIC can handle. LPIs are optional, and, if implemented, can be generated and supported by an Interrupt Translation Service, ITS. See *LPIs on page 6-92*.

Software Generated Interrupts (SGIs)

With the ability of GICv3 to support large-scale systems, the context of an SGI is modified and no longer includes the identity of the source PE. See *Software Generated Interrupts on page 4-55*.

Note

The original SGI format is only available in GIC implementations that support legacy operation.

Shared Peripheral Interrupts (SPIs)

A new set of registers in the Distributor are added to support the setting and clearing of message-based SPIs. See *Shared Peripheral Interrupts on page 4-56*.

System register interface

This interface uses System register instructions in an ARMv8-A or ARMv8-R PE to provide a closely-coupled interface for the CPU interface registers. This interface is used for registers that are associated directly with interrupt handling and priority masking to minimize access latency. For virtualization, the registers that are accessed in this manner include both the registers that are accessed by a VM interrupt handler, and the registers that forward virtual interrupts from a hypervisor to a VM. All other registers are memory-mapped.

For AArch64 state, access to the System register interface is enabled by the following settings:

- ICC_SRE_EL1.SRE == 1.
- ICC_SRE_EL2.SRE == 1.
- ICC_SRE_EL3.SRE == 1.

For AArch32 state, access to the System register interface is enabled by the following settings:

- ICC_SRE.SRE == 1.
- ICC_HSRE.SRE == 1.
- ICC_MSRE.SRE == 1.

Other behavior, which is backwards compatible with GICv2, is described in Chapter 10 *Legacy Operation and Asymmetric Configurations*.

Note

In a GIC that supports legacy operation, memory-mapped access is available for all architected GIC registers.

Unless indicated otherwise, this manual describes the GICv3 architecture in a system with affinity routing, System register access, and two Security states, enabled. This means that:

- GICD_CTLR.ARE_NS == 1.
- GICD_CTLR.ARE_S == 1.
- GICD_CTLR.DS == 0.

For operation in AArch64 state:

- ICC_SRE_EL1.SRE == 1, for both the Secure and the Non-secure copy of this register.
- ICC_SRE_EL2.SRE == 1.
- ICC_SRE_EL3.SRE == 1.

For operation in AArch32 state:

- ICC_SRE.SRE == 1.
- ICC_HSRE.SRE == 1.
- ICC_MSRE.SRE == 1.

From GICv3 onwards, legacy operation with the ARE and SRE control bits set to 0 is deprecated. See Chapter 10 *Legacy Operation and Asymmetric Configurations* for more information about legacy operation.

Changes specific to GICv4

GICv4 adds support for the direct injection of virtual interrupts to a VM, without involving the hypervisor. Direct injections are only supported by systems that implement at least one ITS that translates interrupts into LPIs.

1.2 Terminology

The architecture descriptions in this manual use the same terminology that is used for the ARMv8 architecture. For more information about this terminology, see the introduction to Part A of the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

In addition, the AArch64 System register names are used where appropriate, in preference to listing both the AArch32 and AArch64 System register names. The EL*x* suffix on the AArch64 register name indicates the lowest Exception level at which the register can be accessed. The individual AArch64 System register descriptions contain a reference to the AArch32 System register that provides the same functionality.

The following sections define the architectural terms used in this manual:

- *Interrupt types*.
- *Interrupt states on page 1-20*.
- *Models for handling interrupts on page 1-20*.
- *Additional terms on page 1-21*.

1.2.1 Interrupt types

A device that implements the GIC architecture can control *peripheral interrupts*. Peripheral interrupts are typically asserted by a physical signal to the GIC. The GIC architecture defines the following types of peripheral interrupt:

Locality-specific Peripheral Interrupt (LPI)

An LPI is a targeted peripheral interrupt that is routed to a specific PE within the affinity hierarchy:

- LPIs are always Non-secure Group 1 interrupts, in a system where two Security states are enabled.
- LPIs have edge-triggered behavior.
- LPIs can be routed using an ITS.
- LPIs do not have an active state, and therefore do not require explicit deactivation.
- LPIs are always message-based interrupts.

See *LPIs on page 6-92* for more information.

Private Peripheral Interrupt (PPI)

This is a peripheral interrupt that targets a single, specific PE, and different PEs can use the same interrupt number to indicate different events:

- PPIs can be Group 0 interrupts, Secure Group 1 interrupts, or Non-secure Group 1 interrupts.
- PPIs can support either edge-triggered or level-sensitive behavior.
- PPIs are never routed using an ITS.
- PPIs have an active state and therefore require explicit deactivation.

Note

Commonly, it is expected that PPIs are used by different instances of the same interrupt source on each PE, thereby allowing a common interrupt number to be used for PE specific events, such as the interrupts from a private timer.

Shared Peripheral Interrupt (SPI)

This is a peripheral interrupt that the Distributor can route to a specified PE that can handle the interrupt, or to a PE that is one of a group of PEs in the system that has been configured to accept this type of interrupt:

- SPIs can be Group 0 interrupts, Secure Group 1 interrupts, or Non-secure Group 1 interrupts.
- SPIs can support either edge-triggered or level-sensitive behavior.
- SPIs are never routed using an ITS.
- SPIs have an active state and therefore require explicit deactivation.

See *Shared Peripheral Interrupts on page 4-56* for more information. For more information about the Distributor, see Chapter 2 *GIC Partitioning*.

Software Generated Interrupt (SGI)

SGIs are typically used for inter-processor communication, and are generated by a write to an SGI register in the GIC:

- SGIs can be Group 0 interrupts, Secure Group 1 interrupts, or Non-secure Group 1 interrupts.
- SGIs have edge-triggered behavior.
- SGIs are never routed using an ITS.
- SGIs have an active state and therefore require explicit deactivation.

See *Software Generated Interrupts on page 4-55* for more information.

An interrupt that is edge-triggered has the following property:

- It is asserted on detection of a rising edge of an interrupt signal and then, regardless of the state of the signal, remains asserted until the interrupt is acknowledged by software.

For information about edge-triggered message-based interrupts, see Message-based interrupt.

An interrupt that is level-sensitive has the following properties:

- It is asserted whenever the interrupt signal level is active, and deasserted whenever the level is not active.
- It is explicitly deasserted by software.

1.2.2 Interrupt states

The following states apply at each interface between the GIC and a connected PE:

Inactive An interrupt that is not active or pending.

Pending An interrupt that is recognized as asserted in hardware, or generated by software, and is waiting to be handled by the target PE.

Active An interrupt that has been acknowledged by a PE and is being handled, so that another assertion of the same interrupt is not presented as an interrupt to a PE, until the initial interrupt is no longer active.

LPIs do not have an active state, and transition to the inactive state on being acknowledged by a PE.

Active and pending An interrupt that is active from one assertion of the interrupt, and is pending from a subsequent assertion.

LPIs do not have an active and pending state, and transition to the inactive state on being acknowledged by a PE.

The GIC maintains state for each supported interrupt. The state machine defines the possible transitions between interrupt states, and, for each interrupt type, the conditions that cause a transition. See *Interrupt handling state machine on page 4-50* for more information.

1.2.3 Models for handling interrupts

In a multiprocessor implementation, the following models exist for handling interrupts:

Targeted distribution model

This model applies to all PPIs and to all LPIs. It also applies to:

- SPIs during non-legacy operation, if GICD_IROUTER<n>.Interrupt_Routing_Mode == 0.
- During legacy operation, when GICD_CTLR.ARE_* == 0, if only one bit in the appropriate GICD_ITARGETSR<n> field == 1.

A target PE that has been specified by software receives the interrupt.

Targeted list model

This model applies to SGIs only. Multiple PEs receive the interrupt independently. When a PE acknowledges the interrupt, the interrupt pending state is cleared only for that PE. The interrupt remains pending for each PE independently until it has been acknowledged by the PE.

1 of N model

This model applies to SPIs only. The interrupt is targeted at a specified set of PEs, and is taken on only one PE in that set. The PE that takes the interrupt is selected in an IMPLEMENTATION DEFINED manner. The architecture applies restrictions on which PEs can be selected, see *Enabling the distribution of interrupts on page 4-63*.

Note

- The ARM GIC architecture guarantees that a 1 of N interrupt is presented to only one PE listed in the target PE set.
- A 1 of N interrupt might be presented to a PE where the interrupt is not the highest priority interrupt, or where the interrupt is masked by ICC_PMR_EL1 or within the PE. See *Interrupt lifecycle on page 4-46*.

For SPIs during legacy operation, this model applies when more than one target PE is specified in the target registers.

The hardware implements a mechanism to determine which PE activates the interrupt, if more than one PE can handle the interrupt.

1.2.4 Additional terms

The following additional terms are used throughout this manual:

Idle priority

In GICv3, the idle priority, `0xFF`, is the running priority read from ICC_RPR_EL1 on the CPU interface when no interrupts are active on that interface. During legacy operation, the idle priority, as read from GICC_RPR, is IMPLEMENTATION DEFINED, as in GICv2.

Interrupt Identifier (INTID)

The number space that uniquely identifies an interrupt with an associated event and its source. The interrupt is then routed to one or more PEs for handling. PPI and SGI interrupt numbers are local to each PE. SPIs and LPIs have global interrupt numbers for the physical domain. See *INTIDs on page 3-39* for more information.

Interrupt Routing Infrastructure (IRI)

The Distributor, Redistributors and, optionally, one or more ITSs. See *The GIC logical components on page 2-30* for more information.

Message-based interrupt

A message-based interrupt is an interrupt that is asserted because of a memory write access to an assigned address. Physical interrupts can be converted to message-based interrupts. Message-based interrupts can support either level-sensitive or edge-triggered behavior, although LPIs are always edge-triggered.

GICv3 supports two mechanisms for message-based interrupts:

- A mechanism for communicating an SPI, where the assigned address is held in the Distributor. In this case the message-based interrupt can be either level-sensitive or edge-triggered.
- A mechanism for communicating an LPI, where the assigned address is held in an ITS, if an ITS is implemented, or in the Redistributor.

ARM recommends the use of LPIs to provide support for MSI and MSI-X capabilities in systems that support PCIe. See Chapter 6 *Locality-specific Peripheral Interrupts and the ITS* for more information. GICv3 also includes architected support for signaling SPIs using message-based interrupts, see *Shared Peripheral Interrupts on page 4-56*.

Physical interrupt

An interrupt that targets a physical PE is a physical interrupt. It is signaled to the PE by the physical CPU interface to which the PE is connected.

Running priority

At any given time, the running priority of a CPU interface is either:

- The group priority of the active interrupt, for which there has not been a priority drop on that interface.
- If there is no active interrupt for which there has not been a priority drop on the interface, the running priority is the idle priority `0xFF`.

Sufficient priority

The GIC CPU interface compares the priority of an enabled, pending interrupt with all of the following, to determine whether the interrupt has sufficient priority:

- The Priority Mask Register, ICC_PMR_EL1.
- The preemption settings for the interface, as indicated by ICC_BPR0_EL1 and ICC_BPR1_EL1.
- The current running priority, as indicated by ICC_RPR_EL1 for the CPU interface.

If the interrupt has sufficient priority it is signaled to the connected PE.

Virtual interrupt

An interrupt that targets a VM is a virtual interrupt. It is signaled by the associated virtual CPU interface. See Chapter 5 *Virtual Interrupt Handling and Prioritization* for more information.

Maintenance interrupt

A physical interrupt that signals key events associated with interrupt handling on a VM to allow the hypervisor to track those events. These events are processed by the hypervisor, and include enabling and disabling a particular group of interrupts. See *Maintenance interrupts on page 5-85* for more information.

1.3 Supported configurations and compatibility

In ARMv8-A, EL2 and EL3 are optional, and a PE can support one, both, or neither of these Exception levels. However:

- A PE requires EL3 to support both Secure and Non-secure state.
- A PE requires EL2 to support virtualization.
- If EL3 is not implemented, there is only a single Security state. This Security state is either Secure state or Non-secure state.

GICv3 supports interrupt handling for all of these configurations, and for execution in both AArch32 state and AArch64 state, in accordance with the *interprocessing* rules described in *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

1.3.1 Affinity routing configuration

The GICv3 architecture supports affinity routing. It provides optional support for:

- An asymmetric configuration, where affinity routing is enabled for Non-secure state and disabled for Secure state. This provides support for a Secure legacy environment.
- A legacy-only environment where affinity routing is disabled for both Secure state and Non-secure state.

1.3.2 System register configuration

When affinity routing is enabled for execution in both Security states, the GIC must be configured to use System register access to handle physical interrupts. The architecture does not support having affinity routing enabled for a Security state, and not having System register access configured for that Security state. Configuring the GIC this way results in UNPREDICTABLE behavior. When affinity routing is enabled for execution in Non-secure state, the GIC architecture optionally supports legacy operation for virtual interrupts, that is legacy interrupt handling at Non-secure EL1 under the control of a hypervisor executing at EL2.

1.3.3 GIC control and configuration

Many of the GIC registers are available in different forms, to permit effective interrupt handling:

- For two Security states.
- For different interrupt groups.
- Using System register access for GICv3 or memory-mapped access for legacy operation.

When System register access is enabled, control and configuration of the GIC architecture is handled by architected System registers and the associated accesses that define the GIC programmers' model. See Chapter 8 *Programmers' Model* for more information.

Some registers are always memory-mapped, while others use System register access in GICv3, and memory-mapped access for legacy operations.

Table 1-1 shows the registers that are always memory-mapped.

Table 1-1 Memory-mapped registers

Prefix in short register name	Registers
GICD	Distributor registers
GICR	Redistributor registers[a]
GITS	ITS registers

a. There is one copy of each of the Redistributor registers per PE.

Table 1-2 shows the registers that are memory-mapped for legacy operations, but are replaced by System register access in GICv3 when System register access is enabled.

Table 1-2 Memory-mapped registers for legacy operation

Prefix in short register name	Registers
GICC	Physical CPU interface registers
GICV	Virtual CPU interface registers
GICH	Virtual interface control registers

Note

- An operating system executing at Non-secure EL1 uses either the GICC_* or the GICV_* registers to control interrupts, and is unaware of the difference.
- The GICR_* and GITS_* registers are introduced in GICv3.

Table 1-3 shows the registers that GICv3 supports when System register access is enabled.

Table 1-3 System registers

Prefix in short register name	System registers accessed
ICC	Physical CPU interface registers
ICV	Virtual CPU interface registers
ICH	Virtual interface control registers

The ARMv8 support for virtualization and the Exception level at which a PE is operating determine whether the physical CPU interface registers or the virtual CPU interface registers are accessed.

For more information about register names and the factors that affect which register to use, see *GIC System register access on page 8-160*.

1.3.4 References to the ARMv8 architectural state

Table 1-4 shows the ARMv8 architectural state that is used with or affects the operation of the GIC.

Table 1-4 ARMv8 architectural state affecting GIC operation

AArch64		AArch32		Purpose
State	Field	State	Field	
PSTATE[a]	A	PSTATE[a]	A	SError interrupt mask bit (AArch64 state) Asynchronous Abort mask bit (AArch32 state)
	I		I	IRQ mask bit
	F		F	FIQ mask bit
-	-	DFSR	STATUS/FS	Fault status
	-		ExT	External abort type

Table 1-4 ARMv8 architectural state affecting GIC operation (continued)

AArch64 State	Field	AArch32 State	Field	Purpose
ESR_ELx	EC	HSR	EC	Exception class
	IL		IL	Instruction length for synchronous exceptions
	ISS		ISS	Instruction Specific Syndrome
HCR_EL2	AMO	HCR	AMO	SError interrupt routing (AArch64 state) Asynchronous External Abort interrupt routing (AArch32 state)
	IMO		IMO	Physical IRQ routing
	FMO		FMO	Physical FIQ routing
	RW		RES0	Execution state control for lower Exception levels (AArch64 state)
	VSE		VA	Virtual SError Abort exception (AArch64 state) Virtual Asynchronous Abort exception (AArch32 state)
	VI		VI	Virtual IRQ interrupt
	VF		VF	Virtual FIQ interrupt
HSTR_EL2	T<n>	HSTR	T<n>	Hypervisor system traps
	I		I	IRQ pending
	F		F	FIQ pending
ID_AA64PFR0_EL1	GIC	-	-	System register GIC interface support
ID_PFR1_EL1	GIC	ID_PFR1	GIC	System register GIC CPU interface support
ISR_EL1	A	ISR	A	SError pending (AArch64 state) External Abort pending (AArch32 state)
MPIDR_EL1	Aff3	MPIDR	-	Affinity level 3
	Aff2		Aff2	Affinity level 2
	Aff1		Aff1	Affinity level 1
	Aff0		Aff0	Affinity level 0
SCR_EL3	RW	SCR	RES0	Execution state control for lower Exception levels (AArch64 state only)
	EA		EA	SError interrupt routing (AArch64 state) External Abort interrupt routing (AArch32 state)
	FIQ		FIQ	Physical FIQ routing
	IRQ		IRQ	Physical IRQ routing
	NS		NS	Non-secure bit

a. Process state, PSTATE, is an abstraction of the process state information. For more information, see *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

For more information about these registers and fields, see *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

1.3.5 GICv3 with no legacy operation

In an implementation that does not support legacy operation, affinity routing and System register access are permanently enabled. This means that the associated control bits are RAO/WI. Table 1-5 shows the register fields that are affected by this.

Table 1-5 Control bits for affinity routing and System register access

AArch64 registers	AArch32 registers	Memory-mapped registers
ICC_SRE_EL1.SRE[a]	ICC_SRE.SRE[a]	-
ICC_SRE_EL2.SRE	ICC_HSRE.SRE	-
ICC_SRE_EL3.SRE	ICC_MSRE.SRE	-
-	-	GICD_CTLR.ARE_S
-	-	GICD_CTLR.ARE_NS

a. There is a Secure copy and a Non-secure copy of this register.

1.3.6 GICv3 with legacy operation

Legacy operation is a form of limited backwards compatibility with GICv2 that is provided to allow systems using GICv3 to run code using GICv2, provided that this code meets the restrictions described in this section. Legacy operation is optional in GICv3. See *Legacy support of interrupts and asymmetric configurations on page 10-684*.

In a GICv3 implementation that supports legacy operation, a maximum of eight PEs, whose individual support for a memory-mapped register interface is IMPLEMENTATION DEFINED, are available as physical or virtual interrupt targets within a given VM. It is IMPLEMENTATION DEFINED:

- Whether legacy operation applies to execution in both Security states, or to execution in Secure state only.
- Whether legacy operation is available only in the virtual CPU interface when executing in Non-secure EL1.

In GICv3, the following restrictions apply to legacy operation:

- The GICv2 feature GICC_CTLR.AckCtl was deprecated in GICv2 and is not supported in GICv3. Correspondingly, even in legacy mode, the behavior is as if the GICC_CTLR.AckCtl bit described in GICv2 is RAZ/WI.

  Note

  In a GICv3 implementation that supports legacy operation, a VM is permitted to control Non-secure interrupts when GICV_CTLR.AckCtl set to 1. However, ARM deprecates the use of GICV_CTLR.AckCtl.

- The GICv2 configuration lockdown feature and the associated **CFGSDISABLE** input signal are not supported.
- A hypervisor executing at EL2 can control virtual interrupts only for the PE on which the EL2 software is executing, and cannot control virtual interrupts on other PEs

For legacy operation, an asymmetric configuration is supported where:

- Affinity routing and System register access are enabled in Non-secure state and at EL3.
- Affinity routing and System register access are disabled at Secure EL1.

This allows a secure operating system, running at Secure EL1, to use legacy functionality, provided that it does not configure Non-secure interrupts.

In GICv2 software executing in Secure state could use GICC_AIAR, GICC_AEOIR, GICC_AHPPIR, and GICC_ABPR to control interrupts in Non-secure state. There is no equivalent functionality in asymmetric configurations.

Chapter 2
GIC Partitioning

This chapter describes the GIC logical partitioning. It contains the following sections:

2.1 The GIC logical components

The GICv3 architecture consists of a set of logical components:

- A *Distributor*.
- A *Redistributor* for each PE that is supported.
- 0-16 *Interrupt Translation Service* components (ITS), to support the optional translation of events into LPIs.
- A *CPU interface* for each PE that is supported.

The Distributor, Redistributor and ITS are collectively known as an IRI.

Figure 2-1 shows the IRI.

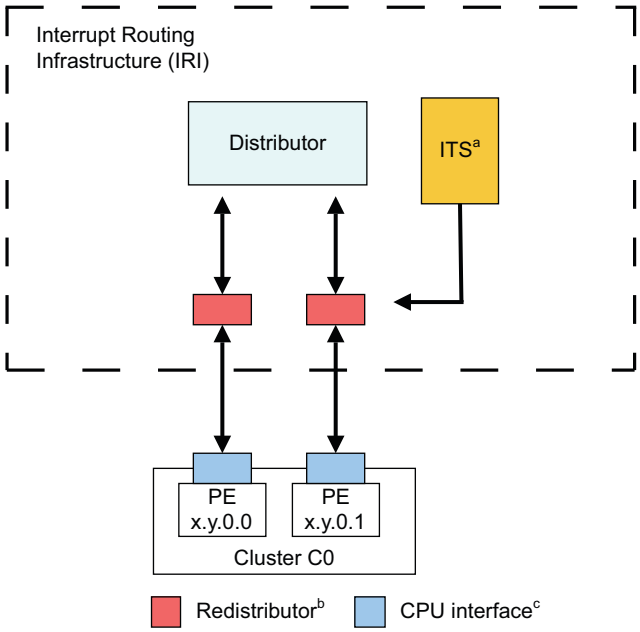

a. The inclusion of an ITS is optional, and there might be more than one ITS in a GIC.
b. There is one Redistributor per PE.
c. There is one CPU interface per PE.

Figure 2-1 Interrupt Routing Infrastructure

The CPU interface handles physical interrupts at all implemented Exception levels:

- Interrupts that are translated into LPIs are optionally routed via the ITS to the Redistributor and the CPU interface.
- PPIs are routed directly from the source to the local Redistributor.
- SPIs are routed from the source through the Distributor to the target Redistributor and the associated CPU interface.
- SGIs are generated by software through the CPU interface and Redistributor. They are then routed through the Distributor to one or more target Redistributors and the associated CPU interfaces.

In GICv3, the ITS is an optional component and translates events into physical LPIs. The architecture also supports direct LPIs that do not require the use of an ITS. Where LPIs are supported, it is IMPLEMENTATION DEFINED whether either:

- Direct LPIs are supported by accessing the registers in the Redistributors.
- LPI support is provided by the ITS.

An implementation must only support one of these methods.

In GICv4, the inclusion of at least one ITS is mandatory to provide support for the direct injection of virtual LPIs.

Figure 2-2 shows the GIC partitioning in an implementation that includes an ITS.

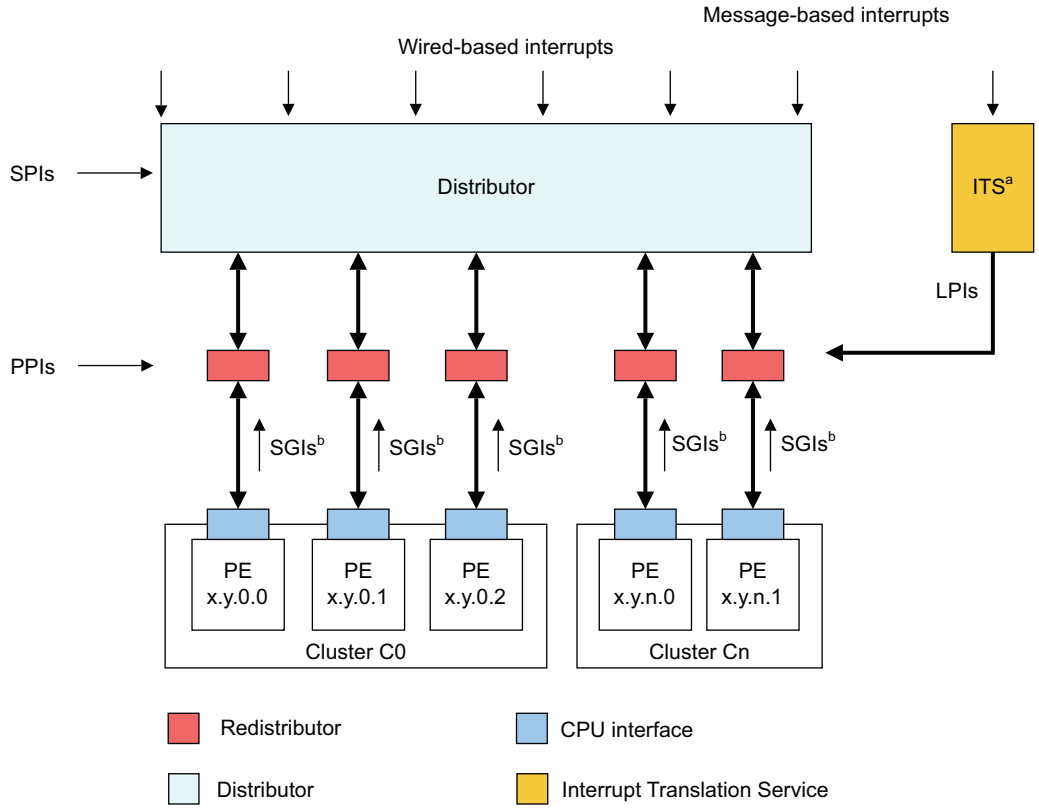

a. The inclusion of an ITS is optional, and there might be more than one ITS in a GIC.
b. SGIs are generated by a PE and routed through the Distributor.

Figure 2-2 GIC logical partitioning with an ITS

The mechanism for communication between the ITS and the Redistributors is IMPLEMENTATION DEFINED.

Figure 2-3 on page 2-32 shows the GIC partitioning in an implementation that does not include an ITS and that supports direct LPIs.

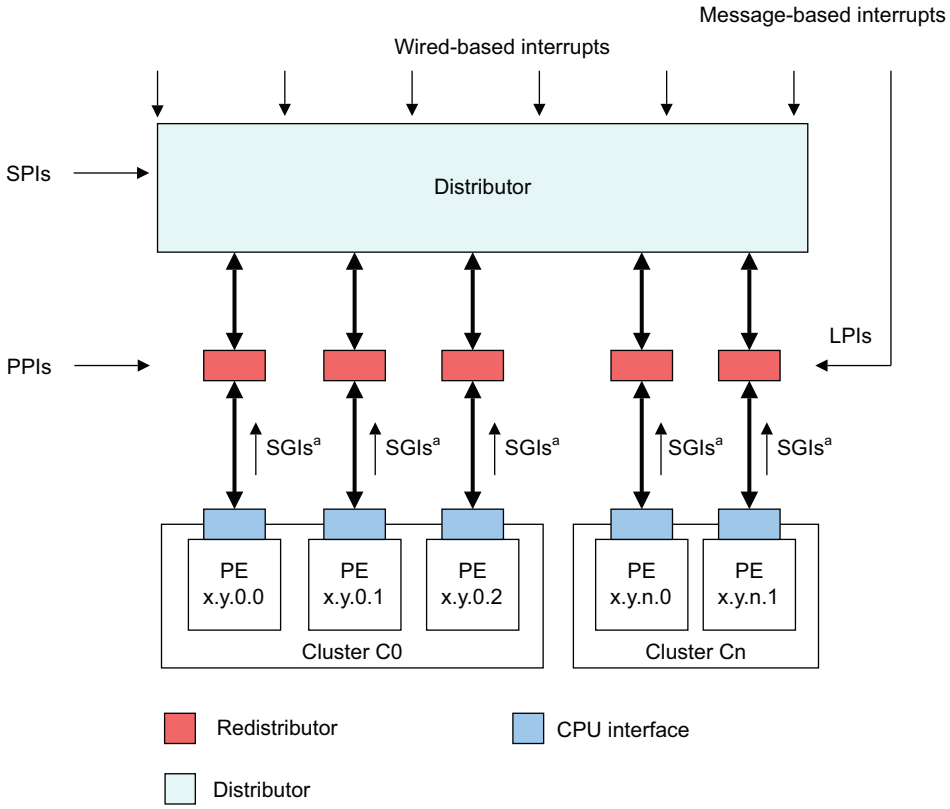

a. SGIs are generated by a PE and routed through the Distributor.

Figure 2-3 GIC logical partitioning without an ITS

The following list describes the components that are depicted in Figure 2-2 on page 2-31 in more detail:

Distributor The Distributor performs interrupt prioritization and distribution of SPIs and SGIs to the Redistributors and CPU interfaces that are connected to the PEs in the system.

GICD_CTLR provides global settings for:

- Enabling affinity routing.
- Disabling security.
- Enabling Secure and Non-secure Group 1 interrupts.
- Enabling Group 0 interrupts.

For SPIs, the Distributor provides a programming interface for:

- Enabling or disabling SPIs.
- Setting the priority level of each SPI.
- Routing information for each SPI.
- Setting each SPI to be level-sensitive or edge-triggered.
- Generating message-based SPIs.
- Assigning each SPI to an interrupt group.
- Controlling the pending and active state of SPIs.

For SGIs, the Redistributor provides a programming interface.

The Distributor registers are identified by the GICD_ prefix.

See Chapter 3 *Distribution and Routing of Interrupts* for more information.

Note

When handling physical interrupts during legacy operation, the Distributor controls the configuration information for PPIs and SGIs. See Chapter 10 *Legacy Operation and Asymmetric Configurations*.

Interrupt translation service, ITS

The ITS is an OPTIONAL hardware mechanism in the GICv3 architecture that routes LPIs to the appropriate Redistributor. Software uses a command queue to configure an ITS. Table structures in memory that are associated with an ITS translate an EventID associated with a device into a pending INTID for a PE.

The ITS is not OPTIONAL in GICv4, and all GICv4 implementations must include at least one ITS.

See *The ITS on page 6-99* for more information.

Redistributor

A Redistributor is the part of the IRI that is connected to the CPU interface of the PE. The Redistributor holds the control, prioritization, and pending information for all physical LPIs using data structures that are held in memory. Two registers in the Redistributor point to these data structures:

- GICR_PROPBASER.
- GICR_PENDBASER.

In GICv4, the Redistributor also includes registers to handle LPIs that are forwarded by an ITS to a Redistributor and directly to a VM, without involving the hypervisor. This is referred to as a *direct injection* of virtual interrupts into a VM.

In GICv4, the Redistributors collectively host the control, prioritization, and pending information for all virtual LPIs using data structures that are held in memory. Two registers in the Redistributor point to these data structures:

- GICR_VPROPBASER.
- GICR_VPENDBASER.

In an implementation that supports LPIs but does not include an ITS, the GICR_* registers contain a simple memory-mapped interface to signal and control physical LPIs.

Redistributors provide a programming interface for:

- Identifying, controlling, and configuring supported features to enable interrupts and interrupt routing of the implementation.
- Enabling or disabling SGIs and PPIs.
- Setting the priority level of SGIs and PPIs.
- Setting each PPI to be level-sensitive or edge-triggered.
- Assigning each SGI and PPI to an interrupt group.
- Controlling the pending state of SGIs and PPIs.
- Controlling the active state of SGIs and PPIs.
- Power management support for the connected PE.
- Where LPIs are supported, base address control for the data structures in memory that support the associated interrupt properties and their pending status.
- Where GICv4 is supported, base address control for the data structures in memory that support the associated virtual interrupt properties and their pending status.

The Redistributor registers are identified by the GICR_ prefix.

See *Affinity routing on page 3-43* and *The Distributor and Redistributors on page 3-38* for more information about the Redistributor.

CPU interface

The GIC architecture supports a CPU interface that provides a register interface to a PE in the system. Each CPU interface provides a programming interface for:

- General control and configuration to enable interrupt handling in accordance with the Security state and legacy support requirements of the implementation.
- Acknowledging an interrupt.
- Performing a priority drop.
- Deactivation of an interrupt.
- Setting an interrupt priority mask for the PE.
- Defining the preemption policy for the PE.
- Determining the highest priority pending interrupt for the PE.

The CPU interface has several components:

- A component that allows a supervisory level of software to control the handling of physical interrupts. The registers that are associated with this are identified by the ICC_ prefix.
- A component that allows a supervisory level of software to control the handling of virtual interrupts. The registers that are associated with this are identified by the ICV_ prefix.
- A component that allows a hypervisor to control the set of pending interrupts. The registers that are associated with this are identified by the ICH_ prefix.

Note

The System registers in the CPU interface are associated with software that is handling interrupts in the physical domain, or with execution at Non-secure EL1 as part of a VM. The configuration of HCR_EL2 determines whether the accesses are to the physical resources or the virtual resources.

The System registers accessible at EL2 that are used for controlling the list of active, pending, and active and pending, virtual interrupts for a PE are identified by the ICH_ prefix.

For more information on handling physical interrupts, see Chapter 4 *Physical Interrupt Handling and Prioritization*.

For more information on handling virtual interrupts, see Chapter 5 *Virtual Interrupt Handling and Prioritization*.

2.2 Interrupt bypass support

In all GIC implementations, a CPU interface optionally includes interrupt signal bypass, so that, when the signaling of an interrupt by the interface is disabled, a legacy interrupt signal is passed to the interrupt request input on the PE, bypassing the GIC functionality.

It is IMPLEMENTATION DEFINED whether bypass is supported.

The controls to determine whether GICv3 FIQ and IRQ outputs or the bypass signals are used differ depending on whether System register access is enabled.

When System register access is enabled, bypass disable is controlled at the highest implemented Exception level using two bits in ICC_SRE_EL1, ICC_SRE_EL2, or ICC_SRE_EL3, as appropriate:

- For FIQ bypass, this is the DFB bit.
- For IRQ bypass, this is the DIB bit.

This bypass mechanism is used when System register access is enabled. For information about bypass support during legacy operation, see *Legacy operation and bypass support on page 10-686*.

The interrupt groups that are supported by the GIC are allocated to FIQs and IRQs, as described in *Interrupt grouping on page 4-58*. Interrupt groups must be disabled at the CPU interface when bypass is enabled, otherwise the behavior of the GICv3 implementation is UNPREDICTABLE. This means that:

- ICC_IGRPEN0_EL1.Enable must have the value 0 when ICC_SRE_ELx.DFB == 0.
- ICC_IGRPEN1_EL1.Enable must have the value 0 when ICC_SRE_ELx.DIB == 0.

For more information about enabling interrupts, see *Enabling the distribution of interrupts on page 4-63*.

For information about the behavior when System register access is not enabled, see Chapter 10 *Legacy Operation and Asymmetric Configurations*.

For FIQs, the following pseudocode determines the source for interrupt signaling to a PE.

```
if ICC_SRE_EL3.SRE == 1 then
    if ICC_SRE_EL3.DFB == 0 then
       if ICC_SRE_EL1.SRE Secure  == 1 then
           BypassFIQsource
       else
        use legacy bypass support
    else
       use GICv3 FIQ output
else
    use legacy bypass support
```

For IRQs, the following pseudocode determines the source for interrupt signaling to a PE.

```
if ICC_SRE_EL3.SRE == 1 then
    if ICC_SRE_EL3.DIB == 0 then
       if ICC_SRE_EL1.SRE Secure  == 1 then
           BypassIRQsource
      else
        use legacy bypass support
    else
       use GICv3 IRQ output
else
    use legacy bypass support
```

Chapter 3
Distribution and Routing of Interrupts

This chapter describes the distribution and routing of interrupts to a target PE using affinity routing, and the assignment of interrupt IDs. It contains the following sections:

3.1 The Distributor and Redistributors

The Distributor provides the routing configuration for SPIs, and holds all the associated routing and priority information.

The Redistributor provides the configuration settings for PPIs.

A Redistributor always presents the pending interrupt with the highest priority to the CPU interface in finite time. For more information about interrupt prioritization, see *Interrupt prioritization on page 4-65*.

The highest priority interrupt might change because:

- The previous highest priority interrupt has been acknowledged.
- The previous highest priority interrupt has been preempted.
- The previous highest priority interrupt is removed and no longer valid.
- The group interrupt enable has been modified.
- The PE is no longer a participating PE. See *Participating nodes on page 3-44*.

3.2 INTIDs

Interrupts are identified using *ID numbers* (INTIDs). The range of INTIDs supported by GICv3 is IMPLEMENTATION DEFINED, according to the following rules:

- For the number of ID bits supported in the Distributor and Redistributor:
  - If LPIs are not supported, the ID space in the Distributor is limited to 10 bits. This is the same as in earlier versions of the GIC architecture.
  - If LPIs are supported, the INTID field is IMPLEMENTATION DEFINED in the range of 14-24 bits, as described in the register description for GICD_TYPER.

  Note

  A Redistributor can be configured through GICR_PROPBASER to use fewer bits than specified by GICD_TYPER.

- For the number of ID bits supported in the ITS:
  - If LPIs are supported, the INTID field is IMPLEMENTATION DEFINED in the range of 14-24 bits.
  - The size of the INTID field is defined by GITS_TYPER.IDbits.

  The ITS must be programmed so that interrupts that are forwarded to a Redistributor are in the range of interrupts that are supported by that Redistributor, otherwise the behavior is UNPREDICTABLE.

- For the number of ID bits supported in the CPU interface:
  - The GICv3 CPU interface supports either a 16-bit or a 24-bit INTID field, the choice being IMPLEMENTATION DEFINED. The number of physical interrupt identifier bits that are supported is indicated by ICC_CTLR_EL1.IDbits and ICC_CTLR_EL3.IDbits.

The valid ID space is governed by the implemented size in the CPU interface and the Distributor. It is a programming error to forward an INTID that is greater than the supported size to a CPU interface.

Unused INTID bits are RAZ. This means that any affected bit field is zero-extended.

Table 3-1 shows how the INTID space is partitioned by interrupt type.

Table 3-1 INTIDs

INTID	Interrupt type	Details	Notes
ID0 – ID15	SGI	These interrupts are local to a CPU interface.	IDs 0-1023 are compatible with earlier versions of the GIC architecture
ID16 – ID31	PPI		
ID32 – ID1019	SPI	Shared peripheral interrupts that the Distributor can route to either a specific PE, or to any one of the PEs in the system that is a participating node, see *Participating nodes on page 3-44*.	
ID1020 – ID1023	Special interrupt number	Interrupt IDs that are reserved for special purposes, as *Special INTIDs on page 3-40* describes.	
ID1024 – ID8191	-	Reserved	-
ID8192 – IMPLEMENTATION DEFINED	LPI	Peripheral hardware interrupts that are routed to a specific PE.	-

Table 3-2 shows the ARM recommended PPI INTID assignments.

Table 3-2 ARM recommended INTIDs for PPIs

INTID	PPI
30	Non-secure physical timer interrupt
29	Secure physical interrupt timer
27	Virtual timer interrupt
26	Hypervisor timer interrupt
25	Virtual CPU Interface Maintenance interrupt
24	Cross Trigger Interface interrupt
23	Performance Monitor Counter Overflow interrupt
22	Debug Communications Channel interrupt

The GICv4 architecture provides a unique ID space for each VM by supporting a vPEID in addition to the INTID space. See *About GIC support for virtualization on page 5-78* for more information about VMs and *The ITS on page 6-99* for more information about vPEIDs.

ARM strongly recommends that implemented interrupts are grouped to use the lowest INTID numbers and as small a range of INTIDs as possible. This reduces the size of the associated tables in memory that must be implemented, and that discovery routines must check.

3.2.1 Special INTIDs

The list of the INTIDs that the GIC architecture reserves for special purposes is as follows:

1020 The GIC returns this value in response to a read of ICC_IAR0_EL1 or ICC_HPPIR0_EL1 at EL3, to indicate that the interrupt being acknowledged is one which is expected to be handled at Secure EL1. This INTID is only returned when the PE is executing at EL3 using AArch64 state, or when the PE is executing in AArch32 state in Monitor mode.

For information about the relation of this value to ICC_CTLR_EL3.RM and legacy operation, see *Legacy support of interrupts and asymmetric configurations on page 10-684*.

1021 The GIC returns this value in response to a read of ICC_IAR0_EL1 or ICC_HPPIR0_EL1 at EL3, to indicate that the interrupt being acknowledged is one which is expected to be handled at Non-secure EL1 or EL2. This INTID is only returned when the PE is executing at EL3 using AArch64 state, or when the PE is executing in AArch32 state in Monitor mode.

For information about the relation of this value to ICC_CTLR_EL3.RM and legacy operation, see *Legacy support of interrupts and asymmetric configurations on page 10-684*.

1022 This value applies to legacy operation only. For more information, see *Use of the special INTID 1022 on page 10-685*.

1023 This value is returned in response to an interrupt acknowledge, if there is no pending interrupt with sufficient priority for it to be signaled to the PE, or if the highest priority pending interrupt is not appropriate for the:

- Current Security state.
- Interrupt group that is associated with the System register.

Note

These INTIDs do not require an end of interrupt or deactivation.

For more information about the use of special INTIDs, see the descriptions for the following registers:

- ICC_IAR0_EL1.
- ICC_IAR1_EL1.
- ICC_HPPIR0_EL1.
- ICC_HPPIR1_EL1.

3.2.2 Implementations with mixed INTD sizes

Implementations might choose to implement different INTID sizes for different parts of the GIC, subject to the following rules:

- PEs might implement either 16 or 24 bits of INTID.

  Note

  A system might include a mixture of PEs that support 16 bits of INTID and PEs that support 24 bits of INTID.

- The Distributor and Redistributors must all implement the same number of INTID bits.
- In systems that support LPIs, the Distributors and all Redistributors must implement at least 14 bits of INTID. The number of bits that is implemented in the Distributor and Redistributors must not exceed the minimum number that is implemented on any PE in the system.

  Note

  Because interrupts might target any PE, each PE must be able to receive the maximum INTID that can be sent by a Redistributor. This means that the INTID size that is supported by the Redistributors cannot exceed the minimum INTID size that is supported by each PE in the system.

- In systems that do not support LPIs, the Distributor and all Redistributors must implement at least 5 bits of INTID and cannot implement more than 10 bits of INTID.
- In systems that include one or more ITSs, an ITS might implement any value up to and including the number of bits that are supported by the Distributor and the Redistributors down to a minimum of 14 bits, which is the minimum number that is required for LPI support.

3.2.3 Valid interrupt ID check pseudocode

The following pseudocode describes how the GIC checks whether an INTID for a physical interrupt is valid:

```
// InterruptIdentifierValid()
// ==========================

boolean InterruptIdentifierValid(bits(64) data, boolean lpiAllowed)

    // First check whether any out of range bits are set
    integer N = CPUInterfaceIDSize();

    if !IsZero(data<63:N>) then
        if GenerateLocalSError() then
            // Reporting of locally generated SEIs is supported
            IMPLEMENTATION_DEFINED “SError INVALID_INTERRUPT_IDENTIFIER”;
            UNPREDICTABLE;

    intID = data<INTID_SIZE-1:0>;

    if !lpiAllowed && IsLPI(intID) then        // LPIs are not supported
        if GenerateLocalSError() then
            // Reporting of locally generated SEIs is supported
            IMPLEMENTATION_DEFINED “SError INVALID_INTERRUPT_IDENTIFIER”;
        UNPREDICTABLE;

    // Now check for special identifiers
    if IsSpecial(intID) then
        return FALSE;                          // It is a special ID

    // All the checks pass so the identifier is valid
```

```
    return TRUE;
```

The following pseudocode describes how the GIC checks whether an INTID for a virtual interrupt is valid:

```
// VirtualIdentifierValid()
// ========================

boolean VirtualIdentifierValid(bits(64) data, boolean lpiAllowed)

    // First check whether any out of range bits are set
    integer N = VIDBits();

    if !IsZero(data<63:N>) then
        if ICH_VTR_EL2.SEIS == '1' then
            // Reporting of locally generated SEIs is supported
            IMPLEMENTATION_DEFINED "SError INVALID_INTERRUPT_IDENTIFIER";
        UNPREDICTABLE;

    intID = data<INTID_SIZE-1:0>;

    if !lpiAllowed && IsLPI(intID) then        // LPIs are not supported
        if ICH_VTR_EL2.SEIS == '1' then
            // Reporting of locally generated SEIs is supported
            IMPLEMENTATION_DEFINED "SError INVALID_INTERRUPT_IDENTIFIER";
        UNPREDICTABLE;

    // Now check for special identifiers
    if IsSpecial(intID) then
        return FALSE;                          // It is a special ID

    // All the checks pass so the identifier is valid
    return TRUE;
```

The following pseudocode describes CPU interface ID size function.

```
// CPUInterfaceIDSize()
// ====================
// Returns the number of Interrupt ID bits implemented at the CPU interface. This value is an
// IMPLEMENTATION DEFINED choice of 16 or 24 and is discoverable from ICC_CTLR_EL1/EL3.IDbits

integer CPUInterfaceIDSize()
    return integer IMPLEMENTATION_DEFINED "CPU interface INTID size 16 or 24";
```

3.3 Affinity routing

Affinity routing is a hierarchical address-based scheme to identify specific PE nodes for interrupt routing.

For a PE, the affinity value is defined in MPIDR_EL1 for AArch64 state, and in MPIDR for AArch32 state:

- Affinity routing is a 32-bit value that is composed of four 8-bit affinity fields. These fields are the nodes *a*, *b*, *c*, and *d*.
- GICv3 using AArch64 state can support:
  - A four level routing hierarchy, `a.b.c.d`.
  - A three level routing hierarchy, `0.b.c.d`.
- GICv3 using AArch32 state only supports three affinity levels.
- ICC_CTLR_EL3.A3V, ICC_CTLR_EL1.A3V, and GICD_TYPER.A3V indicate whether four levels or three levels of affinity are implemented.

Note

An implementation that requires four levels of affinity must only support AArch64 state.

The enumeration notation for specifying nodes in an affinity hierarchy is of the following form, where `Affx` is Affinity level `x`:

```
Aff3.Aff2.Aff1.Aff0
```

Affinity routing for a Security state is enabled in the Distributor, using the *Affinity Routing Enable* (ARE) bits. Affinity routing is enabled:

- For Secure interrupts, if GICD_CTLR.ARE_S is set to 1.
- For Non-secure interrupts, if the GICD_CTLR.ARE_NS bit is set to 1.

GICD_CTLR.ARE_S and GICD_CTLR.ARE_NS are RAO/WI if affinity routing is permanently enabled.

For the handling of physical interrupts when affinity routing is enabled, System register access must also be enabled, see *GIC System register access on page 8-160*. For the other cases, see Chapter 10 *Legacy Operation and Asymmetric Configurations*.

3.3.1 Routing SPIs and SGIs by PE affinity

SPIs are routed using an affinity address and the routing mode information that is held in GICD_IROUTER<n>. SGIs are routed using the affinity address and routing mode information that is written by software when it generates the SGI.

SGIs are generated using the following registers:

- ICC_SGI0R_EL1.
- ICC_SGI1R_EL1.
- ICC_ASGI1R_EL1.

ARM strongly recommends that only values in the range 0-15 are used at affinity level 0 to align with the SGI target list capability. See *Software Generated Interrupts on page 4-55*.

SPIs and SGIs are routed using different registers:

- SPIs are routed using GICD_IROUTER<n>.Interrupt_Routing_Mode:
  - If GICD_IROUTER<n>.Interrupt_Routing_Mode is cleared to 0, SPIs are routed to a single PE specified by `a.b.c.d`.
  - If GICD_IROUTER<n>.Interrupt_Routing_Mode is set to 1, SPIs are routed to any PE defined as a participating node. For more information about participating nodes, see *Participating nodes on page 3-44*.
- SGIs are routed using ICC_SGI0R_EL1.IRM, and ICC_SGI1R_EL1.IRM:
  - If the IRM bit is set to 1, SGIs are routed to all participating PEs in the system, excluding the originating PE.

— If the IRM bit is cleared to 0, SGIs are routed to a group of PEs, specified by `a.b.c.targetlist`. The target list provides a bitfield encoding for affinity level 0 values of 0-15.

3.3.2 Participating nodes

An enabled SPI configured to use the 1 of N distribution model can target a PE when:

- GICR_WAKER.ProcessorSleep == 0 and the interrupt group of the interrupt is enabled on the PE.
- GICD_CTLR.E1NWF == 1.
- GICR_TYPER.DPGS == 1, and, for the interrupt group of the interrupt, GICR_CTLR.{DPG1S, DPG1NS, DPG0} == 0.

For more information about whether a PE can be selected as the target when the 1 of N distribution model is used, see *GICR_CTLR, Redistributor Control Register on page 8-492*.

For more information about enabling interrupts and interrupt groups, see *Enabling the distribution of interrupts on page 4-63*.

3.3.3 Changing affinity routing enables

This manual describes the GICv3 architecture in a system with affinity routing enabled. This means that:

- GICD_CTLR.ARE_NS == 1.
- GICD_CTLR.ARE_S == 1.

If the value of GICD_CTLR.ARE_NS or GICD_CTLR.ARE_S is changed from 1 to 0, the result is UNPREDICTABLE.

When GICD_CTLR. DS == 0, then:

- Changing GICD_CTLR.ARE_S from 0 to 1 is UNPREDICTABLE except when all of the following apply:
  - GICD_CTLR. EnableGrp0 == 0.
  - GICD_CTLR. EnableGrp1S == 0.
  - GICD_CTLR. EnableGrp1NS == 0
- Changing GICD_CTLR.ARE_NS from 0 to 1 is UNPREDICTABLE except when GICD_CTLR. EnableGrp1NS == 0.

When GICD_CTLR. DS == 1, then:

- Changing GICD_CTLR.ARE_S or GICD_CTLR.ARE_NS from 0 to 1 is UNPREDICTABLE except when all of the following apply:
  - GICD_CTLR. EnableGrp0 == 0.
  - GICD_CTLR. EnableGrp1 == 0.

Note

The effect of clearing GICD_CTLR.EnableGrp0, GICD_CTLR.EnableGrp1S, or GICD_CTLR.EnableGrp1NS, as appropriate, must be visible when changing GICD_CTLR.ARE_S or GICD_CTLR.ARE_NS from 0 to 1. Software can poll GICD_CTLR.RWP to check that writes that clear GICD_CTLR.EnableGrp0, GICD_CTLR.EnableGrp1S, or GICD_CTLR.EnableGrp1NS bits have completed.

Chapter 4
Physical Interrupt Handling and Prioritization

This chapter describes the fundamental aspects of GIC interrupt handling and prioritization. It contains the following sections:

4.1 Interrupt lifecycle

GIC interrupt handling is based on the GIC interrupt lifecycle, a series of high-level processes that apply to any interrupt using the GIC architecture. The interrupt lifecycle provides a basis for describing the detailed steps of the interrupt handling process. The GIC also maintains a state machine that controls interrupt state transitions during the lifecycle.

Figure 4-1 shows the GIC interrupt lifecycle for physical interrupts.

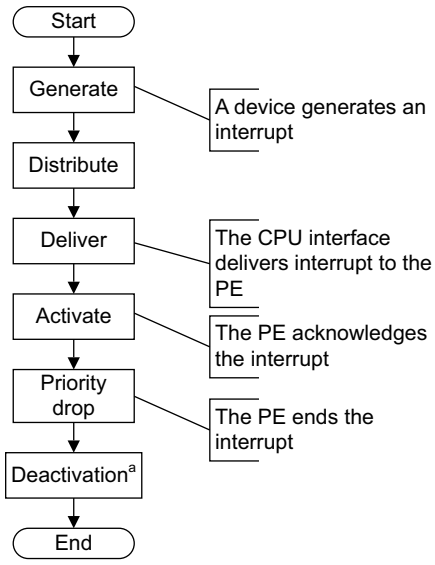

a. This step does not apply to LPIs.

Figure 4-1 Physical interrupt lifecycle

The interrupt lifecycle in Figure 4-1 is as follows:

1. Generate interrupt. An interrupt is generated either by the peripheral or by software.
2. Distribute. The IRI performs interrupt grouping, interrupt prioritization, and controls the forwarding of interrupts to the CPU interfaces.
3. Deliver. A physical CPU interface delivers interrupts to the corresponding PE.
4. Activate. When software running on a PE acknowledges an interrupt, the GIC sets the highest active priority to that of the activated interrupt, and for SPIs, SGIs, and PPIs the interrupt becomes active.
5. Priority drop. Software running on the PE signals to the GIC that the highest priority interrupt has been handled to the point where the running priority can be dropped. The running priority then has the value that it had before the interrupt was acknowledged. This is the point where the end of interrupt is indicated by the interrupt handler. The end of the interrupt can be configured to also perform deactivation of the interrupt.
6. Deactivation. Deactivation clears the active state of the interrupt, and thereby allows the interrupt, when it is pending, to be taken again. Deactivation is not required for LPIs. Deactivation can be configured to occur at the same time as the priority drop, or it can be configured to occur later as the result of an explicit interrupt deactivation operation. This latter approach allows for software architectures where there is an advantage to separating interrupt handling into initial handling and scheduled handling.

4.1.1 Physical CPU interface

A CPU interface provides an interface to a PE that is connected to the GIC. Each CPU interface is connected to a single PE.

A CPU interface receives pending interrupts prioritized by the IRI, and determines whether the interrupt is enabled and has sufficient priority to be signaled to the PE. At any time, the connected PE can determine the:

- INTID of its highest priority pending interrupt, by reading ICC_HPPIR0_EL1 or ICC_HPPIR1_EL1.
- Priority of the highest priority active interrupt, by reading ICC_RPR_EL1.

Note

The priority of the highest priority active interrupt for which there has not been a priority drop is also known as the running priority.

When an LPI is acknowledged, the pending state for the interrupt changes to not pending in the Redistributor. The Redistributor does not maintain an active state for LPIs.

When the PE acknowledges an SGI, a PPI, or an SPI at the CPU interface, the Distributor changes the status of the interrupt to active if:

- It is an edge-triggered interrupt, and another edge has not been detected since the interrupt was acknowledged.
- It is a level-sensitive interrupt, and the level has been deasserted since the interrupt was acknowledged.

When the PE acknowledges an SGI, a PPI, or an SPI at the CPU interface, the Distributor changes the status of the interrupt to active and pending if:

- It is an edge-triggered interrupt, and another edge has been detected since the interrupt was acknowledged.
- It is a level-sensitive interrupt, and the level has not been deasserted since the interrupt was acknowledged.

When the PE acknowledges an SGI, a PPI, or an SPI at the CPU interface, the CPU interface can signal another interrupt to the PE, to preempt interrupts that are active on the PE. If there is no pending interrupt with sufficient priority to be signaled to the PE, the interface deasserts the interrupt request signal to the PE.

The following stages of the interrupt lifecycle are described in the remainder of this section:

- *Activation*.
- *Priority drop* on page 4-48.
- *Deactivation* on page 4-49.

The priority drop and deactivation can be performed as a single operation or can be split, as defined by ICC_CTLR_EL1.EOImode and ICC_CTLR_EL3.EOImode_EL3.

Activation

The interrupt handler reads ICC_IAR0_EL1 for Group 0 interrupts, and ICC_IAR1_EL1 for Group 1 interrupts, in the corresponding CPU interface to acknowledge the interrupt. This read returns either:

- The INTID of the highest priority pending interrupt, if that interrupt is of sufficient priority for it to be signaled to the PE. This is the normal response to an interrupt acknowledge.
- Under certain conditions, an INTID that indicates a special interrupt number, see *INTIDs* on page 3-39.

Whether a read of ICC_IAR0_EL1 and ICC_IAR1_EL1 returns a valid INTID depends on:

- Which of the two registers is accessed.
- The Security state of the PE.
- Whether there is a pending interrupt of sufficient priority to be signaled to the PE, and if so, whether:
  - The highest priority pending interrupt is a Secure Group 1 or a Non-secure Group 1 interrupt.
  - Interrupt signaling is enabled for that interrupt group.
- The Exception level at which the PE is executing.

All interrupts, when acknowledged, modify the *Active Priorities Registers*. See *System register access to the Active Priorities registers* on page 4-69.

In certain circumstances, the value of SCR_EL3.NS affects the value returned when a PE acknowledges an interrupt. That is, when the PE is executing at EL3, a Secure read of ICC_IAR0_EL1 returns a special interrupt number that indicates the required Security state transition for the highest priority pending interrupt. Otherwise, the INTID is returned.

For SGIs in a multiprocessor implementation, the GIC uses the targeted list model, where the acknowledgement of an interrupt by one PE has no effect on the state of the interrupt on other CPU interfaces. When a PE acknowledges the interrupt, the pending state of the interrupt is cleared only for that PE. The interrupt remains pending for the other PEs.

The effects of reading ICC_IAR0_EL1 and ICC_IAR1_EL1 on the state of a returned INTID are not guaranteed to be visible until after the execution of a `DSB`.

Priority drop

After an interrupt has been acknowledged, a valid write to ICC_EOIR0_EL1 for Group 0 interrupts, or a valid write to ICC_EOIR1_EL1 for Group 1 interrupts, results in a priority drop.

A valid write to ICC_EOIR0_EL1 or ICC_EOIR1_EL1 to perform a priority drop is required for each acknowledged interrupt, even for LPIs which do not have an active state. A priority drop must be performed by the same PE that activated the interrupt.

Note

A valid write is a write that is:

- Not UNPREDICTABLE.
- Not ignored.
- Not writing an INTID that is either unsupported or within the range 1020-1023.

For each CPU interface, the GIC architecture requires the order of the valid writes to ICC_EOIR0_EL1 and ICC_EOIR1_EL1 to be the exact reverse of the order of the reads from ICC_IAR0_EL1 and ICC_IAR1_EL1, as shown in Figure 4-2.

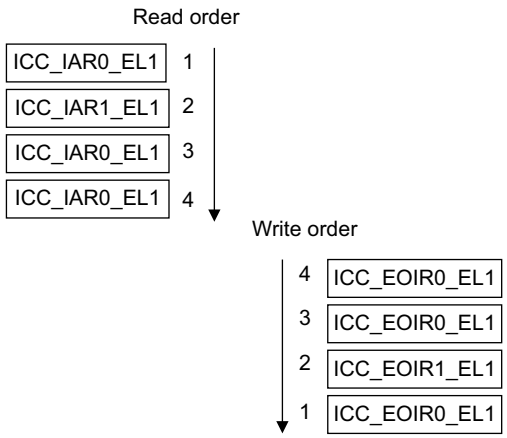

Figure 4-2 Read and write order

On a priority drop, the running priority is reduced from the priority of the interrupt indicated by the write to ICC_EOIR0_EL1 or ICC_EOIR1_EL1 to either:

- The priority of the highest-priority active interrupt for which there has been no write to ICC_EOIR0_EL1 or ICC_EOIR1_EL1.
- The idle priority, `0xFF`, if there is no active interrupt.

Note

For compatibility with possible extensions to the GIC architecture specification, software must preserve the entire register value read from ICC_IAR0_EL1 and ICC_IAR1_EL1 when it acknowledges the interrupt, and use that entire value for the corresponding write to ICC_EOIR0_EL1 and ICC_EOIR1_EL1 by the same PE.

When GICD_CTLR.DS == 0:

- A write to ICC_EOIR0_EL1 performs a priority drop for Group 0 interrupts.

- A write to ICC_EOIR1_EL1 performs a priority drop for Non-secure Group 1 interrupts, if the PE is operating in Non-secure state or at EL3.
- When operating in Secure state, a write to ICC_EOIR1_EL1 performs a priority drop for Secure Group 1 interrupts.

When GICD_CTLR.DS == 1:

- A write to ICC_EOIR0_EL1 performs a priority drop for Group 0 interrupts.
- A write to ICC_EOIR1_EL1 performs a priority drop for Group 1 interrupts.

Deactivation

PPIs, SGIs, and SPIs have an active state in the IRI and must be deactivated.

SGIs and PPIs must be deactivated by the PE that activated the interrupt. SPIs can be deactivated by a different PE.

Interrupt deactivation is required to change the state of an interrupt either:

- From active and pending to pending.
- From active to inactive.

Depending on the Exception level and Security state, ICC_CTLR_EL1.EOImode and ICC_CTLR_EL3.EOImode_EL3 in the appropriate CPU Interface Control Register determine whether priority drop and interrupt deactivation happen together or separately:

- The priority drop and interrupt deactivation happen together when ICC_CTLR_EL1.EOImode or ICC_CTLR_EL3.EOImode_EL3 in the CPU interface is 0, and the PE writes to ICC_EOIR0_EL1 or ICC_EOIR1_EL1. In this case a write to ICC_DIR_EL1 is not required.
- The priority drop and interrupt deactivation are separated when ICC_CTLR_EL1.EOImode or ICC_CTLR_EL3.EOImode_EL3 in the CPU interface is 1, and the PE writes to ICC_EOIR0_EL1 or ICC_EOIR1_EL1. In this case:
  - The priority drop happens when the PE writes to ICC_EOIR0_EL1 or ICC_EOIR1_EL1.
  - Interrupt deactivation happens later, when the PE writes to ICC_DIR_EL1. A valid write to ICC_DIR_EL1 results in interrupt deactivation for a Group 0 or a Group 1 interrupt.

There are no ordering requirements for writes to ICC_DIR_EL1. If software writes to ICC_DIR_EL1 when the following conditions are true, the results are UNPREDICTABLE:

- The appropriate EOIMode bit is cleared to 0.
- The ICC_CTLR_EL1.EOImode or ICC_CTLR_EL3.EOIMode_EL3 is set to 1 and there has not been a corresponding write to ICC_EOIR0_EL1 or ICC_EOIR1_EL1.

When ICC_CTLR_EL1.EOImode or ICC_CTLR_EL3.EOIMode_EL3 == 1 but the interrupt is not active in the Distributor, writes to ICC_DIR_EL1 must be ignored. If supported, an implementation might generate a system error.

Table 4-1 shows how a write to ICC_EOIR0_EL1 or ICC_EOIR1_EL1 affects deactivation.

Table 4-1 Effect of writing to ICC_EOIR0_EL1 or ICC_EOIR1_EL1

Access	ICC_CTLR_EL1.EOImode or ICC_CTLR_EL3.EOImode_EL3	Identified interrupt	Effect
ICC_EOIR1_EL1	0	Group 0	Access ignored
ICC_EOIR0_EL1	0	Group 0	Interrupt deactivated
ICC_EOIR1_EL1	0	Group 1	Interrupt deactivated
ICC_EOIR0_EL1	0	Group 1	Access ignored
-	1	-	Interrupt remains active

When GICD_CTLR.DS == 0, access to certain registers is restricted. See *Interrupt grouping and security* on page 4-59.

The following pseudocode determines whether EOImode is set for the current Exception level and Security state:

```
// EOImodeSet()
// ============

boolean EOImodeSet()

    if HaveEL(EL3) then
        // EL3 is implemented so return the value appropriate to the EL and security state
        if PSTATE.EL == EL3 && ICC_SRE_EL3.SRE == '1' then
            // In EL3
            EOImode = ICC_CTLR_EL3.EOImode_EL3;

        elsif IsSecure() then
            EOImode = ICC_CTLR_EL3.EOImode_EL1S;

        else                                            // Non-secure
            EOImode = ICC_CTLR_EL3.EOImode_EL1NS;
    else
        // No EL3 so return the value from ICC_CTLR_EL1
        EOImode = ICC_CTLR_EL1.EOImode;

    return EOImode == '1';
```

Effect of Security states on writes to ICC_DIR_EL1

The effect of a write to ICC_DIR_EL1 depends on whether the GIC supports one or two Security states:

- If the GIC supports two Security states, a valid:
  - Secure write to ICC_DIR_EL1 deactivates the specified interrupt, regardless of whether that interrupt is a Group 0 or a Group 1 interrupt.
  - Non-secure write to ICC_DIR_EL1 deactivates the specified interrupt only if that interrupt is a Non-secure Group 1 interrupt.
- If the GIC supports only a single Security state, a valid write to ICC_DIR_EL1 deactivates the specified interrupt, regardless of whether that interrupt is a Group 0 or Group 1 interrupt.

Table 4-2 shows the behavior of valid writes to ICC_DIR_EL1. In an implementation that supports only a single Security state, valid writes have the behavior shown for Secure writes to ICC_DIR_EL1.

Table 4-2 Behavior of writes to ICC_DIR_EL1

Security state of writes to ICC_DIR_EL1	Interrupt group	Effect
Non-secure	Non-secure Group 1	Interrupt is deactivated.
Non-secure	Group 0 or Secure Group 1	Write is ignored.
Secure	x	Interrupt is deactivated.

4.1.2 Interrupt handling state machine

The GIC maintains a state machine for each supported interrupt. The possible states of an interrupt are:

- Inactive.
- Pending.
- Active.
- Active and pending.

PPIs, SGIs, and SPIs can have an active and pending state. Interrupts that are active and pending are never signaled to a connected PE.

LPIs have a pending state that is held in memory associated with a Redistributor, and therefore a PE. This also applies to directly injected virtual LPIs, see *Virtual LPI support* on page 5-86.

Note

There is no active or active and pending state for LPIs.

Figure 4-3 shows an instance of the interrupt handling state machine, and the possible state transitions.

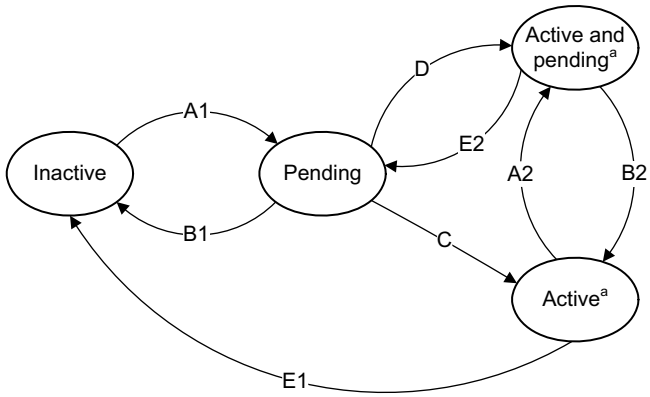

a. Not applicable for LPIs.

Figure 4-3 Interrupt handling state machine

Note

LPIs do not have an active state in the Redistributor, but do require an active priority in the CPU interface. See Chapter 6 *Locality-specific Peripheral Interrupts and the ITS* for more information.

When interrupt forwarding by the Distributor and interrupt signaling by the CPU interface are enabled, the conditions that cause each of the state transitions are as follows:

Transition A1 or A2, add pending state

This transition occurs when the interrupt becomes pending, either as a result of the peripheral generating the interrupt or as result of software generating the interrupt.

Transition B1 or B2, remove pending state

This transition occurs when the interrupt has been deasserted by the peripheral, if the interrupt is a level-sensitive interrupt, or when software has changed the pending state.

For LPIs, it also occurs on acknowledgement of the interrupt.

Transition C, pending to active

This transition occurs on acknowledgement of the interrupt by the PE for edge-triggered SPIs, SGIs, and PPIs.

For SPIs, SGIs, and PPIs, this transition occurs when software reads an INTID value from ICC_IAR0_EL1 or ICC_IAR1_EL1.

Transition D, pending to active and pending

This transition occurs on acknowledgement of the interrupt by the PE for level-sensitive SPIs, SGIs, and PPIs.

Transition E1 or E2, remove active state

This transition occurs when software deactivates an interrupt for SPIs, SGIs, and PPIs.

4.2 Locality-specific Peripheral Interrupts

LPIs are targeted peripheral interrupts that are routed to a specific PE within the affinity hierarchy. In a system where two Security states are enabled, LPIs are always Non-secure Group 1 interrupts. LPIs only support edge-triggered behavior. For more information about LPIs, see *LPIs* on page 6-92.

4.3 Private Peripheral Interrupts

PPIs are interrupts that target a single, specific PE, and different PEs can use the same INTID to indicate different events. PPIs can be Group 0 interrupts, Secure Group 1 interrupts, or Non-secure Group 1 interrupts. They can support either edge-triggered or level-sensitive behavior.

Note

Commonly, ARM expects that PPIs are used by different instances of the same interrupt source on each PE, thereby allowing a common INTID to be used for PE specific events, such as the interrupts from a private timer.

4.4 Software Generated Interrupts

SGIs are typically used for inter-processor communication, and are generated by a write to an SGI register in the GIC. SGIs can be either Group 0 or Group 1 interrupts, and they can support only edge-triggered behavior.

The registers associated with the generation of SGIs are part of the CPU interface:

- A PE generates a Group 1 SGI by writing to ICC_SGI1R_EL1 or ICC_ASGI1R_EL1.
- A PE generates a Group 0 SGI by writing to ICC_SGI0R_EL1.
- Routing information is supplied as the bitfield value in the write to the register that generated the SGI. The SGI can be routed to:
  - The group of PEs specified by `a.b.c.targetlist`. This can include the originating PE.
  - All participating PEs in the system, excluding the originating PE.

  See *Routing SPIs and SGIs by PE affinity* on page 3-43 for more information.

ICC_SGI1R_EL1 allows software executing in a Secure state to generate Secure Group 1 SGIs on a target PE that is executing in Secure state.

ICC_SGI1R_EL1 allows software executing in a Non-secure state to generate Non-secure Group 1 SGIs on a target PE that is executing in Non-secure state.

ICC_ASGI1R_EL1 allows software executing in a Secure state to generate Non-secure Group 1 SGIs.

ICC_ASGI1R_EL1 allows software executing in a Non-secure state to generate Secure Group 1 SGIs, if permitted by the settings of GICR_NSACR in the Redistributor corresponding to the target PE.

ICC_SGI0R_EL1 allows software executing in Secure state, and software executing in Non-secure state, if permitted by the settings of GICR_NSACR in the Redistributor corresponding to the target PE, to generate Secure Group 0 SGIs for each target PE.

For more information about the use of control registers to forward SGIs to a target PE, see Table 8-14 on page 8-171.

4.5 Shared Peripheral Interrupts

SPIs are peripheral interrupts that the Distributor can route to a specified PE that can handle the interrupt, or to a PE that is one of a group of PEs in the system that has been configured to accept this type of interrupt. SPIs can be either Group 0 or Group 1 interrupts, and they can support either edge-triggered or level-sensitive behavior.

SPIs can be either wired-based or message-based interrupts.

Support for message-based SPIs is optional, and can be discovered through GICD_TYPER.MBIS. Message-based SPIs can be:

- Generated by a write to GICD_SETSPI_NSR or GICD_SETSPI_SR
- Cleared by a write to GICD_CLRSPI_NSR orX GICD_CLRSPI_SR.

The effect of a write to these registers depends on whether the targeted SPI is configured to be an edge-triggered or a level-sensitive interrupt:

- For an edge-triggered interrupt, a write to GICD_SETSPI_NSR or GICD_SETSPI_SR sets the interrupt pending. The interrupt is no longer pending when it is activated, or when it is cleared by a write to GICD_CLRSPI_NSR, GICD_CLRSPI_SR, or GICD_ICPENDR<n>.
- For a level-sensitive interrupt, a write to GICD_SETSPI_NSR or GICD_SETSPI_SR sets the interrupt pending. It remains pending until it is deasserted by a write to GICD_CLRSPI_NSR or GICD_CLRSPI_SR. If the interrupt is activated between the time it is asserted by a write to GICD_SETSPI_NSR or GICD_SETSPI_SR and the time it is deactivated by a write to GICD_CLRSPI_NSR or GICD_CLRSPI_SR, then the interrupt becomes active and pending.

  It is IMPLEMENTATION DEFINED for a level-sensitive interrupt whether a write to GICD_ICPENDR<n> has any effect on an interrupt that has been set pending by a write to GICD_SETSPI_NSR or GICD_SETSPI_SR, or whether a write to GICD_CLRSPI_NSR or GICD_CLRSPI_SR has any effect on an interrupt that has been set pending by a write GICD_ISPENDR<n>.
- Changing the configuration of an interrupt from level-sensitive to edge-triggered, or from edge-triggered to level-sensitive, when there is a pending interrupt, leaves the interrupt in an UNKNOWN state.

Figure 4-4 on page 4-57 shows how message-based interrupt requests can trigger SPIs.

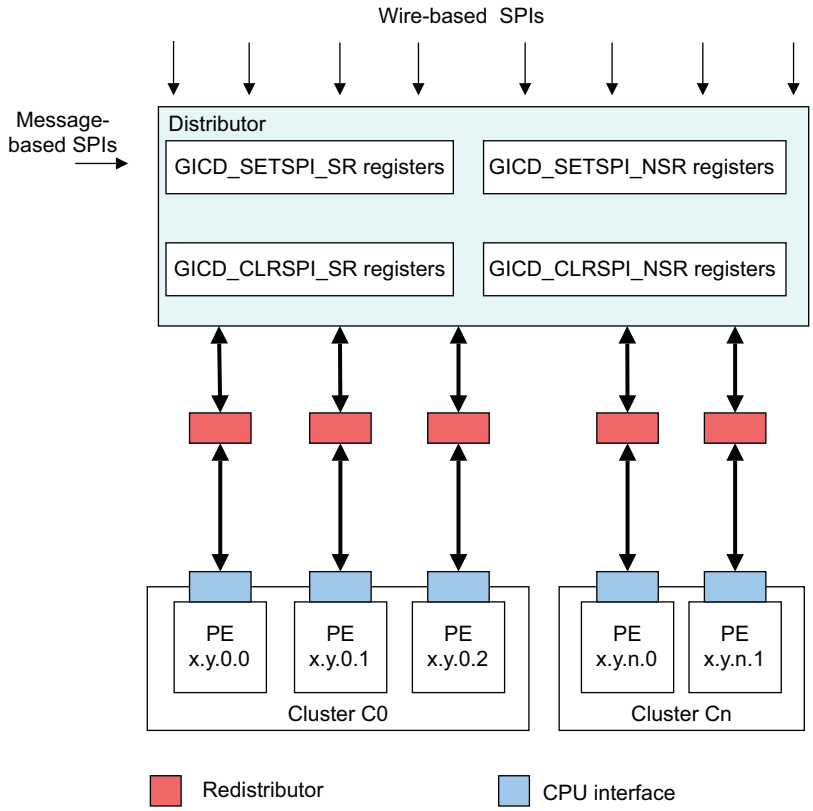

Figure 4-4 Triggering SPIs

4.6 Interrupt grouping

GICv3 uses *interrupt grouping* as a mechanism to align interrupt handling with the ARMv8 Exception model and Security model.

In a system with two Security states, an interrupt is configured as one of the following:

- A Group 0 physical interrupt:
  - ARM expects these interrupts to be handled at EL3.
- A Secure Group 1 physical interrupt:
  - ARM expects these interrupts to be handled at Secure EL1.
- A Non-secure Group 1 physical interrupt:
  - ARM expects these interrupts to be handled at Non-secure EL2 in systems using virtualization, or at Non-secure EL1 in systems not using virtualization

In a system with one Security state, or where specific PEs only operate in one Security state, an interrupt is configured to be either:

- Group 0.
- Group 1.

At the System level, GICD_CTLR.DS indicates if the GIC is configured with one or two Security states. For more information about Security, see *Interrupt grouping and security* on page 4-59.

These interrupt groups are mapped onto the ARMv8 FIQ and IRQ exceptions, see *Interrupt assignment to IRQ and FIQ signals* on page 4-60.

GICD_IGROUPR<n> and GICD_IGRPMODR<n> configure the interrupt group for SPIs. *n* is greater than zero.

GICR_IGROUPR0 and GICR_IGRPMODR0 configure the interrupt group for SGIs and PPIs.

Note

- It is IMPLEMENTATION DEFINED whether the bits associated with implemented SPIs, PPIs, or SGIs are programmable or have a fixed value. See the individual register descriptions for details.
- When GICD_CTLR.DS == 0, LPIs are always Non-secure Group 1 interrupts. When GICD_CTLR.DS == 1, LPIs are always Group 1 interrupts.

System registers control and configure Group 0 and Group 1 interrupts:

- For Group 0 interrupts, software uses:
  - ICC_IAR0_EL1 to read a Group 0 INTID on an interrupt acknowledge.
  - ICC_EOIR0_EL1 to write a Group 0 interrupt completion.
  - ICC_BPR0_EL1 to configure the binary point for Group 0 prioritization.
    This register is also used for Group 1 prioritization when ICC_CTLR_EL1.CBPR == 1.
  - ICC_HPPIR0_EL1 to read the highest Group 0 interrupt that is currently pending.
  - ICC_IGRPEN0_EL1 to enable Group 0 interrupts at the CPU interface.
- For Group 1 interrupts, software uses:
  - ICC_IAR1_EL1 to read a Group 1 INTID on an interrupt acknowledge.
  - ICC_EOIR1_EL1 to write a Group 1 interrupt completion.
  - ICC_BPR1_EL1 to configure the binary point for Group 1 prioritization for the current Security state.
  - ICC_HPPIR1_EL1 to read the highest Group 1 interrupt that is currently pending.
  - ICC_IGRPEN1_EL1 to enable Group 1 interrupts for the target Security state of the interrupt.

In a system with two Security states, Group 0 interrupts are always Secure. For more information about grouping and Security, see *Interrupt grouping and security* on page 4-59.

4.6.1 Interrupt grouping and security

The ARM architecture provides two Security states, each with an associated physical memory address space:

- Secure state.
- Non-secure state.

A software hierarchy of user and privileged code can execute in either state, and software executing in Non-secure state can only access Secure state through a system call to the Secure monitor. The GIC architecture supports the routing and handling of interrupts associated with both Security states.

GICD_CTLR.DS indicates whether a GIC is configured to support the ARMv8-A Security model. This configuration affects:

- Register access, see *GIC System register access* on page 8-160.
- The interrupt groups that are supported.

When GICD_CTLR.DS == 0:

- The GIC supports two Security states, Secure state and Non-secure state.
- The GIC supports three interrupt groups:
  - Group 0.
  - Secure Group 1.
  - Non-secure Group 1.
- Both the Security state and GICR_NSACR determine whether an SGI can be generated.
- The Security state is checked during:
  - Configuration of an interrupt.
  - Acknowledgement of an interrupt.
  - Priority drop.
  - Deactivation.

When GICD_CTLR.DS == 1:

- The GIC supports only a single Security state. This can be either Secure state or Non-secure state.
- The GIC supports two interrupt groups:
  - Group 0.
  - Group 1.
- SGIs can be generated regardless of the settings in GICR_NSACR.
- The Security state is not checked during:
  - Configuration of an interrupt.
  - Acknowledgement of an interrupt.
  - Priority drop.
  - Deactivation.

In a multiprocessor system, one or more PEs within the system might support accesses to resources that are available only in Secure state, or accesses to resources that are available only in Non-secure state. It is a programming error if software configures:

- A Group 0 or Secure Group1 interrupt to be forwarded to a PE that only supports Non-secure state.
- A Non-secure Group1 interrupt to be forwarded to a PE that only supports Secure state.

There is a dedicated register for the priority grouping for each interrupt group, ICC_BPR0_EL1 for Group 0 interrupts and ICC_BPR1_EL1 for Group 1interrupts. However, it is possible to configure a common Binary Point Register for both groups using:

- ICC_CTLR_EL1.CBPR.
- ICC_CTLR_EL3.CBPR_EL1NS and ICC_CTLR_EL3.CBPR_EL1S for an independent common Binary Point Register configuration of Non-secure Group 1 and Secure Group 1 interrupts.

For information about interrupt grouping and legacy operation, see Chapter 10 *Legacy Operation and Asymmetric Configurations*.

4.6.2 Interrupt assignment to IRQ and FIQ signals

A Group 0 physical interrupt, when it is the highest priority pending interrupt and has sufficient priority, is always signaled as an FIQ.

A Group 1 physical interrupt, when it is the highest priority pending interrupt and has sufficient priority, is signaled as an FIQ if either of the following conditions is true, otherwise it is signaled as an IRQ:

- It is an interrupt for the other Security state, that is, the Security state in which the PE is not executing.
- The PE is executing at EL3.

Table 4-3 summarizes the signaling of interrupts when EL3 is using AArch64 state.

Table 4-3 Interrupt signals for two Security states when EL3 is using AArch64 state

Current Exception level	Group 0 interrupts	Group 1 interrupts	
		Secure	Non-secure
Secure EL1or EL0	FIQ	IRQ	FIQ
Non-secure EL1 or EL0	FIQ	FIQ	IRQ
Non-secure EL2	FIQ	FIQ	IRQ
Secure EL3	FIQ	FIQ	FIQ

Table 4-4 summarizes the signaling of interrupts when EL3 is using AArch32 state.

Table 4-4 Interrupt signals for two Security states when EL3 is using AArch32 state

Current Exception level	Group 0 interrupts	Group 1 interrupts	
		Secure	Non-secure
Secure EL0	FIQ	IRQ	FIQ
Non-secure EL1 or EL0	FIQ	FIQ	IRQ
Non-secure EL2	FIQ	FIQ	IRQ
EL3	FIQ	IRQ	FIQ

Table 4-5 summarizes the signaling of interrupts in systems that support only a single Security state, that is where EL3 is not implemented and GICD_CTLR.DS = 1, and where ICC_SRE_EL1.SRE = 0.

Table 4-5 Interrupt signals for a single Security state

Current Exception level	Group 0 interrupts	Group 1 interrupts
Any	FIQ	IRQ

4.6.3 Interrupt routing and System register access

When executing in AArch64 state, interrupt routing to an Exception level is controlled by the following bits:

- SCR_EL3.FIQ, SCR_EL3.NS, and HCR_EL2.FMO control FIQs.
- SCR_EL3.IRQ, SCR_EL3.NS, and HCR_EL2.IMO control IRQs.

This routing also controls the Exception level at which the EL1 CPU interface System registers that control and acknowledge interrupts are accessible. This applies to:

- ICC_IAR0_EL1, ICC_EOIR0_EL1, ICC_HPPIR0_EL1, ICC_BPR0_EL1, ICC_AP0R<n>_EL1 and ICC_IGRPEN0_EL1. These are the registers that are associated with Group 0 interrupts.
- ICC_IAR1_EL1, ICC_EOIR1_EL1, ICC_HPPIR1_EL1, ICC_BPR1_EL1, ICC_AP1R<n>_EL1 and ICC_IGRPEN1_EL1. These are the registers that are associated with Group 1 interrupts.
- ICC_SGI0R_EL1, ICC_SGI1R_EL1, ICC_ASGI1R_EL1, ICC_CTLR_EL1, ICC_DIR_EL1, ICC_PMR_EL1, and ICC_RPR_EL1. These are the Common registers.

When (SCR_EL3.NS == 1 && (HCR_EL2.FMO ==1 || HCR_EL2.IMO == 1)), accesses at EL1 are virtual accesses. Virtual accesses to ICC_SGI0R_EL1, ICC_SGI1R_EL1, and ICC_ASGI1R_EL1 always generate a Trap exception that is taken to EL2.

Where a Distributor supports two Security states a PE might not implement EL2 or EL3. Table 4-6 shows the configurations that are supported in these cases.

Table 4-6 Supported configurations when EL 3 is not implemented

Distributor	EL3	EL2	Security State	Description
Two Security states and GICD_CTLR.DS == 0	No	-	Non-secure	The PE is always Non-secure and can only receive Non-secure Group 1 interrupts. The PE must behave as if software had: • Set ICC_SRE_EL3.Enable to 1 to allow EL2 to use the System registers, if required. • Set ICC_SRE_EL3.DFB to 1. • Set SCR_EL3.FIQ to 1. • Cleared SCR_EL3.IRQ to 0. • Set SCR_EL3.NS to 1. • Cleared ICC_IGRPEN0_EL1.Enable to 0 to disable the signaling of Group 0 interrupts to the PE. • Set the Secure copy of ICC_IGRPEN1_EL1.Enable to 0 to disable the signaling of Secure Group 1 interrupts to this PE.
Two Security states and GICD_CTLR.DS == 0	No	No	Secure	The PE is always Secure and can only receive Group 0 and Secure Group 1 interrupts. The PE must behave as if software had: • Set ICC_SRE_EL3.Enable to 1. • Cleared SCR_EL3.FIQ to 0. • Cleared SCR_EL3.IRQ to 0. • Cleared SCR_EL3.NS to 0. • Cleared the Non-secure copy of ICC_IGRPEN1_EL1.Enable to 0 to disable the signaling of Non-secure Group 1 interrupts to this PE.
One Security state or two Security states and GICD_CTLR.DS == 1	No	-	-	The Distributor and all PEs are always in a single Security state, and can receive Group 0 and Group 1 interrupts. All PEs must behave as if software had: • Set ICC_SRE_EL3.Enable to 1. • Cleared SCR_EL3.FIQ to 0. • Cleared SCR_EL3.IRQ to 0. • Set SCR_EL3.NS to 1.

4.7 Enabling the distribution of interrupts

The following control bits enable and disable the distribution of interrupts:

- GICD_CTLR.EnableGrp1S.
- GICD_CTLR.EnableGrp1NS.
- GICD_CTLR.EnableGrp0.

The following control bits enable and disable the distribution of interrupt groups at the CPU interface:

- ICC_IGRPEN0_EL1.Enable for Group 0 interrupts.
- ICC_IGRPEN1_EL1.Enable for Group 1 interrupts.

  Note

  There is a Secure and a Non-secure copy of this register.

- ICC_IGRPEN1_EL3.{EnableGrp1S, EnableGrp1NS}.

Physical LPIs are enabled by a write to GICR_CTLR.EnableLPIs.

4.7.1 Enabling individual interrupts

PPIs

Individual PPIs can be enabled and disabled by writing to GICD_ISENABLER<n> and GICD_ICENABLER<n>. $n = 0$ for PPIs, if legacy operation for physical interrupts is supported and configured. PPIs can also be enabled and disabled by writing to GICR_ISENABLER0 and GICR_ICENABLER0 when affinity routing is enabled for the Security state of the interrupt.

SPIs

Individual SPIs can be enabled and disabled by writing to GICD_ISENABLER<n> and GICD_ICENABLER<n>. $n > 0$ for SPIs.

SGIs

Individual SGIs can be enabled and disabled by writing to GICD_ISENABLER<n> and GICD_ICENABLER<n>. $n = 0$ for SGIs, if legacy operation for physical interrupts is supported and configured. SGIs can also be enabled and disabled by writing to GICR_ISENABLER0 and GICR_ICENABLER0 when affinity routing is enabled.

Note

Whether SGIs are permanently enabled, or can be enabled and disabled by writes to GICR_ISENABLER0 and GICR_ICENABLER0, is IMPLEMENTATION DEFINED.

LPIs

Individual LPIs can be enabled by setting the enable bits programmed in the LPI Configuration table. For more information about enabling LPIs using the LPI Configuration tables, see *LPI Configuration tables* on page 6-95.

4.7.2 Interaction of group and individual interrupt enables

The GICD_* and GICR_* registers determine, at any moment in time, the highest priority pending interrupt that the hardware is aware of for each target PE. This interrupt is presented to the CPU interface of a PE to evaluate whether it is to be signaled to the PE. The enabling of the interrupts affects this evaluation as follows:

- A pending interrupt that is individually disabled in the GICD_* or GICR_* registers is not one which is considered in the determination of the highest priority pending interrupt, and so cannot be signaled to the PE.
- A pending interrupt that is individually enabled in the GICD_* registers, but is a member of a group that is disabled in GICD_CTLR, is not one that is considered in the determination of the highest priority pending interrupt, and so cannot be signaled to the PE.

- A pending 1 of N interrupt that is individually enabled in the GICD_* registers and is a member of a group that is enabled in GICD_CTLR, but is a member of a group that is disabled in GICC_CTLR for a PE, cannot be selected for that PE. Such an interrupt is not considered in the determination of the highest priority pending interrupt and so cannot be signaled to the PE.
- For a pending direct interrupt that is individually enabled in the GICD_* registers and is a member of a group that is enabled in GICD_CTLR, but is a member of a group that is disabled in GICC_CTLR, it is IMPLEMENTATION DEFINED whether or not the interrupt is considered in the determination of the highest priority pending interrupt. If it is determined to be the highest priority pending interrupt, the interrupt is not signaled to the PE, but will mask a lower priority pending interrupt that is a member of a group that is enabled in GICC_CTLR.

4.7.3 Effect of disabling interrupts

Disabling an interrupt by writing to the appropriate GICD_ICENABLER<n> or to GICR_ICENABLER0, or by writing to the LPI Configuration tables, does not prevent that interrupt from changing state, for example from becoming pending. When GICR_CTLR.EnableLPIs == 0, LPIs are never set pending.

If GICD_CTLR.EnableGrp0, GICD_CTLR.EnableGrp1S, and GICD_CTLR.EnableGrp1NS are all cleared to 0, it is IMPLEMENTATION DEFINED whether:

- An edge-triggered interrupt signal moves the interrupt to the pending state.
- SGIs can be set pending by writing to GICD_SGIR, ICC_SGI0R_EL1, ICC_SGI1R_EL1, or ICC_ASGI1R_EL1.

If an interrupt is pending on a CPU interface when the corresponding GICD_CTLR.EnableGrp0, GICD_CTLR.EnableGrp1NS, or GICD_CTLR.EnableGrp1S bit is written from 1 to 0, then the interrupt must be retrieved from the CPU interface.

Note

This might have no effect on the forwarded interrupt if it has already been activated.

If an interrupt is pending on a CPU interface when software writes ICC_IGRPEN0_EL1.Enable, ICC_IGRPEN0_EL1, ICC_IGRPEN1_EL1.Enable, or ICC_IGRPEN1_EL3.Enable from 1 to 0, the interrupt must be released by the CPU interface to allow the Distributor to forward the interrupt to a different PE.

4.8 Interrupt prioritization

This section describes interrupt prioritization in the GIC architecture. Prioritization describes the:

- Configuration and control of interrupt priority.
- Order of execution of pending interrupts.
- Determination of when interrupts are visible to a target PE, including:
  - Interrupt priority masking.
  - Priority grouping.
  - Preemption of an active interrupt.

Software configures interrupt prioritization in the GIC by assigning a priority value to each interrupt source. Priority values are an 8-bit unsigned binary number. A GIC implementation that supports two Security states must implement a minimum of 32 and a maximum of 256 levels of physical priority. A GIC implementation that supports only a single Security state must implement a minimum of 16 and a maximum of 256 levels of physical priority. If the GIC implements fewer than 256 priority levels, the low-order bits of the priority fields are RAZ/WI. This means that the number of implemented priority field bits is IMPLEMENTATION DEFINED, in the range 4-8. Table 4-7 shows the relation between the priority field bits and the number of priority levels supported by an implementation.

Table 4-7 Effect of not implementing some priority field bits

Implemented priority bits	Possible priority field values	Number of priority levels
[7:0]	`0x00-0xFF` (0-255), all values	256
[7:1]	`0x00-0xFE`, (0-254), even values only	128
[7:2]	`0x00-0xFC` (0-252), in steps of 4	64
[7:3]	`0x00-0xF8` (0-248), in steps of 8	32
[7:4]	`0x00-0xF0` (0-240), in steps of 16	16

In the GIC prioritization scheme, lower numbers have higher priority. This means that the lower the assigned priority value, the higher the priority of the interrupt. Priority field value 0 always indicates the highest possible interrupt priority, and the lowest priority value depends on the number of implemented priority levels.

The GICD_IPRIORITYR<n> registers hold the priority value for each supported SPI. An implementation might reserve an SPI for a particular purpose and assign a fixed priority to that interrupt, meaning the priority value for that interrupt is read-only. For other SPIs the GICD_IPRIORITYR<n> registers can be written by software to set the interrupt priorities. It is IMPLEMENTATION DEFINED whether a write to GICD_IPRIORITYR<n> changes the priority of any active SPI.

In a multiprocessor implementation, the GICR_IPRIORITYR<n> registers define the interrupt priority of each SGI and PPI INTID independently for each target PE. The order in which the CPU interface serializes these SGIs is implementation specific.

LPI Configuration tables and LPI Pending tables in memory store LPI priority information and pending status, see *LPI Configuration tables* on page 6-95 and *LPI Pending tables* on page 6-97.

The GIC security model provides Secure and Non-secure accesses to the interrupt priority settings.The Non-secure accesses can configure interrupts only in the lower priority half of the supported priority values. Therefore, if the GIC implements 32 priority values, Non-secure accesses see only 16 priority values. See *Software accesses of interrupt priority* on page 4-72 for more information.

To determine the number of priority bits implemented for SPIs, software can write `0xFF` to a writable GICD_IPRIORITYR<n> priority field and read back the value stored.

To determine the number of priority bits implemented for SGIs and PPIs, software can write `0xFF` to the GICR_IPRIORITYR<n> priority fields, and read back the value stored.

The GIC architecture does not require all PEs in the system to use the same number of priority bits to control interrupt priority.

In a multiprocessor implementation, ICC_CTLR_EL1.PRIbits and ICC_CTLR_EL3.PRIbits indicate the number of priority bits implemented, independently for each target PE.

Note

ARM recommends that implementations support the same number of priority bits for each PE.

For information about the priority range supported for virtual interrupts, see Chapter 5 *Virtual Interrupt Handling and Prioritization*.

Note

ARM recommends that, before checking the priority range in this way:

- For a peripheral interrupt, software first disables the interrupt.
- For an SGI, software first checks that the interrupt is inactive.

If, on a particular CPU interface, multiple pending interrupts have the same priority, and have sufficient priority for the interface to signal them to the PE, it is implementation specific how the interface selects which interrupt to signal.

The remainder of this section describes:

- *Non-secure accesses to register fields for Secure interrupt priorities*.
- *Priority grouping* on page 4-67.
- *System register access to the Active Priorities registers* on page 4-69.
- *Preemption* on page 4-71.
- *Priority masking* on page 4-72.
- *Software accesses of interrupt priority* on page 4-72.
- *Changing the priority of enabled PPIs, SGIs, and SPIs* on page 4-76.

4.8.1 Non-secure accesses to register fields for Secure interrupt priorities

A GIC that supports two Security states supports the use of:

- Group 0 interrupts as Secure interrupts.
- Secure Group 1 interrupts.
- Non-secure Group 1 interrupts.

In order to support the ARMv8 Security model the register fields associated with Secure interrupts are RAZ/WI for Non-secure accesses. In addition, the following rules apply:

For Non-secure access to a priority field in GICx_IPRIORITYR<n>:

If the priority field corresponds to a Non-secure Group 1 interrupt in *Software accesses of interrupt priority* on page 4-72:

- For SPIs, the priority field is determined by GICD_IPRIORITYR<n>.
- For PPIs and SGIs, the priority field is determined by GICR_IPRIORITYR<n>.

For Non-secure access to ICC_PMR_EL1 and ICC_RPR_EL1:

- If the current priority mask value is in the range of `0x00-0x7F`:
  - A read access returns the value `0x00`.
  - The GIC ignores a write access to ICC_PMR_EL1.
- If the current priority mask value is in the range of `0x80-0xFF`:
  - A read access returns the Non-secure read of the current value.
  - A write access to ICC_PMR_EL1 succeeds, based on the Non-secure read of the priority mask value written to the register.

Note

This means a Non-secure write cannot set a priority mask value in the range of `0x00-0x7F`.

The register descriptions for the following registers provide pseudocode that describes accesses to the registers in a GIC that supports two Security states:

- GICD_IPRIORITYR<n>.
- GICR_IPRIORITYR<n>.
- ICC_PMR_EL1.
- ICC_RPR_EL1.

4.8.2 Priority grouping

Priority grouping uses the following *Binary Point Registers*:

- ICC_BPR0_EL1 for Group 0 interrupts. This register is available in all GIC implementations.
- A Non-secure copy of ICC_BPR1_EL1 for Non-secure Group 1 interrupts. If an implementation supports two Security states, there is a Secure and a Non-secure copy of this register. If an implementation supports only one Security state, there is only one copy of this register
- A Secure copy of ICC_BPR1_EL1 for Secure Group 1 interrupts. This register is available only in a GIC implementation that supports two Security states.

The Binary Point Registers split a priority value into two fields, the *group priority* and the *subpriority*. When determining preemption, all interrupts with the same group priority are considered to have the same priority, regardless of the subpriority.

Where multiple pending interrupts have the same group priority, the GIC uses the subpriority field to resolve the priority within a group. Where two or more pending interrupts in a group have the same subpriority, how the GIC selects between the interrupts is implementation specific.

The GIC uses the group priority field to determine whether a pending interrupt has sufficient priority to preempt execution on a PE, as follows:

- The value of the group priority field for the interrupt must be lower than the value of the running priority of the PE. The running priority is the group priority of the highest priority active interrupt that has not received a priority drop on that PE.
- The value of the priority for the interrupt must be lower than the value of its priority mask.

ICC_BPR0_EL1 determines the priority grouping of Group 0 interrupts:

- When ICC_CTLR_EL3.CBPR_EL1S is set to 1, ICC_BPR0_EL1 also determines the priority grouping of Secure Group 1 interrupts.
- When ICC_CTLR_EL3.CBPR_EL1NS is set to 1, ICC_BPR0_EL1 also determines the priority grouping of Non-secure Group 1 interrupts

ICC_BPR1_EL1 determines the priority of Group 1 interrupts:

- When ICC_CTLR_EL3.CBPR_EL1S is cleared to 0, the Secure copy of ICC_BPR1_EL1 determines the priority grouping of Secure Group 1 interrupts.
- When ICC_CTLR_EL3.CBPR_EL1NS is cleared to 0, the Non-secure copy of ICC_BPR1_EL1 determines the priority grouping of Non-secure Group 1 interrupts.

Table 4-8 shows the priority grouping for Group 1 interrupts when ICC_CTLR_EL3.CBPR_EL1S or ICC_CTLR_EL3.CBPR_EL1NS is cleared to 0.

Table 4-8 Priority grouping for Group 1 interrupts when CBPR is cleared to 0

Value of Binary point field in ICC_BPR1_EL1	Interrupt priority field [7:0]		
	Group priority field	Subpriority field	Field with binary point[a]
0[b]	-	-	-
1	[7:1][c]	[0]	ggggggg.s
2	[7:2][c]	[1:0]	gggggg.ss
3	[7:3][c]	[2:0]	ggggg.sss
4	[7:4][c]	[3:0]	gggg.ssss
5	[7:5][c]	[4:0]	ggg.sssss
6	[7:6][c]	[5:0]	gg.ssssss
7	[7][c]	[6:0]	g.sssssss

a. Group labeling aligns with that shown in Figure 4-8 on page 4-74.
b. Not supported.
c. If a Non-secure write sets the priority value field for a Non-secure interrupt then bit[7] == 1.

Table 4-9 shows the priority grouping for Group 0 interrupts, or for Group 1 interrupts when ICC_CTLR_EL3.CBPR_EL1S or ICC_CTLR_EL3.CBPR_EL1NS is set to 1.

Table 4-9 Priority grouping for Group 1 interrupts when CBPR == 1, or for Group 0 interrupts

Binary point value ICC_BPR0_EL1	Interrupt priority field [7:0]		
	Group priority field	Subpriority field	Field with binary point
0	[7:1][a]	[0]	ggggggg.s
1	[7:2][a]	[1:0]	gggggg.ss
2	[7:3][a]	[2:0]	ggggg.sss
3	[7:4][a]	[3:0]	gggg.ssss
4	[7:5][a]	[4:0]	ggg.sssss
5	[7:6][a]	[5:0]	gg.ssssss
6	[7][a]	[6:0]	g.sssssss
7	No preemption	[7:0]	.sssssss

a. If a Non-secure write sets the priority value field for a Non-secure interrupt then bit[7] == 1.

The minimum binary point value supported depends on the IMPLEMENTATION DEFINED number of priority bits, and is in the range 0 - 3. The number of priority bits implemented is indicated by ICC_CTLR_EL1.PRIBits and ICC_CTLR_EL3.PRIBits.

In a GIC that supports two Security states, when:

- ICC_CTLR_EL3.CBPR_EL1S == 1:
  - Writes to ICC_BPR1_EL1 at Secure EL1 modify ICC_BPR0_EL1.
  - Reads from ICC_BPR1_EL1 at Secure EL1 return the value of ICC_BPR0_EL1.
- ICC_CTLR_EL3.CBPR_EL1NS == 1:
  - Non-secure writes to ICC_BPR1_EL1 modify ICC_BPR0_EL1.
  - Non-secure reads from ICC_BPR1_EL1 return the value of ICC_BPR0_EL1.

Note

- When an interrupt is using ICC_BPR1_EL1, the effective binary point value is that stored in the register, minus one, as shown in Table 4-8 on page 4-68. This means that software with no awareness of the effects of interrupt grouping and where two Security states are supported, sees the same priority grouping mechanism, regardless of whether it is running on a PE that is in Secure state or in Non-secure state.
- Priority grouping always operates on the full priority, which is the value that would be visible to a Secure read. This is different from the value that is visible to a Non-secure read of the priority value corresponding to a Non-secure interrupt. See Figure 4-8 on page 4-74 and Figure 4-9 on page 4-74.

Pseudocode

The following pseudocode indicates the group priority of the interrupt.

```
// GroupBits()
// ===========
// Returns the priority group field for the minimum BPR value for the group

bits(8) GroupBits(bits(8) priority, IntGroup group)
    bit cbpr_G1NS = if HaveEL(EL3) then ICC_CTLR_EL3.CBPR_EL1NS else ICC_CTLR_EL1.CBPR;
    bit cbpr_G1S  = if HaveEL(EL3) then ICC_CTLR_EL3.CBPR_EL1S  else '0';

    if (group == IntGroup_G0 ||
        (group == IntGroup_G1NS && cbpr_G1NS == '1') ||
        (group == IntGroup_G1S  && cbpr_G1S  == '1')) then
        bpr = UInt(ICC_BPR0_EL1.BinaryPoint);
    elsif group == IntGroup_G1S then
        bpr = UInt(ICC_BPR1_EL1S.BinaryPoint);
    else
        bpr = UInt(ICC_BPR1_EL1NS.BinaryPoint);

    mask = Ones(bpr):Zeros(8 - bpr);

    return priority AND mask;
```

4.8.3 System register access to the Active Priorities registers

Physical Group 0 and Group 1 interrupts access different Active Priorities registers, depending on the interrupt group.

For Group 0 interrupts, these registers are ICC_AP0R<n>_EL1, where n = 0-3:

- If 32 or fewer priority levels are implemented, accesses to ICC_AP0R<n>_EL1, where n = 1-3, are UNDEFINED.
- If more than 32 and fewer than 65 priority levels are implemented, accesses to ICC_AP0R<n>_EL1, where n = 2-3, are UNDEFINED.

For Group 1 interrupts, these registers are ICC_AP1R<n>_EL1, where n= 0-3:

- If 32 or fewer priority levels are implemented, accesses to ICC_AP1R<n>_EL1, where n = 1-3, are UNDEFINED.

- If more than 32 and fewer than 65 priority levels are implemented, accesses to ICC_AP1R<n>_EL1, where n = 2-3, are UNDEFINED.

The content of ICC_AP0R<n>_EL1, Secure ICC_AP1R<n>_EL1, and Non-secure ICC_AP1R<n>_EL1 is IMPLEMENTATION DEFINED. However, the value `0x00000000` must be consistent with no priorities being active.

Writing any value other than the last read value, or `0x00000000`, to these registers can cause:

- Interrupts that would otherwise preempt execution to not preempt execution.
- Interrupts that otherwise would not preempt execution to preempt execution.

Writing any value to Non-secure ICC_AP1R<n>_EL1 cannot prevent the correct prioritization and the forwarding of interrupts of higher priority than those in the Non-secure priority range, meaning that this does not create a security hole.

Writes to these registers in any order other than the following can result in UNPREDICTABLE behavior:

1. ICC_AP0R<n>_EL1.
2. Secure ICC_AP1R<n>_EL1.
3. Non-secure ICC_AP1R<n>_EL1.

Note

An `ISB` is not required between each write to ICC_AP0R<n>_EL1, Secure ICC_AP1R<n>_EL1, and Non-secure ICC_AP1R<n>_EL1.

Table 4-10 shows an implementation of ICC_AP0R<n>_EL1.

Table 4-10 Group 0 Active Priorities Register implementation

Minimum value of:		Maximum number of:		
Secure ICC_BPR0_EL1	Non-secure ICC_BPR1_EL1	Group priority bits	Preemption levels	ICC_AP0Rn implementation
3	4	4	16	ICC_AP0R<n>_EL1[15:0], where n = 0
2	3	5	32	ICC_AP0R<n>_EL1[31:0], where n = 0
1	2	6	64	ICC_AP0R<n>_EL1, where n = 0-1
0	1	7	128	ICC_AP0R<n>_EL1, where n = 0-3

Table 4-11 shows an implementation of ICC_AP1R<n>_EL1

Table 4-11 Group 1 Active Priorities Register implementation

Minimum value of:		Maximum number of:		
Secure ICC_BPR0_EL1	Non-secure ICC_BPR1_EL1	Group priority bits	Preemption levels	ICC_AP1Rn implementation
3	4	4	16	ICC_AP1R<n>_EL1[15:0], where n = 0
2	3	5	32	ICC_AP1R<n>_EL1[31:0], where n = 0
1	2	6	64	ICC_AP1R<n>_EL1, where n = 0-1
0	1	7	128	ICC_AP1R<n>_EL1, where n = 0-3

Pseudocode

The following pseudocode indicates the highest active group priority.

```
// GetHighestActiveGroup()
// =======================
// Returns a value indicating the interrupt group of the highest active priority from three
// registers. Returns IntGroup_None if no active priorities.
// Note: having more than one group active at the same priority is UNPREDICTABLE.

IntGroup GetHighestActiveGroup(bits(128) ap0, bits(128) ap1ns, bits(128) ap1s)
```

The following pseudocode indicates the highest active priority.

```
// GetHighestActivePriority()
// ==========================
// Returns the priority of the highest active priority from three registers, expressed as a 7-bit
// unsigned binary number. Returns 0xFF if no bits are active.

bits(8) GetHighestActivePriority(bits(128) ap0, bits(128) ap1ns, bits(128) ap1s)
```

4.8.4 Preemption

A CPU interface supports signaling of higher priority pending interrupts to a target PE before an active interrupt completes. A pending interrupt is only signaled if both:

- Its priority is higher than the priority mask for that CPU interface. See *Priority masking* on page 4-72.
- Its group priority is higher than that of the running priority on the CPU interface. See *Priority grouping* on page 4-67 for more information.

Preemption occurs at the time when the PE takes the new interrupt, and starts handling the new interrupt instead of the previously active interrupt or the currently running process. When this occurs, the initial active interrupt is said to have been *preempted*.

Note

The value of the I or F bit in the Process State, PSTATE, and the Exception level and the interrupt routing controls in software and hardware, determine whether the PE responds to the signaled interrupt by taking the interrupt. For more information, see *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

For more information about enabling interrupts, see *Enabling the distribution of interrupts* on page 4-63.

Preemption level control

ICC_BPR0_EL1 determines whether a Group 0 interrupt is signaled to the PE for possible preemption. In addition:

- When ICC_CTLR_EL3.CBPR_EL1NS == 1, ICC_BPR0_EL1 also determines whether a Non-secure Group 1 interrupt is signaled to the PE for possible preemption.
- When ICC_CTLR_EL3.CBPR_EL1S == 1, ICC_BPR0_EL1 also determines whether a Secure Group 1 interrupt is signaled to the PE for possible preemption.

ICC_BPR1_EL1 determines whether a Group 1 interrupt is signaled to the PE for possible preemption. The Non-secure copy of this register is used for Non-secure Group 1 interrupts. The Secure copy is used for Secure Group 1 interrupts.

When ICC_CTLR_EL3.CBPR_EL1NS is set to 1:

- EL3 can write to ICC_BPR1_EL1(NS).

  When EL3 is using AArch64 state, accesses to ICC_BPR1_EL1(NS) from EL3 are not affected by ICC_CTLR_EL3.CBPR_EL1NS.

  When EL3 is using AArch32 state, accesses to ICC_BPR1_EL1(NS) from Monitor mode are not affected by ICC_CTLR_EL3.CBPR_EL1NS.

- Non-secure writes to ICC_BPR1_EL1 at EL1 or EL2 are ignored.
- Non-secure reads of ICC_BPR1_EL1 at EL1 or EL2 return the value of ICC_BPR0_EL1 +1, saturating at 7.

When ICC_CTLR_EL3.CBPR_EL1S is set to 1:

- Secure reads of ICC_BPR1_EL1 return the value of ICC_BPR0_EL1.
- Secure writes to ICC_BPR1_EL1 update ICC_BPR0_EL1.

4.8.5 Priority masking

The *Priority Mask Register* for a CPU interface, ICC_PMR_EL1, defines a priority threshold for the target PE. The GIC only signals pending interrupts that have a higher priority than this priority threshold to the target PE. A value of zero, the register reset value, masks all interrupts from being signaled to the associated PE. The GIC does not use priority grouping when comparing the priority of a pending interrupt with the priority threshold.

The GIC always masks an interrupt that has the lowest supported priority. This priority is sometimes referred to as the idle priority.

Note

Writing `0xFF` to ICC_PMR_EL1 always sets it to the lowest supported priority. Table 4-7 on page 4-65 shows how the lowest supported priority varies with the number of implemented priority bits.

If the GIC provides support for two Security states, ICC_PMR_EL1 is RAZ/WI to Non-secure accesses, if bit[7] == 0. During normal operation, software executing in Non-secure state does not access ICC_PMR_EL1when it is programmed with such a value.

For information that relates to different GIC configurations, see *Non-secure accesses to register fields for Secure interrupt priorities* on page 4-66.

4.8.6 Software accesses of interrupt priority

This section describes Secure and Non-secure reads of interrupt priority, and the relationship between them. It also describes writes to the priority value fields.

Note

This section applies to any GIC implementation that supports two Security states.

When a PE reads the priority value of a Non-secure Group 1 interrupt, the GIC returns either the Secure or the Non-secure read of that value, depending on whether the access is Secure or Non-secure.

The GIC implements a minimum of 32 and a maximum of 256 priority levels. This means it implements 5-8 bits of the 8-bit priority value fields in the appropriate GICR_IPRIORITYR<n> and GICD_IPRIORITYR<n> register. All of the implemented priority bits can be accessed by a Secure access, and unimplemented low-order bits of the priority fields are RAZ/WI. Figure 4-5 shows the Secure read of a priority value field for an interrupt. The priority value stored in the Distributor is equivalent to the Secure read.

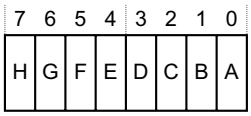

Figure 4-5 Secure read of the priority field for any interrupt

In this view:

- Bits H-D are the bits that the GIC must implement, corresponding to 32 priority levels.
- Bits C-A are the bits the GIC might implement. They are RAZ/WI if not implemented.
- The GIC must implement bits H-A to provide the maximum 256 priority levels.

For Non-secure accesses, the GIC supports half the priority levels it supports for Secure accesses, which means a minimum of 16 priority levels. Figure 4-6 shows the Non-secure view of a priority value field for a Non-secure Group 1 interrupt.

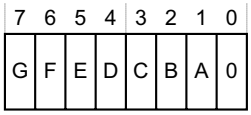

Figure 4-6 Non-secure read of the priority field for a Non-secure Group 1 interrupt

In this read:

- Bits G-D are the bits that the GIC must implement, corresponding to 16 priority levels.
- Bits C-A are the bits the GIC might implement, that are RAZ/WI if not implemented.
- The GIC must implement bits G-A to provide the maximum 128 priority levels.
- Bit [0] is RAZ/WI.

The Non-secure read of a priority value does not show how the value is stored in the registers in the Distributor. For Non-secure writes to a priority field of a Non-secure Group 1 interrupt, before storing the value:

- The value is right-shifted by one bit.
- Bit [7] of the value is set to 1.

This translation means the priority value for the Non-secure Group 1 interrupt is in the bottom half of the priority range.

A Secure read of the priority value for an interrupt returns the value stored in the Distributor. Figure 4-7 shows this Secure read of the priority value field for a Non-secure Group 1 interrupt that has had its priority value field set by a Non-secure access, or has had a priority value with bit[7] == 1 set by a Secure access:

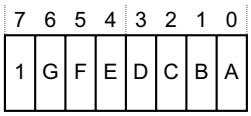

Figure 4-7 Secure read of the priority field for a Non-secure Group 1 interrupt

A Secure write to the priority value field for a Non-secure Group 1 interrupt can set bit [7] to 0. If a Secure write sets bit[7] to 0:

- A Non-secure read returns the value GFEDCBA0.
- A Non-secure write can change the value of the field, but only to a value that has bit [7] set to 1 for the Secure read of the field.

Note

- This behavior of Non-secure accesses applies only to the priority value fields in GICR_IPRIORITYR<n> and GICD_IPRIORITYR<n>, as appropriate:
  - If the Priority field in ICC_PMR_EL1 holds a value with bit [7] == 0, then the field is RAZ/WI for Non-secure accesses.
  - If the Priority field in ICC_RPR_EL1 holds a value with bit [7] == 0, then the field is RAZ for Non-secure reads.
- ARM does not recommend setting bit[7] to 0 for a Non-secure Group 1 interrupt in this way because it places the interrupt in the wrong half of the priority range for maintenance by non-secure code.

Figure 4-8 on page 4-74 shows the relationship between the reads of the priority value fields for Non-secure Group 1 interrupts.

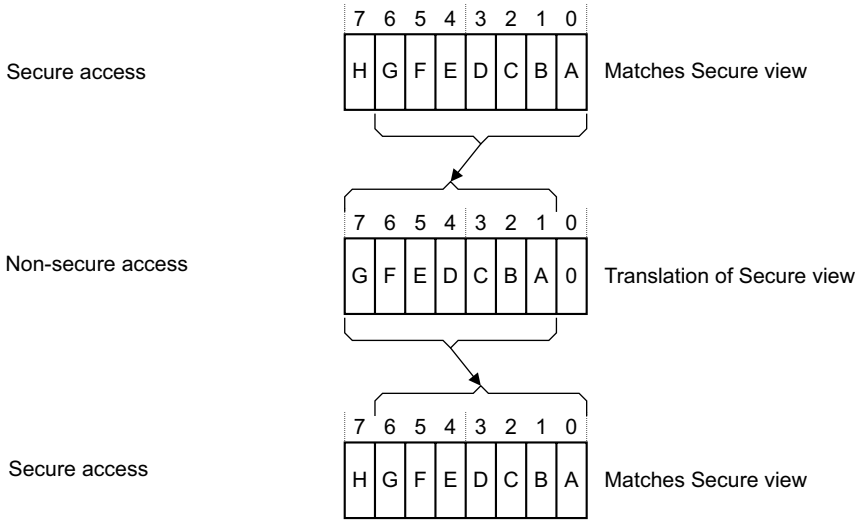

Figure 4-8 Relationship between Secure and Non-secure reads of interrupt priority fields

Figure 4-9 shows how software reads of the interrupt priorities from Secure and Non-secure accesses relate to the priority values held in the Distributor, and to the interrupt values that are visible to Secure and Non-secure accesses. Figure 4-9 applies to a GIC that implements the maximum range of priority values.

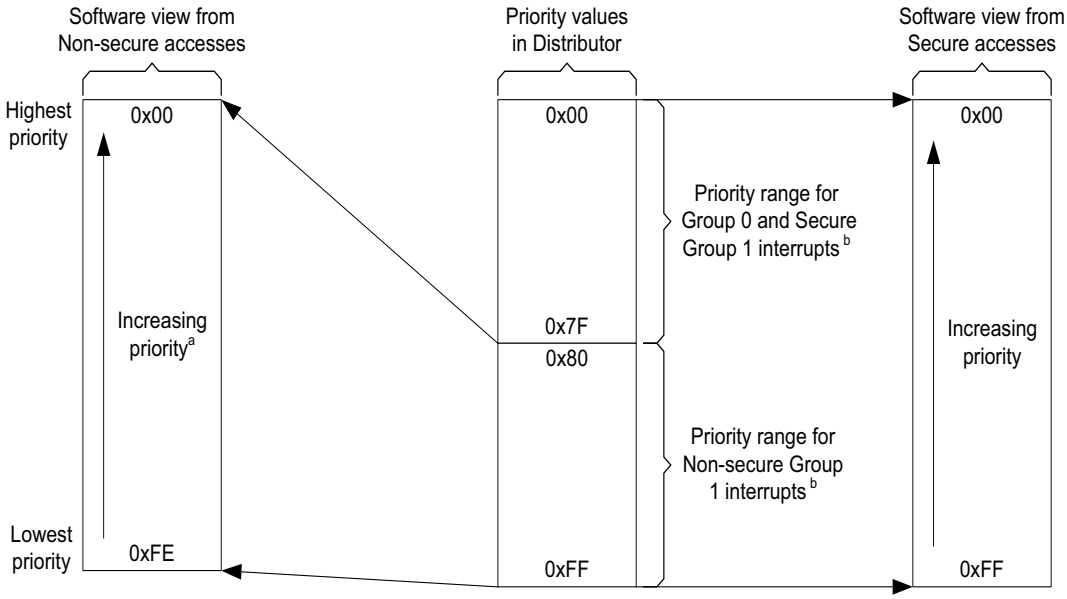

a. All priority values are even (bit [0] == 0) in the software view of Non-secure accesses.

b. Ranges recommended by ARM.

Figure 4-9 Software reads of the priorities of Group 1 and Group 0 interrupts

Table 4-12 on page 4-75 shows how the number of priority value bits implemented by the GIC affects the Secure and Non-secure reads of the priority of a Non-secure Group 1 interrupt.

Note

Software executing in Non-secure state has no visibility of the priority settings of Group 0 interrupts, or where applicable, of Secure Group 1 interrupts.

Table 4-12 Effect of not implementing some priority field bits, two Security states

Implemented priority bits, as seen by a Secure read	Possible priority field values, for a Non-secure Group 1 interrupt	
	Secure read	Non-secure read
[7:0]	0xFF-0x00 (255-0), all values	0xFE-0x00 (254-0), even values only
[7:1]	0xFE-0x00 (254-0), even values only	0xFC-0x00 (252-0), in steps of 4
[7:2]	0xFC-0x00 (252-0), in steps of 4	0xF8-0x00 (248-0), in steps of 8
[7:3]	0xF8-0x00 (248-0), in steps of 8	0xF0-0x00 (240-0), in steps of 16

This model for the presentation of priority values ensures software written to operate with an implementation of this GIC architecture functions as intended regardless of whether the GIC provides support for two Security states. However, programmers must ensure that software assigns the appropriate priority levels to the Group 0 and Group 1 interrupts.

Note

To control priority values, ARM strongly recommends that:

- For a Group 0 interrupt, software sets bit[7] of the priority value field to 0.
- If using a Secure write to set the priority of a Non-secure Group 1 interrupt, software sets bit[7] of the priority value field to 1.

This ensures that all Group 0 and, if applicable, Secure Group 1 interrupts have higher priorities than all Non-secure Group 1 interrupts. However, a system might have requirements that cannot be met with this scheme.

Table 4-13 shows an example priority allocation scheme that ensures:

- Some Group 0 interrupts have higher priority than any other interrupts.
- Some Secure Group 1 interrupts have higher priority than any Non-secure Group 1 interrupt.

Table 4-13 Example priority allocation

Interrupt security configuration	GICR_IPRIORITYR<n>[7:6]
Group 0	0b00
Secure Group 1	0b01
Non-secure Group 1	0b10 0b11

- Software might not be aware that the GIC supports two Security states, and therefore might not know whether it is making Secure or Non-secure accesses to GIC registers. However, for any implemented interrupt, software can write 0xFF to the corresponding GICR_IPRIORITYR<n> priority value field, and then read back the value stored in the field to determine the supported interrupt priority range. ARM recommends that, before checking the priority range in this way:
  - For a peripheral interrupt, software first disables the interrupt.
  - For an SGI, software first checks that the interrupt is inactive.

4.8.7 Changing the priority of enabled PPIs, SGIs, and SPIs

If software writes to the GICD_IPRIORITYR<n> or GICR_IPRIORITYR<n> register of an enabled interrupt while the interrupt is pending, it is IMPLEMENTATION DEFINED whether the GIC uses the old value or the new value. The GIC ensures that no interrupt is handled more than once, and that no interrupt is lost. The effect of the write must be visible in finite time.

Chapter 5
Virtual Interrupt Handling and Prioritization

This chapter describes the fundamental aspects of GIC virtual interrupt handling and prioritization:

5.1 About GIC support for virtualization

An operating system that is executing at EL1 under the control of a hypervisor executing at EL2 is sometimes referred as a *virtual machine* (VM). A VM can support multiprocessing, which means that multiple *virtual PEs* (vPEs), that are scheduled by the hypervisor are executing on one or more physical PEs. When a vPE is executing on a PE, that vPE of the VM is referred to as being scheduled on the physical PE. In ARMv8, when EL2 is implemented and enabled, the GIC CPU interface provides mechanisms to minimize the hypervisor overhead of routing interrupts to a VM. For more information about vPEs, see *Operation overview on page 5-79*.

For more information about EL2 and virtual interrupts, see *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Note

The GIC does not provide additional mechanisms for the virtualization of the GICD_*, GICR_*, and GITS_* registers. To virtualize VM accesses to these registers, the hypervisor must set stage 2 data aborts to those memory locations so that the hypervisor can emulate these effects. For more information about stage 2 data aborts, see *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

When a GIC provides support for virtualization, the VM operates in an environment that has the following features:

- The vPE can be configured to receive virtual Group 0 interrupts.
- The vPE can be configured to receive virtual Group 1 interrupts.
- Virtual Group 0 interrupts are signaled using the virtual FIQ signal to Non-secure EL1.
- Virtual Group 1 interrupts are signaled using the virtual IRQ signal to Non-secure EL1.
- Virtual interrupts can be handled by the vPE as if they were physical interrupts.

Note

This applies when affinity routing and System register access are enabled. For information about support for virtual interrupts in legacy operation, see *Support for legacy operation of VMs on page 10-689*.

EL2 controls the generation of virtual interrupts for a VM. This allows software executing at EL2 to:

- Generate virtual Group 0 and Group 1 interrupts for the vPE.
- Save and restore the interrupt state of the vPE.
- Control the prioritization of the virtual interrupts.
- Change the vPE that is scheduled.

GICv4 introduces the ability to present virtual LPIs from an ITS directly to a vPE, without hypervisor intervention.

Handling virtual interrupts in legacy operation requires a GICV_* memory-mapped interface. See *Support for legacy operation of VMs on page 10-689* for more information.

5.2 Operation overview

GICv3 supports the ARMv8-A virtualization functionality. A hypervisor executing at EL2 uses the ICH_* System register interface to configure and control a virtual PE (vPE) executing at Non-secure EL1. For information about the VM control interface, see *Configuration and control of VMs on page 5-83*. A vPE uses the ICC_*_EL1 System register interface to communicate with the GIC. The configuration of HCR_EL2.{IMO, FMO} determines whether the virtual or the physical interface registers are accessed.

Note

This chapter describes the handling of virtual interrupts in the context of the AArch64 execution state with System register access enabled. The individual AArch64 System register descriptions that are cross-referenced in this chapter contain a reference to the AArch32 System register that provides the same functionality. For information about VMs in legacy operation, see *Support for legacy operation of VMs on page 10-689*.

Software executing at EL3 or EL2 configures the PE to route physical interrupts to EL2. The interrupt can be:

- An interrupt targeting a vPE. The hypervisor sets the corresponding virtual INTID to the pending state on the target vPE and includes the information about the associated physical INTID. When the vPE is not scheduled on a PE, the hypervisor might choose to reschedule the vPE. Otherwise the interrupt is left pending on the vPE for scheduling by the hypervisor at a later time.
- An interrupt targeting the hypervisor. This interrupt might:
  - Have been generated by the system.
  - Be a maintenance interrupt associated with a scheduled VM. See *Maintenance interrupts on page 5-85* for more details.
  - In GICv4, be a doorbell interrupt from an ITS. In GICv4, a virtual interrupt can be presented to a vPE without hypervisor involvement. A doorbell interrupt must be generated when a virtual interrupt is made pending for a vPE but the vPE is not scheduled on a PE.

The hypervisor handles physical interrupts according to the rules described in Chapter 4 *Physical Interrupt Handling and Prioritization* before they are virtualized. For information about the handling of physical interrupts and their virtualization during legacy operation, see Chapter 10 *Legacy Operation and Asymmetric Configurations*.

The GIC virtualization support includes a list of virtual interrupts for a vPE that is stored in hardware List registers, see *Usage model for the List registers on page 5-81*. Each entry in the list corresponds to either a pending or an active interrupt, and the entry describes the virtual interrupt number, the interrupt group, the interrupt state, and the virtual priority of the interrupt. A virtual interrupt described in the list entry can be configured to be associated with a physical SPI or PPI.

The GIC implementation selects the highest priority pending virtual interrupt from the list of interrupts held in the List registers and, if it is of sufficient virtual priority compared to the active virtual interrupts and virtual priority mask, presents it as either a virtual FIQ or a virtual IRQ, depending on the group of the interrupt. The virtual CPU interface controls apply to the virtual interrupt in the same way as the physical interrupt controls apply to the physical interrupt. Therefore, using the virtual CPU interface controls, software executing on the vPE can:

- Set the virtual priority mask.
- Control how the virtual priority is split between the group priority and the subpriority.
- Acknowledge a virtual interrupt.
- Perform a priority drop on the virtual interrupt.
- Deactivate the virtual interrupt.

The virtual CPU interface supports both EOImodes, so that a virtual EOI can perform a priority drop alone, or a combined priority drop and deactivation.

When a virtual interrupt is acknowledged, then the state of the virtual interrupt changes from pending to active in the corresponding List register entry.

When a virtual interrupt is deactivated, then the state of the virtual interrupt changes from active to inactive, or from active and pending to pending, in the corresponding List register entry. If the virtual interrupt is associated with a physical interrupt, then the associated physical interrupt is deactivated.

Virtual interrupts taken to Non-secure EL1 are handled in a similar manner to physical interrupts that are handled in a system with a single Security state, that is where GICD_CTLR.DS is set to 1:

- Group 0 interrupts are signalled using the virtual FIQ signal.
- Group 1 interrupts are signalled using the virtual IRQ signal.
- Group 0 and Group 1 interrupts share an interrupt prioritization and preemption scheme. A minimum of 32 and a maximum of 256 priority levels are supported, as determined by the values in ICH_VTR_EL2.

Note

The priority value is not subject to the shift used for Non-secure physical interrupts. While virtualization supports up to 8 bits of priority, a minimum of 5 and a maximum of 8 bits must be implemented.

Note

For information about the rules governing exception entry on an ARMv8-A PE, see *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Accesses at Non-secure EL1 to Group 0 registers are virtual when HCR_EL2.FMO == 1.

Virtual accesses to the following Group 0 ICC_* registers access the ICV_* equivalents:

- Accesses to ICC_AP0R<n>_EL1 access ICV_AP0R<n>_EL1.
- Accesses to ICC_BPR0_EL1 access ICV_BPR0_EL1.
- Accesses to ICC_EOIR0_EL1 access ICV_EOIR0_EL1.
- Accesses to ICC_HPPIR0_EL1 access ICV_HPPIR0_EL1.
- Accesses to ICC_IAR0_EL1 access ICV_IAR0_EL1.
- Accesses to ICC_IGRPEN0_EL1 access ICV_IGRPEN0_EL1.

Accesses at Non-secure EL1 to Group 1 registers are virtual when HCR_EL2.IMO == 1.

Virtual accesses to the following Group 1 ICC_* registers access the ICV_* equivalents:

- Accesses to ICC_AP1R<n>_EL1 access ICV_AP1R<n>_EL1.
- Accesses to ICC_BPR1_EL1 access ICV_BPR1_EL1.
- Accesses to ICC_EOIR1_EL1 access ICV_EOIR1_EL1.
- Accesses to ICC_HPPIR1_EL1 access ICV_HPPIR1_EL1.
- Accesses to ICC_IAR1_EL1 access ICV_IAR1_EL1.
- Accesses to ICC_IGRPEN1_EL1 access ICV_IGRPEN1_EL1.

Accesses at Non-secure EL1 to the common registers are virtual when either HCR_EL2.IMO == 1 or HCR_EL2.FMO == 1, or both.

Virtual accesses to the following Common ICC_* registers access the ICV_* equivalents:

- Accesses to ICC_RPR_EL1 access ICV_RPR_EL1.
- Accesses to ICC_CTLR_EL1 access ICV_CTLR_EL1.
- Accesses to ICC_DIR_EL1 access ICV_DIR_EL1.
- Accesses to ICC_PMR_EL1 access ICV_PMR_EL1.

A virtual write to ICC_SGI0R_EL1, ICC_SGI1R_EL1, or ICC_ASGI1R_EL1 traps to EL2.

Software executing at EL2 can access some ICV_* register state using ICH_VMCR_EL2 and ICH_VTR_EL2 as follows:

- ICV_PMR_EL1.Priority aliases ICH_VMCR_EL2.VPMR.
- ICV_BPR0_EL1.BinaryPoint aliases ICH_VMCR_EL2.VBPR0.
- ICV_BPR1_EL1.BinaryPoint aliases ICH_VMCR_EL2.VBPR1.
- ICV_CTLR_EL1.EOImode aliases ICH_VMCR_EL2.VEOIM.
- ICV_CTLR_EL1.CBPR aliases ICH_VMCR_EL2.VCBPR.
- ICV_IGRPEN0_EL1aliases ICH_VMCR_EL2.VENG0.
- ICV_IGRPEN1_EL1. aliases ICH_VMCR_EL2.VENG1.
- ICV_CTLR_EL1.PRIbits aliases ICH_VTR_EL2.PRIbits.

- ICV_CTLR_EL1.IDbits aliases ICH_VTR_EL2.IDbits.
- ICV_CTLR_EL1.SEIS aliases ICH_VTR_EL2.SEIS.
- ICV_CTLR_EL1.A3V aliases ICH_VTR_EL2.A3V.

5.2.1 Usage model for the List registers

A fundamental function of an interrupt controller is to develop list of pending interrupts in priority order for each PE, and then to present the highest priority interrupt to the PE if the interrupt is of sufficient priority. For physical interrupts, this task is performed entirely in hardware by the GIC. However, in order to reduce the cost in hardware, the GIC handles virtual interrupts using both hardware and software.

For each physical interrupt received that is targeting a vPE, the hypervisor adds that interrupt to a prioritized list of pending virtual interrupts that is presented to the vPE. The GIC hardware also provides a set of List registers, ICH_LR<n>_EL2, that holds an IMPLEMENTATION DEFINED number of the top entries in the prioritized list for the currently running vPE. Typically, there are at most only a few pending virtual interrupts for that vPE. The interrupts in the List register are then handled by the vPE in hardware, providing the same behavior for the VM as is seen by a non-virtualized operating system handling physical interrupts.

However, the total number of interrupts that are pending, active and pending, or active, can exceed the number of List registers available. In this case, the hypervisor can save one or more active interrupt entries to memory, and later restore them to the List registers based on their priority. In this way, the List registers act as a cache for the list of pending, active, or active and pending interrupts that is controlled by software, for a vPE.

The List registers provide maintenance interrupts for:

- The purpose of signalling when there are no pending interrupts in the List registers to allow the hypervisor to load more pending interrupts to the List registers.
- The purpose of signalling when the List registers are empty or nearly empty to allow the hypervisor to refill the List registers with entries from the list in memory.
- The purpose of signalling when an EOI has been received for an entry that is not in the List registers, which can occur if an active interrupt is held in memory.
- The enabling and disabling of virtual interrupt groups, which might result in a requirement to change the content of the List registers.

For more details on maintenance interrupts, see *Maintenance interrupts on page 5-85*.

Note

Although the List registers might include only active interrupts, with the hypervisor maintaining any pending interrupts in memory, a pending interrupt cannot be signalled to the vPE until the hypervisor adds it to the List registers. Therefore, to minimize interrupt latency and ensure the efficient operation of the vPE, ARM strongly recommends that the List registers contain at least one pending interrupt, if a List register is available for this interrupt.

The List registers form part of the context of the vPE. When there is switch from one vPE running on a PE to another vPE, the hypervisor switches the List registers accordingly.

The number of List registers is IMPLEMENTATION DEFINED, and can be discovered by reading ICH_HCR_EL2

The following pseudocode indicates the number of List registers that are implemented.

```
// NumListRegs()
// =============
// The number of implemented List Registers. This value is IMPLEMENTATION DEFINED.

integer NumListRegs()
    return integer IMPLEMENTATION_DEFINED “Number of List registers”;
```

5.2.2 List register usage resulting in UNPREDICTABLE behavior

The following cases are considered software programming errors and result in UNPREDICTABLE behavior:

- Having two or more interrupts with the same pINTID in the List registers for a single virtual CPU interface.

- Having a List register entry with ICH_LR<n>_EL2.HW= 1, which is associated with a physical interrupt, in active state or in pending state in the List registers if the Distributor does not have the corresponding physical interrupt in either the active state or the active and pending state.
- If ICC_CTLR_EL1.EOImode == 0 or ICC_CTLR_EL3.EOImode_EL3 == 0, then either:
  - Having an active interrupt in the List registers with a priority that is not set in the corresponding Active Priorities Register.
  - Having two interrupts in the List registers in the active state with the same preemption priority.

5.3 Configuration and control of VMs

The virtual GIC works by holding a prioritized list of pending virtual interrupts for each PE. In GICv3 this list is compiled in software and a number of the top entries are held in List registers in hardware. For LPIs, this list can be compiled using tables for each vPE. These tables are controlled by the GICR_* registers.

A hypervisor uses a System register interface that is accessible at EL2 to switch context and to control multiple VMs. The context held in the ICH_* System registers is the context for the scheduled vPE. A vPE is scheduled when:

- ICH_HCR_EL2.En == 1.
- HCR_EL2.FMO == 1, when virtualizing Group 0 interrupts.
- HCR_EL2.IMO == 1, when virtualizing Group 1 interrupts.

When a vPE is scheduled, the ICH_*_EL2 registers affect software executing at Non-secure EL1.

The ICH_*_EL2 registers control and maintain a vPE as follows:

- ICH_HCR_EL2 is used for the top level configuration and control of virtual interrupts.
- Information about the implementation, such as the size of the supported virtual INTIDs and the number of levels of prioritization is read from ICH_VTR_EL2.
- A hypervisor can monitor and provide context for ICC_CTLR_EL1 using ICH_VMCR_EL2.
- A set of List registers, ICH_LR<n>_EL2, are used by the hypervisor to forward a queue of pending interrupts to the PE, see *Usage model for the List registers on page 5-81*. The status of free locations in ICH_LR<n>_EL2 is held in ICH_ELRSR_EL2.
- The end of interrupt status for the List registers is held in ICH_EISR_EL2.
- The VM maintenance interrupt status is held in ICH_MISR_EL2.
- The active priority status is held in:
  - ICH_AP0R<n>_EL2, where n = 0-3.
  - ICH_AP1R<n>_EL2, where n = 0-3.

5.3.1 Association of virtual interrupts with physical interrupts

A virtual interrupt can become pending in response to a physical interrupt, where, for example, the physical interrupt is being used by a peripheral that is owned by a particular VM, or it can be generated for other reasons by the hypervisor where there is no corresponding physical interrupt. This second case can be used, for example, when the hypervisor emulates a virtual peripheral.

To support these two models, for SPIs and PPIs, the GIC List registers provide a mechanism to configure a virtual interrupt be associated with a physical interrupt. The physical interrupt and the virtual interrupt do not necessarily have the same INTID.

Usage model for associating a virtual interrupt with a physical interrupt

A virtual interrupt can be associated with a physical interrupt as follows:

1. The hypervisor configures ICC_CTLR_EL1.EOImode == 1, in this model.
2. On taking a physical PPI or a physical SPI that is targeting a vPE, the interrupt is taken to the hypervisor, and is acknowledged by hypervisor. The makes the physical interrupt active.
3. The hypervisor inserts a virtual interrupt to the list of pending interrupts for the targeted vPE. The hypervisor performs an EOI when it wants to do a priority drop for that interrupt. The hypervisor does not deactivate the interrupt.
4. When this virtual interrupt has a sufficiently high priority in the list of pending interrupts for that vPE, and that vPE is scheduled on the PE, the hypervisor writes this pending virtual interrupt into a List register, and ICH_LR<n>_EL2.HW is set to 1 to indicate that the virtual interrupt is associated with a physical interrupt. The INTID of the associated physical interrupt is held in the same List register.
5. When the vPE is running, it will take the pending virtual interrupt, and acknowledge it in the same way as it would acknowledge a physical interrupt, using the virtual CPU interface. When the interrupt handler running on the vPE has completed its task, and the virtual interrupt is to be deactivated, then the hardware deactivates

both the virtual interrupt and the associated physical interrupt. The virtual interrupt might be deactivated as the result of either an end of interrupt, if ICH_VMCR_EL2.VEOIM== 0, or as the result of a separate deactivation if ICH_VMCR_EL2.VEOIM == 1.

5.3.2 The Active Priorities registers

The active priority is held separately for virtual Group 0 and Group 1 interrupts, using ICH_AP0R<n>_EL2 and ICH_AP1R<n>_EL2, where n = 0-3. The Active Priorities Registers have a bit for each priority group implemented by the implementation. In GICv3, virtualization supports up to 8 bits of priority. However, as a result of interrupt grouping, bit[0] cannot be used for preemption. This means that a maximum of 128 active priority bits are required to maintain context. The number of registers implemented is dependent on the number of group priority bits supported, as shown in Table 5-1.

Table 5-1 Group bit count in the hypervisor Active Priorities Registers

Bits	Register	Number of registers
5	ICH_AP0R<n>_EL2 ICH_AP1R<n>_EL2	n = 0
6	ICH_AP0R<n>_EL2 ICH_AP1R<n>_EL2	n = 0-1
7	ICH_AP0R<n>_EL2 ICH_AP1R<n>_EL2	n = 0-3

If a bit is set to 1 in one of the ICH_AP0R<n>_EL2 registers, the equivalent bit in the ICH_AP1R<n>_EL2 register must be zero when executing in Non-secure EL1 or Non-secure EL0, otherwise the behavior of the GIC is UNPREDICTABLE.

If a bit is set to 1 in one of the ICH_AP1R<n>_EL2 registers, the equivalent bit in the ICH_AP0R<n>_EL2 register must be zero when executing in Non-secure EL1 or Non-secure EL0, otherwise the behavior of the GIC is UNPREDICTABLE.

ICH_AP0R<n>_EL2 provide a list of up to 128 bits where there is a bit for each implemented preemptable priority. If a bit is 1, this indicates that there is a Group 0 interrupt in that priority group which has been acknowledged but has not had a priority drop. If a bit is 0, this indicates that there is no Group 0 interrupt active at that priority, or that all active Group 0 interrupts within that priority group have undergone a priority drop.

Note

Writing to the Link registers does not have an effect on the Active Priorities Registers.

ICH_AP1R<n>_EL2 provide a list of up to 128 bits where there is a bit for each implemented preemptable priority. If a bit is 1, this indicates that there is a Group 1 interrupt in that priority group which has been acknowledged but has not had a priority drop. If a bit is 0, this indicates that there is no Group 1 interrupt active at that priority or that all active Group 1 interrupts within that priority group have undergone a priority drop.

Writing any value other than the last read value of the register, or `0x00000000`, to these registers can cause:

- Virtual interrupts that would otherwise preempt execution to not preempt execution.
- Virtual interrupts that otherwise would not preempt execution to preempt execution at Non-secure EL1 or EL0.

Note

ARM does not expect these registers to be read and written by software for any purpose other than:

- Saving and restoring state, as part of software power management.
- Context switching between vPEs on the same PE.

Writing to the Active Priority Registers in any order other than the following order results in UNPREDICTABLE behavior:

1. ICH_AP0R<n>_EL2.
2. ICH_AP1R<n>_EL2.

Note

An `ISB` is not required between the write to ICH_AP0R<n>_EL2 and the write to ICH_AP1R<n>_EL2.

5.3.3 Maintenance interrupts

Maintenance interrupts can signal key events in the operation of a GIC that implements virtualization. These events are processed by the hypervisor.

Note

- Maintenance interrupts are generated only when the global enable bit for the virtual CPU interface, ICH_HCR_EL2.En, is set to 1.
- ARM strongly recommends that maintenance interrupts are configured to use INTID 25. For more information, see *Server Base System Architecture (SBSA)*.

Maintenance interrupts are level-sensitive interrupts. Configuration bits in ICH_HCR_EL2 can be set to 1 to enable the generation of maintenance interrupts when:

- Group 0 virtual interrupts are enabled.
- Group 1 virtual interrupts are enabled.
- Group 0 virtual interrupts are disabled.
- Group 1 virtual interrupts are disabled.
- There are no pending interrupts in the List registers.
- At least one EOI request occurs with no valid List register entry for the corresponding interrupt.
- There are no valid entries, or there is only one valid entry, in the List registers. This is an underflow condition.
- At least one List register entry has received an EOI request.

See *ICH_MISR_EL2, Interrupt Controller Maintenance Interrupt State Register on page 8-290* for more information about the control and status reporting of maintenance interrupts.

5.4 Virtual LPI support

In GICv3 LPIs can be presented to a virtualized system by the hypervisor, which must be using the System registers. A virtual LPI is generated when the hypervisor writes a vINTID corresponding to the LPI range, that is a vINTID that is greater than 8191, to a List register. Because an LPI does not have an active state, it is not possible to associate a virtual LPI with a physical interrupt.

GICv4 provides support for the direct injection of *virtual LPIs*, vLPIs, in the LPI INTID range. With the direct injection of vLPIs, the GICR_* registers use structures in memory for each vPE to hold LPI configuration and pending information for vLPIs in the same way that they use structures in memory to hold LPI configuration and pending information for physical LPIs. However, the virtual structures are different from the physical structures, with the vLPI tables for the current vPE scheduled on a PE by GICR_VPENDBASER and GICR_VPROPBASER in the Redistributor associated with that PE, For more information about the physical LPI tables, see *LPI Configuration tables on page 6-95* and *LPI Pending tables on page 6-97*.

The Redistributor associated with the PE on which the vPE is scheduled determines the highest priority pending vLPI, and forwards this to the virtual CPU interface of the vPE. This vLPI and the interrupts in the List register are then prioritized together to determine the highest priority pending virtual interrupt for the vPE.

For information about virtual LPIs and the virtual CPU tables, see *The vPE table on page 6-104*.

5.4.1 Direct injection of virtual interrupts

The ITS maps an EventID and a DeviceID to an INTID associated with a PE, see *The ITS on page 6-99* for more information. GICv4 introduces the ability to generate a virtual LPI without involving the hypervisor. In this case an ITS maps the EventID for the interrupt translation using the following mechanism:

- The ITS interruption translation table entry for a vLPI is configured with:
  - A control flag that indicates the EventID is associated with a virtual LPI.
  - A vPEID to index into the ITS vPE table. For more information about vPEID and the vPE table, see *The vPE table on page 6-104*. The vPE table provides:
    1. The base address of the GICR_* registers in the format defined by GITS_TYPER.PTA.
    2. The base address of the virtual LPI Pending table associated with the target VM.
  - A virtual INTID, vINTID, that indicates which vLPI becomes pending.
  - A physical INTID, pINTID, that can be used as a doorbell interrupt to the hypervisor if the vPE is not scheduled on a PE. The value 1023 is used where a doorbell interrupt is not required, otherwise an INTID in the physical LPI range must be provided.

For more information about:

- Physical LPIs, see *LPIs on page 6-92*.
- The ITS and format of an Interrupt Translation Table (ITT), see *The ITS on page 6-99*.
- The commands used to control the handling of virtual LPIs associated with an ITS, see Table 6-6 on page 6-108 and the following commands:
  - *VINVALL on page 6-126*.
  - *VMAPI on page 6-127*.
  - *VMAPP on page 6-128*.
  - *VMAPTI on page 6-129*.
  - *VMOVI on page 6-131*.
  - *VMOVP on page 6-133*.
  - *VSYNC on page 6-135*.

5.4.2 Doorbell interrupts

When an interrupt that targets a vPE becomes pending, it might target a vPE that is not currently scheduled on a PE. Where those interrupts are presented as physical interrupts, the hypervisor can schedule in the vPE as a result of that interrupt. In this case the hypervisor can make the scheduling decisions for the vPE based on the full set of pending virtual interrupts for the vPE.

The equivalent capability is provided in the case of direct injections of vLPIs by the provision of *Doorbell LPI*s.

For a vLPI, the ITS can configure a physical LPI that is sent to a PE when the vLPI becomes pending and the vPE is not scheduled on that PE. This physical LPI is a Doorbell LPI.

The GIC hardware determines whether the vPE is scheduled on a PE when:

- GICR_VPENDBASER.Valid == 1.
- GICR_VPENDBASER.PendingLast holds the same value as defined in the VPT_addr field in the VMAPP command for the vPE that is the target of the vLPI.

If, at the time that a vPE is descheduled from a PE, there are one or more vLPIs pending for the PE, GICR_VPENDBASER.Physical_Address is set to 1. This can be used by the hypervisor to make scheduling decisions.

5.5 Pseudocode

The following pseudocode indicates the number of virtual active priority bits.

```
// ActiveVirtualPRIBits()
// ======================

integer ActiveVirtualPRIBits()
    if VirtualPRIBits() == 8 then
        return 128;
    else
        return 2^(VirtualPREBits() - 1);
```

The following pseudocode indicates the highest active group virtual priority.

```
// GetHighestActiveVGroup()
// ========================
// Returns a value indicating the interrupt group of the highest priority
// bit set from three registers. Returns None if no bits are set.

IntGroup GetHighestActiveVGroup(bits(128) avp0, bits(128) avp1)
    for rval = 0 to ActiveVirtualPRIBits() - 1
        if avp0<rval> != '0' then
            return IntGroup_G0;
        elsif avp1<rval> != '0' then
            return IntGroup_G1NS;

    return IntGroup_None;
```

The following pseudocode indicates the highest active virtual priority.

```
// GetHighestActiveVPriority()
// ===========================
// Returns the index of the highest priority bit set from three registers.

// Returns 0xFF if no bits are set.

bits(8) GetHighestActiveVPriority(bits(128) avp0, bits(128) avp1)
    for rval = 0 to ActiveVirtualPRIBits() - 1
        if avp0<rval> != '0' || avp1<rval> != '0' then
            return rval<7:0>;

    return Ones();
```

The following pseudocode indicates whether any bits are set in the supplied Active Priorities registers.

```
// VPriorityBitsSet()
// ==================
// Returns TRUE if any bit is set in the supplied registers, FALSE otherwise

boolean VPriorityBitsSet(bits(128) avp0, bits(128) avp1)
    for i = 0 to ActiveVirtualPRIBits() - 1
        if avp0<i> != '0' || avp1<i> != '0' then
            return TRUE;

    return FALSE;
```

The following pseudocode clears the highest priority bit in the supplied virtual Active Priorities registers.

```
// VPriorityDrop()
// ===============
// Clears the highest priority bit set in the supplied registers.

VPriorityDrop[bits(128) &avp0, bits(128) &avp1] = bit v
    assert IsZero(v);
    for i = 0 to ActiveVirtualPRIBits() - 1
        if avp0<i> != v then
            avp0<i> = v;
```

```
            return;
        elsif avp1<i> != v then
            avp1<i> = v;
            return;

    return;
```

The following pseudocode determines which active bits are set.

```
// FindActiveVirtualInterrupt()
// ============================
// Find a matching List register. Returns -1 if there is no match.

integer FindActiveVirtualInterrupt(bits(INTID_SIZE) vID)

    for i = 0 to NumListRegs() - 1
        if (!(ICH_LR_EL2[i].State IN {IntState_Active, IntState_ActivePending}) &&
            ICH_LR_EL2[i].VirtualID<INTID_SIZE-1:0> == vID) then
            return i;

    return -1;
```

The following pseudocode indicates the virtual group priority based on the minimum Binary Point register.

```
// VPriorityGroup()
// ================
// Returns the priority group field for the minimum BPR value

bits(8) VPriorityGroup(bits(8) priority, integer group)
    integer vpre_bits = VirtualPREBits();
    mask = Ones(vpre_bits):Zeros(8 - vpre_bits);
    return (priority AND mask);
```

The following pseudocode indicates the virtual group priority based on the appropriate Binary Point register.

```
// VGroupBits()
// ============
// Returns the priority group field for the minimum BPR value for the group

bits(8) VGroupBits(bits(8) priority, bit group)
    bpr = UInt(ICH_VMCR_EL2.VBPR1);

    if group == '0' || ICH_VMCR_EL2.VCBPR == '1' then
        bpr = UInt(ICH_VMCR_EL2.VBPR0);

    mask = Ones(bpr):Zeros(8 - bpr);
    return (priority AND mask);
```

The following pseudocode indicates the number of virtual ID bits.

```
// VIDBits()
// =========

integer VIDBits()
    id_bits = ICH_VTR_EL2.IDbits;
    case id_bits of
        when '000'  return 16;
        when '001'  return 24;
        otherwise   Unreachable();
```

The following pseudocode indicates the number of virtual preemption bits.

```
// VirtualPREBits()
// ================

integer VirtualPREBits()
    return UInt(ICH_VTR_EL2.PREbits) + 1;
```

The following pseudocode indicates the number of virtual priority bits.

```
// VirtualPRIBits()
// ================

integer VirtualPRIBits()
    return UInt(ICH_VTR_EL2.PRIbits) + 1;
```

Chapter 6
Locality-specific Peripheral Interrupts and the ITS

This chapter describes *Locality-specific Peripheral Interrupts* (LPIs) and the *Interrupt Translation Service* (ITS). It contains the following sections:

6.1 LPIs

Locality-specific Peripheral Interrupts (LPIs) are edge-triggered message-based interrupts that can use an *Interrupt Translation Service* (ITS), if it is implemented, to route an interrupt to a specific Redistributor and connected PE. GICv3 provides two types of support for LPIs. LPIs can be supported either:

- Using the ITS to translate an EventID from a device into an LPI INTID. For more information about EventIDs, see *The ITS* on page 6-99.
- By forwarding an LPI INTID directly to the Redistributors, using GICR_SETLPIR.

An implementation must support only one of these methods.

Note

The following registers are mandatory in an implementation that supports LPIs but does not include an ITS. The function of the registers is IMPLEMENTATION DEFINED in implementations that do include an ITS:

- GICR_SETLPIR
- GICR_CLRLPIR.
- GICR_INVLPIR.
- GICR_INVALLR.
- GICR_SYNCR.

These registers control physical LPIs in a system that does not include an ITS.

In an implementation that includes LPIs, at least 8192 LPIs are supported. For this reason, the configuration of each interrupt, and the pending information for each interrupt, is held in tables in memory, rather than in registers, and the tables are pointed to by registers held in the Redistributors.

Note

- ARM expects that an implementation will cache parts of the tables in the Redistributors to reduce latency and memory traffic. The form of these caches is IMPLEMENTATION DEFINED.
- The addresses for the LPI tables are in the Non-secure physical address space.

Figure 6-1 on page 6-93 shows the generation of LPIs in an implementation that includes at least one ITS.

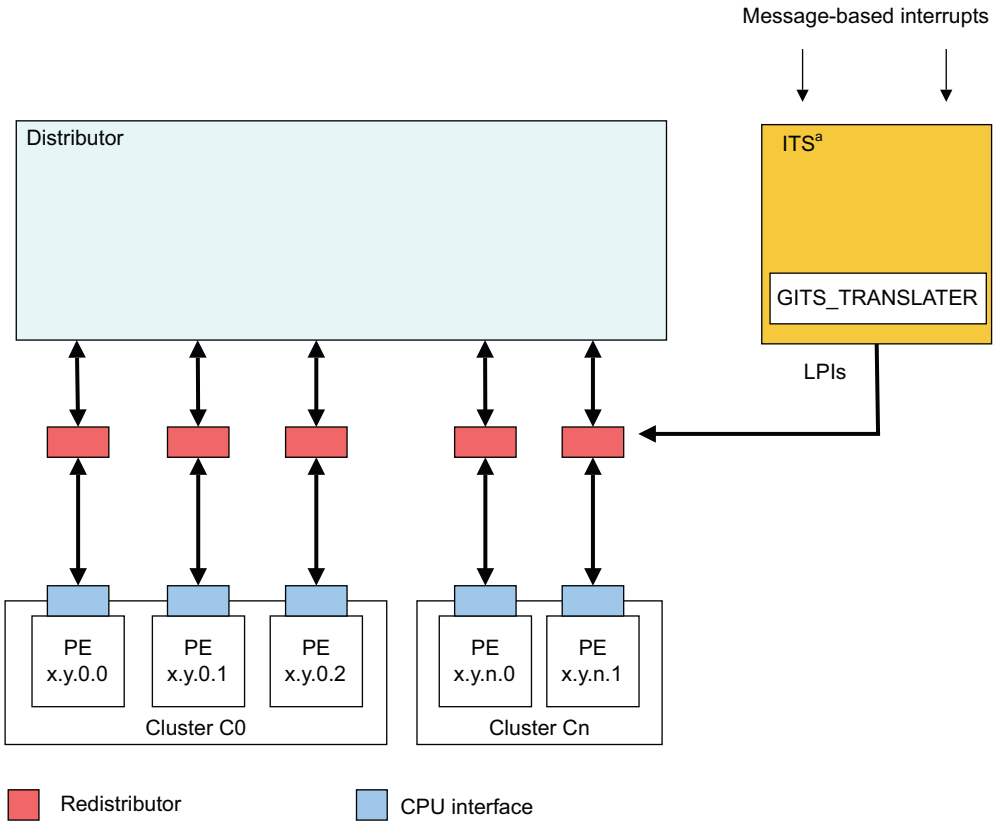

a. There might be zero, one, or more than one ITS in a GIC.

Figure 6-1 Triggering LPIs in an implementation with an ITS

Note

In Figure 6-1, the ITS channel to the Redistributors is IMPLEMENTATION DEFINED.

Figure 6-2 on page 6-94 shows the generation of LPIs in an implementation without an ITS.

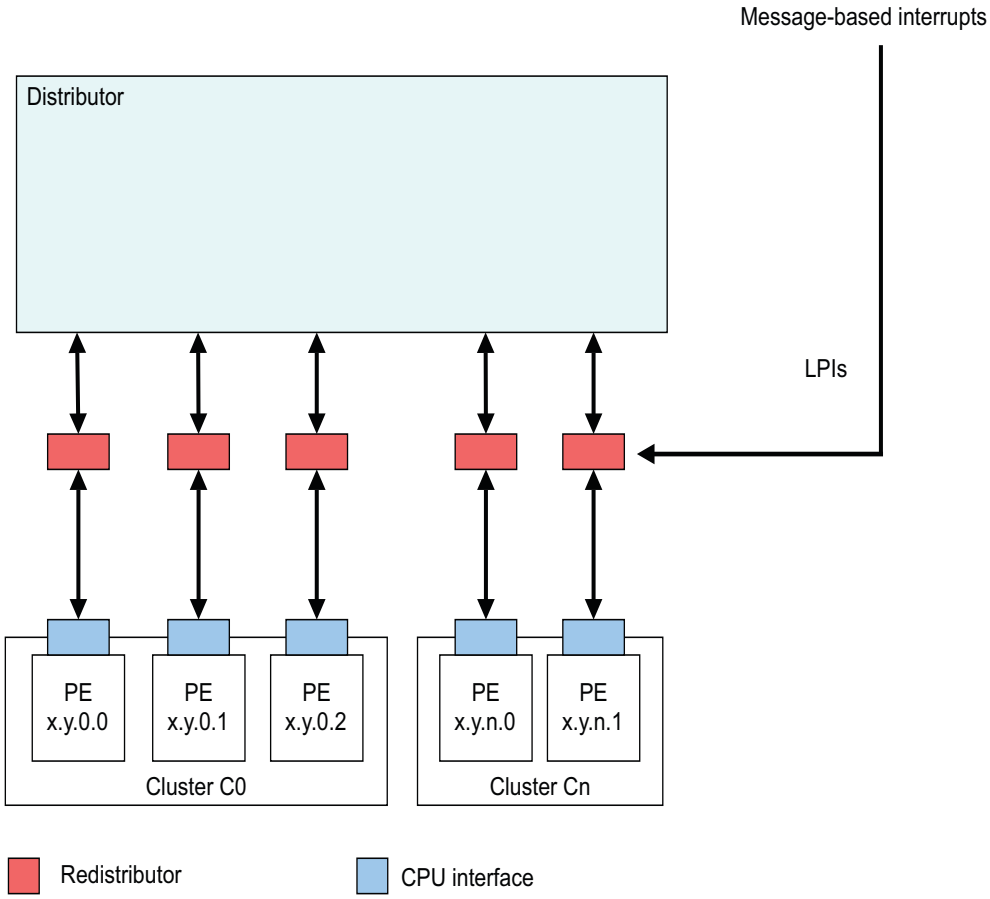

Figure 6-2 Triggering LPIs in an implementation without an ITS

When GICD_CTLR.DS == 0:

- LPIs are only supported when affinity routing is enabled for Non-secure state.
- LPIs are always Non-secure Group 1 interrupts.

When GICD_CTLR.DS == 1:

- LPIs are only supported when affinity routing is enabled.
- LPIs are always Group 1 interrupts.

There is a single global physical LPI space so that LPIs can be moved between all Redistributors. Software programs the size of the single global physical LPI space using GICR_PROPBASER.IDbits.

Note

The size of the physical LPI space is limited to the maximum size that an implementation supports, which is defined in GICD_TYPER.IDbits.

For a given Redistributor, LPI configuration and state are maintained in two tables in memory, described in the following sections:

- *LPI Configuration tables* on page 6-95.
- *LPI Pending tables* on page 6-97.

If a Redistributor supports physical LPIs, it has:

- LPI priority and enable bits programmed in the single LPI Configuration table. The address of the LPI Configuration table is defined by GICR_PROPBASER. If GICR_PROPBASER is updated when GICR_CTLR.EnableLPIs == 1, the effects are UNPREDICTABLE. See *LPI Configuration tables* for more information.
- Memory-backed storage for LPI pending bits in an LPI Pending table. This table is specific to a particular Redistributor. The address of the LPI Pending table is defined by GICR_PENDBASER. If GICR_PENDBASER is updated when GICR_CTLR.EnableLPIs == 1, the effects are UNPREDICTABLE.

GICR_PROPBASER.IDBits sets the size of the ID space, and thereby the number of entries in the LPI Configuration table and the corresponding LPI Pending table.

Physical LPIs are enabled by a write to GICR_CTLR.EnableLPIs.

Note

When LPIs are disabled at the Redistributor interface, that is when GICR_CTLR.EnableLPIs == 0, LPIs cannot become pending. An attempt to make an LPI pending in this situation has no effect, and the LPI is lost. This differs from disabling SGIs, PPIs, and SPIs, which prevents only the signaling of the interrupt to the CPU interface.

GICv4 introduces equivalent tables for handling virtual LPIs with addresses referenced in GICR_VPROPBASER and GICR_VPENDBASER.

In GICv4, virtual LPIs are enabled by a write to GICR_VPENDBASER.Valid.

6.1.1 LPI INTIDs

In both GICv3 and GICv4 software must obey the following rules when mapping physical LPIs.

The behavior of the GIC is UNPREDICTABLE if software:

- Maps multiple EventID-DeviceID combinations to the same physical LPI INTID.
- Assigns doorbell interrupts with the same physical LPI INTID to different physical PEs. This applies to GICv4 only.
- Maps an EventID-DeviceID combination and a doorbell interrupt to the same physical LPI INTID, unless they target the same physical PE. This applies to GICv4 only.

6.1.2 LPI Configuration tables

LPI configuration is global. Whether there are multiple copies of an LPI Configuration table that is pointed at by different Redistributors is IMPLEMENTATION DEFINED.

It is IMPLEMENTATION DEFINED whether GICR_PROPBASER can be set to different values on different Redistributors. GICR_TYPER.CommonLPIAff indicates which Redistributors must have GICR_PROPBASER set to the same value whenever GICR_CTLR.EnableLPIs == 1.

An implementation can treat all copies of GICR_PROPBASER that are required to have the same value as accessing common state.

Setting different values in different copies of GICR_PROPBASER on Redistributors that are required to use a common LPI Configuration table when GICR_CTLR.EnableLPIs == 1 leads to UNPREDICTABLE behavior.

If GICR_PROPBASER is programmed to different values on different Redistributors, it is IMPLEMENTATION DEFINED which copy or copies of GICR_PROPBASER are used when the GIC reads the LPI Configuration tables. However, the copy or copies that are used must correspond to a Redistributor on which GICR_CTLR.EnableLPIs == 1.

To avoid UNPREDICTABLE behavior, software must ensure that all copies of the LPI Configuration tables are identical, and all changes are globally observable, whenever:

- GICR_CTLR.EnableLPIs is written from 0 to 1 on any Redistributor.

- GICR_INVLPIR and GICR_INVALLR are written on any Redistributor with GICR_CTLR.EnableLPIs == 1, if direct LPIs are supported.
- The INV and INVALL command is executed by an ITS, in an implementation that includes at least one ITS.

An LPI Configuration table in memory stores entries containing configuration information for each LPI, where:

- GICR_PROPBASER specifies a 4KB aligned physical address. This is the LPI Configuration table base address.
- For any LPI `N`, the location of the table entry is defined by (base address + (`N` – 8192)).

To change the configuration of an interrupt, software writes to the LPI Configuration tables and then issues the INV or INVALL command. In implementations that do not include an ITS, software writes to GICR_INVALLR or GICR_INVLPIR.

The LPI Configuration table contains an 8-bit entry for each LPI. Figure 6-3 shows the LPI Configuration table entry format.

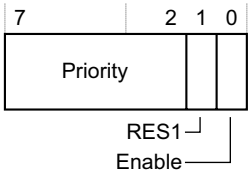

Figure 6-3 LPI Configuration table entry

Table 6-1 shows the LPI Configuration table entry bit assignments.

Table 6-1 LPI Configuration table entry bit assignments

Bits	Name	Function
[7:2]	Priority	The priority of the LPI. These are the most significant bits of the LPI priority. Bits[1:0] of the priority are 0. When GICD_CTLR.DS == 0, this value is shifted in accordance with the security and priority rules specified in *Software accesses of interrupt priority* on page 4-72. This means that LPI priorities are always in the lower half of the priority range. The priority value range is 128-254. In implementations that support a single Security state, the value in this field is not shifted. This means that LPI priorities use the full LPI priority value range of 0-252, in increments of 4. See *Interrupt prioritization* on page 4-65 for more information about interrupt priorities.
[1]	-	RES1.
[0]	Enable	LPI enable. This bit controls whether the LPI is enabled: `0` The LPI is not enabled. `1` The LPI is enabled.

Caching

A Redistributor can cache the information from the LPI Configuration tables pointed to by GICR_PROPBASER, when GICR_CTLR.EnableLPI == 1, subject to all of the following rules:

- Whether or not one or more caches are present is IMPLEMENTATION DEFINED. Where at least one cache is present, the structure and size is IMPLEMENTATION DEFINED.
- An LPI Configuration table entry might be allocated into the cache at any time.
- A cached LPI Configuration table entry is not guaranteed to remain in the cache.
- A cached LPI Configuration table entry is not guaranteed to remain incoherent with memory.

- A change to the LPI configuration is not guaranteed to be visible until an appropriate invalidation operation has completed:
  - If one or more ITS is implemented, invalidation is performed using the INV or INVALL command. A SYNC command completes the INV and INVALL commands.
  - If no ITS is implemented, invalidation is performed by writing to GICR_INVALLR or GICR_INVLPIR.

6.1.3 LPI Pending tables

Software configures the LPI Pending tables, using the implemented range of valid LPI INTIDs, by writing to GICR_PENDBASER. This register provides the base address of the LPI Pending table for physical LPIs.

Each Redistributor maintains entries in a separate LPI Pending table that indicates the pending state of each LPI when GICR_CTLR.EnableLPIs == 1 in the Redistributor:

`0` The LPI is not pending.

`1` The LPI is pending.

For a given LPI:

- The corresponding byte in the LPI Pending table is (base address + (N / 8)).
- The bit position in the byte is (N MOD 8).

An LPI Pending table that contains only zeros, including in the first 1KB, indicates that there are no pending LPIs.

The first 1KB of the LPI Pending table is IMPLEMENTATION DEFINED. However, if the first 1KB of the LPI Pending table and the rest of the table contain only zeros, this must indicate that there are no pending LPIs.

The first 1KB of memory for the LPI Pending tables must contain only zeros on initial allocation, and this must be visible to the Redistributors, or else the effect is UNPREDICTABLE.

During normal operation, the LPI Pending table is maintained solely by the Redistributor.

Behavior is UNPREDICTABLE if software writes to the LPI Pending tables after allocation.

For physical LPIs, when GICR_CTLR.EnableLPIs is changed to 1, the Redistributor must read the pending status of the physical LPIs from the physical LPI Pending table.

Note

If GICR_PENDBASER.PTZ == 1, software guarantees that the LPI Pending table contains only zeros, including in the first 1KB. In this case hardware might not read any part of the table.

For virtual LPIs, when GICR_CTLR.EnableLPIs ==1, and GICR_VPENDBASER.Valid is changed to 1, the Redistributor must read the pending status of the virtual LPIs from the virtual LPI Pending table.

Note

IF GICR_VPENDBASER.IDAI == 0, the software guarantees that the LPI Pending table was written out by the same GIC implementation, meaning that hardware can rely on the first 1KB of the table and might not read the entire table.

6.1.4 Virtual LPI Configuration tables and virtual LPI Pending tables

GICv4 uses the same concept of memory tables to hold the configuration and pending information for virtual LPIs. The format of these tables is the same as for physical LPIs, but the virtual LPI Configuration table is provided by GICR_VPROPBASER and the virtual LPI Pending table is provided by GICR_VPENDBASER, see *Virtual LPI support* on page 5-86.

When scheduling a vPE, GICR_VPENDBASER.IDAI can be cleared to 0:

- When the vPE was last scheduled on a Redistributor on the same GIC.
- When the vPE is scheduled for the first time after the initial allocation, and the entire virtual LPI Pending table contained only zeros on initial allocation.

- In IMPLEMENTATION DEFINED cases.

Clearing GICR_VPENDBASER.IDAI to 0 at any other time results in UNPREDICTABLE behavior.

6.2 The ITS

The ITS translates an input EventID from a device, identified by its DeviceID, and determines:

1. The corresponding INTID for this input.
2. The target Redistributor and, through this, the target PE for that INTID.

For GICv3, the ITS performs this function for events that are translated into physical LPIs. LPIs can be forwarded to a Redistributor either by an ITS or by a direct write to GICR_SETLPIR. An implementation must support only one of these methods.

For GICv4, the ITS also performs this function for interrupts that are directly injected as virtual LPIs.

An ITS has no effect on SGIs, SPIs, or PPIs.

The flow of the ITS translation is as follows:

1. The DeviceID selects a *Device table entry* (DTE) in the *Device table* that describes which *Interrupt translation table* (ITT) to use.
2. The EventID selects an *Interrupt translation entry* (ITE) in the ITT that describes:
   - For physical interrupts:
     - The output physical INTID.
     - The *Interrupt collection number*, ICID.
   - For virtual interrupts, in GICv4:
     - The output virtual INTID.
     - The vPEID.
     - A doorbell to use if the vPE is not scheduled.
3. For physical interrupts, ICID selects a Collection table entry in the *Collection table* (CT) that describes the target Redistributor, and therefore the target PE, to which the interrupt is routed.
4. For virtual interrupts, in GICv4, the vPEID selects a vPE table entry that describes the Redistributor that is currently hosting the target vPE to which the interrupt is routed.

The tables used in the translation process are described in more detail in the following sections:

- *The ITS tables*.
- *The Device table* on page 6-102.
- *The Interrupt translation table* on page 6-103.
- *The Collection table* on page 6-104.
- *The vPE table* on page 6-104.

These tables are created and maintained using the ITS commands described in *ITS commands* on page 6-108. GICv3 and GICv4 do not support direct access to the tables, and the tables must be configured using the ITS commands.

6.2.1 The ITS tables

To allow software to provision memory for the ITS private tables, the GIC provides a set of registers that allow the following features to be discovered:

- The number of private tables that are required.
- The size of each entry in each table.
- The type of each table.

Note

All ITS tables are in the Non-secure physical address space.

The state and configuration of the ITS tables is stored in a set of tables in memory. This memory is allocated by software before enabling the ITS.

GITS_BASER<n> specifies the base address and size of the ITS tables, and must be provisioned before the ITS is enabled.

The ITS tables have either a flat structure or a two-level structure. The structure is determined by GITS_BASER<n>, as follows:

0 Flat table. In this case a contiguous block of memory is allocated for the table. The format of the table is IMPLEMENTATION DEFINED.

Behavior is UNPREDICTABLE if memory that is used for the ITS tables does not contain zeros at the time of the new allocation for use by the ITS.

1 Two-level table. In this case each entry in the level 1 table is 64 bits, and has the following format:

- Bit[63] - Valid:
  - If this bit is cleared to 0, the PhysicalAddress field does not point to the base address of a level 2 table.
  - If this bit is set to 1, the PhysicalAddress field points to the base address of a level 2 table.
- Bits[62:52] - RES 0.
- Bits[51:N] - PhysicalAddress of the level 2 table. N is the number of bits that are required to specify the page size:
  - The size of the level 2 table is determined by GITS_BASER<n>.Page_Size.
- Bits[N-1:0] - RES 0. N is the number of bits that are required to specify the page size.

The level 1 table is indexed by the appropriate ID so that level 1 entry = ID/(Page Size / Entry Size).

Note

This allows software to determine the level 2 table that must be allocated for a given CollectionID, DeviceID, or vPEID.

For level 1 table entries, when Valid == 0:

- If the Type field specifies a valid table type other than an Interrupt collection table, the ITS discards any writes to the level 2 table.
- If the Type field specifies the Collection table, and ICID is greater than or equal to the number indicated by GITS_TYPER.HCC, the ITS discards any writes to the level 2 table.

The format of the level 2 table is IMPLEMENTATION DEFINED.

Behavior is UNPREDICTABLE if:

- Memory that is used for the level 2 tables does not contain zeros at the time of the new allocation for use by the ITS.
- Multiple level 1 table entries with Valid == 1 point to the same level 2 table.

Note

As part of restoring the state of the ITS from powerdown events, the registers that describe the table can point to tables that were previously populated by the ITS, and so might contains values other than zeros. The details of power management of the ITS are IMPLEMENTATION DEFINED. See *ITS power management* on page 6-149.

Figure 6-4 on page 6-101 shows how these tables are used in the translation process.

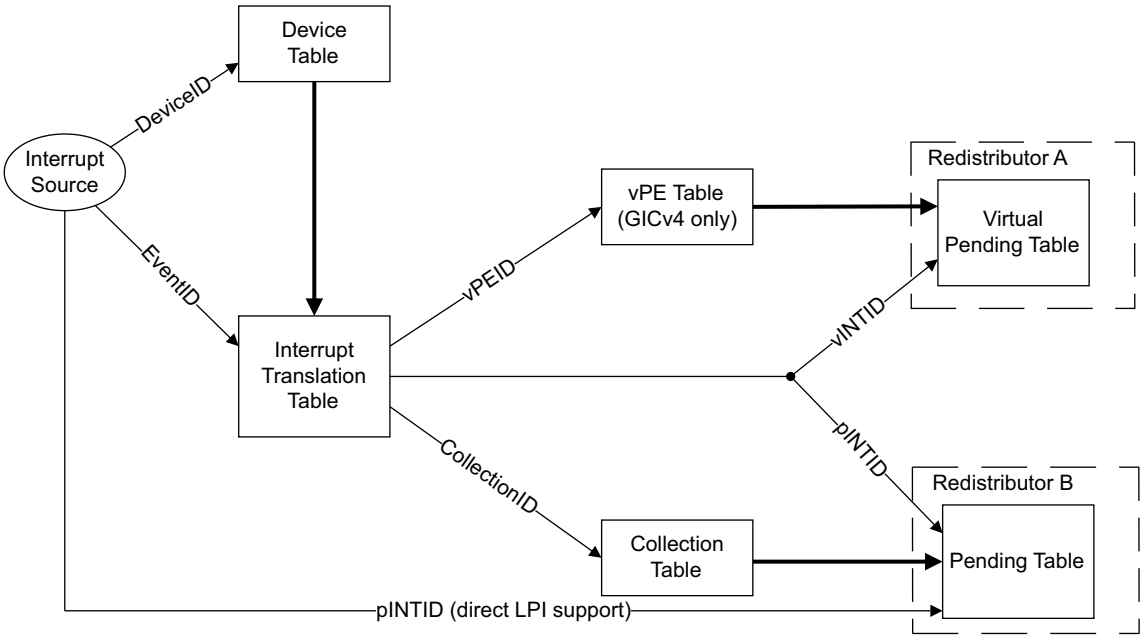

Figure 6-4 ITS tables

When GITS_CTLR.Enabled is written from 0 to 1 behavior is UNPREDICTABLE if any of the following conditions are true:

- GITS_CBASER.Valid == 0.
- GITS_BASER<n>.Valid == 0, for any GITS_BASER<n> register where the Type field indicates Device.
- GITS_BASER<n>.Valid == 0, for any GITS_BASER<n> register where the Type field indicates Interrupt Collection and GITS_TYPER.HCC == 0.
- In GICv4, GITS_BASER<n>.Valid == 0, for any GITS_BASER<n> register where the Type field indicates a vPE.

Software access to the private ITS tables.

If GITS_BASER<n>.Indirect == 0, behavior is UNPREDICTABLE if memory that is used for the ITS tables does not contain all zeros when first allocated to the ITS.

If GITS_CTLR.Indirect == 1, behavior is UNPREDICTABLE if memory that is used for a level 2 table does not contain all zeros when it is first allocated for use by the ITS.

When GITS_CTLR.Enabled = 0 and GITS_CTLR.Quiescent == 1:

- An implementation will not access the tables that are pointed to by any of the GITS_BASER<n> registers.

When GITS_CTLR.Enabled == 1 or GITS_CTLR.Quiescent == 0:

- An implementation will not access a table that is pointed to by any GITS_BASER<n> register for which GITS_BASER<n>.Valid == 0.
- For a table that is pointed to by a GITS_BASER<n> register for which GITS_BASER<n>.Valid == 1 and GITS_BASER<n>.Indirect == 0, behavior is UNPREDICTABLE if the table is written by software.
- For a table that is pointed to by a GITS_BASER<n> register for which GITS_BASER<n>.Valid == 1 and GITS_BASER<n>.Indirect == 1:
  - Behavior is UNPREDICTABLE if any of the level 2 table entries are written by software.
  - An ITS will not cache any entry in the level 1 table where the valid bit is cleared to 0.
  - Behavior is UNPREDICTABLE if any level 1 table entry where the valid bit set to 1 is written by software.

— A write to a level 1 table entry that changes the valid bit from 0 to 1 must be globally visible before software adds a command to the ITS command queue that relies on that entry. Otherwise it is UNKNOWN if the command will succeed or if it will be ignored.

6.2.2 Interrupt collections

In GICv3, the ITS considers all physical LPIs that it generated to be members of *collections*. The data that is associated with a collection can be held in the ITS, in external memory, or in both. The ITS supports collections that are held in memory if any of the GITS_BASER<n>.Type == `0b100`:

- When the ITS supports collections that are held in memory, the total number of collections that is supported is determined by the memory allocated by software:
  - If GITS_BASER<n>.Indirect == 0, the number of collections supported in memory can be calculated using the following formula:

    ((number of pages * page size) / entry size)

    The relevant values for this formula are indicated in GITS_BASER<n>.Size, GITS_BASER<n>.PageSize, and GITS_BASER<n>.EntrySize.
  - If GITS_BASER<n>.Indirect == 1, the number of collections supported in memory can be calculated using the following formula:

    (((number of pages in level 1 table * page size) /8) * (page size/entry size)).

    The relevant values for this formula are indicated in GITS_BASER<n>.Size, GITS_BASER<n>.PageSize, and GITS_BASER<n>.EntrySize.

    Note

    Indirect tables allow sparse allocations, so not all ICIDs in the supported range might be usable.

- Where collections are held in both the ITS and external memory, the total number of collections is indicated by GITS_TYPER.CCT.

When GITS_TYPER.HCC!= 0:

- Collections with identifiers in the range {0... GITS_TYPER.HCC-1} are held in the ITS.
- Collections with identifiers in the range greater than that indicated in GITS_TYPER.HCC are held in external memory, if this is supported.

When GITS_TYPER.HCC == 0:

- The ITS must support collections in external memory, and all collections are held in external memory.

The maximum number of collections that are supported is limited by the size of the ICID:

- If GITS_TYPER.CIL == 0, the ICID is 16 bits.
- If GITS_TYPER.CIL == 1, the ICID is reported by GITS_TYPER.CIDbits.

6.2.3 The Device table

The *Device table* provides a table of *Device table entries* (DTEs). Each DTE describes a mapping between a DeviceID and an ITT base address that points to the memory that the ITS can use to store the translations for the EventID. The ITS uses the ITT to store the translations for every EventID for the specified DeviceID. The DeviceID is a unique identifier assigned to each device that can create a range of EventIDs. For example, ARM expects that the 16-bit Requester ID from a PCIe Root Complex is presented to an ITS as a DeviceID.

The 32-bit DeviceID provides the index value for the table.

Table 6-2 shows an example of the number of bits that might be assigned to each DTE.

Table 6-2 DTE entries

Number of bits	Assignment	Notes
1	Valid	Boolean
40	ITT Address	Base physical address
5	ITT Range	Log2 (number of EventIDs supported by the ITT minus one)

6.2.4 The Interrupt translation table

An *Interrupt translation table* (ITT) is specific to each device that can create numbered events. Each entry in an ITT is referred to as an *Interrupt translation entries* (ITEs).

In GICv3, ITEs are only defined for physical interrupts

In GICv4, ITEs are defined for physical interrupts and for virtual interrupts, and provide a distinction between:

- An entry for a physical LPI and the use of an ICT for routing information.
- An entry for a virtual LPI and the use of a *vPE table*.

An ITT must be assigned a contiguous physical address space that has been zeroed, starting at ITT Address. The size is 2^(DTE.ITT Range + 1)* GITS_TYPER.ITT Size.

For a one-level table, the allocated memory must be zeroed before mapping the device. For a two-level table, the level 1tables must be allocated and initialized, and the level 2 tables must be allocated and zeroed, see GITS_BASER<n>.

ARM expects that devices are mapped to ITTs in such a way that each interrupt request from a particular device maps to a unique ITE. In addition, software must ensure that either:

- Each device maps to a dedicated section of memory. This means that no more than one device can be mapped to a particular section in memory.
- When multiple devices share a section of memory, each ITE within that section is associated with one device only, and all ITS commands that affect an ITE are performed using the associated device.

Note

- For simplification, ARM recommends that each device is mapped to a dedicated section in memory.
- ITS accesses to an ITT use the same Shareability and Cacheability attributes as specified for the Device table.

For physical interrupts, each ITE describes the mapping between the input EventID and:

- The *output physical INTID* (pINTID) that is sent to the target PE.
- The ICID that identifies an entry in the Collection table, that determines the target PE for the LPI. For more information about the Collection table, see *The Collection table* on page 6-104.

For virtual interrupts, each ITE describes the mapping of the EventID as outlined in the preceding list, and:

- The *output virtual INTID* (vINTID) that is sent to the target vPE.
- The *virtual PE number* (vPEID) that identifies an entry in the vPE table to determine the current host Redistributor. For more information about the vPE table, see *The vPE table* on page 6-104.
- A physical LPI that is sent to a physical PE if a virtual interrupt is translated when the target vPE is not currently scheduled on a physical PE.

The EventID provides the index value for the table.

Table 6-3 shows an example of the number of bits that might be stored in an ITE.

Table 6-3 ITE entries

Number of bits	Assignment	Notes
1	Valid	Boolean
1	Interrupt_Type	Boolean, indicates whether the interrupt is physical or virtual
Size of the LPI number space[a]	Interrupt_Number	pINTID or vINTID depending on the interrupt type
Size of the LPI number space[a]	Interrupt_Number HypervisorID	In GICv4 pINTID is used as a doorbell. In GICv3, and in GICv4 when a doorbell is not required, the programmed value is 1023.
16	ICID	Interrupt Collection ID, for physical interrupts only.
16	vPEID	vPE ID, for virtual interrupt only

a. For information about the size of the LPI number space, see *INTIDs* on page 3-39

6.2.5 The Collection table

The Collection table (CT) provides a table of *Collection table entries* (CTEs). For physical LPIs only, each CTE describes a mapping between:

- The ICID generated by the ITT.
- The address of the target Redistributor in the format defined by GITS_TYPER.PTA.

There is a single CT for each ITS, which can be held in registers or in memory, or in a combination of the two. See GITS_BASER<n>.Type and GITS_TYPER.HCC for more information.

The TableID provides the index value for the table. It is derived from ICID.

Table 6-4 shows an example of the number of bits that might be assigned to each CT.

Table 6-4 CT entries

Number of bits	Assignment	Notes
1	Valid	Boolean
Size of RDbase identifier	RDbase	The GIC supports two formats for RDbase, see RDbase

6.2.6 The vPE table

The vPE table consists of vPE table entries that provide a mapping from the vPEID generated by the ITS to:

- The target Redistributor, in the format defined by GITS_TYPER.PTA.
- The base address of the virtual LPI Pending table associated with the target vPE.

An area of memory defined by GITS_BASER<n> holds the vPE table and indicates the size of each entry in the table.

The vPE table describes all the vPEs associated with an ITS. Table 6-5 on page 6-105 shows an example of the number of bits that an implementation might store in a vPE table.

The 16-bit vPEID provides the index value for the table.

Table 6-5 vPE table entries

Number of bits	Assignment	Notes
1	Valid	Boolean
Size of RDbase identifier	RDbase	The GIC supports two formats for RDbase.
Size of address	VPT_addr	VPT_addr locates the LPI Pending table when the VM is not resident in the Redistributor. It is used as the address in GICR_VPENDBASER when the vPE is scheduled in the GICR_* registers associated with RDbase.
5	Size	The size of the vINTID range supported (minus one).

6.2.7 Control and configuration of the ITS

An ITS is controlled and configured using a memory-mapped interface where:

- The version can be read from GITS_IIDR and from GITS_PIDR2.
- GITS_TYPER specifies the features that are supported by an ITS.
- GITS_CTLR controls the operation of an ITS.
- GITS_TRANSLATER receives EventID information. It is IMPLEMENTATION DEFINED how the DeviceID is supplied. See *ITS commands* on page 6-108 for more details.
- GITS_BASER<n> registers provide information about the type, size and access attributes for the architected ITS memory structures.
- GITS_CBASER, GITS_CREADR, and GITS_CWRITER store address information for the ITS command queue interface.

There is an enable bit for each ITS, GITS_CTLR.Enabled.

6.2.8 The ITS command interface

Figure 6-5 on page 6-106 shows how the ITS provides the base address and the size that are used by the ITS command queue.

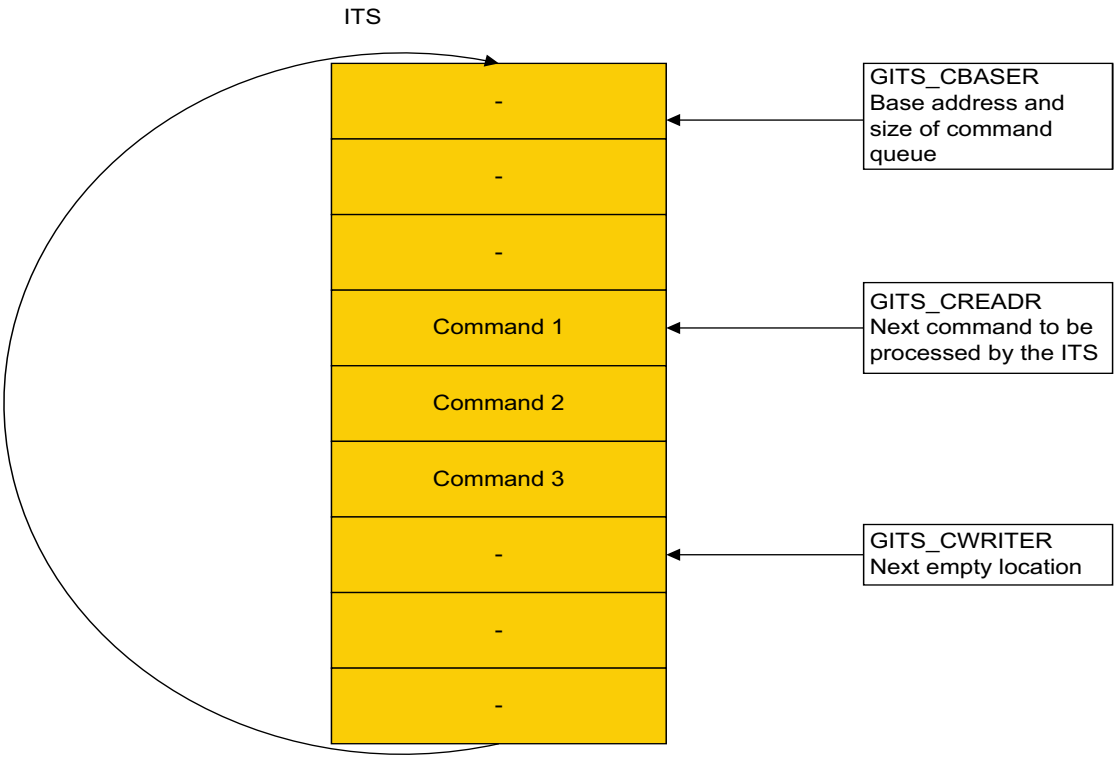

Figure 6-5 The ITS command queue

GITS_CBASER, GITS_CREADR, and GITS_CWRITER define the ITS command queue.

- GITS_CBASER uses the following fields:
  - Valid. This field indicates the allocation of memory for the ITS command queue.
  - Cacheability. This field indicates the cacheability attributes of accesses to the ITS command queue.
  - Shareability. This field indicates the Shareability attributes of accesses to the ITS command queue.
  - Physical address. This field provides the base physical address of the memory containing the ITS command queue.
  - Size. This field indicates the number of 4KB pages of physical memory for the ITS command queue.
- GITS_CREADR specifies the base address offset from which an ITS reads the next command to execute.
- GITS_CWRITER specifies the base address offset of the next free entry to which software writes the next command.

The size of an ITS command queue entry is 32 bytes. This means that there is support for 128 entries in each 4KB page.

The ITS command queue uses a little endian memory order model.

In the ITS command queue:

- The base address is always aligned to 64KB.
- Size is expressed as a multiple of 4KB.
- The address at which the queue wraps is always aligned to 4KB, and is (base address + (Size * 4KB)).

Note

All addresses are Non-secure physical addresses.

When the first command is complete, the ITS starts to process the next command. The read pointer, GITS_CREADR, advances as the ITS processes commands. If GITS_CREADR reaches the top of the memory specified in GITS_CBASER then the pointer wraps back to the base address specified in GITS_CBASER. GITS_CWRITER is controlled by software.

The ITS command queue is empty when GITS_CWRITER and GITS_CREADR specify the same base address offset value.

The ITS command queue is full when GITS_CWRITER points to an address 32 bytes behind GITS_CREADR in the buffer.

When GITS_CREADR.Stalled == 1 no subsequent commands are processed.

The INT ITS command generates an interrupt on execution, and this can generate an interrupt on completion of a particular sequence of commands, see *ITS commands* on page 6-108.

6.2.9 Ordering of translations with the output to ITS commands

Each command queue entry appears to be executed atomically so that a translation request either sees the state of the ITS before a command or the state of the ITS after the command.

A translation request initiated after a SYNC or VSYNC command has completed is translated using an ITS state that is consistent with the state after the command is performed.

In the absence of a SYNC or VSYNC command the ordering of ITS commands and translation requests is not defined by the architecture.

6.3 ITS commands

Table 6-6 provides a summary of all ITS commands.

Table 6-6 ITS commands

Command	Command arguments	Description
CLEAR	DeviceID, EventID	Translates the event defined by EventID and DeviceID into an ICID and pINTID, and instruct the appropriate Redistributor to remove the pending state.
DISCARD	DeviceID, EventID	Translates the event defined by EventID and DeviceID and instructs the appropriate Redistributor to remove the pending state of the interrupt. It also ensures that any caching in the Redistributors associated with a specific EventID is consistent with the configuration held in memory. DISCARD removes the mapping of the DeviceID and EventID from the ITT, and ensures that incoming requests with a particular EventID are silently discarded.
INT	DeviceID, EventID	Translates the event defined by EventID and DeviceID into an ICID and pINTID, and instruct the appropriate Redistributor to set the interrupt pending.
INV	DeviceID, EventID	Specifies that the ITS must ensure that any caching in the Redistributors associated with the specified EventID is consistent with the LPI Configuration tables held in memory.
INVALL	ICID	Specifies that the ITS must ensure any caching associated with the interrupt collection defined by ICID is consistent with the LPI Configuration tables held in memory for all Redistributors.
MAPC	ICID, RDbase	Maps the Collection table entry defined by ICID to the target Redistributor, defined by RDbase.
MAPD	DeviceID, ITT_addr, Size	Maps the Device table entry associated with DeviceID to its associated ITT, defined by ITT_addr and Size.
MAPI	DeviceID, EventID, ICID	Maps the event defined by EventID and DeviceID into an ITT entry with ICID and pINTID = EventID. **Note** • pINTID ≥0x2000 for a valid LPI INTID. • This is equivalent to MAPTI DeviceID, EventID, EventID, ICID
MAPTI[a]	DeviceID, EventID, pINTID, ICID	Maps the event defined by EventID and DeviceID to its associated ITE, defined by ICID and pINTID in the ITT associated with DeviceID. **Note** pINTID ≥0x2000 for a valid LPI INTID.
MOVALL	RDbase1, RDbase2	Instructs the Redistributor specified by RDbase1 to move all of its interrupts to the Redistributor specified by RDbase2.
MOVI	DeviceID, EventID, ICID	Updates the ICID field in the ITT entry for the event defined by DeviceID and EventID. It also translates the event defined by EventID and DeviceID into an ICID and pINTID, and instructs the appropriate Redistributor to move the pending state, if it is set, of the interrupt to the Redistributor defined by the new ICID, and to update the ITE associated with the event to use the new ICID.

Table 6-6 ITS commands (continued)

Command	Command arguments	Description
SYNC	RDbase	Ensures all outstanding ITS operations associated with physical interrupts for the Redistributor specified by RDbase are globally observed before any further ITS commands are executed. Following the execution of a SYNC the effects of all previous commands must apply to subsequent writes to GITS_TRANSLATER. See *Ordering of translations with the output to ITS commands* on page 6-107 for more information.
VINVALL[b]	vPEID	Ensures any cached Redistributor information associated with vPEID is consistent with the associated LPI Configuration tables held in memory.
VMAPI[b]	DeviceID, EventID, Dbell_pINTID, vPEID	Maps the event defined by DeviceID and EventID into an ITT entry with vPEID, vINTID=EventID, and Dbell_PINTID, a doorbell provision. **Note** • vINTID ≥0x2000 for a valid LPI INTID. • This is equivalent to VMAPTI DeviceID, EventID,EventID, pINTID, vPEID • Dbell_pINTID must be either 1023 or Dbell_pINTID ≥0x2000 for a valid LPI INTID.
VMAPP[b]	vPEID, RDbase, VPT_addr, VPT_size	Maps the vPE table entry defined by vPEID to the target RDbase, including an associated virtual LPI Pending table (VPT_addr, VPT_size).
VMAPTI[bc]	DeviceID, EventID, vINTID, Dbell_pINTID, vPEID	Maps the event defined by DeviceID and EventID into an ITT entry with vPEID and vINTID, and Dbell_pINTID, a doorbell provision. **Note** • vINTID ≥0x2000 for a valid LPI INTID. • Dbell_pINTID must be either 1023 or Dbell_pINTID ≥0x2000 for a valid LPI INTID.
VMOVI[b]	DeviceID, EventID, vPEID	Updates the vPEID field in the ITT entry for the event defined by DeviceID and EventID. Translates the event defined by EventID and DeviceID into a vPEID and pINTID, and instructs the appropriate Redistributor to move the pending state, if it is set, of the interrupt to the Redistributor defined by the new vPEID, and updates the ITE associated with the event to use the new vPEID.
VMOVP[b]	vPEID, RDbase, SequenceNumber, ITSList	Updates the vPE table entry defined by vPEID to the target Redistributor specified by RDbase. Software must use SequenceNumber and ITSList to synchronize the execution of VMOVP commands across more than one ITS.
VSYNC[b]	vPEID	Ensures all outstanding ITS operations for the vPEID specified are globally observed before any further ITS commands are executed. Following the execution of a VSYNC the effects of all previous commands must apply to subsequent writes to GITS_TRANSLATER.

a. This command was previously called MAPVI.
b. This command exists in GICv4 only.
c. This command was previously called VMAPVI.

The number of bits of EventID and DeviceID that an implementation supports are discoverable from GITS_TYPER. Unimplemented bits are RES0.

Note

- The INTID of an LPI is in the range of 8192 - maximum number. The maximum number is IMPLEMENTATION DEFINED. See *INTIDs* on page 3-39.
- The following argument names have been changed from those used in preliminary information associated with this GIC specification:
  - Device has been changed to DeviceID.
  - ID has been changed to EventID.
  - pID has been changed to pINTID.
  - vID has been changed to vINTID.
  - pCID has been changed to ICID.
  - target address has been changed to RDbase.
  - VCPU has been changed to vPE.
- The format of the collection target address, RDbase, is indicated by GITS_TYPER.PTA.

6.3.1 IMPLEMENTATION DEFINED sizes in ITS command parameters

Some ITS commands include the following types of parameter that have an IMPLEMENTATION DEFINED size:

DeviceIDs

The maximum number of Device identifiers supported by the associated Device table is determined by the number of bits available, as specified by GITS_TYPER.Devbits.

EventID EventID is limited by the maximum MAPD Size field, which is limited by GITS_TYPER.IDbits.

ICID The number of collections supported is IMPLEMENTATION DEFINED:

- For implementations that do not support Collection tables in external memory, GITS_TYPER.HCC indicates the number of collections.
- For implementations that do support Collection tables in external memory, the number of supported collections is limited by the size of the allocated collection table:
  - The total number of collections supported is calculated as follows:

    GITS_TYPER.HCC + (Size of collection table / Entry size)

    When GITS_TYPER.CIL == 1, the maximum number of collections is limited by GITS_TYPER.CIDbits.

pINTID pINTID is limited by GICR_PROPBASER.IDbits, which is limited by GICD_TYPER.IDbits. This also applies to Dbell_pINTID.

RDbase

RDbase is associated with a Redistributor and is specified in one of two formats:

- The base physical address of RD_base when GITS_TYPER.PTA == 1.

  Note

  Addresses can be up to 52 bits in size and must be 64KB aligned. The RDbase field consists of bits[51:16] of the address.

- A PE number, as indicated in GICR_TYPER.Processor_Number when GITS_TYPER.PTA == 0.

vINTID vINTID can be limited by GICR_VPROPBASER.IDbits, which is limited by GICD_TYPER.IDbits.

vPEID vPEID is limited by the size of the vPE table.

6.3.2 Command errors

If the ITS detects an error in the data provided to a command, the resulting behavior is a CONSTRAINED UNPREDICTABLE choice of:

- Ignoring the command:
  - No action is performed that alters the handling of interrupts.
  - GITS_CREADR is incremented to point to the next command, wrapping if necessary.
  - If GITS_TYPER.SEIS is set to 1, a System error is generated.

    Note

    It is IMPLEMENTATION DEFINED how the System error is recorded and how it is reported to the PE.

- Stalling the ITS command queue:
  - GITS_CREADR is not incremented and continues to point to the entry that triggered the error.
  - GITS_CREADR.Stalled is set to 1.
  - Software can restart the processing of commands by writing 1 to GITS_CWRITER.Retry.
  - If GITS_TYPER.SEIS is set to 1, a System error is generated.

    Note

    It is IMPLEMENTATION DEFINED how the system error is recorded and how it is reported to the PE.

- Treating the data as valid data:
  - The data that generated the error or errors is treated as having a legal value, and the command is processed accordingly.
  - GITS_CREADR is incremented to point to the next command, wrapping if necessary.
  - If GITS_TYPER.SEIS is set to 1 a System error is generated.

    Note

    It is IMPLEMENTATION DEFINED how the System error is recorded and how it is reported to the PE.

See *ITS command error encodings* on page 6-146 for more information.

6.3.3 CLEAR

This command translates the event defined by `EventID and DeviceID` into an `ICID and pINTID`, and instructs the appropriate Redistributor to remove the pending state.

Figure 6-6 shows the format of the `CLEAR` command.

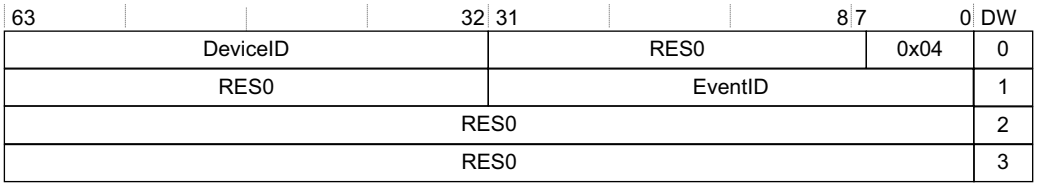

Figure 6-6 CLEAR command format

In Figure 6-6:

- `EventID` identifies the interrupt, associated with a device, for which the pending state is to be cleared.
- `DeviceID` specifies the requesting device.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

The command and its arguments are:

```
CLEAR DeviceID, EventID
```

A command error occurs if any of the following apply:

- DeviceID exceeds the maximum value supported by the ITS.
- The device specified by DeviceID is not mapped to an Interrupt Translation Table, using MAPD.
- EventID exceeds the maximum value allowed by the ITT. This value is specified by the Size field when the MAPD command is issued.
- The EventID for the device is not mapped to a collection, using MAPI or MAPTI.
- The EventID for the device is mapped to a collection that has not been mapped to an RDbase using MAPC.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

The following pseudocode describes the operation of the CLEAR command:

```
// ITS.CLEAR
// =========

ITS.CLEAR(ITSCommand cmd)
    if DeviceOutOfRange(cmd.DeviceID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError CLEAR_DEVICE_OOR";
        IncrementReadPointer();
        return;

    dte = ReadDeviceTable(UInt(cmd.DeviceID));

    if !dte.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError CLEAR_UNMAPPED_DEVICE";
        IncrementReadPointer();
        return;

    if IdOutOfRange(cmd.EventID, dte.ITT_size) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError CLEAR_ID_OOR";
        IncrementReadPointer();
        return;

    InterruptTableEntry ite = ReadTranslationTable(dte.ITT_base, UInt(cmd.EventID));

    if !ite.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError CLEAR_UNMAPPED_INTERRUPT";
        IncrementReadPointer();
        return;

    success = ClearPendingState(ite);

    if !success then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError CLEAR_ITE_INVALID";
        IncrementReadPointer();
        return;

    IncrementReadPointer();
    return;
```

6.3.4 DISCARD

This command translates the event defined by EventID and DeviceID and instructs the appropriate Redistributor to remove the pending state of the interrupt. It also ensures that any caching in the Redistributors associated with a specific EventID is consistent with the configuration held in memory. DISCARD removes the mapping of the DeviceID and EventID from the ITT, and ensures that incoming requests with a particular EventID are silently discarded.

Figure 6-7 on page 6-113 shows the format of the DISCARD command.

63 … 32	31 … 8	7 … 0	DW
DeviceID	RES0	0x0F	0
RES0	EventID		1
RES0			2
RES0			3

Figure 6-7 DISCARD command format

In Figure 6-7:

- `EventID` identifies the interrupt, associated with a device, that is to be discarded.
- `DeviceID` specifies the requesting device.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

The command and its arguments are:

```
DISCARD DeviceID, EventID
```

A command error occurs if any of the following apply:

- `DeviceID` exceeds the maximum value supported by the ITS.
- The device specified by `DeviceID` is not mapped to an ITT, using MAPD.
- `EventID` exceeds the maximum value allowed by the ITT. This value is specified by the `Size` field when the MAPD command is issued.
- The `EventID` for the device is not mapped to a collection, using MAPI or MAPTI.
- The `EventID` for the device is mapped to a collection that has not been mapped to an `RDbase` using MAPC.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

The following pseudocode describes the operation of the `DISCARD` command:

```
// ITS.DISCARD
// ===========

ITS.DISCARD(ITSCommand cmd)
    if DeviceOutOfRange(cmd.DeviceID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError DISCARD_DEVICE_OOR";
        IncrementReadPointer();
        return;

    DeviceTableEntry dte = ReadDeviceTable(UInt(cmd.DeviceID));

    if !dte.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError DISCARD_UNMAPPED_DEVICE";
        IncrementReadPointer();
        return;

    if IdOutOfRange(cmd.EventID, dte.ITT_size) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError DISCARD_ID_OOR";
        IncrementReadPointer();
        return;

    InterruptTableEntry ite = ReadTranslationTable(dte.ITT_base, UInt(cmd.EventID));
    if ite.Valid then
        success = ClearPendingState(ite);
        if !success then
            if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError DISCARD_ITE_INVALID";
            IncrementReadPointer();
            return;
    else
         if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError DISCARD_UNMAPPED_INTERRUPT";
            IncrementReadPointer();
            return;
```

```
    ite.Valid = FALSE;
    WriteTranslationTable(dte.ITT_base, UInt(cmd.EventID), ite);

    IncrementReadPointer();
    return;
```

6.3.5 INT

This command translates the event defined by `EventID and DeviceID` into an `ICID` and `pINTID`, and instructs the appropriate Redistributor to set the interrupt pending.

Figure 6-8 shows the format of the INT command.

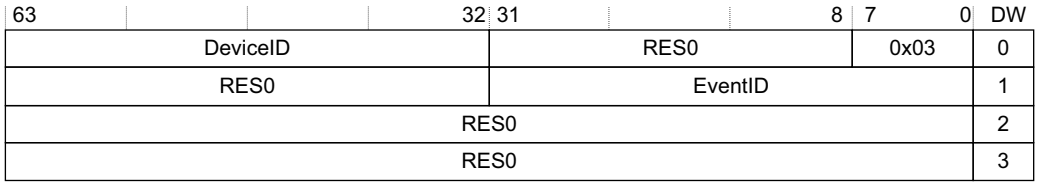

Figure 6-8 INT command format

In Figure 6-8:

- `EventID` identifies an interrupt source associated with a device. The ITS then translates this into an LPI INTID.
- `DeviceID` specifies the requesting device.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

The command and its arguments are:

```
INT DeviceID, EventID
```

A command error occurs if any of the following apply:

- `DeviceID` exceeds the maximum value supported by the ITS.
- The device specified by `DeviceID` is not mapped to an ITT, using MAPD.
- `EventID` exceeds the maximum value allowed by the ITT. This value is specified by the `Size` field when the MAPD command is issued.
- `EventID` is not mapped to a collection, using MAPI or MAPTI.
- The `EventID` for the device is mapped to a collection that has not been mapped to an `RDbase` using MAPC.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

The following pseudocode describes the operation of the INT command:

```
// ITS.INT
// =======

ITS.INT(ITSCommand cmd)
    if DeviceOutOfRange(cmd.DeviceID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError INT_DEVICE_OOR";
        IncrementReadPointer();
        return;

    DeviceTableEntry dte = ReadDeviceTable(UInt(cmd.DeviceID));

    if !dte.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError INT_UNMAPPED_DEVICE";
        IncrementReadPointer();
        return;

    if IdOutOfRange(cmd.EventID, dte.ITT_size) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError INT_ID_OOR";
        IncrementReadPointer();
        return;
```

```
InterruptTableEntry ite = ReadTranslationTable(dte.ITT_base, UInt(cmd.EventID));

if !ite.Valid then
    if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError INT_UNMAPPED_INTERRUPT";
    IncrementReadPointer();
    return;

boolean success = SetPendingState(ite);

if !success then
    if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError INT_ITE_INVALID";
    IncrementReadPointer();
    return;

IncrementReadPointer();
return;
```

6.3.6 INV

This command specifies that the ITS must ensure that any caching in the Redistributors associated with the specified `EventID` is consistent with the LPI Configuration tables held in memory.

Note

The `INV` command performs the same function regardless of whether the interrupt is mapped as a physical interrupt or a virtual interrupt.

Figure 6-9 shows the format of the `INV` command.

63 … 32	31 … 8	7 … 0	DW
DeviceID	RES0	0x0C	0
RES0	EventID		1
RES0			2
RES0			3

Figure 6-9 INV command format

In Figure 6-9:

- `EventID` identifies an interrupt source associated with a device. The ITS then translates this into an LPI INTID.
- `DeviceID` specifies the requesting device.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

The command and its arguments are:

```
INV DeviceID, EventID
```

A command error occurs if any of the following apply:

- `DeviceID` exceeds the maximum value supported by the ITS.
- The device specified by `DeviceID` is not mapped to an ITT, using MAPD.
- `EventID` exceeds the maximum value allowed by the ITT. This value is specified by the `Size` field when the MAPD command is issued.
- `EventID` is not mapped to a collection, using MAPI or MAPTI.
- The `EventID` for the device corresponds to a physical LPI and is mapped to a collection that has not been mapped to an `RDbase` using MAPC.
- The `EventID` for the device corresponds to a virtual LPI associated with a vPE that has not been mapped to a Redistributor using VMAPP.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

The following pseudocode describes the operation of the INV command:

```
// ITS.INV
// =======

ITS.INV(ITSCommand cmd)
    if DeviceOutOfRange(cmd.DeviceID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError INV_DEVICE_OOR";
        IncrementReadPointer();
        return;

    DeviceTableEntry dte = ReadDeviceTable(UInt(cmd.DeviceID));

    if !dte.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError INV_UNMAPPED_DEVICE";
        IncrementReadPointer();
        return;

    if IdOutOfRange(cmd.EventID, dte.ITT_size) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError INV_ID_OOR";
        IncrementReadPointer();
        return;

    InterruptTableEntry ite = ReadTranslationTable(dte.ITT_base, UInt(cmd.EventID));

    if !ite.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError INV_UNMAPPED_INTERRUPT";
        IncrementReadPointer();
        return;

    invalidateByITE(ite);

    IncrementReadPointer();
    return;
```

6.3.7 INVALL

This command specifies that the ITS must ensure any caching associated with the interrupt collection defined by ICID is consistent with the LPI Configuration tables held in memory for all Redistributors.

Figure 6-10 shows the format of the INVALL command.

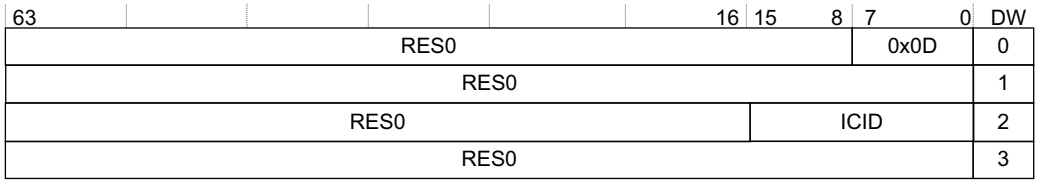

Figure 6-10 INVALL command format

In Figure 6-10:

- ICID specifies the interrupt collection.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

The command and its arguments are:

```
INVALL ICID
```

A command error occurs if any of the following apply:

- The collection specified by ICID exceeds the maximum number supported by the ITS.
- The collection specified by ICID has not been mapped to an RDbase using MAPC.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

The following pseudocode describes the operation of the INVALL command:

```
// ITS.INVALL
// ==========

ITS.INVALL(ITSCommand cmd)
    if (CollectionOutOfRange(cmd.ICID)) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError INVALL_COLLECTION_OOR";
        IncrementReadPointer();
        return;

    CollectionTableEntry cte = ReadCollectionTable(UInt(cmd.ICID));

    if !cte.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError INVALL_UNMAPPED_COLLECTION";
        IncrementReadPointer();
        return;

    // This invalidates any caches containing the configuration data for all interrupts in the
    // collection. Over invalidation is permitted.
    InvalidateCollectionCaches(UInt(cmd.ICID));

    IncrementReadPointer();
    return;
```

6.3.8 MAPC

This command maps the Collection table entry defined by ICID to the target Redistributor, defined by RDbase.

Figure 6-11 shows the format of the MAPC command.

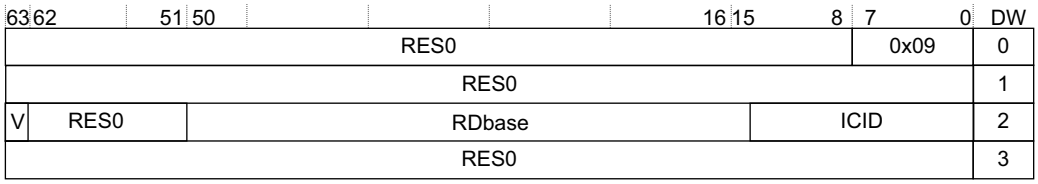

Figure 6-11 MAPC command format

In Figure 6-11:

- V specifies whether RDbase is valid for the collection.
- RDbase specifies the target Redistributor to which interrupts in the collection are forwarded. See *IMPLEMENTATION DEFINED sizes in ITS command parameters* on page 6-110.
- ICID specifies the interrupt collection that is to be mapped.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

If GITS_TYPER.PTA == 1 and a physical address is specified, the target addresses must be 64KB aligned, meaning that only bits[47:16] are required. See *IMPLEMENTATION DEFINED sizes in ITS command parameters* on page 6-110 for more information. In addition, when V is cleared to 0, this field must be written as zero, but hardware might ignore the value.

The command and its arguments are:

```
MAPC ICID, RDbase, V
```

When V is 1:

- Behavior is UNPREDICTABLE if there are interrupts that are mapped to the specified collection and the collection is currently mapped to a Redistributor, unless MAPC is followed by MOVALL to move the pending state for the collection from the old target Redistributor to the new target Redistributor. MOVALL might be issued by a different ITS.
- Behavior is UNPREDICTABLE if RDbase does not specify a valid Redistributor.

When V is 0:

- MAPC removes the mapping of the specified interrupt collection. Interrupts for that are mapped to this collection are ignored.
- Behavior is UNPREDICTABLE if there are interrupts that are mapped to the specified collection, with the restriction that further translation requests from that device are ignored.

A command error occurs if the following applies:

- The collection specified by ICID exceeds the maximum number supported by the ITS.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

Note

When software uses a MAPC command to move a collection from targeting Redistributor A to targeting Redistributor B, it must issue a SYNC command to Redistributor A before issuing the accompanying MOVALL command. Otherwise, interrupts from the collection might still be taken by the PE associated with Redistributor A.

The following pseudocode describes the operation of the MAPC command:

```
// ITS.MAPC
// ========

ITS.MAPC(ITSCommand cmd)
    if CollectionOutOfRange(cmd.ICID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MAPC_COLLECTION_OOR";
        IncrementReadPointer();
        return;


    CollectionTableEntry cte;

    cte.Valid  = cmd.V == '1';
    cte.RDbase = cmd.RDbase;

    WriteCollectionTable(UInt(cmd.ICID), cte);

    IncrementReadPointer();
    return;
```

6.3.9 MAPD

This command maps the Device table entry associated with DeviceID to its associated ITT, defined by ITT_addr and Size.

Software might issue a MAPD command to remap a device that is already mapped, in which case an ITS must invalidate all cached data for the device.

Figure 6-12 shows the format of the MAPD command.

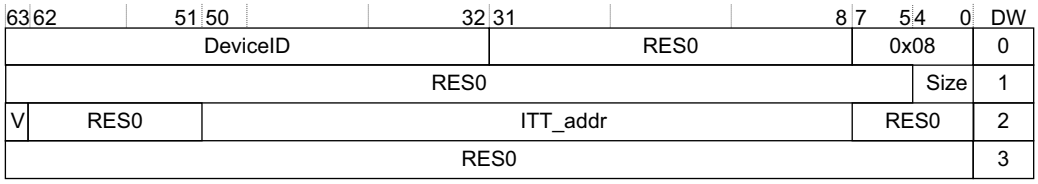

Figure 6-12 MAPD command format

In Figure 6-12:

- DeviceID specifies the device that uses the ITT.

Note

For more information about mapping devices to ITTs, see *The Interrupt translation table* on page 6-103.

- `V` specifies whether the `ITT_addr` and `Size` fields are valid.
- `ITT_addr` specifies bits[51:8] of the physical address of the ITT.
- `Size` is a 5-bit number that specifies the supported number of bits for the device, minus one. The size field enables range checking of EventID for translation requests for this DeviceID.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

If `DeviceID` is mapped to an ITT that contains valid mappings, behavior is UNPREDICTABLE.

The command and its arguments are:

`MAPD DeviceID, ITT_addr, Size, V`

The format of the ITT entries is IMPLEMENTATION DEFINED. A typical example entry size of 8 bytes permits allocation of identifiers to devices in multiples of 32 interrupts.

When `V` is 1:

- `MAPD` associates a `DeviceID` with a 256 byte aligned address of an ITT.

When `V` is 0:

- `MAPD` removes the mapping for the specified `DeviceID`. Translation requests from that device are ignored.
- `MAPD` removes the mapping of the specified `DeviceID`. and interrupt requests from that device are discarded. A subsequent translation for the `DeviceID` does not generate and LPI or VLPI until `DeviceID` has been mapped to the ITT again.

A command error occurs if any of the following apply:

- `DeviceID` exceeds the maximum number of devices supported by an ITS.
- `Size` exceeds the maximum value permitted by the settings of GITS_TYPER.IDbits, when `V` is set to 1.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

Note

- Software might issue a `MAPD` command to re-map a device that has already been mapped and the ITS must invalidate all cached data for that device.
- ITS accesses to an ITT use the same Shareability and Cacheability attributes that are specified for the Device table, see *The Device table* on page 6-102.

The following pseudocode describes the operation of the `MAPD` command:

```
// ITS.MAPD
// ========

ITS.MAPD(ITSCommand cmd)
    if DeviceOutOfRange(cmd.DeviceID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MAPD_DEVICE_OOR";
        IncrementReadPointer();
        return;

    if SizeOutOfRange(cmd.Size) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MAPD_ITTSIZE_OOR";
        IncrementReadPointer();
        return;

    // If a device is Re-mapped software must perform the following actions
    // to ensure the LPI configuration is up to date:
    // 1.  Ensure that the device is quiescent and that all interrupts have
    //     been handled.
    // 2.  Remap the device with the new (empty) ITT
    //
```

```
    DeviceTableEntry dte;

    dte.Valid    = cmd.V == '1';
    dte.ITT_base = cmd.ITT_addr:'00000000';
    dte.ITT_size = cmd.Size;

    WriteDeviceTable(UInt(cmd.DeviceID), dte);

    IncrementReadPointer();
    return;
```

6.3.10 MAPI

This command maps the event defined by `EventID and DeviceID` into an ITT entry with `ICID` and `pINTID = EventID`.

Note

- `pINTID` ≥`0x2000` for a valid LPI INTID.
- This is equivalent to `MAPTI DeviceID, EventID, EventID, ICID`

Figure 6-13 shows the format of the `MAPI` command.

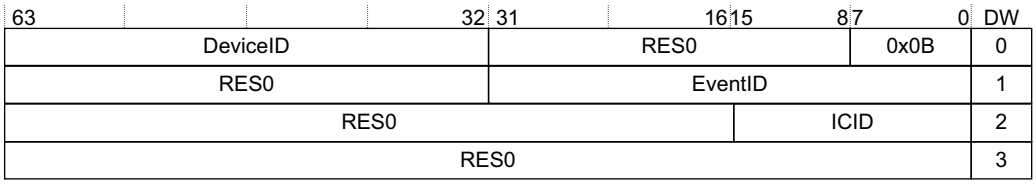

Figure 6-13 MAPI command format

In Figure 6-13:

- `EventID` identifies the interrupt, associated with a device, that is to be mapped.
- `DeviceID` specifies the requesting device.
- `ICID` specifies the interrupt collection that includes the specified interrupt.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

If there is an existing mapping for the `EventID-DeviceID` combination, behavior is UNPREDICTABLE.

The command and its arguments are:

```
MAPI DeviceID, EventID, ICID
```

A command error occurs if any of the following apply:

- `DeviceID` exceeds the maximum value supported by the ITS.
- The device specified by `DeviceID` is not mapped to an ITT, using MAPD.
- `ICID` exceeds the maximum number of interrupt collections supported by an ITS. For more information, see *The Collection table* on page 6-104.
- The `EventID` exceeds the maximum value allowed by the ITT. This value is specified by the `Size` field when the MAPD command is issued.
- `EventID` does not specify a valid LPI identifier. See *INTIDs* on page 3-39.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

The following pseudocode describes the operation of the `MAPI` command:

```
// ITS.MAPI
// ========

ITS.MAPI(ITSCommand cmd)
    if DeviceOutOfRange(cmd.DeviceID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MAPI_DEVICE_OOR";
        IncrementReadPointer();
```

```
        return;

    if CollectionOutOfRange(cmd.ICID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MAPI_COLLECTION_OOR";
        IncrementReadPointer();
        return;

    DeviceTableEntry dte = ReadDeviceTable(UInt(cmd.DeviceID));

    if !dte.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MAPI_UNMAPPED_DEVICE";
        IncrementReadPointer();
        return;

    if IdOutOfRange(cmd.EventID, dte.ITT_size) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MAPI_ID_OOR";
        IncrementReadPointer();
        return;

    if LPIOutOfRange(cmd.EventID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MAPI_ID_OOR";
        IncrementReadPointer();
        return;

    InterruptTableEntry ite = ReadTranslationTable(dte.ITT_base, UInt(cmd.EventID));

    ite.Valid        = TRUE;
    ite.Type         = physical_interrupt;
    ite.OutputID     = cmd.EventID;
    ite.DoorbellID   = ZeroExtend(INTID_SPURIOUS);      // Don't generate a doorbell
    ite.ICID         = cmd.ICID;

    WriteTranslationTable(dte.ITT_base, UInt(cmd.EventID), ite);

    IncrementReadPointer();
    return;
```

6.3.11 MAPTI

This command maps the event defined by `EventID` and `DeviceID` to its associated ITE, defined by `ICID` and `pINTID` in the ITT associated with `DeviceID`.

Figure 6-14 shows the format of the `MAPTI` command.

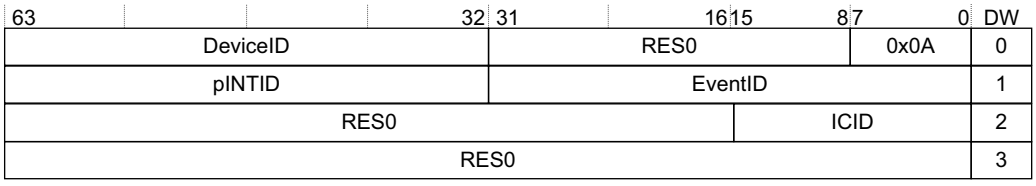

Figure 6-14 MAPTI command format

In Figure 6-14:

- `EventID` identifies the interrupt, associated with a device, that is to be mapped.
- `pINTID` is the INTID of the physical interrupt that is presented to software.
- `DeviceID` specifies the requesting device.
- `ICID` specifies the interrupt collection that includes the specified physical interrupt.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

If there is an existing mapping for the `EventID-DeviceID` combination, behavior is UNPREDICTABLE.

The command and its arguments are:

```
MAPTI DeviceID, EventID, pINTID, ICID
```

A command error occurs if any of the following apply:

- DeviceID exceeds the maximum value supported by the ITS.
- The device specified by DeviceID is not mapped to an ITT using MAPD.
- The number of collections exceeds the maximum number of collections supported by the ITS. For more information, see *The Collection table* on page 6-104.
- The EventID exceeds the maximum value allowed by the ITT. This value is specified by the Size field when the MAPD command is issued.
- pINTID does not specify a valid LPI INTID. For information about the LPI INTID range, see *INTIDs* on page 3-39.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

The following pseudocode describes the operation of the MAPTI command:

```
// ITS.MAPTI
// =========

ITS.MAPTI(ITSCommand cmd)
    if DeviceOutOfRange(cmd.DeviceID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MAPTI_DEVICE_OOR";
        IncrementReadPointer();
        return;

    if CollectionOutOfRange(cmd.ICID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MAPTI_COLLECTION_OOR";
        IncrementReadPointer();
        return;

    DeviceTableEntry dte = ReadDeviceTable(UInt(cmd.DeviceID));

    if !dte.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MAPTI_UNMAPPED_DEVICE";
        IncrementReadPointer();
        return;

    if IdOutOfRange(cmd.EventID, dte.ITT_size) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MAPTI_ID_OOR";
        IncrementReadPointer();
        return;

    if LPIOutOfRange(cmd.pINTID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MAPTI_PHYSICALID_OOR";
        IncrementReadPointer();
        return;

    InterruptTableEntry ite = ReadTranslationTable(dte.ITT_base, UInt(cmd.EventID));

    ite.Valid        = TRUE;
    ite.Type         = physical_interrupt;
    ite.OutputID     = cmd.pINTID;
    ite.DoorbellID   = ZeroExtend(INTID_SPURIOUS);              // Don't generate a doorbell
    ite.ICID         = cmd.ICID;

    WriteTranslationTable(dte.ITT_base, UInt(cmd.EventID), ite);

    IncrementReadPointer();
    return;
```

6.3.12 MOVALL

This command instructs the Redistributor specified by RDbase1 to move all of its interrupts to the Redistributor specified by RDbase2.

Note

Both the mapping of interrupts to collections and the mapping of collections to Redistributors are normally unaffected by this command. Software must ensure that any interrupts that might be affected by this command target the Redistributor specified by RDbase2, otherwise system behavior is UNPREDICTABLE. In particular, an implementation might choose to remap all affected collections to RDbase2.

Figure 6-15 shows the format of the MOVALL command.

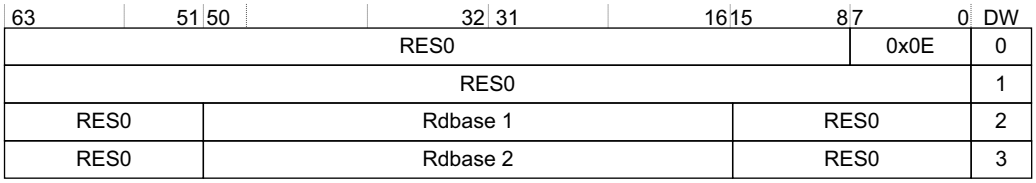

Figure 6-15 MOVALL command format

In Figure 6-15:

- RDbase1 specifies the Redistributor with which the interrupts are currently associated. See *IMPLEMENTATION DEFINED sizes in ITS command parameters* on page 6-110.
- RDbase2 specifies the Redistributor to which the interrupts are to be moved. See *IMPLEMENTATION DEFINED sizes in ITS command parameters* on page 6-110.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

The command and its arguments are:

```
MOVALL RDbase1, RDbase2
```

Behavior is UNPREDICTABLE if RDbase1 and RDbase2 do not specify a valid Redistributor. The format of these fields is specified by GITS_TYPER.PTA.

The following pseudocode describes the operation of the MOVALL command:

```
// ITS.MOVALL
// ==========

ITS.MOVALL(ITSCommand cmd)
    rd1 = cmd.RD1base;
    rd2 = cmd.RD2base;

    if rd1 != rd2 then
        MoveAllPendingState(rd1, rd2);

    IncrementReadPointer();
    return;
```

6.3.13 MOVI

This command updates the ICID field in the ITT entry for the event defined by DeviceID and EventID. It also translates the event defined by EventID and DeviceID into an ICID and pINTID, and instructs the appropriate Redistributor to move the pending state, if it is set, of the interrupt to the Redistributor defined by the new ICID, and to update the ITE associated with the event to use the new ICID.

Figure 6-16 on page 6-124 shows the format of the MOVI command.

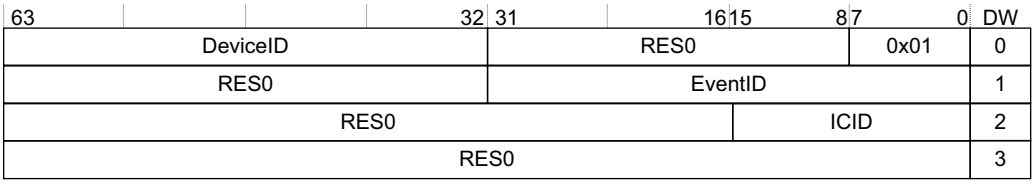

Figure 6-16 MOVI command format

In Figure 6-16:

- `EventID` identifies the interrupt, associated with a device, that is to be redirected.
- `DeviceID` specifies the requesting device.
- `ICID` specifies the new interrupt collection that is to include the specified physical interrupt.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

The command and its arguments are:

```
MOVI DeviceID, EventID, ICID
```

A command error occurs if any of the following apply:

- `DeviceID` exceeds the maximum value supported by the ITS.
- The device specified by `DeviceID` is not mapped to an ITT, using MAPD.
- `ICID` exceeds the maximum number of interrupt collections supported by an ITS.
- `ICID` is not mapped to an `RDbase` using MAPC.
- `EventID` is not mapped to a collection, using MAPI or MAPTI.
- `EventID` corresponds to a virtual LPI.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

Note

If, after using `MOVI` to move an interrupt from collection A to collection B, software moves the same interrupt again from collection B to collection C, a SYNC command must be used before the second MOVI for the Redistributor associated with collection A to ensure correct behavior.

The following pseudocode describes the operation of the `MOVI` command:

```
// ITS.MOVI
// ========

ITS.MOVI(ITSCommand cmd)
    if DeviceOutOfRange(cmd.DeviceID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MOVI_DEVICE_OOR";
        IncrementReadPointer();
        return;

    if CollectionOutOfRange(cmd.ICID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MOVI_COLLECTION_OOR";
        IncrementReadPointer();
        return;

    DeviceTableEntry dte = ReadDeviceTable(UInt(cmd.DeviceID));

    if !dte.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MOVI_UNMAPPED_DEVICE";
        IncrementReadPointer();
        return;

    if IdOutOfRange(cmd.EventID, dte.ITT_size) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MOVI_ID_OOR";
        IncrementReadPointer();
        return;
```

```
InterruptTableEntry ite = ReadTranslationTable(dte.ITT_base, UInt(cmd.EventID));

if !ite.Valid then
    if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MOVI_UNMAPPED_INTERRUPT";
    IncrementReadPointer();
    return;

if ite.Type == virtual_interrupt then
    if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MOVI_ID_IS_VIRTUAL";
    IncrementReadPointer();
    return;

CollectionTableEntry cte1 = ReadCollectionTable(UInt(ite.ICID));

if !cte1.Valid then
    if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MOVI_UNMAPPED_COLLECTION";
    IncrementReadPointer();
    return;

CollectionTableEntry cte2 = ReadCollectionTable(UInt(cmd.ICID));

if !cte2.Valid then
    if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError MOVI_UNMAPPED_COLLECTION";
    IncrementReadPointer();
    return;

bits(32) rd1 = cte1.RDbase;
bits(32) rd2 = cte2.RDbase;

if rd1 != rd2 then
    // Move the move the pending state to rd2 if set taking care of any races where the
    // interrupt has been forwarded to the processor
    MovePendingState(rd1, rd2, ite.OutputID);

ite.ICID = cmd.ICID;

WriteTranslationTable(dte.ITT_base, UInt(cmd.EventID), ite);

IncrementReadPointer();
return;
```

6.3.14 SYNC

This command ensures all outstanding ITS operations associated with physical interrupts for the Redistributor specified by RDbase are globally observed before any further ITS commands are executed. Following the execution of a SYNC, the effects of all previous commands must apply to subsequent writes to GITS_TRANSLATER.

Figure 6-17 shows the format of the SYNC command.

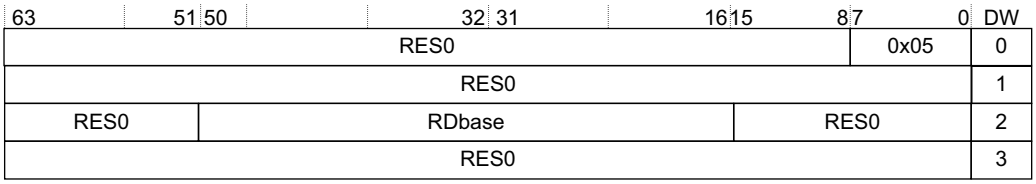

Figure 6-17 SYNC command format

In Figure 6-17:

- RDbase specifies the physical address of the target Redistributor. The format of the target address is determined by GITS_TYPER.PTA. See *IMPLEMENTATION DEFINED sizes in ITS command parameters on page 6-110* for more information.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

The command and its arguments are:

```
SYNC RDbase
```

The following pseudocode describes the operation of the SYNC command:

```
// ITS.SYNC
// ========

ITS.SYNC(ITSCommand cmd)
    // Wait for external effects of any physical comamnds to be observable by all redistributors
    // and ensure the internal effects of any previous commands affect any subsequent interrupt
    // requests or commands
    WaitForCompletion(cmd.RDbase);

    IncrementReadPointer();
```

6.3.15 VINVALL

This command ensures that any cached Redistributor information associated with vPEID is consistent with the associated LPI Configuration tables held in memory.

This command is provided only in GICv4.

Figure 6-18 shows the format of the VINVALL command.

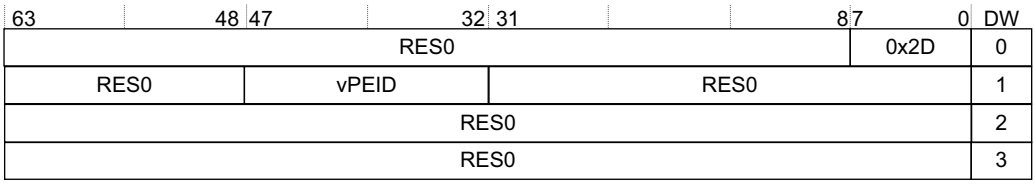

Figure 6-18 VINVALL command format

In Figure 6-18:

- vPEID specifies the vPE.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

The command and its arguments are:

```
VINVALL vPEID
```

A command error occurs if any of the following apply:

- vPEID exceeds the maximum number supported by the ITS, as defined by GITS_BASER<n>.
- The PE specified by vPEID is not mapped to a Redistributor using VMAPP.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

The following pseudocode describes the operation of the VINVALL command:

```
// ITS.VINVALL
// ===========

ITS.VINVALL(ITSCommand cmd)
    if VCPUOutOfRange(cmd.VCPUID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VINVALL_VCPU_OOR";
        IncrementReadPointer();
        return;

    VCPUTableEntry vte = ReadVCPUTable(UInt(cmd.VCPUID));

    if !vte.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VINVALL_VCPU_INVALID";
        IncrementReadPointer();
        return;

    InvalidateVCPUCaches(UInt(cmd.VCPUID));
```

```
    IncrementReadPointer();
    return;
```

6.3.16 VMAPI

This command maps the event defined by DeviceID and EventID into an ITT entry with vPEID, vINTID=EventID, and Dbell_PINTID, a doorbell provision.

Note

- vINTID ≥0x2000 for a valid LPI INTID.
- This is equivalent to VMAPTI DeviceID, EventID,EventID, pINTID, vPEID.
- Dbell_pINTID must be either 1023 or Dbell_pINTID ≥0x2000 for a valid LPI INTID.

This command is provided only in GICv4.

Figure 6-19 shows the format of the VMAPI command.

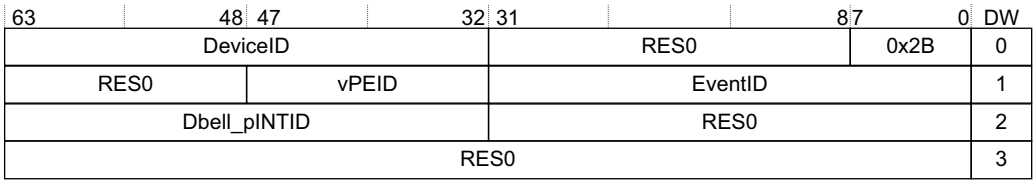

Figure 6-19 VMAPI command format

In Figure 6-19:

- EventID identifies the interrupt, associated with a device, that is to be presented to the VM.
- DeviceID specifies the requesting device.
- vPEID specifies the vPE.
- Dbell_pINTID specifies the ID that is presented to the hypervisor if the vPE is not scheduled.

  Note

  If Dbell_pINTID indicates a spurious interrupt, then no physical interrupt is generated.

- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

If there is an existing mapping for the EventID-DeviceID combination, behavior is UNPREDICTABLE.

The command and its arguments are:

```
VMAPI DeviceID, EventID, Dbell_pINTID, vPEID
```

A command error occurs if any of the following apply:

- DeviceID exceeds the maximum value supported by the ITS.
- vPEID exceeds the maximum number supported by the ITS, as defined by GITS_BASER<n>.
- The device specified by DeviceID is not mapped to an ITT, using MAPD.
- EventID exceeds the maximum value allowed by the ITT. This value is specified by the Size field when the MAPD command is issued.
- EventID does not specify a valid LPI INTID. For information about valid LPI INTIDs, see *INTIDs* on page 3-39.
- Dbell_pINTID does not specify a valid doorbell INTID, where a valid INTID is either:
  - 1023.
  - Within the supported range for LPIs.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

The following pseudocode describes the operation of the VMAPI command:

```
// ITS.VMAPI
// =========

ITS.VMAPI(ITSCommand cmd)
    if DeviceOutOfRange(cmd.DeviceID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMAPI_DEVICE_OOR";
        IncrementReadPointer();
        return;

    if VCPUOutOfRange(cmd.VCPUID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMAPI_VCPU_OOR";
        IncrementReadPointer();
        return;

    DeviceTableEntry dte = ReadDeviceTable(UInt(cmd.DeviceID));

    if !dte.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMAPI_UNMAPPED_DEVICE";
        IncrementReadPointer();
        return;

    if IdOutOfRange(cmd.EventID, dte.ITT_size) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMAPI_ID_OOR";
        IncrementReadPointer();
        return;

    if LPIOutOfRange(cmd.EventID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMAPI_ID_OOR";
        IncrementReadPointer();
        return;

    if LPIOutOfRange(cmd.Dbell_pINTID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMAPI_PHYSICALID_OOR";
        IncrementReadPointer();
        return;

    InterruptTableEntry ite = ReadTranslationTable(dte.ITT_base, UInt(cmd.EventID));

    ite.Valid        = TRUE;
    ite.Type         = virtual_interrupt;
    ite.OutputID     = cmd.EventID;
    ite.DoorbellID   = cmd.Dbell_pINTID;
    ite.VCPUID       = cmd.VCPUID;

    WriteTranslationTable(dte.ITT_base, UInt(cmd.EventID), ite);

    IncrementReadPointer();
    return;
```

6.3.17 VMAPP

This command maps the vPE table entry defined by `vPEID` to the target `RDbase`, including an associated virtual LPI Pending table (`VPT_addr`, `VPT_size`).

Figure 6-20 shows the format of the `VMAPP` command.

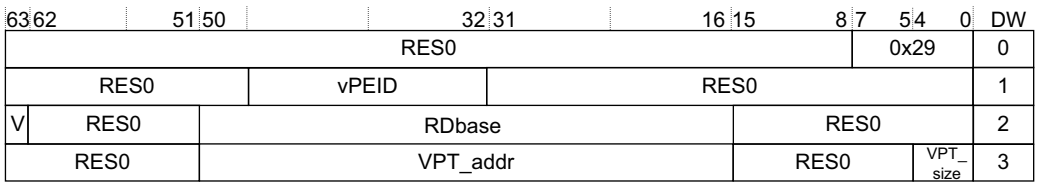

Figure 6-20 VMAPP command format

In Figure 6-20 on page 6-128:

- vPEID specifies the vPE.
- V specifies whether the RDbase and VPT_addr are valid for the vPE.
- RDbase specifies the target Redistributor that owns the vPE and to which the ITS directs commands for that PE. See *IMPLEMENTATION DEFINED sizes in ITS command parameters* on page 6-110.
- VPT_addr specifies bits [51:16] of the physical address of the virtual LPI Pending table for the vPE.

  Note

  The target addresses must be 64KB aligned, meaning that only bits [51:16] are required. Bits[15:0] of the physical address are 0.

- VPT_size specifies the number of vINTID bits that the vPE supports, minus one.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

The command and its arguments are:

```
VMAPP vPEID, RDbase, VPT_addr, VPT_size, V
```

When V is 0:

- VMAPP removes the mapping for the specified vPE. Interrupts that are mapped to this vPE are discarded.

When V is 1:

- Behavior is UNPREDICTABLE if RDbase does not specify a valid Redistributor.

A command error occurs if any of the following apply:

- vPEID exceeds the maximum number supported by the ITS, as defined by GITS_BASER<n>.
- Size exceeds the maximum value permitted by the settings of GITS_TYPER.IDbits.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

The following pseudocode describes the operation of the VMAPP command:

```
// ITS.VMAPP
// =========

ITS.VMAPP(ITSCommand cmd)
    if VCPUOutOfRange(cmd.VCPUID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMAPP_VCPU_OOR";
        IncrementReadPointer();
        return;

    if SizeOutOfRange(cmd.VPT_size) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMAPP_VPTSIZE_OOR";
        IncrementReadPointer();
        return;

    VCPUTableEntry vte;

    vte.Valid    = cmd.V == '1';
    vte.RDbase   = cmd.RDbase;
    vte.VPT_base = cmd.VPT_addr:Zeros(16);
    vte.VPT_size = cmd.VPT_size;

    WriteVCPUTable(UInt(cmd.VCPUID), vte);

    IncrementReadPointer();
    return;
```

6.3.18 VMAPTI

This command maps the event defined by DeviceID and EventID into an ITT entry with vPEID and vINTID, and Dbell_pINTID, a doorbell provision.

Note

- vINTID ≥0x2000 for a valid LPI INTID.
- Dbell_pINTID must be either 1023 or Dbell_pINTID ≥0x2000 for a valid LPI INTID.

This command is provided only in GICv4.

Figure 6-21 shows the format of the VMAPTI command.

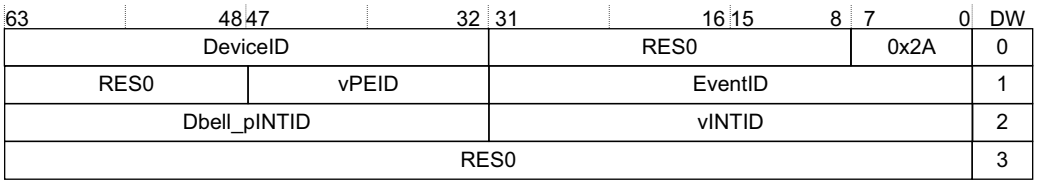

Figure 6-21 VMAPTI command format

In Figure 6-21:

- vPEID specifies the vPE.
- DeviceID specifies a device owned by the vPE.
- vINTID specifies the INTID presented to the vPE that controls the device that DeviceID specifies.
- Dbell_pINTID specifies the pINTID that is presented to the PE if the vPE is not scheduled.

  Note

  If Dbell_pINTID is 1023 then no physical interrupt is generated.

- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

If there is an existing mapping for the EventID-DeviceID combination, behavior is UNPREDICTABLE.

The command and its arguments are:

```
VMAPTI DeviceID, EventID, vINTID, Dbell_pINTID, vPEID
```

A command error occurs if any of the following apply:

- DeviceID exceeds the maximum value supported by the ITS.
- vPEID exceeds the maximum number supported by the ITS, as defined by GITS_BASER<n>.
- The device specified by DeviceID is not mapped to an ITT, using MAPD.
- The EventID exceeds the maximum value allowed by the ITT. This value is specified by the Size field when the MAPD command is issued.
- vINTID does not specify a valid LPI INTID, see *INTIDs* on page 3-39.
- Dbell_pINTID does not specify a valid doorbell INTID, where a valid INTID is either:
  - 1023.
  - Within the supported range for LPIs.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

The following pseudocode describes the operation of the VMAPTI command:

```
// ITS.VMAPTI
// ==========

ITS.VMAPTI(ITSCommand cmd)
    if DeviceOutOfRange(cmd.DeviceID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMAPTI_DEVICE_OOR";
        IncrementReadPointer();
        return;

    if VCPUOutOfRange(cmd.VCPUID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMAPTI_VCPU_OOR";
        IncrementReadPointer();
```

```
    return;

DeviceTableEntry dte = ReadDeviceTable(UInt(cmd.DeviceID));

if !dte.Valid then
    if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMAPTI_UNMAPPED_DEVICE";
    IncrementReadPointer();
    return;

if IdOutOfRange(cmd.EventID, dte.ITT_size) then
    if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMAPTI_ID_OOR";
    IncrementReadPointer();
    return;

if LPIOutOfRange(cmd.vINTID) then
    if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMAPTI_VIRTUALID_OOR";
    IncrementReadPointer();
    return;

if LPIOutOfRange(cmd.pINTID) then
    if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMAPTI_PHYSICALID_OOR";
    IncrementReadPointer();
    return;

InterruptTableEntry ite = ReadTranslationTable(dte.ITT_base, UInt(cmd.EventID));

ite.Valid        = TRUE;
ite.Type         = virtual_interrupt;
ite.OutputID     = cmd.vINTID;
ite.DoorbellID   = cmd.Dbell_pINTID;
ite.VCPUID       = cmd.VCPUID;

WriteTranslationTable(dte.ITT_base, UInt(cmd.EventID), ite);

IncrementReadPointer();
return;
```

6.3.19 VMOVI

This command updates the vPEID field in the ITT entry for the event defined by DeviceID and EventID. It also translates the event defined by EventID and DeviceID into a vPEID and pINTID, and instructs the appropriate Redistributor to move the pending state of the interrupt to the Redistributor defined by the new vPEID, and updates the ITE associated with the event to use the new vPEID.

This command is provided only in GICv4.

Figure 6-22 shows the format of the VMOVI command.

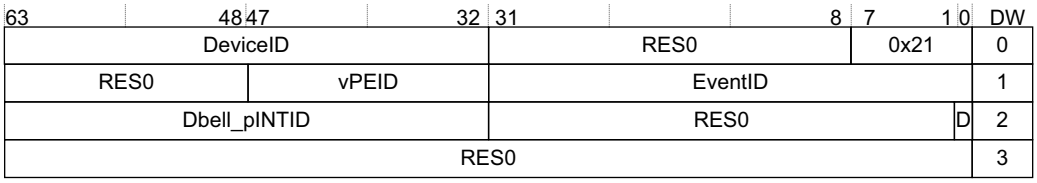

Figure 6-22 VMOVI command format

In Figure 6-22:

- vPEID specifies the vPE.
- EventID identifies the interrupt, associated with a device and already mapped by the ITS, that is to be moved to a new target specified by vPEID.
- D specifies whether the Dbell_pINTID field is valid.
- DeviceID specifies the device that generates the interrupt.
- Dbell_pINTID specifies the ID that is presented to the hypervisor if the vPE is not scheduled.

Note

If Dbell_pINTID is 1023 then no physical interrupt is generated.

- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

The command and its arguments are:

VMOVI DeviceID, EventID, vPEID, [Dbell_pINTID]

A command error occurs if any of the following apply:

- DeviceID exceeds the maximum value supported by the ITS.
- vPEID exceeds the maximum number supported by the ITS, as defined by GITS_BASER<n>.
- The device specified by DeviceID is not mapped to an ITT, using MAPD.
- The EventID exceeds the maximum value allowed by the ITT. This value is specified by the Size field when the MAPD command is issued.
- The vPE is not mapped to a Redistributor, using VMAPP.
- EventID corresponds to a physical LPI.
- If D is 1 and pINTID does not specify a valid doorbell INTID, where a valid INTID is either:
  - 1023.
  - Within the supported range for LPIs.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

Note

If, after using VMOVI to move an interrupt from vPE A to vPE B, software moves the same interrupt again, a VSYNC command must be issued to vPE A between the moves to ensure correct behavior.

The following pseudocode describes the operation of the VMOVI command:

```
// ITS.VMOVI
// =========

ITS.VMOVI(ITSCommand cmd)

    if DeviceOutOfRange(cmd.DeviceID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMOVI_DEVICE_OOR";
        IncrementReadPointer();
        return;

    if VCPUOutOfRange(cmd.VCPUID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMOVI_COLLECTION_OOR";
        IncrementReadPointer();
        return;

    if ( cmd.V == '1' && LPIOutOfRange(cmd.pINTID) && cmd.pINTID != '1023') then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMOVI_PHYSICALID_OOR";
        IncrementReadPointer();
        Return;

    DeviceTableEntry dte = ReadDeviceTable(UInt(cmd.DeviceID));

    if !dte.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMOVI_UNMAPPED_DEVICE";
        IncrementReadPointer();
        return;

    if IdOutOfRange(cmd.EventID, dte.ITT_size) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMOVI_ID_OOR";
        IncrementReadPointer();
        return;

    InterruptTableEntry ite = ReadTranslationTable(dte.ITT_base, UInt(cmd.EventID));
```

```
    if !ite.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMOVI_UNMAPPED_INTERRUPT";
        IncrementReadPointer();
        return;

    if ite.Type == physical_interrupt then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMOVI_ID_IS_PHYSICAL";
        IncrementReadPointer();
        return;

    VCPUTableEntry vte1 = ReadVCPUTable(UInt(ite.VCPUID));
    VCPUTableEntry vte2 = ReadVCPUTable(UInt(cmd.VCPUID));

    if !vte1.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMOVI_ITEVCPU_INVALID";
        IncrementReadPointer();
        return;

    if !vte2.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMOVI_CMDVCPU_INVALID";
        IncrementReadPointer();
        return;


    bits(32) rd1 = vte1.RDbase;
    Address vpt1 = vte1.VPT_base;
    bits(32) rd2 = vte2.RDbase;
    Address vpt2 = vte2.VPT_base;

    ite.VCPUID = cmd.VCPUID;

    if cmd.V == '1' then
        ite.DoorbellID = cmd.pINTID;


    WriteTranslationTable(dte.ITT_base, UInt(cmd.EventID), ite);
    // From this point new interrupts sent to the new VCPU move the pending state to rd2 if set taking
care of any races where the interrupt
    // has been forwarded to the processor
    MoveVirtualPendingState(rd1, vpt1, vpt2, ite.OutputID);


    IncrementReadPointer();
    return;
```

6.3.20 VMOVP

This command updates the vPE table entry defined by `vPEID` to the target `RDbase`. Software must use `SequenceNumber` and `ITSList` to synchronize the execution of VMOVP commands across more than one ITS.

This command is provided only in GICv4.

Software must ensure that this command is not executed with a `vPEID` that is scheduled on the target Redistributor, otherwise system behavior is UNPREDICTABLE.

Figure 6-23 shows the format of the `VMOVP` command.

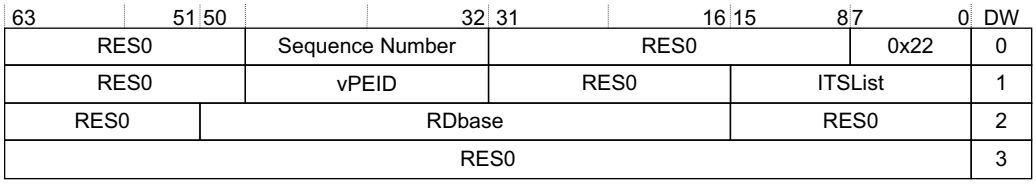

Figure 6-23 VMOVP command format

In Figure 6-23 on page 6-133:

- vPEID specifies the vPE.
- RDbase specifies the Redistributor to which interrupts are forwarded. See *IMPLEMENTATION DEFINED sizes in ITS command parameters* on page 6-110.
- Sequence Number specifies the identity of the synchronization point that every ITS included in ITS List uses.

  When GITS_TYPER.VMOVP == 0 Sequence Number specifies the identity of the synchronization point that is used by all ITSs that are included in ITSList.

  When GITS_TYPER.VMOVP == 1 Sequence Number is RES0.

  For more information, see *Use of the sequence number field*.
- ITSList specifies the ITS instances that are included in the synchronization operation, where:
  - Each bit in ITS List identifies an ITS where bit[n] corresponds to ITS n.
  - An ITS is included if the corresponding bit is set to 1.

  When GITS_TYPER.VMOVP == 0 ITSList specifies which ITSs are included in the synchronization operation. Each bit of ITSList corresponds to an ITS, for example bit[0] of ITSList corresponds to ITS 0, bit[1] to ITS 1.

  When GITS_TYPER.VMOVP ==1 ITSList is RES0.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

The command and its arguments are:

```
VMOVP, vPEID, RDbase, SequenceNumber, ITSList
```

A command error occurs if any of the following apply:

- If the PE specified by vPEID is not mapped to a Redistributor, using VMAPP.
- vPEID exceeds the maximum number supported by the ITS, as defined by GITS_BASER<n>.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

The following pseudocode describes the operation of the VMOVP command:

```
// ITS.VMOVP
// =========

ITS.VMOVP(ITSCommand cmd)
    if VCPUOutOfRange(cmd.VCPUID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMOVP_VCPU_OOR";
        IncrementReadPointer();
        return;

    VCPUTableEntry vte = ReadVCPUTable(UInt(cmd.VCPUID));

    if !vte.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VMOVP_VCPU_INVALID";
        IncrementReadPointer();
        return;

    vte.RDbase = cmd.RDbase;

    WriteVCPUTable(UInt(cmd.VCPUID), vte);

    IncrementReadPointer();
    return;
```

Use of the sequence number field

Where more than one ITS controls interrupts for the same vPE, moving this vPE must be co-ordinated between the different ITSs. This is controlled by software, as follows:

- The VMOVP command must be issued for each ITS that controls interrupts for the vPE that is being moved. Each of these commands must have a common sequence number. That sequence number cannot be used for other VMOVP commands until all commands that previously used that sequence number have been processed by all ITSs.

- The VMOVP command issued for each ITS contains a list of all the ITSs that are affected by moving the vPE. This is the ITS List.
- Each ITS must have the sequence numbers presented to it in the same order in that they are presented to the other ITSs.

Not following this approach results in UNPREDICTABLE behavior.

6.3.21 VSYNC

This command ensures all outstanding ITS operations for the vPEID specified are globally observed before any further ITS commands are executed. Following the execution of a VSYNC the effects of all previous commands must apply to subsequent writes to GITS_TRANSLATER.

This command is provided only in GICv4.

Figure 6-24 shows the format of the VSYNC command.

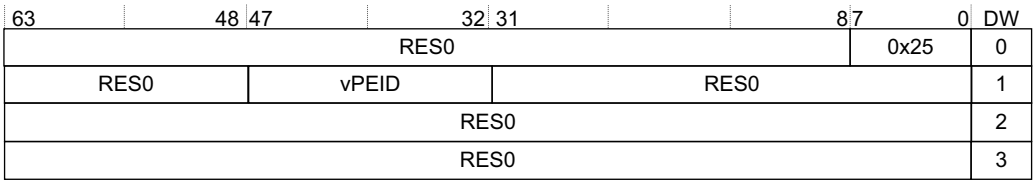

Figure 6-24 VSYNC command format

In Figure 6-24:

- vPEID specifies the vPE for which commands must be synchronized.
- DW is the doubleword offset within a 32 byte, or four doubleword, ITS command packet.

The command and its arguments are:

```
VSYNC vPEID
```

A command error occurs if any of the following apply:

- If the PE specified by vPEID is not mapped to a Redistributor, using VMAPP.
- vPEID exceeds the maximum number supported by the ITS, as defined by GITS_BASER<n>.

In this case, the ITS must take the actions described in *Command errors* on page 6-111.

The following pseudocode describes the operation of the VSYNC command:

```
// ITS.VSYNC
// =========

ITS.VSYNC(ITSCommand cmd)
    if VCPUOutOfRange(cmd.VCPUID) then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VSYNC_VCPU_OOR";
        IncrementReadPointer();
        return;

    VCPUTableEntry vte = ReadVCPUTable(UInt(cmd.VCPUID));

    if !vte.Valid then
        if GITS_TYPER.SEIS == '1' then IMPLEMENTATION_DEFINED "SError VSYNC_VCPU_INVALID";
        IncrementReadPointer();
        return;

    bits(32) rd_base = vte.RDbase;

    // Wait for the external effects of any virtual commands to be observable by all redistributors
    // and ensure the internal effects of any previous commands affect any subsequent interrupt
    // requests or commands
    WaitForVirtualCompletion(rd_base);
```

```
                IncrementReadPointer();
                return;
```

6.4 Common ITS pseudocode functions

The following terminology appears in some of the pseudocode functions in this section:

Interrupt Translation
: An action that causes the ITS to attempt to set a particular pending bit in a particular table.

Pending Interrupt
: A particular pending bit is set in a particular table.

The pseudocode functions in this section are based on the following assumptions:

- Each ITS function must be performed as an atomic operation. Implementations must ensure that the observed behavior is consistent with that from a strictly atomic implementation.
- Where the pseudocode issues a sequence of read and write operations to a particular Redistributor, these operations must be performed in the order in which they were generated.
- Where the pseudocode issues a write to a particular Redistributor, operation does not need to wait for the completion of the write.
- Where the pseudocode issues writes to memory to update a table, operation does not need to wait for the completion of these writes. A write to memory might never become visible to an external observer. However, the effect of any such writes must be ordered by any subsequent ITS operations, including the handling of interrupt translations. There are no ordering rules other than the standard rule that memory must appear as if writes to each location occurred in program order.
- Because each ITS function is performed as an atomic operation, any new interrupt translation that occurs after the function must be subject to the effects of that function.
- The effects of caching of the Redistributor LPI Configuration and LPI Pending tables are specified explicitly in the pseudocode.
- An interrupt translation might set a pending bit and pending bits remain set until handled by the PE. While an interrupt is pending it might be affected by an interrupt translation that is updated by a subsequent ITS function.

Note

Some variable names used in the pseudocode differ from those used in the body text. For a list of the affected variables, see *Pseudocode terminology* on page B-734.

The following pseudocode invalidates any associated caching for the LPI configuration in the Redistributor for the specified translation.

```
// InvalidateByITE
// ===============

boolean InvalidateByITE(InterruptTableEntry ite)
    if ite.Type == physical_interrupt then
        CollectionTableEntry cte = ReadCollectionTable(UInt(ite.ICID));

        if !cte.Valid then
            return FALSE;

        InvalidateInterruptCaches(ite.ICID, ite.OutputID);
    else
        VCPUTableEntry vte = ReadVCPUTable(UInt(ite.VCPUID));

        if !vte.Valid then
            return FALSE;

        InvalidateVirtualInterruptCaches(ite.VCPUID, ite.OutputID);

    return TRUE;
```

The following pseudocode describes moving a pending interrupt.

```
// MovePendingState()
// ==================
```

```
MovePendingState(bits(32) rd1, bits(32) rd2, bits(32) ID)
    if IsPending(GICR_PENDBASER[rd1], ID) then
        // The interrupt is pending in the source redistributor

        // Make sure the interrupt is released or taken by the processor for
        // example by sending a clear and waiting for the response
        EnsureInterruptNotPendingOnProcessor(rd1, ID);

        if IsPending(GICR_PENDBASER[rd1], ID) then
            // The CPU released the interrupt and it is still pending
            // Note: the following must be done without any possibility of the
            // source redistributor re-forwarding the interrupt to the processor
            ClearPendingStateLocal(GICR_PENDBASER[rd1], ID);
            SetPendingStateLocal(GICR_PENDBASER[rd2], ID);
```

The following pseudocode describes moving a pending virtual interrupt.

```
// MoveVirtualPendingState()
// =========================

MoveVirtualPendingState(bits(32) rd_base, Address vpt1, Address vpt2, bits(32) ID)
    if IsPending(vpt1, ID) then
        // The interrupt is pending in the source redistributor

        // Make sure the interrupt is released or taken by the processor for example by sending a
        // VClear and waiting for the response
        EnsureVirtualInterruptNotPendingOnProcessor(rd_base, vpt1, ID);

        if IsPending(vpt1, ID) then
            // The CPU released the interrupt and it is still pending
            // Note: the following must be done without any possibility of the source redistributor
            // re-forwarding the interrupt to the processor
            ClearVirtualPendingStateLocal(vpt1, ID);
            SetVirtualPendingStateLocal(vpt2, ID);
    return;
```

6.4.1 ITS helper functions

This subsection describes the ITS helper functions. These functions are placeholder functions for behavior that is not architected and that is IMPLEMENTATION DEFINED.

The functions are indicated by the hierarchical path names, for example `shared/gic/its/its_helper`:

shared/gic/its/its_helper/Address

```
// Address()
// ============

type Address = bits(48);
```

shared/gic/its/its_helper/ClearPendingState

```
// ClearPendingState()
// ===================

boolean ClearPendingState(InterruptTableEntry ite);
```

shared/gic/its/its_helper/ClearPendingStateLocal

```
// ClearPendingStateLocal()
// ============================

// Clears the pending state of the physical interrupt specified by INTID
// for the redistributor which owns the LPI pending table specified by PendBase

ClearPendingStateLocal(PBType PendBase, bits(32) INTID);


// ClearVirtualPendingStateLocal()
// ============================

// Clears the pending state of the virtual interrupt specified by vINTID
// in the LPI pending table  specified by base

ClearPendingStateLocal(Address base, bits(32) vINTID);
```

shared/gic/its/its_helper/CollectionOutOfRange

```
// CollectionOutOfRange()
// ======================

// Returns TRUE if the value supplied  has bits above the implemented range
// or if the value exceeds the total number of collections supported in
// hardware and external memory

boolean CollectionOutOfRange(bits(16) collection);
```

shared/gic/its/its_helper/CollectionTableEntry

```
//CollectionTableEntry()
// ======================

type CollectionTableEntry is (
    boolean Valid,
    bits(32) RDbase
)
```

shared/gic/its/its_helper/DeviceOutOfRange

```
// DeviceOutOfRange()
// ==================

// Returns TRUE if the value supplied has bits above the implemented range

boolean DeviceOutOfRange(bits(32) device);
```

shared/gic/its/its_helper/DeviceTableEntry

```
// DeviceTableEntry()
// ==================

type DeviceTableEntry is (
    boolean Valid,
    Address ITT_base,
    bits(5) ITT_size
)
```

shared/gic/its/its_helper/EndOfCommand

```
// EndOfCommand()
// ==============

// Terminate processing of the current command without incrementing the read pointer.
// This means the command will be run again.

EndOfCommand();
```

shared/gic/its/its_helper/EnsureInterruptNotPendingOnProcessor

```
// EnsureInterruptNotPendingOnProcessor()
// ======================================

// Returns when the physical interrupt specified by ID is not pending on
// the CPU interface connected to the redistributor specified by rd1

EnsureInterruptNotPendingOnProcessor(bits(32) rd1, bits(32) ID);
```

shared/gic/its/its_helper/EnsureVirtualInterruptNotPendingOnProcessor

```
// EnsureVirtualInterruptNotPendingOnProcessor()
// =============================================

// Returns when the virtual interrupt specified by ID is not pending on
// the CPU interface connected to the redistributor specified by rd1

EnsureVirtualInterruptNotPendingOnProcessor(bits(32) rd1, Address vpt, bits(32) ID);
```

shared/gic/its/its_helper/IdOutOfRange

```
// IdOutOfRange()
// ==============================

// Returns TRUE if the value supplied has bits above the implemented size or above the ITT_size

boolean IdOutOfRange(bits(32) ID, bits(5) ITT_size);
```

shared/gic/its/its_helper/IncrementReadPointer

```
// IncrementReadPointer()
// ======================

//Increments GITS_CREADR, wrapping if appropriate

IncrementReadPointer();
```

shared/gic/its/its_helper/InterruptTableEntry

```
// InterruptTableEntry()
// =====================

type InterruptTableEntry is (
    boolean Valid,
    InterruptType Type,
    bits(32) OutputID,
    bits(32) DoorbellID,
    bits(16) ICID,
    bits(16) VCPUID
)
```

shared/gic/its/its_helper/InterruptType

```
// InterruptType
// =====================

enumeration InterruptType { virtual_interrupt, physical_interrupt };
```

shared/gic/its/its_helper/InvalidateCollectionCaches

```
//InvalidateCollectionCaches()
// ============================

// Invalidates any caching of configuration for interrupts which are
// members of the collection specified by “collection”

InvalidateCollectionCaches(integer collection);
```

shared/gic/its/its_helper/InvalidateInterruptCaches

```
// InvalidateInterruptCaches()
// ===============================
```

```
// Invalidates any caching of configuration for the physical
// interrupt specified by  interruptID , which is a member of
// the collection specified by collection

InvalidateInterruptCaches(bits(16) collection, bits(32) interruptID);
```

shared/gic/its/its_helper/InvalidateInterruptConfigurationCaches

```
// InvalidateInterruptConfigurationCaches()
// ========================================

InvalidateInterruptConfigurationCaches(bits(32) ID, integer collection);
```

shared/gic/its/its_helper/InvalidateVCPUCaches

```
// InvalidateVCPUCaches()
// ==========================

// Invalidates any caching of configuration for the vPE specificed by vcpu_id

InvalidateVCPUCaches(integer vcpu_id);
```

shared/gic/its/its_helper/InvalidateVirtualConfigurationCaches

```
// InvalidateVirtualConfigurationCaches
// ================================

InvalidateVirtualConfigurationCaches(bits(32) ID, bits(16) VCPU);
```

shared/gic/its/its_helper/InvalidateVirtualInterruptCaches

```
// InvalidateVirtualInterruptCaches()
// ==================================

// Invalidates any caching of configuration for the virtual interrpt specified
// by the interruptID for the vPE specified by vcpi_id

InvalidateVirtualInterruptCaches(bits(16) vcpu_id, bits(32) interruptID);
```

shared/gic/its/its_helper/IsPending

```
// IsPending()
// ========================

// Returns TRUE if the physical interrupt specified by interrupt ID
// is pending for the Redistributor which owns the LPI pending table
// specified by PendBase

boolean IsPending(PBType PendBase, bits(32) interruptID);
```

shared/gic/its/its_helper/IsPending

```
// IsPending()
// ========================

// Returns TRUE if the virtual interrupt specified by interruptID
// is pending in the LPI pending table specified by base

boolean IsPending(Address base, bits(32) interruptID);
```

shared/gic/its/its_helper/LPIOutOfRange

```
//LPIOutOfRange()
// ========================

// Returns TRUE if the value supplied is larger than that permitted by GICD_TYPER.IDbits or not in the
// LPI range and is not 1023

boolean LPIOutOfRange(bits(32) ID);
```

shared/gic/its/its_helper/MoveAllPendingState

```
// MoveAllPendingState()
// ====================

// Moves the pending state of all interrupts from the Redistributor specified by rd1
// to the Redistributor specified by rd2

MoveAllPendingState(bits(32) rd1, bits(32) rd2);
```

shared/gic/its/its_helper/ReadCollectionTable

```
// ReadCollectionTableEntry()
========================

// Reads a collection table entry from memory

CollectionTableEntry ReadCollectionTable(integer index);
```

shared/gic/its/its_helper/ReadDeviceTable

```
// ReadDevicePointer()
// ==========================

// Reads a device table entry from memory

DeviceTableEntry ReadDeviceTable(integer index);
```

shared/gic/its/its_helper/ReadTranslationTable

```
// ReadTranslationTable()
// ======================

// Reads an ITT table entry from memory

InterruptTableEntry ReadTranslationTable(Address base, integer index);
```

shared/gic/its/its_helper/ReadVCPUTable

```
//ReadVCPUTable()
// ===========================

// Reads a VCPU table entry from memory

VCPUTableEntry ReadVCPUTable(integer index);
```

shared/gic/its/its_helper/RetargetVirtualInterrupt

```
// RetargetVirtualInterrupt()
// ======================

RetargetVirtualInterrupt(integer device, bits(32) ID, integer vcpu);
```

shared/gic/its/its_helper/SetPendingState

```
// SetPendingState()
// =================

boolean SetPendingState(InterruptTableEntry ite);
```

shared/gic/its/its_helper/SetPendingStateLocal

```
// SetPendingStateLocal()
// ======================

// Sets the pending state of the physical interrupt specified by INTID
//  for the Redistributor that owns  the LPI pending table specified by PendBase

SetPendingStateLocal(PBType PendBase, bits(32) INTID);



// SetVirtualPendingStateLocal()
// =============================

// Sets the pending state of the virtual interupt specified by INTID
// in the LPI pending table specified by base

SetPendingStateLocal(Address base, bits(32)INTID);
```

shared/gic/its/its_helper/SizeOutOfRange

```
// SizeOutOfRange()
// ================

// Returns TRUE if the value supplied exceeds the maximum allowed by GITS_TYPER.IDbits

boolean SizeOutOfRange(bits(5) ITT_size);
```

shared/gic/its/its_helper/VCPUOutOfRange

```
// VCPUOutOfRange()
// ================

// Returns TRUE if the value supplied has bits above the implemented range or
// if the value supplied exceeds the maximum configured size in the
// appropriate GITS_BASERn

boolean VCPUOutOfRange(bits(16) vcpu);
```

shared/gic/its/its_helper/VCPUTableEntry

```
//VCPUTableEntry()
// ================

type VCPUTableEntry is (
    boolean Valid,
    bits(32) RDbase,
    Address VPT_base,
    bits(5) VPT_size
)
```

shared/gic/its/its_helper/WaitForCompletion

```
// WaitForCompletion()
// ===================
```

```
// Returns when all external effects of any phsical commands are observable
//  by all Redistributors and the internal effects of any previous
// commands affect any subsequent interrupt requests or commands

WaitForCompletion(bits(32) RDbase);
```

shared/gic/its/its_helper/WaitForVirtualCompletion

```
// WaitForVirtualCompletion()
// ==========================

WaitForVirtualCompletion(bits(32) RDbase);
```

shared/gic/its/its_helper/WriteCollectionTable

```
//WriteCollectionTable()
// ==========================

// Writes a collection table entry to memory

WriteCollectionTable(integer index, CollectionTableEntry cte);
```

shared/gic/its/its_helper/WriteDeviceTable

```
// WriteDeviceTable()
// ==========================

// Writes a device table entry to memory

WriteDeviceTable(integer index, DeviceTableEntry dte);
```

shared/gic/its/its_helper/WriteTranslationTable

```
// WriteTranslationTable()
// ===================================

// Writes an ITT table entry to memory

WriteTranslationTable(Address base, integer index, InterruptTableEntry cte);
```

shared/gic/its/its_helper/WriteVCPUTable

```
// WriteVCPUTable()
// ===============================

// Writes a VCPU table entry to memory

WriteVCPUTable(integer index, VCPUTableEntry vte);
```

6.5 ITS command error encodings

When an ITS supports system errors, that is when GITS_TYPER.SEIS == 1, ITS command errors can be reported to software. It is IMPLEMENTATION DEFINED how these errors are recorded and reported.

Table 6-7 shows the ITS command error encodings.

Table 6-7 ITS command error encodings

Encoding	Error mnemonic	Command	Error description
0x01_0801	MAPD_DEVICE_OOR	MAPD	Out of range
0x01_0802	MAPD_ITTSIZE_OOR		
0x01_0903	MAPC_COLLECTION_OOR	MAPC	Out of range
0x01_0B01	MAPI_DEVICE_OOR	MAPI	
0x01_0B03	MAPI_COLLECTION_OOR		
0x01_0B04	MAPI_UNMAPPED_DEVICE		Unmapped device
0x01_0B05	MAPI_ID_OOR		Out of range
0x01_0A01	MAPTI_DEVICE_OOR	MAPTI	
0x01_0A03	MAPTI_COLLECTION_OOR		
0x01_0A04	MAPTI_UNMAPPED_DEVICE		Unmapped device
0x01_0A05	MAPTI_ID_OOR		Out of range
0x01_0A06	MAPTI_PHYSICALID_OOR		
0x01_0101	MOVI_DEVICE_OOR	MOVI	
0x01_0103	MOVI_COLLECTION_OOR		
0x01_0104	MOVI_UNMAPPED_DEVICE		Unmapped device
0x01_0105	MOVI_ID_OOR		Out of range
0x01_0107	MOVI_UNMAPPED_INTERRUPT		Unmapped interrupt
0x01_0108	MOVI_ID_IS_VIRTUAL		vINTID specified
0x01_0109	MOVI_UNMAPPED_COLLECTION		Unmapped interrupt collection
0x01_0F01	DISCARD_DEVICE_OOR	DISCARD	Out of range
0x01_0F04	DISCARD_UNMAPPED_DEVICE		Unmapped device
0x01_0F05	DISCARD_ID_OOR		Out of range
0x01_0F07	DISCARD_UNMAPPED_INTERRUPT		Unmapped interrupt
0x01_0F10	DISCARD_ITE_INVALID		Invalid translation table entry

Table 6-7 ITS command error encodings (continued)

Encoding	Error mnemonic	Command	Error description
0x01_0C01	INV_DEVICE_OOR	INV	Out of range
0x01_0C04	INV_UNMAPPED_DEVICE		Unmapped device
0x01_0C05	INV_ID_OOR		Out of range
0x01_0C07	INV_UNMAPPED_INTERRUPT		Unmapped interrupt
0x01_0C10	INV_ITE_INVALID		Invalid translation table entry
0x01_0D03	INVALL_COLLECTION_OOR	INVALL	Out of range
0x01_0D09	INVALL_UNMAPPED_COLLECTION		Unmapped interrupt collection
0x01_0301	INT_DEVICE_OOR	INT	Out of range
0x01_0304	INT_UNMAPPED_DEVICE		Unmapped device
0x01_0305	INT_ID_OOR		Out of range
0x01_0307	INT_UNMAPPED_INTERRUPT		Unmapped interrupt
0x01_0310	INT_ITE_INVALID		Invalid translation table entry
0x01_0501	CLEAR_DEVICE_OOR	CLEAR	Out of range
0x01_0504	CLEAR_UNMAPPED_DEVICE		Unmapped device
0x01_0505	CLEAR_ID_OOR		Out of range
0x01_0507	CLEAR_UNMAPPED_INTERRUPT		Unmapped interrupt
0x01_0510	CLEAR_ITE_INVALID		Invalid translation table entry
0x01_2911	VMAPP_VCPU_OOR	VMAPP	Out of range
0x01_2912	VMAPP_VPTSIZE_OOR		
0x01_2b01	VMAPI_DEVICE_OOR	VMAPI	Out of range
0x01_2b11	VMAPI_VCPU_OOR		
0x01_2b04	VMAPI_UNMAPPED_DEVICE		Unmapped device
0x01_2b05	VMAPI_ID_OOR		Out of range
0x01_2b06	VMAPI_PHYSICALID_OOR		
0x01_2a01	VMAPTI_DEVICE_OOR	VMAPTI	
0x01_2a11	VMAPTI_VCPU_OOR		
0x01_2a04	VMAPTI_UNMAPPED_DEVICE		Unmapped device
0x01_2a05	VMAPTI_ID_OOR		Out of range
0x01_2a13	VMAPTI_VIRTUALID_OOR		
0x01_2a06	VMAPTI_PHYSICALID_OOR		
0x01_2d11	VINVALL_VCPU_OOR	VINVALL	
0x01_2d14	VINVALL_VCPU_INVALID		Invalid vPE specified

Table 6-7 ITS command error encodings (continued)

Encoding	Error mnemonic	Command	Error description
0x01_2511	VSYNC_VCPU_OOR	VSYNC	Out of range
0x01_2514	VSYNC_VCPU_INVALID		Invalid vPE specified
0x01_2211	VMOVP_VCPU_OOR	VMOVP	Out of range
0x01_2214	VMOVP_VCPU_INVALID		Invalid vPE specified
0x01_2101	VMOVI_DEVICE_OOR	VMOVI	Out of range
0x01_2103	VMOVI_COLLECTION_OOR		
0x01_2104	VMOVI_UNMAPPED_DEVICE		Unmapped device
0x01_2105	VMOVI_ID_OOR		Out of range
0x01_2106	VMOVI_PHYSICALID_OOR		
0x01_2107	VMOVI_UNMAPPED_INTERRUPT		Unmapped interrupt
0x01_2115	VMOVI_ID_IS_PHYSICAL		pINTID specified
0x01_2116	VMOVI_ITEVCPU_INVALID		Invalid translation table entry
0x01_2117	VMOVI_CMDVCPU_INVALID		Invalid vPE specified

6.6 ITS power management

This subsection describes the software sequences for enabling and disabling an ITS. It contains the following sections:

- *Enabling an ITS*.
- *Disabling an ITS*.

6.6.1 Enabling an ITS

On power up, an ITS is reset to the quiescent state where GITS_CTLR.Quiescent == 1 and GITS_CTLR.Enable == 0. To enable an ITS, software must:

1. Ensure any memory structures required to support the device, interrupt translation, interrupt collection, or virtual CPU tables are initialized or restored.
2. Ensure that the ITS command queue has been provisioned.
3. Set GITS_CTLR.Enable to 1.
4. Configure the ITS as required using the appropriate ITS commands. For more information about the ITS commands, see *ITS commands* on page 6-108.
5. Wait for GITS_CTLR.Quiescent == 0.

6.6.2 Disabling an ITS

To disable an ITS, software must:

1. Ensure that all interrupts that target the ITS that is being powered down are either redirected or disabled.
2. Disable the ITS by clearing GITS_CTLR.Enable to 0. The disabled ITS completes all outstanding operations and then sets GITS_CTLR.Quiescent to 1.
3. Ensure the ITS is quiescent by polling until GITS_CTLR.Quiescent == 1.

When GITS_CTLR.Enable == 0, write accesses to GITS_TRANSLATER are ignored. When GITS_CTLR.Quiescent == 1, all operations have completed and memory backed state is committed. The ITS can then be powered down to an IMPLEMENTATION DEFINED state.

Chapter 7
Power Management

This chapter describes power management. It contains the following section:

- *Power management on page 7-152*.

7.1 Power management

In an implementation compliant with the GICv3 architecture, the CPU interface and the PE must be in the same power domain, but this does not have to be the same power domain as that within which the associated Redistributor is located. This means that it is possible to have a situation where the PE and its CPU interface are powered down, and the Redistributor, Distributor, and ITS, are powered up. In this situation, the GIC architecture supports the use of interrupts targeted at the PE to signal a powerup event to the PE and CPU interface.

Note

ARM strongly recommends that the GIC is not configured in such a way that an interrupt can cause wake-up of a particular PE, if on waking software on that PE cannot handle the interrupt.

GICv3 provides power management to control this situation, because the architecture is designed to allow the Redistributors designed by one organization to be used with PEs and CPU interfaces that have been designed by a different organization.

All other aspects of power management for the GIC are IMPLEMENTATION DEFINED.

Before powering down the CPU interface and the PE when the Redistributor is powered up, software must put the interface between the CPU interface and the Redistributor into the quiescent state or the system will become UNPREDICTABLE. The transition to the quiescent state is initiated by setting GICR_WAKER.ProcessorSleep to 1. When the interface is quiescent, GICR_WAKER.ChildrenAsleep is also set to 1.

GICR_WAKER.ProcessorSleep == 1 has the following effects:

- The Redistributor does not forward any interrupts for the PE to the CPU interface. If there is a pending interrupt for the PE that would otherwise be forwarded to the PE, a hardware signal, **WakeRequest**, is asserted to indicate that the PE is to have its power restored. In a GICv4 implementation, this applies to virtual LPIs in addition to any other interrupts.
- The Distributor does not select this PE as a candidate for selection for a 1 of N interrupt, unless GICD_CTLR.E1NWF == 1, and the PE has been selected by an IMPLEMENTATION DEFINED mechanism:
  - For a 1 of N interrupt that causes wake-up, the GIC is not required to select a new target PE if the PE that received the **WakeRequest** does not handle the interrupt on waking.

When the interface between the Redistributor and the CPU interface is in a quiescent state, the following architectural state of the CPU interface can be saved as part of saving the state within the power domain of the CPU interface and the PE:

- The CPU interface state related to physical interrupts of the connected PE.
- The CPU interface state related to virtual interrupts that is part of the vPE that is scheduled on the associated PE.

Setting GICR_WAKER.ProcessorSleep to 1 when the physical group enables in the CPU interface are set to 1 results in UNPREDICTABLE behavior.

When GICR_WAKER.ProcessorSleep == 1 or GICR_WAKER.ChildrenAsleep == 1 then a write to any GICC_*, GICV_*, GICH_*, ICC_*, or ICH_* registers, other than those in the following list, is UNPREDICTABLE:

- ICC_SRE_EL1.
- ICC_SRE_EL2.
- ICC_SRE_EL3.

Chapter 8
Programmers' Model

This chapter provides information about the GIC register interfaces and describes all of the GIC registers. It contains the following sections:

8.1 About the programmers' model

The GIC is partitioned into several logical components, as defined in Chapter 2 *GIC Partitioning*, and each component supports one or more programming interfaces. Software uses these programming interfaces to access the programmers' model and control the GIC. The interfaces are either memory-mapped or support System register accesses as follows:

- The Distributor, Redistributor, and ITS programming interfaces are always memory-mapped.
- The CPU interfaces for physical and virtual interrupt handling, and the virtual machine control interface used by the hypervisor use:
  - System register interfaces for the operation of GICv3 and GICv4.
  - Memory-mapped interfaces for legacy operation.

    Note

    Support for legacy operation is optional. Implementations are allowed to support legacy operation for virtual interrupts only, meaning that the GICV_* registers are the only memory-mapped CPU interface registers that are provided. In these implementations, GICC_* registers and GICH_* registers are not provided. GICC_* and GICH_* registers are only required to support legacy operation by physical interrupts.

When accessing a System register, the register content accessed depends on:

- The Exception level at which the PE is executing.
- Whether the access is Secure or Non-secure.
- For a Non-secure access at EL1, whether the Exception level is configured by HCR_EL2 when executing in AArch64 state, or by HCR when executing in AArch32 state, to handle virtual or physical interrupts.

8.1.1 GIC register names

All of the GIC registers have names that provide a short mnemonic for the function of the register:

- Memory-mapped registers are prefixed by one of the following:
  - GICC, to indicate a CPU interface register.
  - GICD, to indicate a Distributor register.
  - GICH, to indicate a virtual interface control register, typically accessed by a hypervisor.
  - GICR, to indicate a Redistributor register.
  - GICV, to indicate a virtual CPU interface register.
  - GITS, to indicate an ITS register.
- System registers are prefixed by:
  - ICC, to indicate a physical GIC CPU interface System register.
  - ICV, to indicate a virtual GIC CPU interface System register.
  - ICH, to indicate a virtual interface control System register.
- The remaining letters are a mnemonic for the register, for example the GIC Distributor Control Register is called GICD_CTLR.

Figure 8-1 on page 8-155 shows the interfaces that the programmer can use for the different logical components when affinity routing and System register access are enabled for all Exception levels.

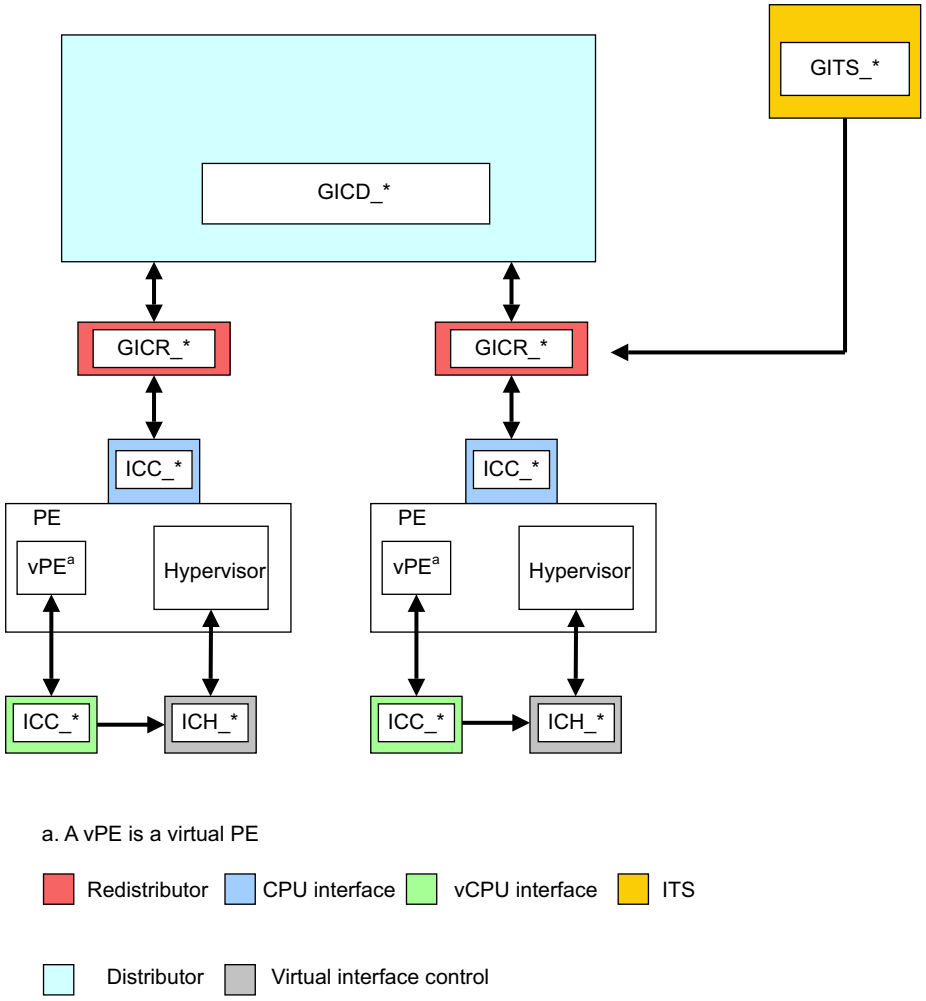

Figure 8-1 Register interfaces without legacy support (GICv3 only)

A System register might be accessible from different Exception levels. In AArch64 state, a register suffix defines the lowest Exception level at which the register is accessible. That is, any access to ICC_*_ELx must be from Exception level ELx or higher.

8.1.2 Relation between System registers and memory-mapped registers

The GIC architecture permits, but does not require, that the same registers can be shared between memory-mapped registers and the equivalent System registers. This means that if the memory-mapped registers have been accessed while ICC_SRE_ELx.SRE == 0, the System registers might be modified. Therefore, ARM recommends that software only relies on the reset values of the System registers if there has been no use of the GIC functionality while the memory-mapped registers are in use, otherwise ARM recommends that the values are treated as UNKNOWN.

Table 8-1 shows the registers that are shared between the memory-mapped registers and the System registers.

Table 8-1 Relation between System registers and memory-mapped registers

System registers[a]		Memory-mapped CPU interface registers	Memory-mapped virtual CPU interface registers
AArch64	**AArch32**		
ICC_AP0R<n>_EL1	ICC_AP0R<n>	GICC_APR<n>, GICC_NSAPR<n>	GICV_APR<n>
ICC_AP1R<n>_EL1	ICC_AP1R<n>		
ICC_BPR0_EL1	ICC_BPR0	GICC_BPR, GICC_ABPR[bc]	GICV_BPR
ICC_BPR1_EL1	ICC_BPR1		GICV_ABPR
ICC_CTLR_EL1	ICC_CTLR	GICC_CTLR	GICV_CTLR
ICC_CTLR_EL3	ICC_MCTLR		
ICC_DIR_EL1	ICC_DIR	GICC_DIR	GICV_DIR
ICC_EOIR0_EL1	ICC_EOIR0	GICC_EOIR, GICC_AEOIR	GICV_EOIR
ICC_EOIR1_EL1	ICC_EOIR1		GICV_AEOIR
ICC_HPPIR0_EL1	ICC_HPPIR0	GICC_HPPIR, GICC_AHPPIR	GICV_HPPIR
ICC_HPPIR1_EL1	ICC_HPPIR1		GICV_AHPPIR
ICC_IAR0_EL1	ICC_IAR0	GICC_IAR, GICC_AIAR[d]	GICV_IAR
ICC_IAR1_EL1	ICC_IAR0		GICV_AIAR
ICC_IGRPEN0_EL1	ICC_IGRPEN0	GICC_CTLR	GICV_CTLR
ICC_IGRPEN1_EL1	ICC_IGRPEN1		
ICC_IGRPEN1_EL3	ICC_MGRPEN1		
ICC_PMR_EL1	ICC_PMR	GICC_PMR	GICV_PMR
ICC_RPR_EL1	ICC_RPR	GICC_RPR	GICV_RPR
ICH_AP0R<n>_EL2	ICH_AP0R<n>	GICH_APR<n>	-
ICH_AP1R<n>_EL2	ICH_AP1R<n>		-
ICH_EISR_EL2	ICH_EISR	GICH_EISR	-
ICH_ELRSR_EL2	ICH_ELRSR	GICH_ELRSR	-
ICH_HCR_EL2	ICH_HCR	GICH_HCR	-
ICH_LR<n>_EL2	ICH_LR<n>	GICH_LR<n>	-
	ICH_LRC<n>		-
ICH_MISR_EL2	ICH_MISR	GICH_MISR	-
ICH_VMCR_EL2	ICH_VMCR	GICH_VMCR	-
ICH_VTR_EL2	ICH_VTR	GICH_VTR	-

a. There are also System registers prefixed with ICV, rather than ICC, and these are the virtual GIC CPU interface System registers, see *AArch64 System register descriptions of the virtual registers* on page 8-238 and *AArch32 System register descriptions of the virtual registers* on page 8-365.
b. This register is an alias of the Non-secure copy of GICC_BPR.
c. If ICC_CTLR_EL3.CBPR_EL1NS == 1, Secure accesses to this register access (and might modify) ICC_BPR0_EL1.
d. In GIC implementations that support two Security states, this register is an alias of the Non-secure view of GICC_IAR.

8.1.3 GIC memory-mapped register access

Access to the following registers must be supported:

- Single copy atomic 32-bit accesses to:
  - All GICC_*, GICV_* and GICH_* registers, where implemented.
  - All GITS_* registers.
  - All GICD_* registers.
  - All GICR_* registers.

  For the GITS_*, GICD_* and GICR_* registers, the upper 32 bits and the lower 32 bits can be accessed independently, unless the register requires a 64 bit access.
- Single copy atomic 64-bit accesses to:
  - All 64-bit GITS_* registers.
  - All 64-bit GICD_* registers.
  - All 64-bit GICR_* registers.
- Byte accesses to:
  - GICD_IPRIORITYR<n>.
  - GICD_ITARGETSR<n>.
  - GICD_SPENDSGIR<n>.
  - GICD_CPENDSGIR<n>.

ARM does not expect the following registers to be accessed directly by software, but single-copy atomic 16-bit and 32-bit accesses to these registers must be supported:

- GITS_TRANSLATER.
- GICD_SETSPI_NSR.
- GICD_CLRSPI_NSR.
- GICD_SETSPI_SR.
- GICD_CLRSPI_SR.

All other accesses to these registers result in UNPREDICTABLE behavior.

In the GIC architecture, all registers that are doubleword-accessible, halfword-accessible, or byte-accessible use a little endian memory order model.

The following accesses are not supported:

- Byte access to registers other than those specifically listed in this section.
- Unaligned word accesses. These accesses are not word single-copy atomic.
- Unaligned doubleword accesses. These accesses are not doubleword single-copy atomic.
- Word accesses for registers marked as requiring a 64-bit access.
- Doubleword accesses, other than those specifically listed in this section.
- Quadword or higher.
- Exclusive accesses.

For each of these access types, it is UNPREDICTABLE whether:

- The access generates an external abort or not.
- The defined side-effects of a read occur or not. A read returns UNKNOWN data.

- A write is ignored or sets the accessed register or registers to UNKNOWN values.

For memory-mapped accesses by a PE that complies with the ARM architecture, the single-copy atomicity rules for the instruction, the type of instruction, and the type of memory accessed, determine the size of the access made by the instruction. Example 8-1 shows this.

Example 8-1 Access sizes for memory-mapped accesses

Two Load Doubleword instructions made to consecutive doubleword-aligned locations generate a pair of single-copy atomic doubleword reads. However, if the accesses are made to Normal memory or Device-GRE memory they might appear as a single quadword access that is not supported by the peripheral.

ARMv8 does not require the size of each element accessed by a multi-register load or store instruction to be identifiable by the memory system beyond the PE. Any memory-mapped access to a GIC is defined to be beyond the PE.

Software must use a Device-nGRE or stronger memory-type, and use only single register load and store instructions, to create memory accesses that are supported by the peripheral.

Reads and writes of the memory-mapped registers complete in the order in which they arrive at the GIC. For access to different register locations, software must create this order by:

- Marking the memory as Device-nGnRnE or Device-nGnRE.
- Using the appropriate memory barriers.

Software must be able to guarantee completion of a write, for example by:

- Marking the memory as Device-nGnRnE and executing a DSB barrier, if the system supports this property.
- Reading back the value written.

For more information on endianness, memory ordering, and barrier instructions, see *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile.*

The access type definitions for the memory-mapped register interface are:

RW Read and write.

RO Read only. Writes are ignored.

WO Write only. Reads return an UNKNOWN value.

8.1.4 Access to memory-mapped registers when System register access is enabled

Because memory-mapped accesses and System register accesses might not access the same state, and are not guaranteed to be synchronized when System registers access is enabled for a particular Exception level and Security state, ARM recommends that the System registers be used instead of the memory-mapped registers that provide the same functionality.

In implementations that include the GICC_* registers, and where the Secure copy of ICC_SRE_EL1.SRE is programmable, the following state must be shared between System register access and memory-mapped access to ensure the correct operation of preemption:

- GICC_PMR and ICC_PMR_EL1 or ICC_PMR must access the same state.
- GICC_APR<n> and ICC_AP0R<n>_EL1 must access the same state.
- GICC_NSAPR<n> and ICC_AP1R<n>_EL1(NS) must access the same state.
- GICC_CTLR.CBPR and ICC_CTLR_EL3(NS).CBPR must access the same state.
- Secure accesses to GICC_BPR and ICC_BPR0_EL1 must access the same state when GICC_CTLR.CBPR == 0.
- Secure accesses to GICC_ABPR and ICC_BPR1_EL1 must access the same state when GICC_CTLR.CBPR == 0.

Note

- Software must follow the rules specified in *GIC System register access* on page 8-160 when changing the setting of the SRE fields.
- *Relation between System registers and memory-mapped registers* on page 8-155 specifies the relationship between memory-mapped registers and System registers. State can only be shared between registers that perform the same function, and the registers listed in Table 8-1 on page 8-156 might share state.

When changing from a state where the registers are required to access the same state to a state where the registers are not required to access the same state, or when changing from a state where the registers are not required to access the same state to a state where the registers are required to access the same state, the content of the registers becomes UNKNOWN.

Note

The priority bits implemented for memory-mapped and System register state must be the same, as must the minimum value of the Binary Point Register for Group 0 interrupts for both Secure and Non-secure views.

Accesses to the GICC_* registers might be affected by whether System register access is enabled or not, depending on the implementation:

- If the Secure copy of ICC_SRE_EL1.SRE == 1, then the GICC_* registers might not be accessible or might be RAZ/WI.

  Note

  When EL3 is configured to use AArch32 state, Secure EL1 is not accessible but software must still set the Secure copy of ICC_SRE.SRE to 1, to enable support for Secure Group 1 interrupts, otherwise the system is UNPREDICTABLE.

- If ICC_SRE_EL2.SRE ==1, then the GICH_* registers might not be accessible or might be RAZ/WI.
- If the Non-secure copy of ICC_SRE_EL1.SRE == 1, then the GICV_* registers might not be accessible or might be RAZ/WI.

Note

In implementations where the Non-secure copy of ICC_SRE_EL1.SRE is programmable, that is, it is not RAO/WI, the GICV_* register interface must still be provided.

An implementation might be able to detect accesses to memory-mapped registers that must not be accessed because an SRE bit is 1, and report them in an IMPLEMENTATION DEFINED manner.

8.1.5 Execution state

The ARMv8-A architecture has two Execution states:

- AArch64 state.
- AArch32 state.

To see the mapping between the AArch64 System registers and the AArch32 System registers, see:

- Table 8-3 on page 8-162.
- Table 8-4 on page 8-163.

8.1.6 Observability of the effects of accesses to the GIC registers

The PE and CPU interface logic must ensure that:

- Writes to ICC_PMR_EL1 are self-synchronizing.

Note

This ensures that no interrupts with a priority lower than the priority value in ICC_PMR_EL1 are taken after a write to ICC_PMR_EL1 is architecturally executed.

- Reads of ICC_IAR0_EL1 and ICC_IAR1_EL1 are self-synchronizing when interrupts are masked by the PE, that is when PSTATE.F == 1, for reads of ICC_IAR0_EL1, and when PSTATE.I == 1, for reads of ICC_IAR1_EL1.

  Note

  This ensures that the effect of activating an interrupt on the signaling of an interrupt exception is observed when a read of ICC_IAR0_EL1 and ICC_IAR1_EL1 is architecturally executed. This means that no spurious interrupt exception occurs if interrupts are unmasked by an instruction immediately following the read.

- Instructions that change the current Exception level from EL3 to a lower Exception level, for example the `ERET` instruction, must be synchronized with any corresponding change in the allocation of interrupts as FIQs and IRQs, so that no spurious FIQ is taken after the architectural execution of the instruction, see *Interrupt assignment to IRQ and FIQ signals* on page 4-60.
- Architectural execution of a `DSB` instruction guarantees that
  - The last value written to ICC_PMR_EL1 or GICC_PMR is observed by the associated Redistributor.
  - The last value written to ICC_SGI0R_EL1 or ICC_SGI1R_EL1 is observed by the associated Redistributor.
  - The last value written to ICC_ASGI1R_EL1 is observed by the associated Redistributor.
  - The last value written to ICC_IGRPEN0_EL1, ICC_IGRPEN1_EL1, ICC_IGRPEN1_EL3 or GICC_CTLR.{EnableGrp0, EnableGrp1}is observed by the associated Redistributor.
  - The last value written to ICH_VMCR_EL2.{VENG0, VENG1}, or GICV_CTLR.{EnableGrp0, EnableGrp1} is observed by the associated Redistributor.
  - The last SPI INTID read from ICC_IAR0_EL1, ICC_IAR1_EL1, GICC_IAR or GICC_AIAR is observed by the Distributor and by accesses from any PE to the Distributor.
  - The last SPI, PPI or LPI INTID read from ICC_IAR0_EL1, ICC_IAR1_EL1, GICC_IAR or GICC_AIAR is observed by the associated Redistributor and by accesses from the PE to the associated Redistributor.
  - The last Deactivate command for an SPI generated by a write to ICC_EOIR0_EL1, ICC_EOIR1_EL1, GICC_AEOIR, GICC_EOIR, ICC_DIR_EL1 or GICC_DIR is observed by the Distributor and by accesses from any PE to the Distributor.
  - The last Deactivate command for an SGI or PPI generated by a write to ICC_EOIR0_EL1, ICC_EOIR1_EL1, GICC_AEOIR, GICC_EOIR, ICC_DIR_EL1 or GICC_DIR is observed by the Redistributor and by accesses from any PE to the Redistributor.

Note

An `ISB` or other context synchronization operation must precede the `DSB` to ensure visibility of System register writes.

8.1.7 GIC System register access

The GIC System register interface is managed by Exception level, using the following AArch64 System registers:

- ICC_SRE_EL3, if EL3 is implemented.
- ICC_SRE_EL2, if EL2 is implemented.
- ICC_SRE_EL1.

Table 8-2 shows the permitted ICC_SRE_ELx.SRE settings.

Table 8-2 Permitted ICC_SRE_ELx.SRE settings

ICC_SRE_EL1(S)	ICC_SRE_EL1(NS)	ICC_SRE_EL2	ICC_SRE_EL3	Notes
0	0	0	0	Legacy, see Chapter 10 *Legacy Operation and Asymmetric Configurations*
0	0	0	1	Supported only when EL3 is using AArch64
0	0	1	1	Supported only when EL3 is using AArch64 and virtual interrupts are enabled
0	1	1	1	Supported only when EL3 is using AArch64
1	0	1	1	Supported only when virtual interrupts are enabled
1	1	1	1	Fully supported System register access

All combinations of ICC_SRE_ELx.SRE settings not listed in Table 8-2 result in UNPREDICTABLE behavior.

All settings other than ICC_SRE_ELx.SRE == 1 are deprecated.

Note

- When HCR_EL2 is configured so that virtualization at EL1 is enabled, it is IMPLEMENTATION DEFINED whether a Non-secure access to ICC_SRE_EL1.SRE or ICC_SRE.SRE is programmable to support a legacy VM.
- ARM expects that when ICC_SRE_EL3.SRE == 1 and ICC_SRE_EL1(S).SRE == 0, then ICC_CTLR_EL3.RM == 1.

The following changes to ICC_SRE_ELx result in UNPREDICTABLE behavior:

- Changing the value of ICC_SRE_EL3.SRE from 1 to 0.
- Changing the value of ICC_SRE_EL2.SRE from 1 to 0.
- Changing the value of ICC_SRE_EL1(S).SRE from 1 to 0.

Note

ICC_SRE_EL1(NS) can be changed from 1 to 0 to allow different VMs to have different ICC_SRE_EL1 values.

Each ICC_SRE_ELx register listed in this section provides:

- An SRE bit to enable the ICC_* System register interface at that Exception level. For EL2 and EL3, the SRE bit also enables access to all ICH_* registers.
- DIB and DFB bits to support interrupt bypass for the Exception level hierarchy. For more information about bypass, see *Interrupt bypass support* on page 2-35.

In addition:

- ICC_SRE_EL3.Enable controls EL1 access to ICC_SRE_EL1, and EL2 access to ICC_SRE_EL1 and ICC_SRE_EL2.
- ICC_SRE_EL2.Enable controls Non-secure EL1 accesses to ICC_SRE_EL1 if EL3 is not present or ICC_SRE_EL3.Enable == 1.

Note

The ICC_SRE_ELx register associated with the highest implemented Exception level is always accessible to allow software executing at that Exception level to configure the System register at different Exception levels.

The System register interface can be used for execution in both AArch32 state and AArch64 state.

For AArch32 state, accesses to GIC registers that are visible in the System register interface use the following instructions:

- The `MRC` instruction for 32-bit read accesses.
- The `MCR` instruction for32-bit write accesses.
- The `MCRR` instruction for 64-bit write accesses to ICC_SGI0R_EL1, ICC_SGI1R_EL1 and ICC_ASGI1R_EL1.

See the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for information about the form of the `MRC`, `MCR`, and `MCRR` instructions.

System registers support 32-bit or 64-bit accesses. See the individual register description for the associated access size.

The access type definitions for the System register interface are:

RW Read and write.

RO Read only. Writes result in an UNDEFINED exception.

WO Write only. Reads result in an UNDEFINED exception.

Note

For more information about UNDEFINED exceptions, see *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Table 8-3 shows the AArch64 and AArch32 register mappings for System register accesses by the GIC CPU interface.

Table 8-3 System register accesses for GIC CPU interface registers

Name of System register accessed	
AArch64	**AArch32**
ICC_IAR0_EL1[a]	ICC_IAR0
ICC_IAR1_EL1[a]	ICC_IAR1
ICC_EOIR0_EL1[a]	ICC_EOIR0
ICC_EOIR1_EL1[a]	ICC_EOIR1
ICC_HPPIR0_EL1[a]	ICC_HPPIR0
ICC_HPPIR1_EL1[a]	ICC_HPPIR1
ICC_BPR0_EL1[a]	ICC_BPR0
ICC_BPR1_EL1[a]	ICC_BPR1
ICC_DIR_EL1[a]	ICC_DIR
ICC_PMR_EL1	ICC_PMR
ICC_RPR_EL1	ICC_RPR
ICC_AP0R<n>_EL1[a]	ICC_AP0R<n>

Table 8-3 System register accesses for GIC CPU interface registers (continued)

Name of System register accessed	
AArch64	**AArch32**
ICC_AP1R<n>_EL1[a]	ICC_AP1R<n>
ICC_CTLR_EL1	ICC_CTLR
ICC_CTLR_EL3	ICC_MCTLR
ICC_IGRPEN0_EL1	ICC_IGRPEN0
ICC_IGRPEN1_EL1	ICC_IGRPEN1
ICC_IGRPEN1_EL3	ICC_MGRPEN1
ICC_SGI1R_EL1	ICC_SGI1R
ICC_ASGI1R_EL1	ICC_ASGI1R
ICC_SGI0R_EL1	ICC_SGI0R
ICC_SRE_EL1	ICC_SRE
ICC_SRE_EL2	ICC_HSRE
ICC_SRE_EL3	ICC_MSRE

a. In addition to ICC_SRE_EL*.SRE, ICC_SRE.SRE, ICC_HSRE.SRE, and ICC_MSRE.SRE, SCR_EL3 and HCR_EL2 control accessibility to these registers.

The GIC virtual interface control registers are accessible when ICC_SRE_EL2.SRE == 1.

Table 8-4 shows the AArch64 and AArch32 System register mappings for the GIC virtual interface control registers.

Table 8-4 System register mappings for GIC virtual interface control registers

Name of System register accessed	
AArch64	**AArch32**
ICH_HCR_EL2	ICH_HCR
ICH_VTR_EL2	ICH_VTR
ICH_MISR_EL2	ICH_MISR
ICH_EISR_EL2	ICH_EISR
ICH_ELRSR_EL2	ICH_ELRSR
ICH_AP0R<n>_EL2[a]	ICH_AP0R<n>[a]
ICH_AP1R<n>_EL2[a]	ICH_AP1R<n>[a]

Table 8-4 System register mappings for GIC virtual interface control registers (continued)

Name of System register accessed	
AArch64	**AArch32**
ICH_LR<n>_EL2[63:32][b]	ICH_LRC<n>[b]
ICH_LR<n>_EL2[31:0][b]	ICH_LR<n>[b]
ICH_VMCR_EL2	ICH_VMCR

a. n = 0-3
b. n = 0-15.

AArch64 System register access instruction encodings

Table 8-5 shows the format of the A64 `MSR` and `MRS` instructions to access the physical and virtual CPU interface.

Table 8-5 Mapping of MSR and MRS to physical and virtual CPU interface registers, AArch64 state

System register	Access	opc0	opc1	CRn	CRm	opc2
ICC_AP0R<n>_EL1[a]	RW	3	0	c12	c8	4-7
ICC_AP1R<n>_EL1[a]	RW	3	0	c12	c9	0-3
ICC_ASGI1R_EL1	WO	3	0	c12	c11	6
ICC_BPR0_EL1	RW	3	0	c12	c8	3
ICC_BPR1_EL1[b]	RW	3	0	c12	c12	3
ICC_CTLR_EL1[b]	RW	3	0	c12	c12	4
ICC_CTLR_EL3	RW	3	6	c12	c12	4
ICC_DIR_EL1	WO	3	0	c12	c11	1
ICC_EOIR0_EL1	WO	3	0	c12	c8	1
ICC_EOIR1_EL1	WO	3	0	c12	c12	1
ICC_HPPIR0_EL1	RO	3	0	c12	c8	2
ICC_HPPIR1_EL1	RO	3	0	c12	c12	2
ICC_IAR0_EL1	RO	3	0	c12	c8	0
ICC_IAR1_EL1	RO	3	0	c12	c12	0
ICC_IGRPEN0_EL1	RW	3	0	c12	c12	6
ICC_IGRPEN1_EL1[b]	RW	3	0	c12	c12	7
ICC_IGRPEN1_EL3	RW	3	6	c12	c12	7
ICC_PMR_EL1	RW	3	0	c4	c6	0
ICC_RPR_EL1	RO	3	0	c12	c11	3
ICC_SGI0R_EL1	WO	3	0	c12	c11	7
ICC_SGI1R_EL1	WO	3	0	c12	c11	5

Table 8-5 Mapping of MSR and MRS to physical and virtual CPU interface registers, AArch64 state

System register	Access	opc0	opc1	CRn	CRm	opc2
ICC_SRE_EL1	RW	3	0	c12	c12	5
ICC_SRE_EL2	RW	3	4	c12	c9	5
ICC_SRE_EL3	RW	3	6	c12	c12	5

a. n = 0-3.

b. There is a Secure copy and a Non-secure copy of this register.

Table 8-6 shows the format of the A64 `MSR` and `MRS` instructions that access the virtual interface control registers.

Table 8-6 Mapping of MSR and MRS to virtual interface control registers, AArch64 state

System register	Access	opc0	opc1	CRn	CRm	opc2
ICH_AP0R<n>_EL2	RW	3	4	c12	c8	0-3
ICH_AP1R<n>_EL2	RW	3	4	c12	c9	0-3
ICH_HCR_EL2	RW	3	4	c12	c11	0
ICH_VTR_EL2	RO	3	4	c12	c11	1
ICH_MISR_EL2	RO	3	4	c12	c11	2
ICH_EISR_EL2	RO	3	4	c12	c11	3
ICH_ELRSR_EL2	RO	3	4	c12	c11	5
ICH_VMCR_EL2	RW	3	4	c12	c11	7
ICH_LR<n>_EL2[a]	RW	3	4	c12	c12, c13	0-7

a. n = 0-15

For more information about the A64 instructions, see the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

AArch32 System register access instruction encodings

Table 8-7 shows the format of the A32 and T32 `MCR` and `MRC` instructions that access the physical and virtual CPU interface.

Table 8-7 Mapping of MCR and MRC to physical and virtual CPU interface registers, AArch32 state

System register	Access	opc1	CRn	CRm	opc2	Notes
ICC_AP0R<n>	RW	0	c12	c8	4-7	-
ICC_AP1R<n>[a]	RW	0	c12	c9	0-3	-
ICC_ASGI1R	WO	1	-	c12	-	Accessed using the `MCRR` and `MRRC` instructions
ICC_BPR0	RW	0	c12	c8	3	-
ICC_BPR1[a]	RW	0	c12	c12	3	-

Table 8-7 Mapping of MCR and MRC to physical and virtual CPU interface registers, AArch32 state (continued)

System register	Access	opc1	CRn	CRm	opc2	Notes
ICC_CTLR[a]	RW	0	c12	c12	4	-
ICC_DIR	WO	0	c12	c11	1	-
ICC_EOIR0	WO	0	c12	c8	1	-
ICC_EOIR1	WO	0	c12	c12	1	-
ICC_HPPIR0	RO	0	c12	c8	2	-
ICC_HPPIR1	RO	0	c12	c12	2	-
ICC_HSRE	RW	4	c12	c9	5	-
ICC_IAR0	RO	0	c12	c8	0	-
ICC_IAR1	RO	0	c12	c12	0	-
ICC_IGRPEN0	RW	0	c12	c12	6	-
ICC_IGRPEN1[a]	RW	0	c12	c12	7	-
ICC_MCTLR	RW	6	c12	c12	4	-
ICC_MGRPEN1	RW	6	c12	c12	7	-
ICC_MSRE	RW	6	c12	c12	5	-
ICC_PMR	RW	0	c4	c6	0	-
ICC_RPR	RO	0	c12	c11	3	-
ICC_SGI0R	WO	2	-	c12	-	Accessed using the MCRR and MRRC instructions
ICC_SGI1R	WO	0	-	c12	-	Accessed using the MCRR and MRRC instructions
ICC_SRE	RW	0	c12	c12	5	-

a. There is a Secure copy and a Non-secure copy of this register.

Table 8-8 shows the format of the A32 and T32 MCR and MRC instructions that access the virtual interface control registers.

Table 8-8 Mapping of MCR and MRC to virtual interface control registers, AArch32 state

System register	Access	opc1	CRn	CRm	opc2
ICH_AP0R<n>[a]	RW	4	c12	c8	0-3
ICH_AP1R<n>[a]	RW	4	c12	c9	0-3
ICH_HCR	RW	4	c12	c11	0
ICH_VTR	RO	4	c12	c11	1
ICH_MISR	RO	4	c12	c11	2
ICH_EISR	RO	4	c12	c11	3

Table 8-8 Mapping of MCR and MRC to virtual interface control registers, AArch32 state

System register	Access	opc1	CRn	CRm	opc2
ICH_ELRSR	RO	4	c12	c11	5
ICH_VMCR	RW	4	c12	c11	7
ICH_LR<n>[a]	RW	4	c12	c12, c13	0-7
ICH_LRC<n>[b]	RW	4	c12	c14, c15	0-7

a. n = 0-3.
b. n = 0-15

For more information about the T32 and A32 instructions, see the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Implementations with fixed System register enables

GICv3 implementations that are not required to be backwards compatible with GICv2 might have some System register enable bits that are RAO/WI. GICv3 supports the following options:

- ICC_SRE_EL3.SRE might be RAO/WI. This means that software executing at EL3 must always access the GIC using the System registers, but lower Exception levels might use the memory-mapped registers to access the GIC.
- ICC_SRE_EL2.SRE might be RAO/WI if ICC_SRE_EL3.SRE is also RAO/WI, or if EL3 is not implemented. This means that software executing at EL2 must always access the GIC using the System registers and software executing at Non-secure EL1 might use the memory-mapped registers to access the GIC.
- The Non-secure copy of ICC_SRE_EL1.SRE might be RAO/WI if:
  - EL2 and EL3 are not implemented.
  - Only EL2 is implemented, and ICC_SRE_EL2.SRE is RAO/WI.
  - Only EL3 is implemented, and ICC_SRE_EL3.SRE is RAO/WI.
  - Both EL2 and EL3 are implemented, and both ICC_SRE_EL2.SRE and ICC_SRE_EL3.SRE are RAO/WI.

  This means that software executing at Non-secure EL1 must always access the GIC using the System registers.
- The Secure copy of ICC_SRE_EL1.SRE might be RAO/WI if:
  - EL3 is not implemented.
  - EL3 is implemented, EL2 is not implemented, and ICC_SRE_EL3.SRE and the Non-secure copy of ICC_SRE_EL1 are RAO/WI.
  - Both EL2 and EL3 are implemented, and both ICC_SRE_EL2.SRE and ICC_SRE_EL3.SRE are RAO/WI.

  ICC_SRE_EL3.SRE and ICC_SRE_EL2.SRE are also RAO/WI. This means that software executing in Secure EL1 must access the GIC using the System registers.

8.1.8 Access to Common registers

When System register access is enabled for interrupts at Non-secure EL1, Group 0 and Group 1 interrupts are virtualized separately. This means that a VM operating at EL1 might control both physical interrupts and virtual interrupts. For example, a VM might be configured to handle:

- Virtual Group 0 interrupts by setting SCR_EL3.NS and HCR_EL2.FMO to 1.
- Physical Group 1 interrupts by setting SCR_EL3.NS to 1, and clearing SCR_EL3.IRQ and HCR_EL2.IMO to 0.

For most operations, this separate virtualization is achieved by using different registers to handle Group 0 and Group 1 interrupts. However, a number of registers are common to both Group 0 and Group 1 interrupts. These Common registers are:

- ICC_SGI0R_EL1, ICC_SGI1R_EL1, ICC_ASGI1R_EL1.
- ICC_CTLR_EL1.
- ICC_DIR_EL1.
- ICC_PMR_EL1.
- ICC_RPR_EL1.

The rules governing whether accesses to the Common registers are physical accesses, virtual accesses, or whether they generate a Trap exception, are as follows:

- When ICH_HCR_EL2.TC == 1, Non-secure accesses at EL1 generate a Trap Exception that is taken to EL2.
- When ICH_HCR_EL2.TDIR == 1, Non-secure writes at EL1 to ICC_DIR_EL1 generate a Trap exception that is taken to EL2.
- When HCR_EL2.FMO == 1 || HCR_EL2.IMO == 1 Non-secure accesses at EL1 are virtual accesses:
  - Accesses to all ICC_* registers that are accessible at EL1, other than ICC_SRE_EL1*, access the equivalent ICV_* registers instead.
  - Virtual accesses to ICC_SGI0R_EL1, ICC_SGI1R_EL1 and ICC_ASGI1R_EL1 always generate a Trap exception that is taken to EL2.

Otherwise, the lowest Exception level at which the Common registers can be accessed is the lowest Exception level that is either:

- Specified by SCR_EL3.FIQ, SCR_EL3.NS, and HCR_EL2.FMO.
- Specified by SCR_EL3.IRQ, SCR_EL3.NS, and HCR_EL2.IMO.

This means that the Common registers can be accessed at:

- EL1, without trapping, when (SCR_EL3.FIQ == 0 || SCR_EL3.IRQ == 0) && (SCR_EL3.NS == 0 || HCR_EL2.FMO == 0 || HCR_EL2. IMO == 0).
- EL2, without trapping, when (SCR_EL3.FIQ == 0 || SCR_EL3.IRQ == 0) && SCR_EL3.NS == 1.

Note

ARM expects that software configures a GIC so that:

- ICH_HCR_EL2.TC == 1 when Group 0 and Group 1 are configured asymmetrically and therefore access different states, for example one group accesses the virtualized state and the other group accesses the physical state.
- ICH_HCR_EL2.TC == 0 when the configuration is symmetric, and accesses to ICC_DIR_EL1 access the physical state or the virtualized state for both Group 0 and Group 1.

8.1.9 Traps and enables for the ICC_SRE_ELx registers

The read/write behavior of ICC_SRE_ELx.SRE is controlled as follows:

- ICC_SRE_EL1(NS) is controlled by ICC_SRE_EL2.{SRE, Enable} and ICC_SRE_EL3.{SRE, Enable}:
  - If ICC_SRE_EL2.SRE == 0 or ICC_SRE_EL3.SRE == 0, then ICC_SRE_EL1.SRE(NS) is RAZ/WI.
  - If ICC_SRE_EL2.Enable == 0, then accesses to ICC_SRE_EL1(NS) are trapped to EL2.
  - If ICC_SRE_EL3.Enable == 0, then accesses to ICC_SRE_EL1(NS) are trapped to EL3.
- ICC_SRE_EL1(S) is controlled by ICC_SRE_EL3.{SRE, Enable}:
  - If ICC_SRE_EL3.SRE == 0, then ICC_SRE_EL1(S) is RAZ/WI.
  - If ICC_SRE_EL3.Enable == 0, then accesses to ICC_SRE_EL1(S) are trapped to EL3.
- ICC_SRE_EL2 is controlled by ICC_SRE_EL3.{SRE, Enable}:
  - If ICC_SRE_EL3.SRE == 0, then ICC_SRE_EL2.SRE is RAZ/WI.
  - If ICC_SRE_EL2.SRE == 0, then ICC_SRE_EL2.Enable is treated as 1 for all purposes, other than reading/writing the register.

— If ICC_SRE_EL3.Enable == 0, then accesses to ICC_SRE_EL2 trap to EL3.

- If ICC_SRE_EL3.SRE == 0, then ICC_SRE_EL3.Enable is treated as 1 for all purposes, other than reading/writing the register.
- In an implementation that includes EL3, if ICC_SRE_EL1(S).SRE == 1 and ICC_SRE_EL3.SRE == 1, then ICC_SRE_EL2.SRE == 0 leads to UNPREDICTABLE behavior.

In the following tables:

x Indicates that the bit can be either 0 or 1.

- Indicates that this access is not applicable.

[0] RAZ/WI. Reads return 0 and writes are ignored. This bits is treated as 0.

{1} This bit is treated as 1 for all purposes, other than reading/writing the register.

(1) This bit must be set to 1, otherwise behavior is UNPREDICTABLE.

NS Indicates the value of SCR_EL3.NS.

RW Indicates that a read/write access is allowed.

T(EL2) Generates a Trap exception that is taken to EL2.

T(EL3) Generates a Trap exception that is taken to EL3. When EL3 is using AArch32, this is replaced by an Undefined exception that is taken to the current Exception level.

UND Generates an UNDEFINED exception or a trap to the current Exception level.

Table 8-9 shows the conditions under which ICC_SRE_EL3 can be accessed.

Table 8-9 ICC_SRE_EL3 access

ICC_SRE_EL3		ICC_SRE_EL2		ICC_SRE_EL1		EL1		EL2	EL3
SRE	Enable	SRE	Enable	SRE NS=0	SRE NS=1	NS = 0	NS = 1	NS = 1	
x	x	x	x	x	x	UND	UND	UND	RW

Table 8-10 shows the conditions under which ICC_SRE_EL2 can be accessed.

Table 8-10 ICC_SRE_EL2 access

ICC_SRE_EL3		ICC_SRE_EL2		ICC_SRE_EL1		EL1		EL2	EL3	
SRE	Enable	SRE	Enable	SRE NS=0	SRE NS=1	NS=0	NS=1	NS = 1	NS=0	NS=1
0	{1}	[0]	{1}	[0]	[0]	UND	UND	RW	UND	RW
1	0	x	x	0	x	UND	UND	T(EL3)	UND	RW
1	0	(1)	x	1	x	UND	UND	T(EL3)	UND	RW
1	1	x	x	0	x	UND	UND	RW	UND	RW
1	1	(1)	x	1	x	UND	UND	RW	UND	RW

Table 8-11 shows the conditions under which ICC_SRE_EL1(S) can be accessed when EL3 is implemented.

Table 8-11 ICC_SRE_EL1(S) access

ICC_SRE_EL3		ICC_SRE_EL2		ICC_SRE_EL1		EL1		EL2	EL3	
SRE	Enable	SRE	Enable	SRE NS=0	SRE NS=1	NS=0	NS=1	NS = 1	NS=0	NS=1
0	{1}	[0]	{1}	[0]	[0]	RW	N/A	N/A	RW	N/A
1	0	x	x	x	x	T(EL3)	N/A	N/A	RW	N/A
1	1	x	x	x	x	RW	N/A	N/A	RW	N/A

Table 8-12 shows the conditions under which ICC_SRE_EL1(NS) can be accessed when EL3 is implemented.

Table 8-12 ICC_SRE_EL1(NS) access

ICC_SRE_EL3		ICC_SRE_EL2		ICC_SRE_EL1		EL1		EL2	EL3	
SRE	Enable	SRE	Enable	SRE NS=0	SRE NS=1	NS=0	NS=1	NS = 1	NS=0	NS=1
0	{1}	[0]	{1}	[0]	[0]	N/A	RW	RW	N/A	RW
1	0	0	{1}	0	[0]	N/A	T(EL3)	T(EL3)	N/A	RW
1	1	0	{1}	0	[0]	N/A	RW	RW	N/A	RW
1	0	(1)	0	1	x	N/A	T(EL2)	T(EL3)	N/A	RW
1	1	(1)	0	1	x	N/A	T(EL2)	RW	N/A	RW
1	0	(1)	1	1	x	N/A	T(EL3)	T(EL3)	N/A	RW
1	1	(1)	1	1	x	N/A	RW	RW	N/A	RW
1	0	1	0	0	x	N/A	T(EL2)	T(EL3)	N/A	RW
1	1	1	0	0	x	N/A	T(EL2)	RW	N/A	RW
1	0	1	1	0	x	N/A	T(EL3)	T(EL3)	N/A	RW
1	1	1	1	0	x	N/A	RW	RW	N/A	RW

Table 8-13 shows the conditions under which the single copy of ICC_SRE_EL1 can be accessed when EL3 is not implemented.

Table 8-13 ICC_SRE_EL1 access

ICC_SRE_EL2		ICC_SRE_EL1	EL1	EL2
SRE	Enable			
0	{1}	[0]	RW	RW
0	1	[0]	RW	RW
1	0	x	T(EL2)	RW
1	1	x	RW	RW

Accesses that are not described in these tables are not possible.

8.1.10 Use of control registers for SGI forwarding

Table 8-14 shows the conditions that determine which SGI register is accessed, and whether an SGI is forwarded to a specified target CPU interface when affinity routing is enabled.

Table 8-14 Forwarding an SGI to a target PE

Access	SGI register accessed		Configuration of specified SGI on target PE	Signal SGI?
	AArch64	AArch32		
Secure EL1 EL3	ICC_SGI1R_EL1	ICC_SGI1R	Secure Group 0	Yes, provided GICD_CTLR.DS == 1
			Secure Group 1	Yes
			Non-secure Group 1	No
	ICC_ASGI1R_EL1	ICC_ASGI1R	Secure Group 0	No
			Secure Group 1	No
			Non-secure Group 1	Yes
	ICC_SGI0R_EL1	ICC_SGI0R	Secure Group 0	Yes
			Secure Group 1	No
			Non-secure Group 1	No

Table 8-14 Forwarding an SGI to a target PE (continued)

Access	SGI register accessed		Configuration of specified SGI on target PE	Signal SGI?
	AArch64	AArch32		
Non-secure EL1 EL2	ICC_SGI1R_EL1	ICC_SGI1R	Secure Group 0	Yes, provided either that: • This is permitted by the corresponding field in GICR_NSACR at each target PE. • GICD_CTLR.DS == 1.
			Secure Group 1	Yes, if permitted by the corresponding field in GICR_NSACR at each target PE
			Non-secure Group 1	Yes
	ICC_ASGI1R_EL1	ICC_ASGI1R	Secure Group 0	Yes, provided either that: • This is permitted by the corresponding field in GICR_NSACR at each target PE. • GICD_CTLR.DS == 1.
			Secure Group 1	If permitted by the corresponding field in GICR_NSACR.
			Non-secure Group 1	No
	ICC_SGI0R_EL1	ICC_SGI0R	Secure Group 0	Yes, provided either that: • This is permitted by the corresponding field in GICR_NSACR at each target PE. • GICD_CTLR.DS == 1.
			Secure Group 1	No
			Non-secure Group 1	No

Note

- When System register access is not enabled for Secure EL1, or when GICD_CTLR.DS == 1, the Distributor treats Secure Group 1 interrupts as Group 0 interrupts. When Table 8-14 on page 8-171 indicates that a Secure Group 1 interrupt is generated, the Distributor must send a Secure Group 0 interrupt to the CPU interface.
- Generating SGIs for the other Security state is only supported when affinity routing is enabled for both Security states.

8.1.11 GIC Security States

When a GIC supports two Security states, the behavior of PE accesses to the GIC registers depends on whether the access is Secure or Non-secure. Except where this document explicitly indicates otherwise, when accessing GIC registers:

- A Non-secure read of a register field holding state information for a Secure interrupt returns zero.

- The GIC ignores any Non-secure write to a register field holding state information for a Secure interrupt.

The ARM architecture defines the following register types:

Banked The device implements Secure and Non-secure copies of the register. See *Register banking* for more information.

Secure The register is accessible only from a Secure access. The address of a Secure register is RAZ/WI to any Non-secure access.

Common The register is accessible from both Secure and Non-secure accesses. The access permissions of some or all fields in the register might depend on whether the access is Secure or Non-secure.

8.1.12 Register banking

Register banking refers to providing multiple copies of a register. The GIC banks registers in the following cases:

- If a GIC supports two Security states, some registers are Banked to provide separate Secure and Non-secure copies of the registers. The Secure and Non-secure register bit assignments can differ. A Secure access to the register address accesses the Secure copy of the register, and a Non-secure access accesses the Non-secure copy.
- If the GIC is implemented as part of a multiprocessor system:
  - Some registers are Banked to provide a separate copy for each connected PE. These include the registers associated with PPIs and SGIs, and GICD_NSACR<n>, where n=0, when implemented.
  - The GIC implements the CPU interface registers independently for each CPU interface, and each connected PE accesses the registers for the interface to which it connects.

The following GIC System registers are banked by Security state:

- ICC_AP1R<n>_EL1.
- ICC_BPR1_EL1.
- ICC_CTLR_EL1.
- ICC_IGRPEN1_EL1.
- ICC_SRE_EL1.

Note

These are the only ARMv8 AArch64 System registers that are banked by Security state.

Where legacy operation supports physical interrupts, the following GICC_* memory-mapped registers are banked by Security state:

- GICC_CTLR.
- GICC_BPR.

8.1.13 Identification registers

Register offsets `0xFD0-0xFFC` are defined as read-only identification register space. For ARM implementations of the GIC architecture, the assignment of this register space, and the naming of registers in this space, is consistent with the ARM identification scheme for CoreLink and CoreSight components. ARM strongly recommends that other implementers also use this scheme to provide a consistent software discovery model.

The architecture specification defines offsets 0xFD0 - 0x0FFC in the Distributor register map as identification register space, as Table 8-15 shows

Table 8-15 The GIC identification register space, Distributor register map

Offset	Name	Type	Description
0xFFD0-0xFFE4	-	RO	IMPLEMENTATION DEFINED registers
0xFFE8	GICD_PIDR2	RO	Distributor Peripheral ID2 Register
0xFFEC-0xFFFC	-	RO	IMPLEMENTATION DEFINED registers

The architecture specification defines offsets 0xFD0 - 0x0FFC in the Redistributor register map as identification register space, as Table 8-16 shows.

Table 8-16 The GIC identification register space, Redistributor register map

Offset	Name	Type	Description
0xFFD0-0xFFE4	-	RO	IMPLEMENTATION DEFINED registers
0xFFE8	GICR_PIDR2	RO	Redistributor Peripheral ID2 Register
0xFFEC-0xFFFC	-	RO	IMPLEMENTATION DEFINED registers

The architecture specification defines offsets 0xFD0 - 0x0FFC in the ITS register map as identification register space, as Table 8-17 shows.

Table 8-17 The GIC identification register space, ITS register map

Offset	Name	Type	Description
0xFFD0-0xFEF4	-	RO	IMPLEMENTATION DEFINED registers
0xFEF8	GITS_PIDR2	RO	ITS Peripheral ID2 Register
0xFFEC-0xFFFC	-	RO	IMPLEMENTATION DEFINED registers

ARM generic ID registers can be used in the IMPLEMENTATION DEFINED register space.

GICD_PIDR2, Peripheral ID2 Register

Where an implementation implements the CoreLink and CoreSight ID scheme described in *Identification registers on page 8-173*, the GICD_PIDR2 characteristics are:

Purpose This register provides a four-bit architecturally-defined architecture revision field. The remaining bits of the register are IMPLEMENTATION DEFINED.

Usage constraints Bits[31:8] of the register are reserved, RAZ.

Configurations This register is available in all configurations of the GIC.

Attributes See the register summary in Table 8-15.

Figure 8-2 on page 8-175 shows the GICD_PIDR2 bit assignments.

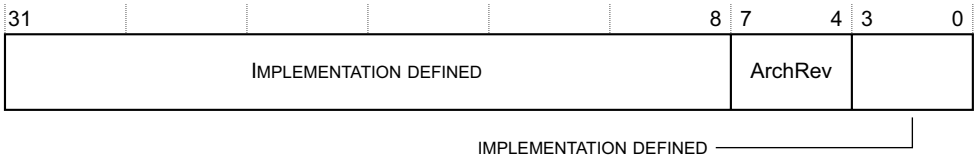

Figure 8-2 GICD_PIDR2 bit assignments

Table 8-18 shows the GICD_PIDR2 bit assignments.

Table 8-18 GICD_PIDR2 bit assignments

Bits	Name	Function
[31:8]	-	IMPLEMENTATION DEFINED. The CoreLink and CoreSight Peripheral ID Registers scheme requires these bits to be reserved, RES0, and ARM strongly recommends that implementations follow this scheme.
[7:4]	ArchRev	Revision field for the GIC architecture. The value of this field depends on the GIC architecture version that applies to the Distributor or Redistributor: • `0x1`. GICv1. • `0x2`. GICv2. • `0x3`. GICv3. • `0x4`. GICv4. • `0x5-0xF`. Reserved
[3:0]	-	IMPLEMENTATION DEFINED.

GICR_PIDR2, Redistributor Peripheral ID2 Register

Where an implementation implements the CoreLink and CoreSight ID scheme described in *Identification registers on page 8-173*, the GICR_PIDR2 characteristics are:

Purpose This register provides a four-bit architecturally-defined architecture revision field. The remaining bits of the register are IMPLEMENTATION DEFINED.

Usage constraints Bits[31:8] of the register are reserved, RAZ.

Configurations This register is available in all configurations of the GIC.

Attributes See the register summary in Table 8-16 on page 8-174.

The GICR_PIDR2 bit assignments are the same as those for GICD_PIDR2.

GITS_PIDR2, Redistributor Peripheral ID2 Register

Where an implementation implements the CoreLink and CoreSight ID scheme described in *Identification registers on page 8-173*, the GITS_PIDR2 characteristics are:

Purpose This register provides a four-bit architecturally-defined architecture revision field. The remaining bits of the register are IMPLEMENTATION DEFINED.

Usage constraints Bits[31:8] of the register are reserved, RAZ.

Configurations This register is available in all configurations of the GIC.

Attributes See the register summary in Table 8-17 on page 8-174.

The GITS_PIDR2 bit assignments are the same as those for GICD_PIDR2.

The ARM implementation of the GIC identification registers

Note

- The ARM implementation of these registers is consistent with the identification scheme for CoreLink and CoreSight components. This implementation identifies the device as a GIC that implements this architecture. It does not identify the designer or manufacturer of the GIC implementation. For information about the designer and manufacturer of a GIC implementation, see the descriptions for GICD_IIDR and GICC_IIDR.
- In other contexts, this identification scheme identifies a component in a system. The GIC use of the scheme is different. It identifies only that the device is an implementation of a version of the GIC architecture defined by this specification. Software must read GICD_IIDR and GICC_IIDR to discover, for example, the implementer and version of the GIC hardware.

All component classes require the implementation of the Component and Peripheral Identification registers, as described in:

- *Component Identification Registers, CIDR0-CIDR3*.
- *Peripheral Identification Registers, PIDR0 - PIDR7*.

Component Identification Registers, CIDR0-CIDR3

Table 8-19 shows the Component Identification Registers.

Table 8-19 Component Identification Registers

Name	Offset	Bits	Field	Value	Description
CIDR3	0xFFFC	[7:0]	PRMBL_3	0xB1	Preamble
CIDR2	0xFFF8	[7:0]	PRMBL_2	0x05	Preamble
CIDR1	0xFFF4	[7:4]	CLASS	0xF	Component Class
		[3:0]	PRMBL_1	0x0	Preamble
CIDR0	0xFFF0	[7:0]	PRMBL_0	0x0D	Preamble

Peripheral Identification Registers, PIDR0 - PIDR7

Table 8-20 shows the Peripheral Identification Registers.

Table 8-20 Peripheral Identification Registers

Name	Offset	Bits	Field	Value	Description
PIDR7	0xFFDC	[7:0]	-	RES0	Reserved
PIDR6	0xFFD8	[7:0]	-	RES0	Reserved
PIDR5	0XFFD4	[7:0]	-	RES0	Reserved
PIDR4	0xFFD0	[7:4]	SIZE	0x4	64 KB software visible page
		[3:0]	DES_2	0x4	ARM implementation

Table 8-20 Peripheral Identification Registers (continued)

Name	Offset	Bits	Field	Value	Description
PIDR2	0xFFE8	[7:4]	ARCHREV	IMP DEF	• 0x1. GICv1. • 0x2. GICv2. • 0x3. GICv3. • 0x4. GICv4. • 0x5 - 0xF. Reserved.
		[3]	JEDEC	0x1	JEP code
		[2:0]	DES_1	0x3	JEP106 identification code, bits[6:4]
PIDR1	0xFFE4	[7:4]	DES_0	0xB	JEP106 identification code, bits[3:0]
		[3:0]	PART_1	0x4	Part number, bits[11:8]
PIDR0	0xFFE0	[7:0]	PART_0	0x92	Part number, bits[7:0]

A component is uniquely identified by the following fields:

- JEP106 continuation code.
- JEP106 identification code.
- Part Number.
- ArchRev.
- Customer Modified.
- RevAnd.

The meaning of the fields is as follows:

JEP106 continuation code, JEP106 identification code (DES_2, DES_1, DES_0)

These indicate the designer of the component and not the implementer, except where the two are the same. To obtain a number, or to see the assignment of these codes, contact JEDEC at http://www.jedec.org.

A JEDEC code takes the following form:

- A sequence of zero or more bytes, all of the value 0x7F.
- A following 8-bit number, that is not 0x7F, and where bit[7] is an odd parity bit.

For example, ARM Limited is assigned the code 0x7F 0x7F 0x7F 0x7F 0x3B.

The encoding used in the Peripheral Identification Registers is as follows:

- The continuation code is the number of times 0x7F appears before the final number. for example, for ARM Limited this is 0x4.
- The identification code is bits[6:0] of the final number. For example, for ARM Limited this is 0x3B.

Part number (PART_1, PART_0)

This is selected by the designer of the component.

ArchRev In GICv3, this field is ArchRev, see *GICD_PIDR2, Peripheral ID2 Register* on page 8-174.

Customer Modified (CMOD) Where the component is reusable IP, this value indicates if the customer has modified the behavior of the component. In most cases this field is zero.

RevAnd (REVAND) The RevAnd field is an incremental value starting at 0x0 for the first design of a component. This only increases by 1 for both major and minor revisions, and is simply used as a look-up to establish the exact major and minor revision.

4KB Count (SIZE) This is a 4-bit value that indicates the total contiguous size of the memory block used by this component in powers of 2 from the standard 4KB. If a component only requires a single 4KB then this must read as log to the base of 2 of the number of 4KB blocks.

8.2 AArch64 System register descriptions

This section describes each of the physical AArch64 GIC System registers in register name order. The ICC prefix indicates a GIC CPU interface System register. Each AArch64 System register description contains a reference to the AArch32 register that provides the same functionality.

Unless otherwise stated, the bit assignments for the GIC System registers are the same as those for the equivalent GICC_* and GICV_* memory-mapped registers.

The ICC prefix is used by the System register access mechanism to select the physical or virtual interface System registers according to the setting of HCR_EL2. The equivalent memory-mapped physical registers are described in *The GIC CPU interface register descriptions* on page 8-548.The equivalent virtual interface memory-mapped registers are described in *The GIC virtual CPU interface register descriptions* on page 8-587.

Table 8-21 shows the encodings for the AArch64 System registers.

Table 8-21 Encodings for the AArch64 System registers

Register	Width (bits)	Access instruction encoding					Notes
		Op0	Op1	CRn	CRm	Op2	
ICC_PMR_EL1	32	3	0	4	6	0	RW
ICC_IAR0_EL1	32			12	8	0	RO
ICC_EOIR0_EL1	32					1	WO
ICC_HPPIR0_EL1	32					2	RO
ICC_BPR0_EL1	32					3	RW
ICC_AP0R<n>_EL1	32					4-7	RW, <n> = `Op2-4`
ICC_AP1R<n>_EL1	32				9	0-3	RW, <n> = `Op2`
ICC_DIR_EL1	32				11	1	WO
ICC_RPR_EL1	32					3	RO
ICC_SGI1R_EL1	64					5	WO
ICC_ASGI1R_EL1	64					6	WO
ICC_SGI0R_EL1	64					7	WO
ICC_IAR1_EL1	32				12	0	RO
ICC_EOIR1_EL1	32					1	WO
ICC_HPPIR1_EL1	32					2	RO
ICC_BPR1_EL1	32					3	RW
ICC_CTLR_EL1	32					4	RW
ICC_SRE_EL1	32					5	RW
ICC_IGRPEN0_EL1	32					6	RW
ICC_IGRPEN1_EL1	32					7	RW
ICC_SRE_EL2	32	3	4	12	9	5	RW

Table 8-21 Encodings for the AArch64 System registers (continued)

Register	Width (bits)	Access instruction encoding					Notes
		Op0	Op1	CRn	CRm	Op2	
ICC_CTLR_EL3	32	3	6	12	12	4	RW
ICC_SRE_EL3	32					5	RW
ICC_IGRPEN1_EL1	32					7	RW

The following access encodings are IMPLEMENTATION DEFINED.

op0	op1	CRn	CRm	op2
11	000	1100	1101	000

8.2.1 ICC_AP0R<n>_EL1, Interrupt Controller Active Priorities Group 0 Registers, n = 0 - 3

The ICC_AP0R<n>_EL1 characteristics are:

Purpose

Provides information about Group 0 active priorities.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	RW	RW	RW	RW

The ICC_AP0R<n>_EL1 registers are only accessible at Non-secure EL1 when HCR_EL2.FMO is set to 0.

Note

When HCR_EL2.FMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_AP0R<n>_EL1 results in an access to ICV_AP0R<n>_EL1.

Writing to these registers with any value other than the last read value of the register (or `0x00000000` when there are no Group 0 active priorities) might result in UNPREDICTABLE behavior of the interrupt prioritization system, causing:

- Interrupts that should preempt execution to not preempt execution.
- Interrupts that should not preempt execution to preempt execution.

ICC_AP0R1_EL1 is only implemented in implementations that support 6 or more bits of priority. ICC_AP0R2_EL1 and ICC_AP0R3_EL1 are only implemented in implementations that support 7 bits of priority. Unimplemented registers are UNDEFINED.

Writing to the active priority registers in any order other than the following order will result in UNPREDICTABLE behavior:

- ICC_AP0R<n>_EL1.
- Secure ICC_AP1R<n>_EL1.
- Non-secure ICC_AP1R<n>_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.FIQ==1, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR_EL2.FMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_AP0R<n>_EL1 is architecturally mapped to AArch32 System register ICC_AP0R<n>.

Attributes

ICC_AP0R<n>_EL1 is a 32-bit register.

Field descriptions

The ICC_AP0R<n>_EL1 bit assignments are:

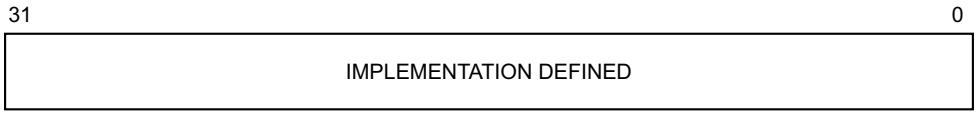

IMPLEMENTATION DEFINED, bits [31:0]

IMPLEMENTATION DEFINED.

When this register has an architecturally-defined reset value, this field resets to `0`.

The contents of these registers are IMPLEMENTATION DEFINED with the one architectural requirement that the value `0x00000000` is consistent with no interrupts being active.

Accessing the ICC_AP0R<n>_EL1:

To access the ICC_AP0R<n>_EL1:

```
MRS <Xt>, ICC_AP0R<n>_EL1 ; Read ICC_AP0R<n>_EL1 into Xt, where n is in the range 0 to 3
MSR ICC_AP0R<n>_EL1, <Xt> ; Write Xt to ICC_AP0R<n>_EL1, where n is in the range 0 to 3
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1000	1:n<1:0>

When HCR_EL2.FMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_AP0R<n>_EL1.

8.2.2 ICC_AP1R<n>_EL1, Interrupt Controller Active Priorities Group 1 Registers, n = 0 - 3

The ICC_AP1R<n>_EL1 characteristics are:

Purpose

Provides information about Group 1 active priorities.

Usage constraints

ICC_AP1R<n>_EL1(S) is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	-	-	RW

ICC_AP1R<n>_EL1(NS) is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	-	RW	RW	-

The ICC_AP1R<n>_EL1 registers are only accessible at Non-secure EL1 when HCR_EL2.IMO is set to 0.

Note

When HCR_EL2.IMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_AP1R<n>_EL1 results in an access to ICV_AP1R<n>_EL1.

Writing to these registers with any value other than the last read value of the register (or `0x00000000` when there are no Group 1 active priorities) might result in UNPREDICTABLE behavior of the interrupt prioritization system, causing:

- Interrupts that should preempt execution to not preempt execution.
- Interrupts that should not preempt execution to preempt execution.

ICC_AP1R1_EL1 is only implemented in implementations that support 6 or more bits of priority. ICC_AP1R2_EL1 and ICC_AP1R3_EL1 are only implemented in implementations that support 7 bits of priority. Unimplemented registers are UNDEFINED.

Writing to the active priority registers in any order other than the following order will result in UNPREDICTABLE behavior:

- ICC_AP0R<n>_EL1.
- Secure ICC_AP1R<n>_EL1.
- Non-secure ICC_AP1R<n>_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.IRQ==1, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR_EL2.IMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_AP1R<n>_EL1(S) is architecturally mapped to AArch32 System register ICC_AP1R<n> (S).

AArch64 System register ICC_AP1R<n>_EL1(NS) is architecturally mapped to AArch32 System register ICC_AP1R<n> (NS).

Attributes

ICC_AP1R<n>_EL1 is a 32-bit register.

Field descriptions

The ICC_AP1R<n>_EL1 bit assignments are:

IMPLEMENTATION DEFINED, bits [31:0]

IMPLEMENTATION DEFINED.

When this register has an architecturally-defined reset value, this field resets to `0`.

The contents of these registers are IMPLEMENTATION DEFINED with the one architectural requirement that the value `0x00000000` is consistent with no interrupts being active.

Accessing the ICC_AP1R<n>_EL1:

To access the ICC_AP1R<n>_EL1:

```
MRS <Xt>, ICC_AP1R<n>_EL1 ; Read ICC_AP1R<n>_EL1 into Xt, where n is in the range 0 to 3
MSR ICC_AP1R<n>_EL1, <Xt> ; Write Xt to ICC_AP1R<n>_EL1, where n is in the range 0 to 3
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1001	0:n<1:0>

When HCR_EL2.IMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_AP1R<n>_EL1.

8.2.3 ICC_ASGI1R_EL1, Interrupt Controller Alias Software Generated Interrupt Group 1 Register

The ICC_ASGI1R_EL1 characteristics are:

Purpose

Generates Group 1 SGIs for the Security state that is not the current Security state.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	WO	WO	WO	WO	WO

This register allows software executing in a Secure state to generate Non-secure Group 1 SGIs. It will also allow software executing in a Non-secure state to generate Secure Group 1 SGIs, if permitted by the settings of GICR_NSACR in the Redistributor corresponding to the target PE.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If HCR_EL2.FMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HCR_EL2.IMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICC_SRE_EL1.SRE==0, write accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, write accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, write accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, Secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, write accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR_EL2.IMO==0, and HCR_EL2.FMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_ASGI1R_EL1 performs the same function as AArch32 System operation ICC_ASGI1R.

Under certain conditions a write to ICC_ASGI1R_EL1 can generate Group 0 interrupts, see Table 8-14 on page 8-171.

Attributes

ICC_ASGI1R_EL1 is a 64-bit register.

Field descriptions

The ICC_ASGI1R_EL1 bit assignments are:

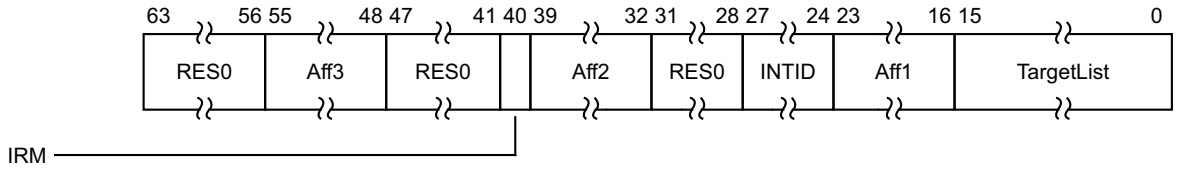

Bits [63:56]

Reserved, RES0.

Aff3, bits [55:48]

The affinity 3 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

Bits [47:41]

Reserved, RES0.

IRM, bit [40]

Interrupt Routing Mode. Determines how the generated interrupts should be distributed to PEs. Possible values are:

`0` Interrupts routed to the PEs specified by Aff3.Aff2.Aff1.<target list>.

`1` Interrupts routed to all PEs in the system, excluding "self".

Aff2, bits [39:32]

The affinity 2 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

Bits [31:28]

Reserved, RES0.

INTID, bits [27:24]

The INTID of the SGI.

Aff1, bits [23:16]

The affinity 1 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

TargetList, bits [15:0]

Target List. The set of PEs for which SGI interrupts will be generated. Each bit corresponds to the PE within a cluster with an Affinity 0 value equal to the bit number.

If a bit is 1 and the bit does not correspond to a valid target PE, the bit must be ignored by the Distributor. It is IMPLEMENTATION DEFINED whether, in such cases, a Distributor can signal a system error.

Note

This restricts a system to sending targeted SGIs to PE with an affinity 0 number of less than 16.

If SRE is set only for secure EL3, software executing at EL3 might use the System register interface to generate SGIs. Hence, the Distributor must always be able to receive and acknowledge Generate SGI packets received from CPU interface regardless of the ARE settings for a Security state. However, the Distributor might discard such packets.

If the IRM bit is 1, this field is RES0.

Accessing the ICC_ASGI1R_EL1:

To access the ICC_ASGI1R_EL1:

```
MSR ICC_ASGI1R_EL1, <Xt> ; Write Xt to ICC_ASGI1R_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1011	110

8.2.4 ICC_BPR0_EL1, Interrupt Controller Binary Point Register 0

The ICC_BPR0_EL1 characteristics are:

Purpose

Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines Group 0 interrupt preemption.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	RW	RW	RW	RW

ICC_BPR0_EL1 is only accessible at Non-secure EL1 when HCR_EL2.FMO is set to 0.

Note

When HCR_EL2.FMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_BPR0_EL1 results in an access to ICV_BPR0_EL1.

The minimum binary point value is derived from the number of implemented priority bits. The number of priority bits is IMPLEMENTATION DEFINED, and reported by ICC_CTLR_EL1.PRIbits and ICC_CTLR_EL3.PRIbits.

An attempt to program the binary point field to a value less than the minimum value sets the field to the minimum value. On a reset, the binary point field is UNKNOWN.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.FIQ==1, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR_EL2.FMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_BPR0_EL1 is architecturally mapped to AArch32 System register ICC_BPR0.

Virtual accesses to this register update ICH_VMCR_EL2.VBPR0.

Attributes

ICC_BPR0_EL1 is a 32-bit register.

Field descriptions

The ICC_BPR0_EL1 bit assignments are:

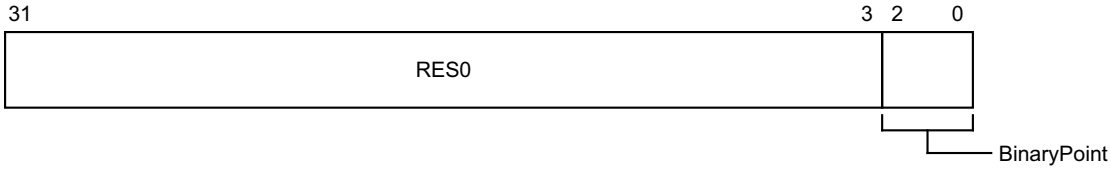

Bits [31:3]

Reserved, RES0.

BinaryPoint, bits [2:0]

The value of this field controls how the 8-bit interrupt priority field is split into a group priority field, that determines interrupt preemption, and a subpriority field. This is done as follows:

Binary point value	Group priority field	Subpriority field	Field with binary point
0	[7:1]	[0]	ggggggg.s
1	[7:2]	[1:0]	gggggg.ss
2	[7:3]	[2:0]	ggggg.sss
3	[7:4]	[3:0]	gggg.ssss
4	[7:5]	[4:0]	ggg.sssss
5	[7:6]	[5:0]	gg.ssssss
6	[7]	[6:0]	g.sssssss
7	No preemption	[7:0]	.sssssssss

Accessing the ICC_BPR0_EL1:

To access the ICC_BPR0_EL1:

```
MRS <Xt>, ICC_BPR0_EL1 ; Read ICC_BPR0_EL1 into Xt
MSR ICC_BPR0_EL1, <Xt> ; Write Xt to ICC_BPR0_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1000	011

When HCR_EL2.FMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_BPR0_EL1.

8.2.5 ICC_BPR1_EL1, Interrupt Controller Binary Point Register 1

The ICC_BPR1_EL1 characteristics are:

Purpose

Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines Group 1 interrupt preemption.

Usage constraints

ICC_BPR1_EL1(S) is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	-	-	RW

ICC_BPR1_EL1(NS) is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	-	RW	RW	-

ICC_BPR1_EL1 is only accessible at Non-secure EL1 when HCR_EL2.IMO is set to 0.

Note

When HCR_EL2.IMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_BPR1_EL1 results in an access to ICV_BPR1_EL1.

On a reset, the binary point field is UNKNOWN.

An attempt to program the binary point field to a value less than the minimum value sets the field to the minimum value.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*M.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, accesses to this register from EL1 generate an Undefined exception that is taken to EL1.

If ICC_SRE_EL2.SRE==0, accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.IRQ==1, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR_EL2.IMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_BPR1_EL1(S) is architecturally mapped to AArch32 System register ICC_BPR1 (S).

AArch64 System register ICC_BPR1_EL1(NS) is architecturally mapped to AArch32 System register ICC_BPR1 (NS).

Virtual accesses to this register update IICH_VMCR_EL2.VBPR1.

Attributes

ICC_BPR1_EL1 is a 32-bit register.

Field descriptions

The ICC_BPR1_EL1 bit assignments are:

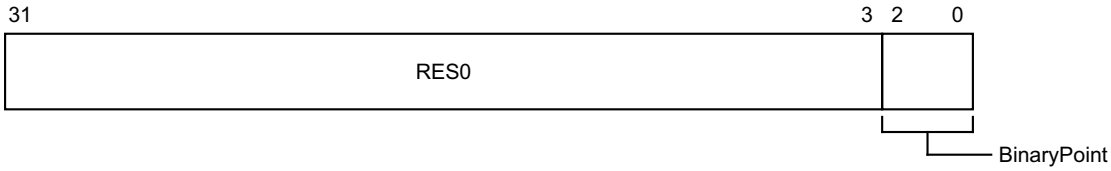

Bits [31:3]

Reserved, RES0.

BinaryPoint, bits [2:0]

If the GIC is configured to use separate binary point fields for Group 0 and Group 1 interrupts, the value of this field controls how the 8-bit interrupt priority field is split into a group priority field, that determines interrupt preemption, and a subpriority field. This is done for the Secure ICC_BPR1_EL1.BinaryPoint as follows:

Binary point value	Group priority field	Subpriority field	Field with binary point
0	[7:1]	[0]	ggggggg.s
1	[7:2]	[1:0]	gggggg.ss
2	[7:3]	[2:0]	ggggg.sss
3	[7:4]	[3:0]	gggg.ssss
4	[7:5]	[4:0]	ggg.sssss
5	[7:6]	[5:0]	gg.ssssss
6	[7]	[6:0]	g.sssssss
7	No preemption	[7:0]	.ssssssss

For the Non-secure ICC_BPR1_EL1.BinaryPoint, the split of the interrupt priority field is as follows:

Binary point value	Group priority field	Subpriority field	Field with binary point
0	-	-	-
1	[7:1]	[0]	ggggggg.s
2	[7:2]	[1:0]	gggggg.ss
3	[7:3]	[2:0]	ggggg.sss
4	[7:4]	[3:0]	gggg.ssss

Binary point value	Group priority field	Subpriority field	Field with binary point
5	[7:5]	[4:0]	ggg.sssss
6	[7:6]	[5:0]	gg.ssssss
7	[7]	[6:0]	g.sssssss

Writing 0 to this field will set this field to its reset value, which is IMPLEMENTATION DEFINED and non-zero.

If EL3 is implemented and ICC_CTLR_EL3.CBPR_EL1S is 1:

- Writing to this register at Secure EL1 modifies ICC_BPR0_EL1.
- Reading this register at Secure EL1 returns the value of ICC_BPR0_EL1.

If EL3 is implemented and ICC_CTLR_EL3.CBPR_EL1NS is 1, Non-secure accesses to this register at EL1 or EL2 behave as follows, depending on the values of HCR_EL2.IMO and SCR_EL3.IRQ:

HCR_EL2.IMO	SCR_EL3.IRQ	Behavior
0	0	Non-secure EL1 and EL2 reads return ICC_BPR0_EL1 + 1 saturated to 0b111. Non-secure EL1 and EL2 writes are ignored.
0	1	Non-secure EL1 and EL2 accesses trap to EL3.
1	0	Non-secure EL1 accesses affect virtual interrupts. Non-secure EL2 reads return ICC_BPR0_EL1 + 1 saturated to 0b111. Non-secure EL2 writes are ignored.
1	1	Non-secure EL1 accesses affect virtual interrupts. Non-secure EL2 accesses trap to EL3.

If EL3 is not implemented and ICC_CTLR_EL1.CBPR is 1, Non-secure accesses to this register at EL1 or EL2 behave as follows, depending on the values of HCR_EL2.IMO:

HCR_EL2.IMO	Behavior
0	Non-secure EL1 and EL2 reads return ICC_BPR0_EL1 + 1 saturated to 0b111. Non-secure EL1 and EL2 writes are ignored.
1	Non-secure EL1 accesses affect virtual interrupts. Non-secure EL2 reads return ICC_BPR0_EL1 + 1 saturated to 0b111. Non-secure EL2 writes are ignored.

Accessing the ICC_BPR1_EL1:

To access the ICC_BPR1_EL1:

```
MRS <Xt>, ICC_BPR1_EL1 ; Read ICC_BPR1_EL1 into Xt
MSR ICC_BPR1_EL1, <Xt> ; Write Xt to ICC_BPR1_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	011

When HCR_EL2.IMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_BPR1_EL1.

8.2.6 ICC_CTLR_EL1, Interrupt Controller Control Register (EL1)

The ICC_CTLR_EL1 characteristics are:

Purpose

Controls aspects of the behavior of the GIC CPU interface and provides information about the features implemented.

Usage constraints

ICC_CTLR_EL1(S) is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	-	-	RW

ICC_CTLR_EL1(NS) is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	-	RW	RW	-

ICC_CTLR_EL1 is only accessible at Non-secure EL1 when HCR_EL2.{FMO, IMO} == {0, 0}.

Note

When HCR_EL2.{FMO, IMO} != {0, 0}, at Non-secure EL1, the instruction encoding used to access ICC_CTLR_EL1 results in an access to ICV_CTLR_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TC==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR_EL2.IMO==0, and HCR_EL2.FMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_CTLR_EL1(S) is architecturally mapped to AArch32 System register ICC_CTLR (S).

AArch64 System register ICC_CTLR_EL1(NS) is architecturally mapped to AArch32 System register ICC_CTLR (NS).

Attributes

ICC_CTLR_EL1 is a 32-bit register.

Field descriptions

The ICC_CTLR_EL1 bit assignments are:

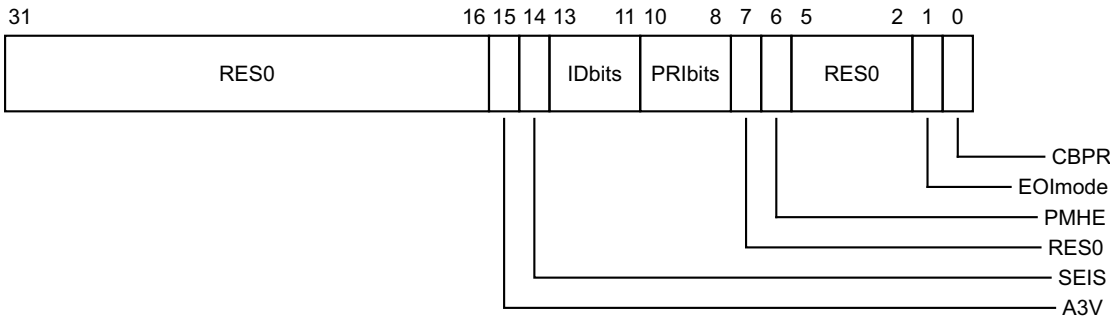

Bits [31:16]

Reserved, RES0.

A3V, bit [15]

Affinity 3 Valid. Read-only and writes are ignored. Possible values are:

`0` The CPU interface logic only supports zero values of Affinity 3 in SGI generation System registers.

`1` The CPU interface logic supports non-zero values of Affinity 3 in SGI generation System registers.

If EL3 is implemented, this bit is an alias of ICC_CTLR_EL3.A3V.

SEIS, bit [14]

SEI Support. Read-only and writes are ignored. Indicates whether the CPU interface supports local generation of SEIs:

`0` The CPU interface logic does not support local generation of SEIs.

`1` The CPU interface logic supports local generation of SEIs.

If EL3 is implemented, this bit is an alias of ICC_CTLR_EL3.SEIS.

IDbits, bits [13:11]

Identifier bits. Read-only and writes are ignored. The number of physical interrupt identifier bits supported:

`000` 16 bits.

`001` 24 bits.

All other values are reserved.

If EL3 is implemented, this field is an alias of ICC_CTLR_EL3.IDbits.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

PRIbits, bits [10:8]

Priority bits. Read-only and writes are ignored. The number of priority bits implemented, minus one.

An implementation that supports two Security states must implement at least 32 levels of physical priority (5 priority bits).

An implementation that supports only a single Security state must implement at least 16 levels of physical priority (4 priority bits).

Note

This field always returns the number of priority bits implemented, regardless of the Security state of the access or the value of GICD_CTLR.DS.

For physical accesses, this field determines the minimum value of ICC_BPR0_EL1.

If EL3 is implemented, physical accesses return the value from ICC_CTLR_EL3.PRIbits.

If EL3 is not implemented, physical accesses return the value from this field.

Bit [7]

Reserved, RES0.

PMHE, bit [6]

Priority Mask Hint Enable. Controls whether the priority mask register is used as a hint for interrupt distribution:

0 Disables use of ICC_PMR_EL1 as a hint for interrupt distribution.

1 Enables use of ICC_PMR_EL1 as a hint for interrupt distribution.

If EL3 is implemented, this bit is an alias of ICC_CTLR_EL3.PMHE. Whether this bit can be written as part of an access to this register depends on the value of GICD_CTLR.DS:

- If GICD_CTLR.DS == 0, this bit is read-only.
- If GICD_CTLR.DS == 1, this bit is read/write.

If EL3 is not implemented, it is IMPLEMENTATION DEFINED whether this bit is read-only or read-write:

- If this bit is read-only, an implementation can choose to make this field RAZ/WI or RAO/WI.
- If this bit is read/write, it resets to zero.

Bits [5:2]

Reserved, RES0.

EOImode, bit [1]

EOI mode for the current Security state. Controls whether a write to an End of Interrupt register also deactivates the interrupt:

0 ICC_EOIR0_EL1 and ICC_EOIR1_EL1 provide both priority drop and interrupt deactivation functionality. Accesses to ICC_DIR_EL1 are UNPREDICTABLE.

1 ICC_EOIR0_EL1 and ICC_EOIR1_EL1 provide priority drop functionality only. ICC_DIR_EL1 provides interrupt deactivation functionality.

The Secure ICC_CTLR_EL1.EOIMode is an alias of ICC_CTLR_EL3.EOImode_EL1S.

The Non-secure ICC_CTLR_EL1.EOIMode is an alias of ICC_CTLR_EL3.EOImode_EL1NS

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

CBPR, bit [0]

Common Binary Point Register. Controls whether the same register is used for interrupt preemption of both Group 0 and Group 1 interrupts:

0 ICC_BPR0_EL1 determines the preemption group for Group 0 interrupts only.
ICC_BPR1_EL1 determines the preemption group for Group 1 interrupts.

1 ICC_BPR0_EL1 determines the preemption group for both Group 0 and Group 1 interrupts.

If EL3 is implemented:

- This bit is an alias of ICC_CTLR_EL3.CBPR_EL1{S,NS} where S or NS corresponds to the current Security state.
- If GICD_CTLR.DS == 0, this bit is read-only.

- If GICD_CTLR.DS == 1, this bit is read/write.

If EL3 is not implemented, this bit is read/write.

Accessing the ICC_CTLR_EL1:

To access the ICC_CTLR_EL1:

```
MRS <Xt>, ICC_CTLR_EL1 ; Read ICC_CTLR_EL1 into Xt
MSR ICC_CTLR_EL1, <Xt> ; Write Xt to ICC_CTLR_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	100

When HCR_EL2.{FMO, IMO} != {0, 0}, execution of this encoding at Non-secure EL1 results in an access to ICV_CTLR_EL1.

8.2.7 ICC_CTLR_EL3, Interrupt Controller Control Register (EL3)

The ICC_CTLR_EL3 characteristics are:

Purpose

Controls aspects of the behavior of the GIC CPU interface and provides information about the features implemented.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	-	RW	RW

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

Configurations

AArch64 System register ICC_CTLR_EL3 can be mapped to AArch32 System register ICC_MCTLR, but this is not architecturally mandated.

Attributes

ICC_CTLR_EL3 is a 32-bit register.

Field descriptions

The ICC_CTLR_EL3 bit assignments are:

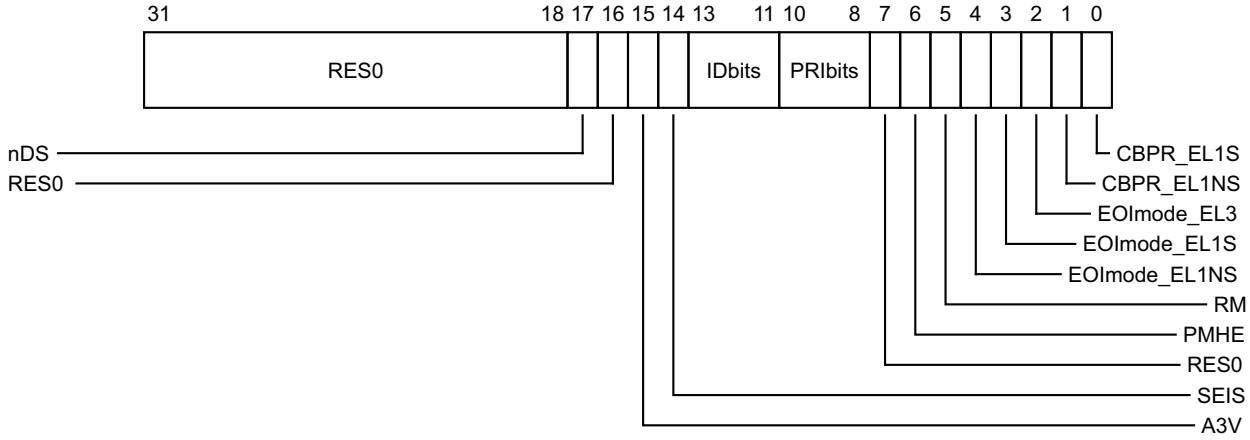

Bits [31:18]

Reserved, RES0.

nDS, bit [17]

Disable Security not supported. Read-only and writes are ignored. Possible values are:

`0` The CPU interface logic supports disabling of security.

1 The CPU interface logic does not support disabling of security, and requires that security is not disabled.

Bit [16]

Reserved, RES0.

A3V, bit [15]

Affinity 3 Valid. Read-only and writes are ignored. Possible values are:

0 The CPU interface logic does not support non-zero values of the Aff3 field in SGI generation System registers.

1 The CPU interface logic supports non-zero values of the Aff3 field in SGI generation System registers.

SEIS, bit [14]

SEI Support. Read-only and writes are ignored. Indicates whether the CPU interface supports generation of SEIs:

0 The CPU interface logic does not support generation of SEIs.

1 The CPU interface logic supports generation of SEIs.

IDbits, bits [13:11]

Identifier bits. Read-only and writes are ignored. The number of physical interrupt identifier bits supported:

000 16 bits.

001 24 bits.

All other values are reserved.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

PRIbits, bits [10:8]

Priority bits. Read-only and writes are ignored. The number of priority bits implemented, minus one.

An implementation that supports two Security states must implement at least 32 levels of physical priority (5 priority bits).

An implementation that supports only a single Security state must implement at least 16 levels of physical priority (4 priority bits).

Note

This field always returns the number of priority bits implemented, regardless of the value of SCR_EL3.NS or the value of GICD_CTLR.DS.

The division between group priority and subpriority is defined in the binary point registers ICC_BPR0_EL1 and ICC_BPR1_EL1.

This field determines the minimum value of ICC_BPR0_EL1.

Bit [7]

Reserved, RES0.

PMHE, bit [6]

Priority Mask Hint Enable.

0 Disables use of the priority mask register as a hint for interrupt distribution.

1 Enables use of the priority mask register as a hint for interrupt distribution.

Software must write ICC_PMR_EL1 to 0xFF before clearing this field to 0.

- An implementation might choose to make this field RAO/WI if priority based routing is always used.
- An implementation might choose to make this field RAZ/WI if priority based routing is never used.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

RM, bit [5]

Routing Modifier. For legacy operation of EL1 software with GICC_CTLR.FIQen set to 1, this bit indicates whether interrupts can be acknowledged or observed as the Highest Priority Pending Interrupt, or whether a special INTID value is returned.

Possible values of this bit are:

0 Secure Group 0 and Non-secure Group 1 interrupts can be acknowledged and observed as the highest priority interrupt at the Secure Exception level where the interrupt is taken.

1 When accessed at EL3 in AArch64 state:

- Secure Group 0 interrupts return a special INTID value of 1020. This affects accesses to ICC_IAR0_EL1 and ICC_HPPIR0_EL1.
- Non-secure Group 1 interrupts return a special INTID value of 1021. This affects accesses to ICC_IAR1_EL1 and ICC_HPPIR1_EL1.

Note

The Routing Modifier bit is supported in AArch64 only. In systems without EL3 the behavior is as if the value is 0.

Software must ensure this bit is 0 when the Secure copy of ICC_SRE_EL1.SRE is 1, otherwise system behavior is UNPREDICTABLE.

In systems without EL3 or where the secure copy of ICC_SRE_EL1 is RES1, this bit is RES0.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

EOImode_EL1NS, bit [4]

EOI mode for interrupts handled at Non-secure EL1 and EL2. Controls whether a write to an End of Interrupt register also deactivates the interrupt:

0 ICC_EOIR0_EL1 and ICC_EOIR1_EL1 provide both priority drop and interrupt deactivation functionality. Accesses to ICC_DIR_EL1 are UNPREDICTABLE.

1 ICC_EOIR0_EL1 and ICC_EOIR1_EL1 provide priority drop functionality only. ICC_DIR_EL1 provides interrupt deactivation functionality.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

EOImode_EL1S, bit [3]

EOI mode for interrupts handled at Secure EL1. Controls whether a write to an End of Interrupt register also deactivates the interrupt:

0 ICC_EOIR0_EL1 and ICC_EOIR1_EL1 provide both priority drop and interrupt deactivation functionality. Accesses to ICC_DIR_EL1 are UNPREDICTABLE.

1 ICC_EOIR0_EL1 and ICC_EOIR1_EL1 provide priority drop functionality only. ICC_DIR_EL1 provides interrupt deactivation functionality.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

EOImode_EL3, bit [2]

EOI mode for interrupts handled at EL3. Controls whether a write to an End of Interrupt register also deactivates the interrupt:

0 ICC_EOIR0_EL1 and ICC_EOIR1_EL1 provide both priority drop and interrupt deactivation functionality. Accesses to ICC_DIR_EL1 are UNPREDICTABLE.

1 ICC_EOIR0_EL1 and ICC_EOIR1_EL1 provide priority drop functionality only. ICC_DIR_EL1 provides interrupt deactivation functionality.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

CBPR_EL1NS, bit [1]

Common Binary Point Register, EL1 Non-secure. Controls whether the same register is used for interrupt preemption of both Group 0 and Group 1 Non-secure interrupts at EL1 and EL2:

0 ICC_BPR0_EL1 determines the preemption group for Group 0 interrupts only.
ICC_BPR1_EL1 determines the preemption group for Non-secure Group 1 interrupts.

1 ICC_BPR0_EL1 determines the preemption group for Group 0 interrupts and Non-secure Group 1 interrupts. Non-secure accesses to GICC_BPR and ICC_BPR1_EL1 access the state of ICC_BPR0_EL1.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

CBPR_EL1S, bit [0]

Common Binary Point Register, EL1 Secure. Controls whether the same register is used for interrupt preemption of both Group 0 and Group 1 Secure interrupts at EL1:

0 ICC_BPR0_EL1 determines the preemption group for Group 0 interrupts only.
ICC_BPR1_EL1 determines the preemption group for Secure Group 1 interrupts.

1 ICC_BPR0_EL1 determines the preemption group for Group 0 interrupts and Secure Group 1 interrupts. Secure accesses to ICC_BPR1_EL1 access the state of ICC_BPR0_EL1.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the ICC_CTLR_EL3:

To access the ICC_CTLR_EL3:

```
MRS <Xt>, ICC_CTLR_EL3 ; Read ICC_CTLR_EL3 into Xt
MSR ICC_CTLR_EL3, <Xt> ; Write Xt to ICC_CTLR_EL3
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	110	1100	1100	100

8.2.8 ICC_DIR_EL1, Interrupt Controller Deactivate Interrupt Register

The ICC_DIR_EL1 characteristics are:

Purpose

When interrupt priority drop is separated from interrupt deactivation, a write to this register deactivates the specified interrupt.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	WO	WO	WO	WO	WO

The ICC_DIR_EL1 register is only accessible at Non-secure EL1 when HCR_EL2.{FMO, IMO} == {0, 0}.

Note

At Non-secure EL1, the instruction encoding used to access ICC_DIR_EL1 results in an access to ICV_DIR_EL1 in the following cases:

- When HCR_EL2.FMO is set to 1.
- When HCR_EL2.IMO is set to 1.

There are two cases when writing to ICC_DIR_EL1 that were UNPREDICTABLE for a corresponding GICv2 write to GICC_DIR:

- When EOImode == '0'. GICv3 implementations must ignore such writes. In systems supporting system error generation, an implementation might generate an SEI.
- When EOImode == '1' but no EOI has been issued. The interrupt will be de-activated by the Distributor, however the active priority in the CPU interface for the interrupt will remain set (because no EOI was issued).

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*M.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, write accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, write accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, write accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TDIR==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, Secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, write accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR_EL2.IMO==0, and HCR_EL2.FMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_DIR_EL1 performs the same function as AArch32 System operation ICC_DIR.

Attributes

ICC_DIR_EL1 is a 32-bit register.

Field descriptions

The ICC_DIR_EL1 bit assignments are:

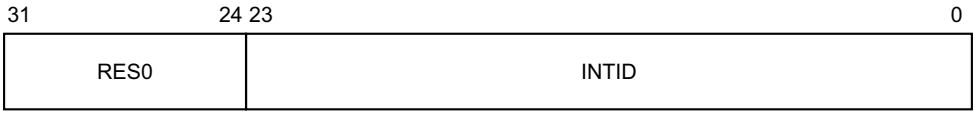

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the interrupt to be deactivated.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICC_CTLR_EL1.IDbits and ICC_CTLR_EL3.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICC_DIR_EL1:

To access the ICC_DIR_EL1:

```
MSR ICC_DIR_EL1, <Xt> ; Write Xt to ICC_DIR_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1011	001

This encoding results in an access to ICV_DIR_EL1 at Non-secure EL1 in the following cases:

- When HCR_EL2.FMO is set to 1, and the INTID field refers to a Group 0 interrupt.
- When HCR_EL2.IMO is set to 1, and the INTID field refers to a Group 1 interrupt.

8.2.9 ICC_EOIR0_EL1, Interrupt Controller End Of Interrupt Register 0

The ICC_EOIR0_EL1 characteristics are:

Purpose

A PE writes to this register to inform the CPU interface that it has completed the processing of the specified Group 0 interrupt.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	WO	WO	WO	WO	WO

ICC_EOIR0_EL1 is only accessible at Non-secure EL1 when HCR_EL2.FMO is set to 0.

Note

When HCR_EL2.FMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_EOIR0_EL1 results in an access to ICV_EOIR0_EL1.

A write to this register must correspond to the most recent valid read from an Interrupt Acknowledge Register for which there has not been a priority drop and that this identifier was read from ICC_IAR0_EL1 while operating in the same Security state as that in which the write occurs, otherwise the system behavior is UNPREDICTABLE. A valid read is a read that returns a valid interrupt ID, that is not a special INTID.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, write accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, write accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, write accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TALL0==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.FIQ==1, Secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, write accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR_EL2.FMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_EOIR0_EL1 performs the same function as AArch32 System operation ICC_EOIR0.

Attributes

ICC_EOIR0_EL1 is a 32-bit register.

Field descriptions

The ICC_EOIR0_EL1 bit assignments are:

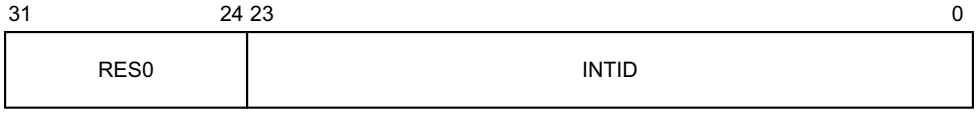

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID from the corresponding ICC_IAR0_EL1 access.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICC_CTLR_EL1.IDbits and ICC_CTLR_EL3.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

If the EOImode bit for the current Exception level and Security state is 0, a write to this register drops the priority for the interrupt, and also deactivates the interrupt.

If the EOImode bit for the current Exception level and Security state is 1, a write to this register only drops the priority for the interrupt. Software must write to ICC_DIR_EL1 to deactivate the interrupt.

The EOImode bit for the current Exception level and Security state is determined as follows:

- If EL3 is not implemented, the appropriate bit is ICC_CTLR_EL1.EOImode.
- If EL3 is implemented and the software is executing at EL3, the appropriate bit is ICC_CTLR_EL3.EOImode_EL3.
- If EL3 is implemented and the software is not executing at EL3, the bit depends on the current Security state:
  - If the software is executing in Secure state, the bit is ICC_CTLR_EL3.EOImode_EL1S.
  - If the software is executing in Non-secure state, the bit is ICC_CTLR_EL3.EOImode_EL1NS.

Accessing the ICC_EOIR0_EL1:

To access the ICC_EOIR0_EL1:

```
MSR ICC_EOIR0_EL1, <Xt> ; Write Xt to ICC_EOIR0_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1000	001

When HCR_EL2.FMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_EOIR0_EL1.

8.2.10 ICC_EOIR1_EL1, Interrupt Controller End Of Interrupt Register 1

The ICC_EOIR1_EL1 characteristics are:

Purpose

A PE writes to this register to inform the CPU interface that it has completed the processing of the specified Group 1 interrupt.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	WO	WO	WO	WO	WO

ICC_EOIR1_EL1 is only accessible at Non-secure EL1 when HCR_EL2.IMO is set to 0.

Note

When HCR_EL2.IMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_EOIR1_EL1 results in an access to ICV_EOIR1_EL1.

A write to this register must correspond to the most recent valid read from an Interrupt Acknowledge Register for which there has not been a priority drop and that this identifier was read from ICC_IAR1_EL1 while operating in the same Security state as that in which the write occurs, otherwise the system behavior is UNPREDICTABLE. A valid read is a read that returns a valid interrupt ID, that is not a special INTID.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, write accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, write accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, write accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TALL1==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.IRQ==1, Secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, write accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR_EL2.IMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_EOIR1_EL1 performs the same function as AArch32 System operation ICC_EOIR1.

Attributes

ICC_EOIR1_EL1 is a 32-bit register.

Field descriptions

The ICC_EOIR1_EL1 bit assignments are:

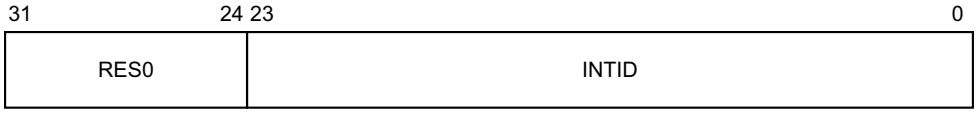

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID from the corresponding ICC_IAR1_EL1 access.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICC_CTLR_EL1.IDbits and ICC_CTLR_EL3.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

If the EOImode bit for the current Exception level and Security state is 0, a write to this register drops the priority for the interrupt, and also deactivates the interrupt.

If the EOImode bit for the current Exception level and Security state is 1, a write to this register only drops the priority for the interrupt. Software must write to ICC_DIR_EL1 to deactivate the interrupt.

The EOImode bit for the current Exception level and Security state is determined as follows:

- If EL3 is not implemented, the appropriate bit is ICC_CTLR_EL1.EOImode.
- If EL3 is implemented and the software is executing at EL3, the appropriate bit is ICC_CTLR_EL3.EOImode_EL3.
- If EL3 is implemented and the software is not executing at EL3, the bit depends on the current Security state:
  - If the software is executing in Secure state, the bit is ICC_CTLR_EL3.EOImode_EL1S.
  - If the software is executing in Non-secure state, the bit is ICC_CTLR_EL3.EOImode_EL1NS.

Accessing the ICC_EOIR1_EL1:

To access the ICC_EOIR1_EL1:

```
MSR ICC_EOIR1_EL1, <Xt> ; Write Xt to ICC_EOIR1_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	001

When HCR_EL2.IMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_EOIR1_EL1.

8.2.11 ICC_HPPIR0_EL1, Interrupt Controller Highest Priority Pending Interrupt Register 0

The ICC_HPPIR0_EL1 characteristics are:

Purpose

Indicates the highest priority pending Group 0 interrupt on the CPU interface.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RO	RO	RO	RO	RO

ICC_HPPIR0_EL1 is only accessible at Non-secure EL1 when HCR_EL2.FMO is set to 0.

Note

When HCR_EL2.FMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_HPPIR0_EL1 results in an access to ICV_HPPIR0_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, read accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, read accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, read accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TALL0==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.FIQ==1, Secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, read accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR_EL2.FMO==0, Non-secure read accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_HPPIR0_EL1 performs the same function as AArch32 System operation ICC_HPPIR0.

Attributes

ICC_HPPIR0_EL1 is a 32-bit register.

Field descriptions

The ICC_HPPIR0_EL1 bit assignments are:

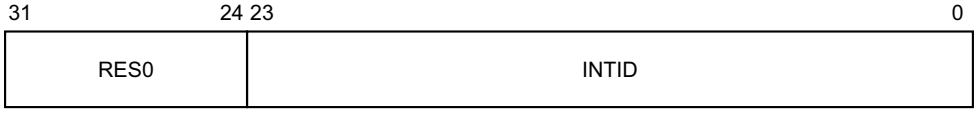

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the highest priority pending interrupt, if that interrupt is observable at the current Security state and Exception level.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. These special INTIDs can be one of: 1020, 1021, or 1023. See *Special INTIDs on page 3-40*, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICC_CTLR_EL1.IDbits and ICC_CTLR_EL3.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICC_HPPIR0_EL1:

To access the ICC_HPPIR0_EL1:

```
MRS <Xt>, ICC_HPPIR0_EL1 ; Read ICC_HPPIR0_EL1 into Xt
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1000	010

When HCR_EL2.FMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_HPPIR0_EL1.

8.2.12 ICC_HPPIR1_EL1, Interrupt Controller Highest Priority Pending Interrupt Register 1

The ICC_HPPIR1_EL1 characteristics are:

Purpose

Indicates the highest priority pending Group 1 interrupt on the CPU interface.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RO	RO	RO	RO	RO

ICC_HPPIR1_EL1 is only accessible at Non-secure EL1 when HCR_EL2.IMO is set to 0.

Note

When HCR_EL2.IMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_HPPIR1_EL1 results in an access to ICV_HPPIR1_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, read accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, read accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, read accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TALL1==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.IRQ==1, Secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, read accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR_EL2.IMO==0, Non-secure read accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_HPPIR1_EL1 performs the same function as AArch32 System operation ICC_HPPIR1.

Attributes

ICC_HPPIR1_EL1 is a 32-bit register.

Field descriptions

The ICC_HPPIR1_EL1 bit assignments are:

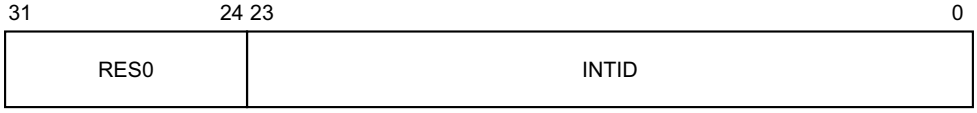

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the highest priority pending interrupt, if that interrupt is observable at the current Security state and Exception level.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. These special INTIDs can be one of: 1020, 1021, or 1023. See *Special INTIDs on page 3-40*, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICC_CTLR_EL1.IDbits and ICC_CTLR_EL3.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICC_HPPIR1_EL1:

To access the ICC_HPPIR1_EL1:

```
MRS <Xt>, ICC_HPPIR1_EL1 ; Read ICC_HPPIR1_EL1 into Xt
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	010

When HCR_EL2.IMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_HPPIR1_EL1.

8.2.13 ICC_IAR0_EL1, Interrupt Controller Interrupt Acknowledge Register 0

The ICC_IAR0_EL1 characteristics are:

Purpose

The PE reads this register to obtain the INTID of the signaled Group 0 interrupt. This read acts as an acknowledge for the interrupt.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RO	RO	RO	RO	RO

ICC_IAR0_EL1 is only accessible at Non-secure EL1 when HCR_EL2.FMO is set to 0.

Note

When HCR_EL2.FMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_IAR0_EL1 results in an access to ICV_IAR0_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, read accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, read accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, read accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TALL0==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.FIQ==1, Secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, read accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR_EL2.FMO==0, Non-secure read accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_IAR0_EL1 performs the same function as AArch32 System operation ICC_IAR0.

To allow software to ensure appropriate observability of actions initiated by GIC register accesses, the PE and CPU interface logic must ensure that reads of this register are self-synchronising when interrupts are masked by the PE (that is when PSTATE.{I,F} == {0,0}). This ensures that the effect of activating an interrupt on the signaling of interrupt exceptions is observed when a read of this register is architecturally executed so that no spurious interrupt exception occurs if interrupts are unmasked by an instruction immediately following the read. See *Observability of the effects of accesses to the GIC registers* on page 8-159, for more information.

Attributes

ICC_IAR0_EL1 is a 32-bit register.

Field descriptions

The ICC_IAR0_EL1 bit assignments are:

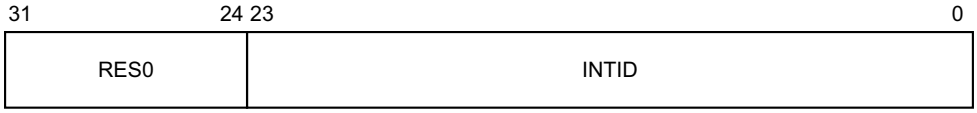

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled interrupt.

This is the INTID of the highest priority pending interrupt, if that interrupt is of sufficient priority for it to be signaled to the PE, and if it can be acknowledged at the current Security state and Exception level.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. These special INTIDs can be one of: 1020, 1021, or 1023. See *Special INTIDs on page 3-40*, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICC_CTLR_EL1.IDbits and ICC_CTLR_EL3.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICC_IAR0_EL1:

To access the ICC_IAR0_EL1:

```
MRS <Xt>, ICC_IAR0_EL1 ; Read ICC_IAR0_EL1 into Xt
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1000	000

When HCR_EL2.FMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_IAR0_EL1.

8.2.14 ICC_IAR1_EL1, Interrupt Controller Interrupt Acknowledge Register 1

The ICC_IAR1_EL1 characteristics are:

Purpose

The PE reads this register to obtain the INTID of the signaled Group 1 interrupt. This read acts as an acknowledge for the interrupt.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RO	RO	RO	RO	RO

ICC_IAR1_EL1 is only accessible at Non-secure EL1 when HCR_EL2.IMO is set to 0.

Note

When HCR_EL2.IMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_IAR1_EL1 results in an access to ICV_IAR1_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, read accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, read accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, read accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TALL1==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.IRQ==1, Secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, read accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR_EL2.IMO==0, Non-secure read accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_IAR1_EL1 performs the same function as AArch32 System operation ICC_IAR1.

To allow software to ensure appropriate observability of actions initiated by GIC register accesses, the PE and CPU interface logic must ensure that reads of this register are self-synchronising when interrupts are masked by the PE (that is when PSTATE.{I,F} == {0,0}). This ensures that the effect of activating an interrupt on the signaling of interrupt exceptions is observed when a read of this register is architecturally executed so that no spurious interrupt exception occurs if interrupts are unmasked by an instruction immediately following the read. See *Observability of the effects of accesses to the GIC registers* on page 8-159, for more information.

Attributes

ICC_IAR1_EL1 is a 32-bit register.

Field descriptions

The ICC_IAR1_EL1 bit assignments are:

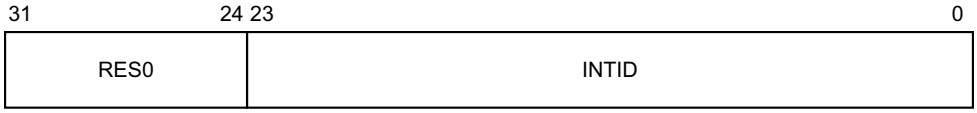

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled interrupt.

This is the INTID of the highest priority pending interrupt, if that interrupt is of sufficient priority for it to be signaled to the PE, and if it can be acknowledged at the current Security state and Exception level.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. These special INTIDs can be one of: 1020, 1021, or 1023. See *Special INTIDs on page 3-40*, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICC_CTLR_EL1.IDbits and ICC_CTLR_EL3.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICC_IAR1_EL1:

To access the ICC_IAR1_EL1:

```
MRS <Xt>, ICC_IAR1_EL1 ; Read ICC_IAR1_EL1 into Xt
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	000

When HCR_EL2.IMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_IAR1_EL1.

8.2.15 ICC_IGRPEN0_EL1, Interrupt Controller Interrupt Group 0 Enable register

The ICC_IGRPEN0_EL1 characteristics are:

Purpose

Controls whether Group 0 interrupts are enabled or not.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	RW	RW	RW	RW

ICC_IGRPEN0_EL1 is only accessible at Non-secure EL1 when HCR_EL2.FMO is set to 0.

Note

When HCR_EL2.FMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_IGRPEN0_EL1 results in an access to ICV_IGRPEN0_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.FIQ==1, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR_EL2.FMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_IGRPEN0_EL1 is architecturally mapped to AArch32 System register ICC_IGRPEN0.

Attributes

ICC_IGRPEN0_EL1 is a 32-bit register.

Field descriptions

The ICC_IGRPEN0_EL1 bit assignments are:

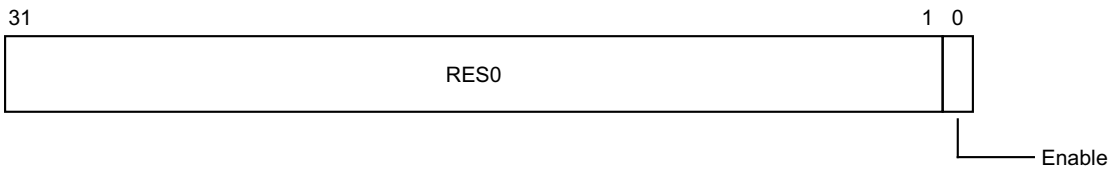

Bits [31:1]

Reserved, RES0.

Enable, bit [0]

Enables Group 0 interrupts.

0 Group 0 interrupts are disabled.

1 Group 0 interrupts are enabled.

Virtual accesses to this register update ICH_VMCR_EL2.VENG0.

If the highest priority pending interrupt for that PE is a Group 0 interrupt using 1 of N targeting, then the interrupt will be targeted to another PE as a result of the Enable bit changing from 1 to 0.

When this register has an architecturally-defined reset value, this field resets to 0.

Accessing the ICC_IGRPEN0_EL1:

To access the ICC_IGRPEN0_EL1:

```
MRS <Xt>, ICC_IGRPEN0_EL1 ; Read ICC_IGRPEN0_EL1 into Xt
MSR ICC_IGRPEN0_EL1, <Xt> ; Write Xt to ICC_IGRPEN0_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	110

When HCR_EL2.FMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_IGRPEN0_EL1.

8.2.16 ICC_IGRPEN1_EL1, Interrupt Controller Interrupt Group 1 Enable register

The ICC_IGRPEN1_EL1 characteristics are:

Purpose

Controls whether Group 1 interrupts are enabled for the current Security state.

Usage constraints

ICC_IGRPEN1_EL1(S) is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	-	-	RW

ICC_IGRPEN1_EL1(NS) is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	-	RW	RW	-

ICC_IGRPEN1_EL1 is only accessible at Non-secure EL1 when HCR_EL2.IMO is set to 0.

Note

When HCR_EL2.IMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_IGRPEN1_EL1 results in an access to ICV_IGRPEN1_EL1.

If EL3 is present and this register is accessed at EL3, the copy of this register appropriate to the current setting of SCR_EL3.NS is accessed.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.IRQ==1, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR_EL2.IMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_IGRPEN1_EL1(S) is architecturally mapped to AArch32 System register ICC_IGRPEN1 (S).

AArch64 System register ICC_IGRPEN1_EL1(NS) is architecturally mapped to AArch32 System register ICC_IGRPEN1 (NS).

Attributes

ICC_IGRPEN1_EL1 is a 32-bit register.

Field descriptions

The ICC_IGRPEN1_EL1 bit assignments are:

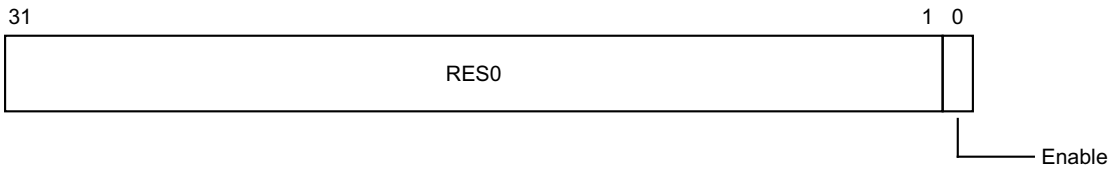

Bits [31:1]

Reserved, RES0.

Enable, bit [0]

Enables Group 1 interrupts for the current Security state.

`0` Group 1 interrupts are disabled for the current Security state.

`1` Group 1 interrupts are enabled for the current Security state.

Virtual accesses to this register update ICH_VMCR_EL2.VENG1.

If EL3 is present:

- The Secure ICC_IGRPEN1_EL1.Enable bit is a read/write alias of the ICC_IGRPEN1_EL3.EnableGrp1S bit.
- The Non-secure ICC_IGRPEN1_EL1.Enable bit is a read/write alias of the ICC_IGRPEN1_EL3.EnableGrp1NS bit.

If the highest priority pending interrupt for that PE is a Group 1 interrupt using 1 of N model, then the interrupt will target another PE as a result of the Enable bit changing from 1 to 0.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the ICC_IGRPEN1_EL1:

To access the ICC_IGRPEN1_EL1:

```
MRS <Xt>, ICC_IGRPEN1_EL1 ; Read ICC_IGRPEN1_EL1 into Xt
MSR ICC_IGRPEN1_EL1, <Xt> ; Write Xt to ICC_IGRPEN1_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	111

When HCR_EL2.IMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_IGRPEN1_EL1.

8.2.17 ICC_IGRPEN1_EL3, Interrupt Controller Interrupt Group 1 Enable register (EL3)

The ICC_IGRPEN1_EL3 characteristics are:

Purpose

Controls whether Group 1 interrupts are enabled or not.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	-	RW	RW

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

Configurations

AArch64 System register ICC_IGRPEN1_EL3 can be mapped to AArch32 System register ICC_MGRPEN1, but this is not architecturally mandated.

Attributes

ICC_IGRPEN1_EL3 is a 32-bit register.

Field descriptions

The ICC_IGRPEN1_EL3 bit assignments are:

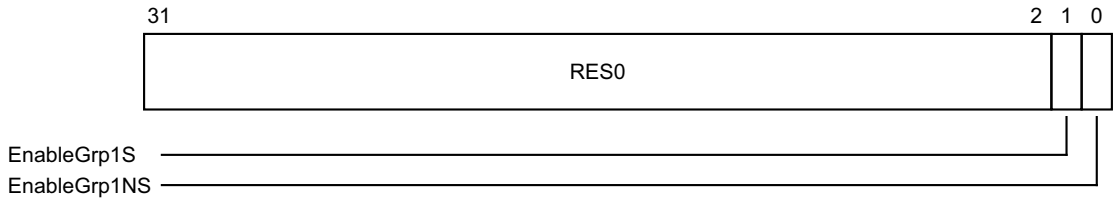

Bits [31:2]

Reserved, RES0.

EnableGrp1S, bit [1]

Enables Group 1 interrupts for the Secure state.

`0` Secure Group 1 interrupts are disabled.

`1` Secure Group 1 interrupts are enabled.

If EL3 is present, the Secure ICC_IGRPEN1_EL1.Enable bit is a read/write alias of the ICC_IGRPEN1_EL3.EnableGrp1S bit.

If the highest priority pending interrupt for that PE is a Group 1 interrupt using 1 of N model, then the interrupt will target another PE as a result of the Enable bit changing from 1 to 0.

When this register has an architecturally-defined reset value, this field resets to `0`.

EnableGrp1NS, bit [0]

Enables Group 1 interrupts for the Non-secure state.

`0` Non-secure Group 1 interrupts are disabled.

`1` Non-secure Group 1 interrupts are enabled.

If EL3 is present, the Non-secure ICC_IGRPEN1_EL1.Enable bit is a read/write alias of the ICC_IGRPEN1_EL3.EnableGrp1NS bit.

If the highest priority pending interrupt for that PE is a Group 1 interrupt using 1 of N model, then the interrupt will target another PE as a result of the Enable bit changing from 1 to 0.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the ICC_IGRPEN1_EL3:

To access the ICC_IGRPEN1_EL3:

```
MRS <Xt>, ICC_IGRPEN1_EL3 ; Read ICC_IGRPEN1_EL3 into Xt
MSR ICC_IGRPEN1_EL3, <Xt> ; Write Xt to ICC_IGRPEN1_EL3
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	110	1100	1100	111

8.2.18 ICC_PMR_EL1, Interrupt Controller Interrupt Priority Mask Register

The ICC_PMR_EL1 characteristics are:

Purpose

Provides an interrupt priority filter. Only interrupts with a higher priority than the value in this register are signaled to the PE.

Writes to this register must be high performance and must ensure that no interrupt of lower priority than the written value occurs after the write, without requiring an ISB or an exception boundary.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	RW	RW	RW	RW

ICC_PMR_EL1 is only accessible at Non-secure EL1 when HCR_EL2.{FMO, IMO} == {0, 0}.

Note

When HCR_EL2.{FMO, IMO} != {0, 0}, at Non-secure EL1, the instruction encoding used to access ICC_PMR_EL1 results in an access to ICV_PMR_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TC==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR_EL2.IMO==0, and HCR_EL2.FMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_PMR_EL1 is architecturally mapped to AArch32 System register ICC_PMR.

To allow software to ensure appropriate observability of actions initiated by GIC register accesses, the PE and CPU interface logic must ensure that writes to this register are self-synchronising. This ensures that no interrupts below the written PMR value will be taken after a write to this register is architecturally executed. See *Observability of the effects of accesses to the GIC registers* on page 8-159, for more information.

Attributes

ICC_PMR_EL1 is a 32-bit register.

Field descriptions

The ICC_PMR_EL1 bit assignments are:

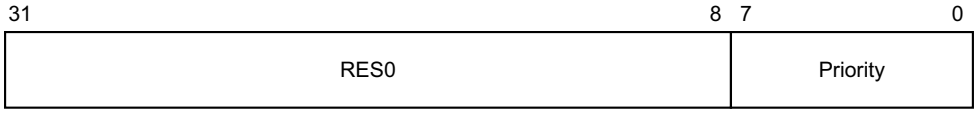

Bits [31:8]

Reserved, RES0.

Priority, bits [7:0]

The priority mask level for the CPU interface. If the priority of an interrupt is higher than the value indicated by this field, the interface signals the interrupt to the PE.

The possible priority field values are as follows:

Implemented priority bits	Possible priority field values	Number of priority levels
[7:0]	0x00-0xFF (0-255), all values	256
[7:1]	0x00-0xFE (0-254), even values only	128
[7:2]	0x00-0xFC (0-252), in steps of 4	64
[7:3]	0x00-0xF8 (0-248), in steps of 8	32
[7:4]	0x00-0xF0 (0-240), in steps of 16	16

Unimplemented priority bits are RAZ/WI.

When this register has an architecturally-defined reset value, this field resets to 0.

Accessing the ICC_PMR_EL1:

To access the ICC_PMR_EL1:

```
MRS <Xt>, ICC_PMR_EL1 ; Read ICC_PMR_EL1 into Xt
MSR ICC_PMR_EL1, <Xt> ; Write Xt to ICC_PMR_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	0100	0110	000

When HCR_EL2.{FMO, IMO} != {0, 0}, execution of this encoding at Non-secure EL1 results in an access to ICV_PMR_EL1.

8.2.19 ICC_RPR_EL1, Interrupt Controller Running Priority Register

The ICC_RPR_EL1 characteristics are:

Purpose

Indicates the Running priority of the CPU interface.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RO	RO	RO	RO	RO

ICC_RPR_EL1 is only accessible at Non-secure EL1 when HCR_EL2.{FMO, IMO} == {0, 0}.

Note

When HCR_EL2.{FMO, IMO} != {0, 0}, at Non-secure EL1, the instruction encoding used to access ICC_RPR_EL1 results in an access to ICV_RPR_EL1.

Software cannot determine the number of implemented priority bits from a read of this register.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, read accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, read accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, read accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TC==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, Secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, read accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR_EL2.IMO==0, and HCR_EL2.FMO==0, Non-secure read accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_RPR_EL1 performs the same function as AArch32 System operation ICC_RPR.

Attributes

ICC_RPR_EL1 is a 32-bit register.

Field descriptions

The ICC_RPR_EL1 bit assignments are:

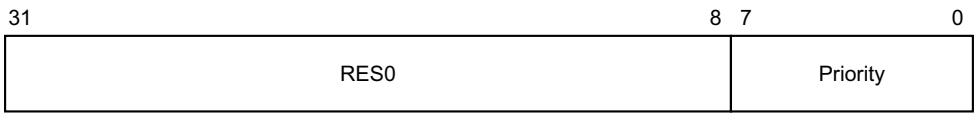

Bits [31:8]

Reserved, RES0.

Priority, bits [7:0]

The current running priority on the CPU interface. This is the group priority of the current active interrupt.

If there are no active interrupts on the CPU interface, or all active interrupts have undergone a priority drop, the value returned is the Idle priority.

The priority returned is the group priority as if the BPR for the current Exception level and Security state was set to the minimum value of BPR for the number of implemented priority bits.

Note

If 8 bits of priority are implemented the group priority is bits[7:1] of the priority.

Accessing the ICC_RPR_EL1:

To access the ICC_RPR_EL1:

```
MRS <Xt>, ICC_RPR_EL1 ; Read ICC_RPR_EL1 into Xt
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1011	011

When HCR_EL2.{FMO, IMO} != {0, 0}, execution of this encoding at Non-secure EL1 results in an access to ICV_RPR_EL1.

8.2.20 ICC_SGI0R_EL1, Interrupt Controller Software Generated Interrupt Group 0 Register

The ICC_SGI0R_EL1 characteristics are:

Purpose

Generates Secure Group 0 SGIs, including from the Non-secure state when permitted by GICR_NSACR.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	WO	WO	WO	WO	WO

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If HCR_EL2.FMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HCR_EL2.IMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICC_SRE_EL1.SRE==0, write accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, write accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, write accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, Secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, write accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR_EL2.IMO==0, and HCR_EL2.FMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_SGI0R_EL1 performs the same function as AArch32 System operation ICC_SGI0R.

Attributes

ICC_SGI0R_EL1 is a 64-bit register.

Field descriptions

The ICC_SGI0R_EL1 bit assignments are:

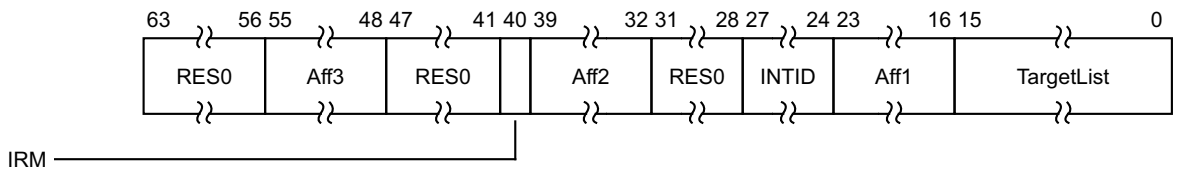

Bits [63:56]

Reserved, RES0.

Aff3, bits [55:48]

The affinity 3 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

Bits [47:41]

Reserved, RES0.

IRM, bit [40]

Interrupt Routing Mode. Determines how the generated interrupts should be distributed to PEs. Possible values are:

`0` Interrupts routed to the PEs specified by Aff3.Aff2.Aff1.<target list>.

`1` Interrupts routed to all PEs in the system, excluding "self".

Aff2, bits [39:32]

The affinity 2 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

Bits [31:28]

Reserved, RES0.

INTID, bits [27:24]

The INTID of the SGI.

Aff1, bits [23:16]

The affinity 1 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

TargetList, bits [15:0]

Target List. The set of PEs for which SGI interrupts will be generated. Each bit corresponds to the PE within a cluster with an Affinity 0 value equal to the bit number.

If a bit is 1 and the bit does not correspond to a valid target PE, the bit must be ignored by the Distributor. It is IMPLEMENTATION DEFINED whether, in such cases, a Distributor can signal a system error.

Note

This restricts a system to sending targeted SGIs to PEs with an affinity 0 number that is less than 16.

If SRE is set only for secure EL3, software executing at EL3 might use the System register interface to generate SGIs. Hence, the Distributor must always be able to receive and acknowledge Generate SGI packets received from CPU interface regardless of the ARE settings for a Security state. However, the Distributor might discard such packets.

If the IRM bit is 1, this field is RES0.

Accessing the ICC_SGI0R_EL1:

To access the ICC_SGI0R_EL1:

```
MSR ICC_SGI0R_EL1, <Xt> ; Write Xt to ICC_SGI0R_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1011	111

8.2.21 ICC_SGI1R_EL1, Interrupt Controller Software Generated Interrupt Group 1 Register

The ICC_SGI1R_EL1 characteristics are:

Purpose

Generates Group 1 SGIs for the current Security state.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	WO	WO	WO	WO	WO

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If HCR_EL2.FMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HCR_EL2.IMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICC_SRE_EL1.SRE==0, write accesses to this register from EL1 are trapped to EL1.

If ICC_SRE_EL2.SRE==0, write accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, write accesses to this register from EL3 are trapped to EL3.

If ICH_HCR_EL2.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, Secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, write accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR_EL2.IMO==0, and HCR_EL2.FMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

Configurations

AArch64 System register ICC_SGI1R_EL1 performs the same function as AArch32 System operation ICC_SGI1R.

Under certain conditions a write to ICC_SGI1R_EL1 can generate Group 0 interrupts, see Table 8-14 on page 8-171.

Attributes

ICC_SGI1R_EL1 is a 64-bit register.

Field descriptions

The ICC_SGI1R_EL1 bit assignments are:

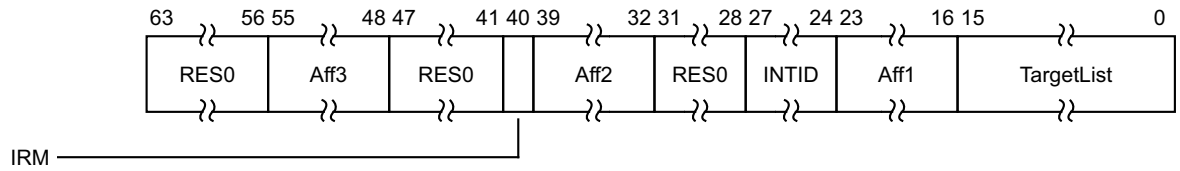

Bits [63:56]

Reserved, RES0.

Aff3, bits [55:48]

The affinity 3 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

Bits [47:41]

Reserved, RES0.

IRM, bit [40]

Interrupt Routing Mode. Determines how the generated interrupts should be distributed to PEs. Possible values are:

`0` Interrupts routed to the PEs specified by Aff3.Aff2.Aff1.<target list>.

`1` Interrupts routed to all PEs in the system, excluding "self".

Aff2, bits [39:32]

The affinity 2 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

Bits [31:28]

Reserved, RES0.

INTID, bits [27:24]

The INTID of the SGI.

Aff1, bits [23:16]

The affinity 1 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

TargetList, bits [15:0]

Target List. The set of PEs for which SGI interrupts will be generated. Each bit corresponds to the PE within a cluster with an Affinity 0 value equal to the bit number.

If a bit is 1 and the bit does not correspond to a valid target PE, the bit must be ignored by the Distributor. It is IMPLEMENTATION DEFINED whether, in such cases, a Distributor can signal a system error.

Note

This restricts a system to sending targeted SGIs to PEs with an affinity 0 number that is less than 16.

If SRE is set only for secure EL3, software executing at EL3 might use the System register interface to generate SGIs. Hence, the Distributor must always be able to receive and acknowledge Generate SGI packets received from CPU interface regardless of the ARE settings for a Security state. However, the Distributor might discard such packets.

If the IRM bit is 1, this field is RES0.

Accessing the ICC_SGI1R_EL1:

To access the ICC_SGI1R_EL1:

```
MSR ICC_SGI1R_EL1, <Xt> ; Write Xt to ICC_SGI1R_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1011	101

8.2.22 ICC_SRE_EL1, Interrupt Controller System Register Enable register (EL1)

The ICC_SRE_EL1 characteristics are:

Purpose

Controls whether the System register interface or the memory-mapped interface to the GIC CPU interface is used for EL1.

Usage constraints

ICC_SRE_EL1(S) is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	-	-	RW

ICC_SRE_EL1(NS) is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	-	RW	RW	-

Execution with ICC_SRE_EL1.SRE set to 0 might make some System registers UNKNOWN.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL2.Enable==0, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE_EL3.Enable==0, accesses to this register from EL1 and EL2 are trapped to EL3.

Configurations

AArch64 System register ICC_SRE_EL1(S) is architecturally mapped to AArch32 System register ICC_SRE (S).

AArch64 System register ICC_SRE_EL1(NS) is architecturally mapped to AArch32 System register ICC_SRE (NS).

Attributes

ICC_SRE_EL1 is a 32-bit register.

Field descriptions

The ICC_SRE_EL1 bit assignments are:

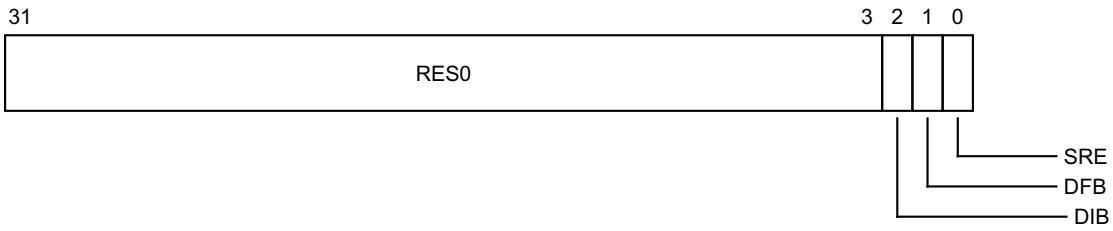

Bits [31:3]

Reserved, RES0.

DIB, bit [2]

Disable IRQ bypass.

0 IRQ bypass enabled.

1 IRQ bypass disabled.

If EL3 is implemented and GICD_CTLR.DS == 0, this field is a read-only alias of ICC_SRE_EL3.DIB.

If EL3 is implemented and GICD_CTLR.DS == 1, and EL2 is not implemented, this field is a read-write alias of ICC_SRE_EL3.DIB.

If EL3 is not implemented or GICD_CTLR.DS == 1, and EL2 is implemented, this field is a read-only alias of ICC_SRE_EL2.DIB.

In systems that do not support IRQ bypass, this field is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

DFB, bit [1]

Disable FIQ bypass.

0 FIQ bypass enabled.

1 FIQ bypass disabled.

If EL3 is implemented and GICD_CTLR.DS == 0, this field is a read-only alias of ICC_SRE_EL3.DFB.

If EL3 is implemented and GICD_CTLR.DS == 1, and EL2 is not implemented, this field is a read-write alias of ICC_SRE_EL3.DFB.

If EL3 is not implemented or GICD_CTLR.DS == 1, and EL2 is implemented, this field is a read-only alias of ICC_SRE_EL2.DFB.

In systems that do not support FIQ bypass, this field is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

SRE, bit [0]

System Register Enable.

0 The memory-mapped interface must be used. Access at EL1 to any ICC_* System register other than ICC_SRE_EL1 is trapped to EL1.

1 The System register interface for the current Security state is enabled.

Virtual accesses modify ICH_VMCR_EL2.VSRE.

If software changes this bit from 1 to 0 in the Secure instance of this register, the results are UNPREDICTABLE.

If an implementation supports only a System register interface to the GIC CPU interface, this bit is RAO/WI.

If EL3 is implemented and ICC_SRE_EL3.SRE==0 the Secure copy of this bit is RAZ/WI.

If EL2 is implemented and ICC_SRE_EL2.SRE==0 the Non-secure copy of this bit is RAZ/WI.

If EL3 is implemented and ICC_SRE_EL3.SRE==0 the Non-secure copy of this bit is RAZ/WI.

GICv3 implementations that do not require GICv2 compatibility might choose to make this bit RAO/WI. The following options are supported:

- The non-secure copy of ICC_SRE_EL1.SRE may be RAO/WI if ICC_SRE_EL2.SRE is also RAO/WI. This means all Non-secure software, including VMs using only virtual interrupts, must access the GIC using System registers.

- The secure copy of ICC_SRE_EL1.SRE might be RAO/WI if ICC_SRE_EL3.SRE and ICC_SRE_EL2.SRE are also RAO/WI. This means that all Secure software must access the GIC using System registers and all Non-secure accesses to registers for physical interrupts must use system registers.

Note

A VM using only virtual interrupts might still use memory-mapped access if the Non-secure copy of ICC_SRE_EL1.SRE is not RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

Accessing the ICC_SRE_EL1:

To access the ICC_SRE_EL1:

```
MRS <Xt>, ICC_SRE_EL1 ; Read ICC_SRE_EL1 into Xt
MSR ICC_SRE_EL1, <Xt> ; Write Xt to ICC_SRE_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	101

8.2.23 ICC_SRE_EL2, Interrupt Controller System Register Enable register (EL2)

The ICC_SRE_EL2 characteristics are:

Purpose

Controls whether the System register interface or the memory-mapped interface to the GIC CPU interface is used for EL2.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RW	RW	-

Execution with ICC_SRE_EL2.SRE set to 0 might make some System registers UNKNOWN.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL3.Enable==0, accesses to this register from EL2 are trapped to EL3.

Configurations

AArch64 System register ICC_SRE_EL2 is architecturally mapped to AArch32 System register ICC_HSRE.

Attributes

ICC_SRE_EL2 is a 32-bit register.

Field descriptions

The ICC_SRE_EL2 bit assignments are:

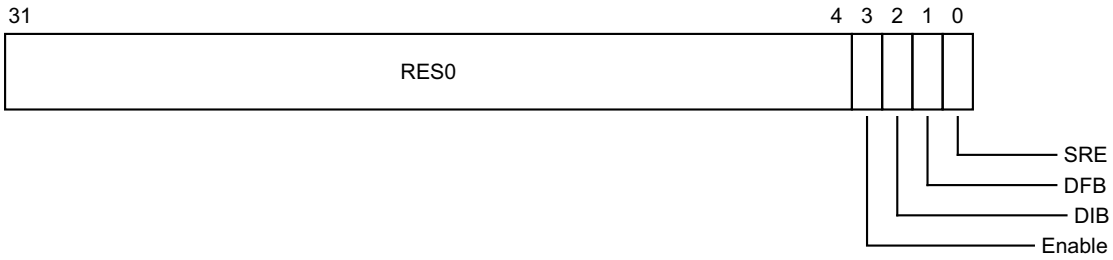

Bits [31:4]

Reserved, RES0.

Enable, bit [3]

Enable. Enables lower Exception level access to ICC_SRE_EL1.

`0` Non-secure EL1 accesses to ICC_SRE_EL1 trap to EL2.

`1` Non-secure EL1 accesses to ICC_SRE_EL1 do not trap to EL2.

If ICC_SRE_EL2.SRE is RAO/WI, an implementation is permitted to make the Enable bit RAO/WI.

If ICC_SRE_EL2.SRE is 0, the Enable bit behaves as 1 for all purposes other than reading the value of the bit.

DIB, bit [2]

Disable IRQ bypass.

0 IRQ bypass enabled.

1 IRQ bypass disabled.

If EL3 is implemented and GICD_CTLR.DS is 0, this field is a read-only alias of ICC_SRE_EL3.DIB.

If EL3 is implemented and GICD_CTLR.DS is 1, this field is a read-write alias of ICC_SRE_EL3.DIB.

In systems that do not support IRQ bypass, this bit is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

DFB, bit [1]

Disable FIQ bypass.

0 FIQ bypass enabled.

1 FIQ bypass disabled.

If EL3 is implemented and GICD_CTLR.DS is 0, this field is a read-only alias of ICC_SRE_EL3.DFB.

If EL3 is implemented and GICD_CTLR.DS is 1, this field is a read-write alias of ICC_SRE_EL3.DFB.

In systems that do not support FIQ bypass, this bit is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

SRE, bit [0]

System Register Enable.

0 The memory-mapped interface must be used. Access at EL2 to any ICH_* or ICC_* register other than ICC_SRE_EL1 or ICC_SRE_EL2, is trapped to EL2.

1 The System register interface to the ICH_* registers and the EL1 and EL2 ICC_* registers is enabled for EL2.

If software changes this bit from 1 to 0, the results are UNPREDICTABLE.

If an implementation supports only a System register interface to the GIC CPU interface, this bit is RAO/WI.

If EL3 is implemented and ICC_SRE_EL3.SRE==0 this bit is RAZ/WI.

GICv3 implementations that do not require GICv2 compatibility might choose to make this bit RAO/WI, but this is only allowed if ICC_SRE_EL3.SRE is also RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

Accessing the ICC_SRE_EL2:

To access the ICC_SRE_EL2:

```
MRS <Xt>, ICC_SRE_EL2 ; Read ICC_SRE_EL2 into Xt
MSR ICC_SRE_EL2, <Xt> ; Write Xt to ICC_SRE_EL2
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	100	1100	1001	101

8.2.24 ICC_SRE_EL3, Interrupt Controller System Register Enable register (EL3)

The ICC_SRE_EL3 characteristics are:

Purpose

Controls whether the System register interface or the memory-mapped interface to the GIC CPU interface is used for EL3.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	-	RW	RW

This register is always System register accessible.

Traps and Enables

There are no traps or enables affecting this register.

Configurations

AArch64 System register ICC_SRE_EL3 can be mapped to AArch32 System register ICC_MSRE, but this is not architecturally mandated.

Attributes

ICC_SRE_EL3 is a 32-bit register.

Field descriptions

The ICC_SRE_EL3 bit assignments are:

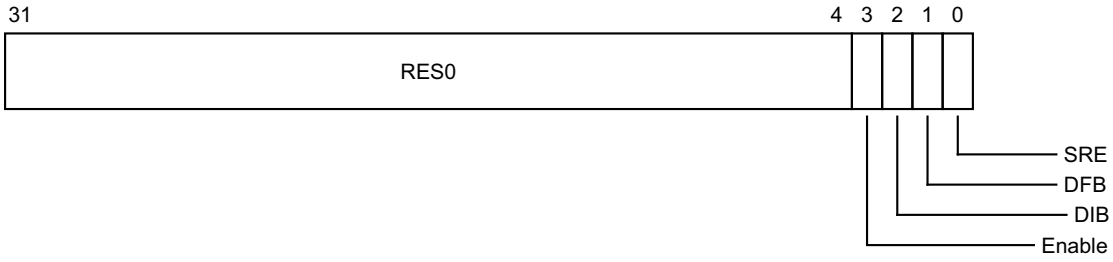

Bits [31:4]

Reserved, RES0.

Enable, bit [3]

Enable. Enables lower Exception level access to ICC_SRE_EL1 and ICC_SRE_EL2.

0 Secure EL1 accesses to Secure ICC_SRE_EL1 trap to EL3.
EL2 accesses to Non-secure ICC_SRE_EL1 and ICC_SRE_EL2 trap to EL3.
Non-secure EL1 accesses to ICC_SRE_EL1 trap to EL3, unless these accesses are trapped to EL2 as a result of ICC_SRE_EL3.Enable == 0.

1 Secure EL1 accesses to Secure ICC_SRE_EL1 do not trap to EL3.
EL2 accesses to Non-secure ICC_SRE_EL1 and ICC_SRE_EL2 do not trap to EL3.
Non-secure EL1 accesses to ICC_SRE_EL1 do not trap to EL3.

If ICC_SRE_EL3.SRE is RAO/WI, an implementation is permitted to make the Enable bit RAO/WI.

If ICC_SRE_EL3.SRE is 0, the Enable bit behaves as 1 for all purposes other than reading the value of the bit.

DIB, bit [2]

Disable IRQ bypass.

0 IRQ bypass enabled.

1 IRQ bypass disabled.

In systems that do not support IRQ bypass, this bit is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

DFB, bit [1]

Disable FIQ bypass.

0 FIQ bypass enabled.

1 FIQ bypass disabled.

In systems that do not support FIQ bypass, this bit is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

SRE, bit [0]

System Register Enable.

0 The memory-mapped interface must be used. Access at EL3 to any ICH_* or ICC_* register other than ICC_SRE_EL1, ICC_SRE_EL2, or ICC_SRE_EL3 is trapped to EL3

1 The System register interface to the ICH_* registers and the EL1, EL2, and EL3 ICC_* registers is enabled for EL3.

If software changes this bit from 1 to 0, the results are UNPREDICTABLE.

GICv3 implementations that do not require GICv2 compatibility might choose to make this bit RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

Accessing the ICC_SRE_EL3:

To access the ICC_SRE_EL3:

```
MRS <Xt>, ICC_SRE_EL3 ; Read ICC_SRE_EL3 into Xt
MSR ICC_SRE_EL3, <Xt> ; Write Xt to ICC_SRE_EL3
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	110	1100	1100	101

8.3 AArch64 System register descriptions of the virtual registers

This section describes each of the virtual AArch64 GIC System registers in register name order. The ICV prefix indicates a virtual GIC CPU interface System register. Each AArch64 System register description contains a reference to the AArch32 register that provides the same functionality.

Unless otherwise stated, the bit assignments for the GIC System registers are the same as those for the equivalent GICC_* and GICV_* memory-mapped registers.

The ICV_* registers are only accessible at Non-secure EL1. Whether an access encoding maps to an ICC_* register or the equivalent ICV_* register is determined by HCR_EL2, see Chapter 5 *Virtual Interrupt Handling and Prioritization*. The equivalent virtual interface memory-mapped registers are described in *The GIC virtual CPU interface register descriptions* on page 8-587.

The encodings for the virtual registers are the same as for the physical registers, see Table 8-21 on page 8-179.

8.3.1 ICV_AP0R<n>_EL1, Interrupt Controller Virtual Active Priorities Group 0 Registers, n = 0 - 3

The ICV_AP0R<n>_EL1 characteristics are:

Purpose

Provides information about virtual Group 0 active priorities.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	-	-	-	-

The ICV_AP0R<n>_EL1 registers are only accessible at Non-secure EL1 when HCR_EL2.FMO is set to 1.

Note

When HCR_EL2.FMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_AP0R<n>_EL1 results in an access to ICC_AP0R<n>_EL1.

Writing to these registers with any value other than the last read value of the register (or `0x00000000` when there are no Group 0 active priorities) might result in UNPREDICTABLE behavior of the virtual interrupt prioritization system, causing:

- Interrupts that should preempt execution to not preempt execution.
- Interrupts that should not preempt execution to preempt execution.

ICV_AP0R1_EL1 is only implemented in implementations that support 6 or more bits of priority. ICV_AP0R2_EL1 and ICV_AP0R3_EL1 are only implemented in implementations that support 7 bits of priority. Unimplemented registers are UNDEFINED.

Writing to the active priority registers in any order other than the following order might result in UNPREDICTABLE behavior of the interrupt prioritization system:

- ICV_AP0R<n>_EL1.
- ICV_AP1R<n>_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_AP0R<n>_EL1 is architecturally mapped to AArch32 System register ICV_AP0R<n>.

Attributes

ICV_AP0R<n>_EL1 is a 32-bit register.

Field descriptions

The ICV_AP0R<n>_EL1 bit assignments are:

IMPLEMENTATION DEFINED, bits [31:0]

IMPLEMENTATION DEFINED.

When this register has an architecturally-defined reset value, this field resets to `0`.

The contents of these registers are IMPLEMENTATION DEFINED with the one architectural requirement that the value `0x00000000` is consistent with no interrupts being active.

Accessing the ICV_AP0R<n>_EL1:

To access the ICV_AP0R<n>_EL1 when HCR_EL2.FMO is set to 1, and executing at Non-secure EL1:

```
MRS <Xt>, ICC_AP0R<n>_EL1 ; Read ICV_AP0R<n>_EL1 into Xt, where n is in the range 0 to 3
MSR ICC_AP0R<n>_EL1, <Xt> ; Write Xt to ICV_AP0R<n>_EL1, where n is in the range 0 to 3
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1000	1:n<1:0>

When HCR_EL2.FMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_AP0R<n>_EL1.

8.3.2 ICV_AP1R<n>_EL1, Interrupt Controller Virtual Active Priorities Group 1 Registers, n = 0 - 3

The ICV_AP1R<n>_EL1 characteristics are:

Purpose

Provides information about virtual Group 1 active priorities.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	-	-	-	-

The ICV_AP1R<n>_EL1 registers are only accessible at Non-secure EL1 when HCR_EL2.IMO == 1.

Note

When HCR_EL2.IMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_AP1R<n>_EL1 results in an access to ICC_AP1R<n>_EL1.

Writing to these registers with any value other than the last read value of the register (or `0x00000000` when there are no Group 1 active priorities) might result in UNPREDICTABLE behavior of the virtual interrupt prioritization system, causing:

- Interrupts that should preempt execution to not preempt execution.
- Interrupts that should not preempt execution to preempt execution.

ICV_AP1R1_EL1 is only implemented in implementations that support 6 or more bits of priority. ICV_AP1R2_EL1 and ICV_AP1R3_EL1 are only implemented in implementations that support 7 bits of priority. Unimplemented registers are UNDEFINED.

Writing to the active priority registers in any order other than the following order might result in UNPREDICTABLE behavior of the interrupt prioritization system:

- ICV_AP0R<n>_EL1.
- ICV_AP1R<n>_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_AP1R<n>_EL1 is architecturally mapped to AArch32 System register ICV_AP1R<n>.

Attributes

ICV_AP1R<n>_EL1 is a 32-bit register.

Field descriptions

The ICV_AP1R<n>_EL1 bit assignments are:

IMPLEMENTATION DEFINED, bits [31:0]

IMPLEMENTATION DEFINED.

When this register has an architecturally-defined reset value, this field resets to `0`.

The contents of these registers are IMPLEMENTATION DEFINED with the one architectural requirement that the value `0x00000000` is consistent with no interrupts being active.

Accessing the ICV_AP1R<n>_EL1:

To access the ICV_AP1R<n>_EL1 when HCR_EL2.IMO is set to 1, and executing at Non-secure EL1:

```
MRS <Xt>, ICC_AP1R<n>_EL1 ; Read ICV_AP1R<n>_EL1 into Xt, where n is in the range 0 to 3
MSR ICC_AP1R<n>_EL1, <Xt> ; Write Xt to ICV_AP1R<n>_EL1, where n is in the range 0 to 3
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1001	0:n<1:0>

When HCR_EL2.IMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_AP1R<n>_EL1.

8.3.3 ICV_BPR0_EL1, Interrupt Controller Virtual Binary Point Register 0

The ICV_BPR0_EL1 characteristics are:

Purpose

Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines virtual Group 0 interrupt preemption.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	-	-	-	-

ICV_BPR0_EL1 is only accessible at Non-secure EL1 when HCR_EL2.FMO is set to 1.

Note

When HCR_EL2.FMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_BPR0_EL1 results in an access to ICC_BPR0_EL1.

The minimum binary point value is derived from the number of implemented priority bits. The number of priority bits is IMPLEMENTATION DEFINED, and reported by ICV_CTLR_EL1.PRIbits.

An attempt to program the binary point field to a value less than the minimum value sets the field to the minimum value. On a reset, the binary point field is set to the minimum supported value.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_BPR0_EL1 is architecturally mapped to AArch32 System register ICV_BPR0.

Attributes

ICV_BPR0_EL1 is a 32-bit register.

Field descriptions

The ICV_BPR0_EL1 bit assignments are:

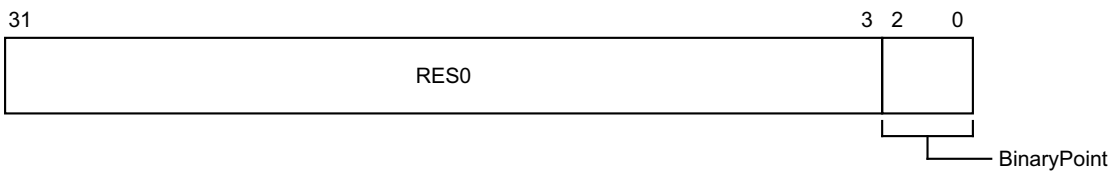

Bits [31:3]

Reserved, RES0.

BinaryPoint, bits [2:0]

The value of this field controls how the 8-bit interrupt priority field is split into a group priority field, that determines interrupt preemption, and a subpriority field. This is done as follows:

Binary point value	Group priority field	Subpriority field	Field with binary point
0	[7:1]	[0]	ggggggg.s
1	[7:2]	[1:0]	gggggg.ss
2	[7:3]	[2:0]	ggggg.sss
3	[7:4]	[3:0]	gggg.ssss
4	[7:5]	[4:0]	ggg.sssss
5	[7:6]	[5:0]	gg.ssssss
6	[7]	[6:0]	g.sssssss
7	No preemption	[7:0]	.ssssssss

Accessing the ICV_BPR0_EL1:

To access the ICV_BPR0_EL1 when HCR_EL2.FMO is set to 1, and executing at Non-secure EL1:

```
MRS <Xt>, ICC_BPR0_EL1 ; Read ICV_BPR0_EL1 into Xt
MSR ICC_BPR0_EL1, <Xt> ; Write Xt to ICV_BPR0_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1000	011

When HCR_EL2.FMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_BPR0_EL1.

8.3.4 ICV_BPR1_EL1, Interrupt Controller Virtual Binary Point Register 1

The ICV_BPR1_EL1 characteristics are:

Purpose

Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines virtual Group 1 interrupt preemption.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	-	-	-	-

ICV_BPR1_EL1 is only accessible at Non-secure EL1 when HCR_EL2.IMO is set to 1.

Note

When HCR_EL2.IMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_BPR1_EL1 results in an access to ICC_BPR1_EL1.

The reset value is IMPLEMENTATION DEFINED, but is equal to the minimum value of ICV_BPR0_EL1 plus one.

An attempt to program the binary point field to a value less than the reset value sets the field to the reset value.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_BPR1_EL1 is architecturally mapped to AArch32 System register ICV_BPR1.

Attributes

ICV_BPR1_EL1 is a 32-bit register.

Field descriptions

The ICV_BPR1_EL1 bit assignments are:

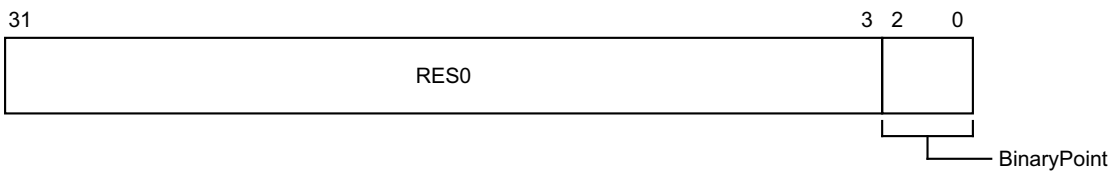

Bits [31:3]

Reserved, RES0.

BinaryPoint, bits [2:0]

If the GIC is configured to use separate binary point fields for virtual Group 0 and virtual Group 1 interrupts, the value of this field controls how the 8-bit interrupt priority field is split into a group priority field, that determines interrupt preemption, and a subpriority field. This is done as follows:

Binary point value	Group priority field	Subpriority field	Field with binary point
0	-	-	-
1	[7:1]	[0]	ggggggg.s
2	[7:2]	[1:0]	gggggg.ss
3	[7:3]	[2:0]	ggggg.sss
4	[7:4]	[3:0]	gggg.ssss
5	[7:5]	[4:0]	ggg.sssss
6	[7:6]	[5:0]	gg.ssssss
7	[7]	[6:0]	g.sssssss

Writing 0 to this field will set this field to its reset value, which is IMPLEMENTATION DEFINED and non-zero.

If ICV_CTLR_EL1.CBPR is set to 1, Non-secure EL1 reads return ICV_BPR0_EL1 + 1 saturated to `0b111`. Non-secure EL1 writes are ignored.

Accessing the ICV_BPR1_EL1:

To access the ICV_BPR1_EL1 when HCR_EL2.IMO is set to 1, and executing at Non-secure EL1:

```
MRS <Xt>, ICC_BPR1_EL1 ; Read ICV_BPR1_EL1 into Xt
MSR ICC_BPR1_EL1, <Xt> ; Write Xt to ICV_BPR1_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	011

When HCR_EL2.IMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_BPR1_EL1.

8.3.5 ICV_CTLR_EL1, Interrupt Controller Virtual Control Register

The ICV_CTLR_EL1 characteristics are:

Purpose

Controls aspects of the behavior of the GIC virtual CPU interface and provides information about the features implemented.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	-	-	-	-

ICV_CTLR_EL1 is only accessible at Non-secure EL1 when HCR_EL2.{FMO, IMO} != {0, 0}.

Note

When HCR_EL2.{FMO, IMO} == {0, 0}, at Non-secure EL1, the instruction encoding used to access ICV_CTLR_EL1 results in an access to ICC_CTLR_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TC==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_CTLR_EL1 is architecturally mapped to AArch32 System register ICV_CTLR.

Attributes

ICV_CTLR_EL1 is a 32-bit register.

Field descriptions

The ICV_CTLR_EL1 bit assignments are:

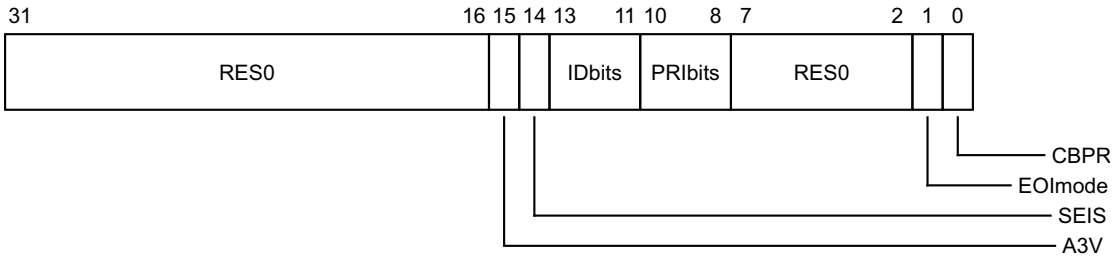

Bits [31:16]

Reserved, RES0.

A3V, bit [15]

Affinity 3 Valid. Read-only and writes are ignored. Possible values are:

0 The virtual CPU interface logic only supports zero values of Affinity 3 in SGI generation System registers.

1 The virtual CPU interface logic supports non-zero values of Affinity 3 in SGI generation System registers.

SEIS, bit [14]

SEI Support. Read-only and writes are ignored. Indicates whether the virtual CPU interface supports local generation of SEIs:

0 The virtual CPU interface logic does not support local generation of SEIs.

1 The virtual CPU interface logic supports local generation of SEIs.

IDbits, bits [13:11]

Identifier bits. Read-only and writes are ignored. The number of virtual interrupt identifier bits supported:

000 16 bits.

001 24 bits.

All other values are reserved.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

PRIbits, bits [10:8]

Priority bits. Read-only and writes are ignored. The number of priority bits implemented, minus one.

An implementation must implement at least 32 levels of physical priority (5 priority bits).

Note

This field always returns the number of priority bits implemented.

The division between group priority and subpriority is defined in the binary point registers ICV_BPR0_EL1 and ICV_BPR1_EL1.

Bits [7:2]

Reserved, RES0.

EOImode, bit [1]

Virtual EOI mode. Controls whether a write to an End of Interrupt register also deactivates the virtual interrupt:

0 ICV_EOIR0_EL1 and ICV_EOIR1_EL1 provide both priority drop and interrupt deactivation functionality. Accesses to ICV_DIR_EL1 are UNPREDICTABLE.

1 ICV_EOIR0_EL1 and ICV_EOIR1_EL1 provide priority drop functionality only. ICV_DIR_EL1 provides interrupt deactivation functionality.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

CBPR, bit [0]

Common Binary Point Register. Controls whether the same register is used for interrupt preemption of both virtual Group 0 and virtual Group 1 interrupts:

0 ICV_BPR0_EL1 determines the preemption group for virtual Group 0 interrupts only.
ICV_BPR1_EL1 determines the preemption group for virtual Group 1 interrupts.

1 ICV_BPR0_EL1 determines the preemption group for both virtual Group 0 and virtual Group 1 interrupts.

Reads of ICV_BPR1_EL1 return ICV_BPR0_EL1 plus one, saturated to 0b111. Writes to ICV_BPR1_EL1 are ignored.

Accessing the ICV_CTLR_EL1:

To access the ICV_CTLR_EL1 when HCR_EL2.{FMO, IMO} != {0, 0}, and executing at Non-secure EL1:

```
MRS <Xt>, ICC_CTLR_EL1 ; Read ICV_CTLR_EL1 into Xt
MSR ICC_CTLR_EL1, <Xt> ; Write Xt to ICV_CTLR_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	100

When HCR_EL2.{FMO, IMO} == {0, 0}, execution of this encoding at Non-secure EL1 results in an access to ICC_CTLR_EL1.

8.3.6 ICV_DIR_EL1, Interrupt Controller Deactivate Virtual Interrupt Register

The ICV_DIR_EL1 characteristics are:

Purpose

When interrupt priority drop is separated from interrupt deactivation, a write to this register deactivates the specified virtual interrupt.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	WO	-	-	-	-

The ICV_DIR_EL1 register is only accessible at Non-secure EL1 in the following cases:

- When HCR_EL2.FMO is set to 1.
- When HCR_EL2.IMO is set to 1.

Note

At Non-secure EL1, the instruction encoding used to access ICV_DIR_EL1 results in an access to ICC_DIR_EL1 when HCR_EL2.{FMO, IMO} == {0, 0}.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure write accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TDIR==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_DIR_EL1 performs the same function as AArch32 System operation ICV_DIR.

Attributes

ICV_DIR_EL1 is a 32-bit register.

Field descriptions

The ICV_DIR_EL1 bit assignments are:

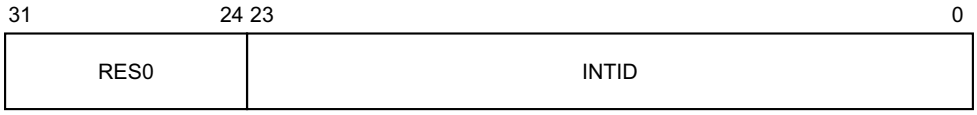

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the virtual interrupt to be deactivated.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICV_CTLR_EL1.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICV_DIR_EL1:

To access the ICV_DIR_EL1:

```
MSR ICC_DIR_EL1, <Xt> ; Write Xt to ICV_DIR_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1011	001

This encoding results in an access to ICV_DIR_EL1 at Non-secure EL1 in the following cases:

- When HCR_EL2.FMO is set to 1, and the INTID field refers to a Group 0 interrupt.
- When HCR_EL2.IMO is set to 1, and the INTID field refers to a Group 1 interrupt.

This encoding results in an access to ICC_DIR_EL1 at Non-secure EL1 in the following cases:

- When HCR_EL2.{FMO, IMO} == {0, 0}.
- When HCR_EL2.FMO is set to 1, and the INTID field does not refer to a Group 0 interrupt.
- When HCR_EL2.IMO is set to 1, and the INTID field does not refer to a Group 1 interrupt.

8.3.7 ICV_EOIR0_EL1, Interrupt Controller Virtual End Of Interrupt Register 0

The ICV_EOIR0_EL1 characteristics are:

Purpose

A PE writes to this register to inform the CPU interface that it has completed the processing of the specified virtual Group 0 interrupt.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	WO	-	-	-	-

ICV_EOIR0_EL1 is only accessible at Non-secure EL1 when HCR_EL2.FMO is set to 1.

Note

When HCR_EL2.FMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_EOIR0_EL1 results in an access to ICC_EOIR0_EL1.

A write to this register must correspond to the most recent valid read from an Interrupt Acknowledge Register and that this identifier was read from ICV_IAR0_EL1 while operating in the same Security state as that in which the write occurs, otherwise the system behavior is UNPREDICTABLE. A valid read is a read that returns a valid interrupt ID that is not a special INTID.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure write accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TALL0==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_EOIR0_EL1 performs the same function as AArch32 System operation ICV_EOIR0.

Attributes

ICV_EOIR0_EL1 is a 32-bit register.

Field descriptions

The ICV_EOIR0_EL1 bit assignments are:

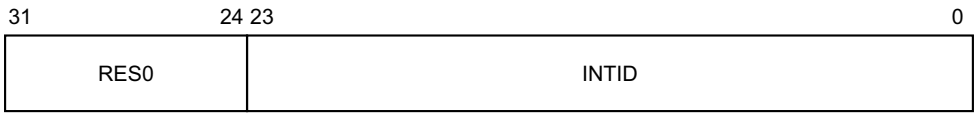

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID from the corresponding ICV_IAR0_EL1 access.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICV_CTLR_EL1.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

If the ICV_CTLR.EOImode bit is 0, a write to this register drops the priority for the virtual interrupt, and also deactivates the virtual interrupt.

If the ICV_CTLR.EOImode bit is 1, a write to this register only drops the priority for the virtual interrupt. Software must write to ICV_DIR_EL1 to deactivate the virtual interrupt.

Accessing the ICV_EOIR0_EL1:

To access the ICV_EOIR0_EL1 when HCR_EL2.FMO is set to 1, and executing at Non-secure EL1:

```
MSR ICC_EOIR0_EL1, <Xt> ; Write Xt to ICV_EOIR0_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1000	001

When HCR_EL2.FMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_EOIR0_EL1.

8.3.8 ICV_EOIR1_EL1, Interrupt Controller Virtual End Of Interrupt Register 1

The ICV_EOIR1_EL1 characteristics are:

Purpose

A PE writes to this register to inform the CPU interface that it has completed the processing of the specified virtual Group 1 interrupt.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	WO	-	-	-	-

ICV_EOIR1_EL1 is only accessible at Non-secure EL1 when HCR_EL2.IMO is set to 1.

Note

When HCR_EL2.IMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_EOIR1_EL1 results in an access to ICC_EOIR1_EL1.

A write to this register must correspond to the most recent valid read from an Interrupt Acknowledge Register and that this identifier was read from ICV_IAR1_EL1 while operating in the same Security state as that in which the write occurs, otherwise the system behavior is UNPREDICTABLE. A valid read is a read that returns a valid interrupt ID, that is not a special INTID.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure write accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TALL1==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_EOIR1_EL1 performs the same function as AArch32 System operation ICV_EOIR1.

Attributes

ICV_EOIR1_EL1 is a 32-bit register.

Field descriptions

The ICV_EOIR1_EL1 bit assignments are:

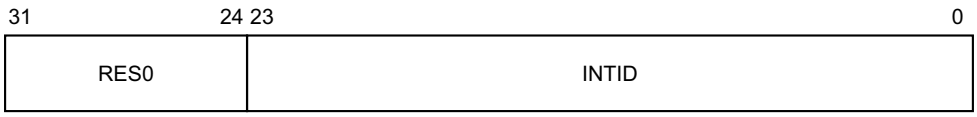

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID from the corresponding ICV_IAR1_EL1 access.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICV_CTLR_EL1.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

If the ICV_CTLR.EOImode bit is 0, a write to this register drops the priority for the virtual interrupt, and also deactivates the virtual interrupt.

If the ICV_CTLR.EOImode bit is 1, a write to this register only drops the priority for the virtual interrupt. Software must write to ICV_DIR_EL1 to deactivate the virtual interrupt.

Accessing the ICV_EOIR1_EL1:

To access the ICV_EOIR1_EL1 when HCR_EL2.IMO is set to 1, and executing at Non-secure EL1:

```
MSR ICC_EOIR1_EL1, <Xt> ; Write Xt to ICV_EOIR1_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	001

When HCR_EL2.IMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_EOIR1_EL1.

8.3.9 ICV_HPPIR0_EL1, Interrupt Controller Virtual Highest Priority Pending Interrupt Register 0

The ICV_HPPIR0_EL1 characteristics are:

Purpose

Indicates the highest priority pending virtual Group 0 interrupt on the virtual CPU interface.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RO	-	-	-	-

ICV_HPPIR0_EL1 is only accessible at Non-secure EL1 when HCR_EL2.FMO is set to 1.

Note

When HCR_EL2.FMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_HPPIR0_EL1 results in an access to ICC_HPPIR0_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure read accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TALL0==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_HPPIR0_EL1 performs the same function as AArch32 System operation ICV_HPPIR0.

Attributes

ICV_HPPIR0_EL1 is a 32-bit register.

Field descriptions

The ICV_HPPIR0_EL1 bit assignments are:

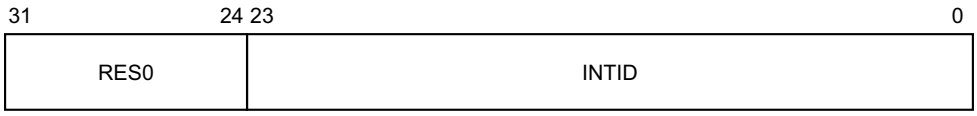

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the highest priority pending virtual interrupt.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. This special INTID can take the value 1023 only. See *Special INTIDs* on page 3-40, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICV_CTLR_EL1.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICV_HPPIR0_EL1:

To access the ICV_HPPIR0_EL1 when HCR_EL2.FMO is set to 1, and executing at Non-secure EL1:

```
MRS <Xt>, ICC_HPPIR0_EL1 ; Read ICV_HPPIR0_EL1 into Xt
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1000	010

When HCR_EL2.FMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_HPPIR0_EL1.

8.3.10 ICV_HPPIR1_EL1, Interrupt Controller Virtual Highest Priority Pending Interrupt Register 1

The ICV_HPPIR1_EL1 characteristics are:

Purpose

Indicates the highest priority pending virtual Group 1 interrupt on the virtual CPU interface.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RO	-	-	-	-

ICV_HPPIR1_EL1 is only accessible at Non-secure EL1 when HCR_EL2.IMO is set to 1.

Note

When HCR_EL2.IMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_HPPIR1_EL1 results in an access to ICC_HPPIR1_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure read accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TALL1==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_HPPIR1_EL1 performs the same function as AArch32 System operation ICV_HPPIR1.

Attributes

ICV_HPPIR1_EL1 is a 32-bit register.

Field descriptions

The ICV_HPPIR1_EL1 bit assignments are:

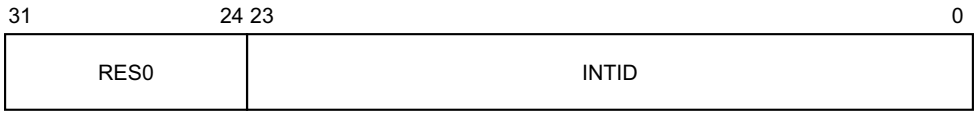

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the highest priority pending virtual interrupt.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. This special INTID can take the value 1023 only. See *Special INTIDs* on page 3-40, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICV_CTLR_EL1.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICV_HPPIR1_EL1:

To access the ICV_HPPIR1_EL1 when HCR_EL2.IMO is set to 1, and executing at Non-secure EL1:

```
MRS <Xt>, ICC_HPPIR1_EL1 ; Read ICV_HPPIR1_EL1 into Xt
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	010

When HCR_EL2.IMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_HPPIR1_EL1.

8.3.11 ICV_IAR0_EL1, Interrupt Controller Virtual Interrupt Acknowledge Register 0

The ICV_IAR0_EL1 characteristics are:

Purpose

The PE reads this register to obtain the INTID of the signaled virtual Group 0 interrupt. This read acts as an acknowledge for the interrupt.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RO	-	-	-	-

ICV_IAR0_EL1 is only accessible at Non-secure EL1 when HCR_EL2.FMO is set to 1.

Note

When HCR_EL2.FMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_IAR0_EL1 results in an access to ICC_IAR0_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure read accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TALL0==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_IAR0_EL1 performs the same function as AArch32 System operation ICV_IAR0.

To allow software to ensure appropriate observability of actions initiated by GIC register accesses, the PE and CPU interface logic must ensure that reads of this register are self-synchronising when interrupts are masked by the PE (that is when PSTATE.{I,F} == {0,0}). This ensures that the effect of activating an interrupt on the signaling of interrupt exceptions is observed when a read of this register is architecturally executed so that no spurious interrupt exception occurs if interrupts are unmasked by an instruction immediately following the read. See *Observability of the effects of accesses to the GIC registers* on page 8-159, for more information.

Attributes

ICV_IAR0_EL1 is a 32-bit register.

Field descriptions

The ICV_IAR0_EL1 bit assignments are:

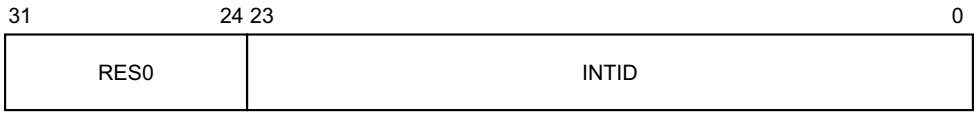

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled virtual interrupt.

This is the INTID of the highest priority pending virtual interrupt, if that interrupt is of sufficient priority for it to be signaled to the PE, and if it can be acknowledged.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. This special INTID can take the value 1023 only. See *Special INTIDs* on page 3-40, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICV_CTLR_EL1.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICV_IAR0_EL1:

To access the ICV_IAR0_EL1 when HCR_EL2.FMO is set to 1, and executing at Non-secure EL1:

```
MRS <Xt>, ICC_IAR0_EL1 ; Read ICV_IAR0_EL1 into Xt
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1000	000

When HCR_EL2.FMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_IAR0_EL1.

8.3.12 ICV_IAR1_EL1, Interrupt Controller Virtual Interrupt Acknowledge Register 1

The ICV_IAR1_EL1 characteristics are:

Purpose

The PE reads this register to obtain the INTID of the signaled virtual Group 1 interrupt. This read acts as an acknowledge for the interrupt.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RO	-	-	-	-

ICV_IAR1_EL1 is only accessible at Non-secure EL1 when HCR_EL2.IMO is set to 1.

Note

When HCR_EL2.IMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_IAR1_EL1 results in an access to ICC_IAR1_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure read accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TALL1==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_IAR1_EL1 performs the same function as AArch32 System operation ICV_IAR1.

To allow software to ensure appropriate observability of actions initiated by GIC register accesses, the PE and CPU interface logic must ensure that reads of this register are self-synchronising when interrupts are masked by the PE (that is when PSTATE.{I,F} == {0,0}. This ensures that the effect of activating an interrupt on the signaling of interrupt exceptions is observed when a read of this register is architecturally executed so that no spurious interrupt exception occurs if interrupts are unmasked by an instruction immediately following the read. See *Observability of the effects of accesses to the GIC registers* on page 8-159, for more information.

Attributes

ICV_IAR1_EL1 is a 32-bit register.

Field descriptions

The ICV_IAR1_EL1 bit assignments are:

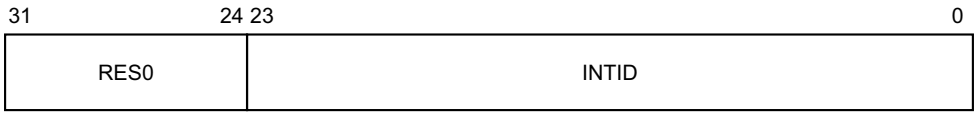

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled virtual interrupt.

This is the INTID of the highest priority pending virtual interrupt, if that interrupt is of sufficient priority for it to be signaled to the PE, and if it can be acknowledged.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. This special INTID can take the value 1023 only. See *Special INTIDs* on page 3-40, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICV_CTLR_EL1.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICV_IAR1_EL1:

To access the ICV_IAR1_EL1 when HCR_EL2.IMO is set to 1, and executing at Non-secure EL1:

```
MRS <Xt>, ICC_IAR1_EL1 ; Read ICV_IAR1_EL1 into Xt
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	000

When HCR_EL2.IMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_IAR1_EL1.

8.3.13 ICV_IGRPEN0_EL1, Interrupt Controller Virtual Interrupt Group 0 Enable register

The ICV_IGRPEN0_EL1 characteristics are:

Purpose

Controls whether virtual Group 0 interrupts are enabled or not.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	-	-	-	-

ICV_IGRPEN0_EL1 is only accessible at Non-secure EL1 when HCR_EL2.FMO is set to 1.

Note

When HCR_EL2.FMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_IGRPEN0_EL1 results in an access to ICC_IGRPEN0_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_IGRPEN0_EL1 is architecturally mapped to AArch32 System register ICV_IGRPEN0.

Attributes

ICV_IGRPEN0_EL1 is a 32-bit register.

Field descriptions

The ICV_IGRPEN0_EL1 bit assignments are:

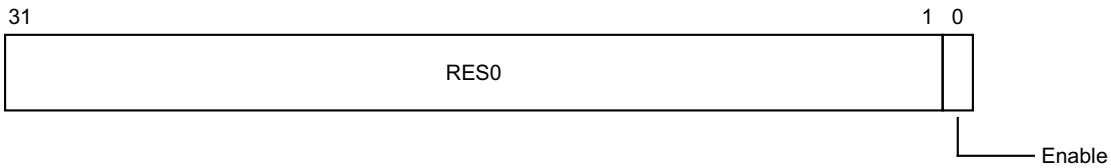

Bits [31:1]

Reserved, RES0.

Enable, bit [0]

Enables virtual Group 0 interrupts.

0 Virtual Group 0 interrupts are disabled.

1 Virtual Group 0 interrupts are enabled.

When this register has an architecturally-defined reset value, this field resets to 0.

Accessing the ICV_IGRPEN0_EL1:

To access the ICV_IGRPEN0_EL1 when HCR_EL2.FMO is set to 1, and executing at Non-secure EL1:

```
MRS <Xt>, ICC_IGRPEN0_EL1 ; Read ICV_IGRPEN0_EL1 into Xt
MSR ICC_IGRPEN0_EL1, <Xt> ; Write Xt to ICV_IGRPEN0_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	110

When HCR_EL2.FMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_IGRPEN0_EL1.

8.3.14 ICV_IGRPEN1_EL1, Interrupt Controller Virtual Interrupt Group 1 Enable register

The ICV_IGRPEN1_EL1 characteristics are:

Purpose

Controls whether virtual Group 1 interrupts are enabled for the current Security state.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	-	-	-	-

ICV_IGRPEN1_EL1 is only accessible at Non-secure EL1 when HCR_EL2.IMO is set to 1.

Note

When HCR_EL2.IMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_IGRPEN1_EL1 results in an access to ICC_IGRPEN1_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_IGRPEN1_EL1 is architecturally mapped to AArch32 System register ICV_IGRPEN1.

Attributes

ICV_IGRPEN1_EL1 is a 32-bit register.

Field descriptions

The ICV_IGRPEN1_EL1 bit assignments are:

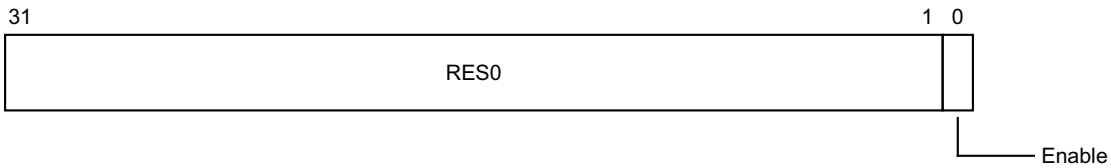

Bits [31:1]

Reserved, RES0.

Enable, bit [0]

Enables virtual Group 1 interrupts.

0 Virtual Group 1 interrupts are disabled.

1 Virtual Group 1 interrupts are enabled.

When this register has an architecturally-defined reset value, this field resets to 0.

Accessing the ICV_IGRPEN1_EL1:

To access the ICV_IGRPEN1_EL1 when HCR_EL2.IMO is set to 1, and executing at Non-secure EL1:

```
MRS <Xt>, ICC_IGRPEN1_EL1 ; Read ICV_IGRPEN1_EL1 into Xt
MSR ICC_IGRPEN1_EL1, <Xt> ; Write Xt to ICV_IGRPEN1_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1100	111

When HCR_EL2.IMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_IGRPEN1_EL1.

8.3.15 ICV_PMR_EL1, Interrupt Controller Virtual Interrupt Priority Mask Register

The ICV_PMR_EL1 characteristics are:

Purpose

Provides a virtual interrupt priority filter. Only virtual interrupts with a higher priority than the value in this register are signaled to the PE.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RW	-	-	-	-

ICV_PMR_EL1 is only accessible at Non-secure EL1 when HCR_EL2.{FMO, IMO} != {0, 0}.

Note

When HCR_EL2.{FMO, IMO} == {0, 0}, at Non-secure EL1, the instruction encoding used to access ICV_PMR_EL1 results in an access to ICC_PMR_EL1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TC==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_PMR_EL1 is architecturally mapped to AArch32 System register ICV_PMR.

To allow software to ensure appropriate observability of actions initiated by GIC register accesses, the PE and CPU interface logic must ensure that writes to this register are self-synchronising. This ensures that no interrupts below the written PMR value will be taken after a write to this register is architecturally executed. See *Observability of the effects of accesses to the GIC registers* on page 8-159, for more information.

Attributes

ICV_PMR_EL1 is a 32-bit register.

Field descriptions

The ICV_PMR_EL1 bit assignments are:

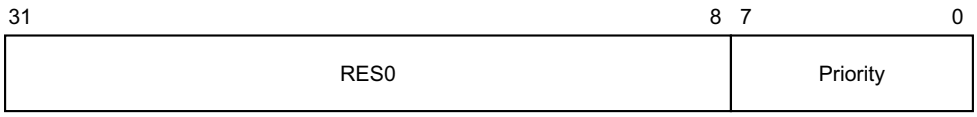

Bits [31:8]

Reserved, RES0.

Priority, bits [7:0]

The priority mask level for the virtual CPU interface. If the priority of a virtual interrupt is higher than the value indicated by this field, the interface signals the virtual interrupt to the PE.

The possible priority field values are as follows:

Implemented priority bits	Possible priority field values	Number of priority levels
[7:0]	0x00-0xFF (0-255), all values	256
[7:1]	0x00-0xFE (0-254), even values only	128
[7:2]	0x00-0xFC (0-252), in steps of 4	64
[7:3]	0x00-0xF8 (0-248), in steps of 8	32
[7:4]	0x00-0xF0 (0-240), in steps of 16	16

Unimplemented priority bits are RAZ/WI.

When this register has an architecturally-defined reset value, this field resets to 0.

Accessing the ICV_PMR_EL1:

To access the ICV_PMR_EL1 when HCR_EL2.{FMO, IMO} != {0, 0}, and executing at Non-secure EL1:

```
MRS <Xt>, ICC_PMR_EL1 ; Read ICV_PMR_EL1 into Xt
MSR ICC_PMR_EL1, <Xt> ; Write Xt to ICV_PMR_EL1
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	0100	0110	000

When HCR_EL2.{FMO, IMO} == {0, 0}, execution of this encoding at Non-secure EL1 results in an access to ICC_PMR_EL1.

8.3.16 ICV_RPR_EL1, Interrupt Controller Virtual Running Priority Register

The ICV_RPR_EL1 characteristics are:

Purpose

Indicates the Running priority of the virtual CPU interface.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	RO	-	-	-	-

ICV_RPR_EL1 is only accessible at Non-secure EL1 when HCR_EL2.{FMO, IMO} != {0, 0}.

Note

When HCR_EL2.{FMO, IMO} == {0, 0}, at Non-secure EL1, the instruction encoding used to access ICV_RPR_EL1 results in an access to ICC_RPR_EL1.

If there are no active interrupts on the virtual CPU interface, or all active interrupts have undergone a priority drop, the value returned is the Idle priority.

Software cannot determine the number of implemented priority bits from a read of this register.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL1.SRE==0, Non-secure read accesses to this register from EL1 are trapped to EL1.

If ICH_HCR_EL2.TC==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

Configurations

AArch64 System register ICV_RPR_EL1 performs the same function as AArch32 System operation ICV_RPR.

Attributes

ICV_RPR_EL1 is a 32-bit register.

Field descriptions

The ICV_RPR_EL1 bit assignments are:

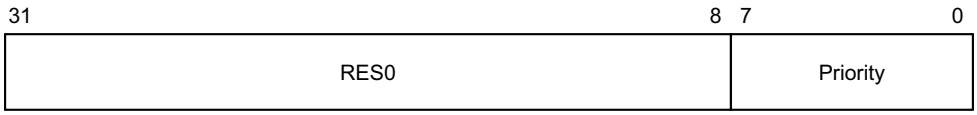

Bits [31:8]

Reserved, RES0.

Priority, bits [7:0]

The current running priority on the virtual CPU interface. This is the group priority of the current active virtual interrupt.

If there are no active interrupts on the virtual CPU interface, or all active interrupts have undergone a priority drop, the value returned is the Idle priority.

The priority returned is the group priority as if the BPR for the current Exception level and Security state was set to the minimum value of BPR for the number of implemented priority bits.

Note

If 8 bits of priority are implemented the group priority is bits[7:1] of the priority.

Accessing the ICV_RPR_EL1:

To access the ICV_RPR_EL1 when HCR_EL2.{FMO, IMO} != {0, 0}, and executing at Non-secure EL1:

```
MRS <Xt>, ICC_RPR_EL1 ; Read ICV_RPR_EL1 into Xt
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	000	1100	1011	011

When HCR_EL2.{FMO, IMO} == {0, 0}, execution of this encoding at Non-secure EL1 results in an access to ICC_RPR_EL1.

8.4 AArch64 virtualization control System registers

This section describes each of the virtualization control AArch64 GIC System registers in register name order. The ICH prefix indicates a virtual interface control System register. Each AArch64 System register description contains a reference to the AArch32 register that provides the same functionality.

Unless otherwise stated, the bit assignments for the GIC System registers are the same as those for the equivalent GICH_* memory-mapped registers. See *The GIC virtual interface control register descriptions* on page 8-619.

Table 8-22 shows the encodings for the AArch64 virtualization control System registers.

Table 8-22 Encodings for AArch64 virtualization control System registers

Register	Width (bits)	Access instruction encoding					Notes
		Op0	Op1	CRn	CRm	Op2	
ICH_AP0R<n>_EL2	32	3	4	12	8	0-3	RW, <n>=Op2.
ICH_AP1R<n>_EL2	32				9	0-3	RW, <n>=Op2.
ICH_HCR_EL2	32				11	0	RW
ICH_VTR_EL2	32					1	RO
ICH_MISR_EL2	32					2	RO
ICH_EISR_EL2	32					3	RO
ICH_ELRSR_EL2	32					5	RO
ICH_VMCR_EL2	32					7	RW
ICH_LR<n>_EL2	64				12, 13	0-7	RW: • For CRm==12, <n>=Op2. • For CRm==13, <n>=Op2+8.

8.4.1 ICH_AP0R<n>_EL2, Interrupt Controller Hyp Active Priorities Group 0 Registers, n = 0 - 3

The ICH_AP0R<n>_EL2 characteristics are:

Purpose

Provides information about Group 0 virtual active priorities for EL2.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RW	RW	RW

ICH_AP0R1_EL2 is only implemented in implementations that support 6 or more bits of priority. ICH_AP0R2_EL2 and ICH_AP0R3_EL2 are only implemented in implementations that support 7 bits of priority. Unimplemented registers are UNDEFINED.

Writing to these registers with any value other than the last read value of the register (or `0x00000000` for a newly set up virtual machine) can result in UNPREDICTABLE behavior of the virtual interrupt prioritization system allowing either:

- Virtual interrupts that should preempt execution to not preempt execution.
- Interrupts that should not preempt execution to preempt execution at Non-secure EL1 or EL0.

Writing to the active priority registers in any order other than the following order will result in UNPREDICTABLE behavior:

- ICH_AP0R<n>_EL2.
- ICH_AP1R<n>_EL2.

Having the bit corresponding to a priority set in both ICH_AP0R<n>_EL2 and ICH_AP1R<n>_EL2 can result in UNPREDICTABLE behavior of the interrupt prioritization system for virtual interrupts.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL2.SRE==0, accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

Configurations

AArch64 System register ICH_AP0R<n>_EL2 is architecturally mapped to AArch32 System register ICH_AP0R<n>.

Attributes

ICH_AP0R<n>_EL2 is a 32-bit register.

Field descriptions

The ICH_AP0R<n>_EL2 bit assignments are:

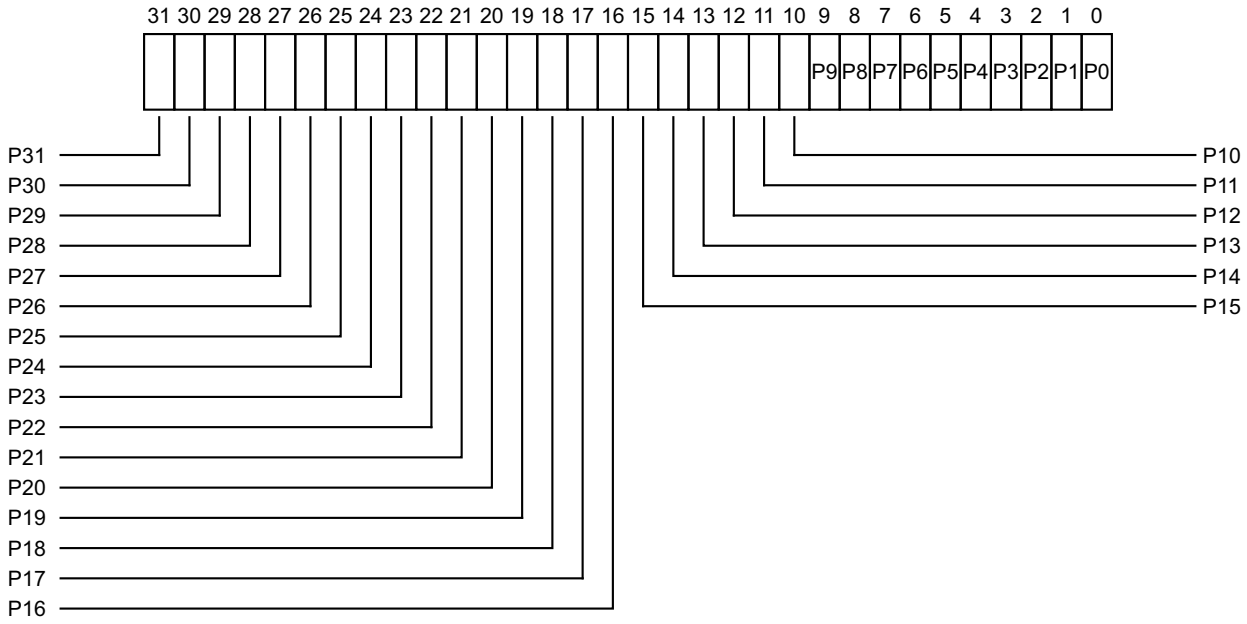

P<x>, bit [x], for x = 0 to 31

Provides the access to the virtual active priorities for Group 0 interrupts. Possible values of each bit are:

`0` There is no Group 0 interrupt active with this priority level, or all active Group 0 interrupts with this priority level have undergone priority-drop.

`1` There is a Group 0 interrupt active with this priority level which has not undergone priority drop.

The correspondence between priority levels and bits depends on the number of bits of priority that are implemented.

If 5 bits of priority are implemented (bits [7:3] of priority), then there are 32 priority levels, and the active state of these priority levels are held in ICH_AP0R0_EL2 in the bits corresponding to Priority[7:3].

If 6 bits of priority are implemented (bits [7:2] of priority), then there are 64 priority levels, and:

- The active state of priority levels 0 - 124 are held in ICH_AP0R0_EL2 in the bits corresponding to 0:Priority[6:2].
- The active state of priority levels 128 - 252 are held in ICH_AP0R1_EL2 in the bits corresponding to 1:Priority[6:2].

If 7 bits of priority are implemented (bits [7:1] of priority), then there are 128 priority levels, and:

- The active state of priority levels 0 - 62 are held in ICH_AP0R0_EL2 in the bits corresponding to 00:Priority[5:1].
- The active state of priority levels 64 - 126 are held in ICH_AP0R1_EL2 in the bits corresponding to 01:Priority[5:1].
- The active state of priority levels 128 - 190 are held in ICH_AP0R2_EL2 in the bits corresponding to 10:Priority[5:1].
- The active state of priority levels 192 - 254 are held in ICH_AP0R3_EL2 in the bits corresponding to 11:Priority[5:1].

Note

Having the bit corresponding to a priority set to 1 in both ICH_AP0R<n>_EL2 and ICH_AP1R<n>_EL2 might result in UNPREDICTABLE behavior of the interrupt prioritization system for virtual interrupts.

When this register has an architecturally-defined reset value, this field resets to `0`.

Software must ensure that ICH_AP0R<n>_EL2 is zero for VMs that use memory-mapped access to the GIC otherwise the effects are UNPREDICTABLE.

The active priorities for Group 0 and Group 1 interrupts for memory mapped guests (i.e. when non-secure EL1 is not using system registers) are held in ICH_AP1R<n>_EL2 and virtual non-secure EL1 (i.e. guest) reads and writes to GICV_APR access ICH_AP1R<n>_EL2. This means this register is inaccessible to memory mapped guests. It is recommended that EL2 software ensures this register is written to zero for memory mapped guests.

Accessing the ICH_AP0R<n>_EL2:

To access the ICH_AP0R<n>_EL2:

```
MRS <Xt>, ICH_AP0R<n>_EL2 ; Read ICH_AP0R<n>_EL2 into Xt, where n is in the range 0 to 3
MSR ICH_AP0R<n>_EL2, <Xt> ; Write Xt to ICH_AP0R<n>_EL2, where n is in the range 0 to 3
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	100	1100	1000	0:n<1:0>

8.4.2 ICH_AP1R<n>_EL2, Interrupt Controller Hyp Active Priorities Group 1 Registers, n = 0 - 3

The ICH_AP1R<n>_EL2 characteristics are:

Purpose

Provides information about Group 1 virtual active priorities for EL2.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RW	RW	RW

ICH_AP1R1_EL2 is only implemented in implementations that support 6 or more bits of priority. ICH_AP1R2_EL2 and ICH_AP1R3_EL2 are only implemented in implementations that support 7 bits of priority. Unimplemented registers are UNDEFINED.

Writing to these registers with any value other than the last read value of the register (or `0x00000000` for a newly set up virtual machine) can result in UNPREDICTABLE behavior of the virtual interrupt prioritization system allowing either:

- Virtual interrupts that should preempt execution to not preempt execution.
- Interrupts that should not preempt execution to preempt execution at Non-secure EL1 or EL0.

Writing to the active priority registers in any order other than the following order will result in UNPREDICTABLE behavior:

- ICH_AP0R<n>_EL2.
- ICH_AP1R<n>_EL2.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL2.SRE==0, accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

Configurations

AArch64 System register ICH_AP1R<n>_EL2 is architecturally mapped to AArch32 System register ICH_AP1R<n>.

Attributes

ICH_AP1R<n>_EL2 is a 32-bit register.

Field descriptions

The ICH_AP1R<n>_EL2 bit assignments are:

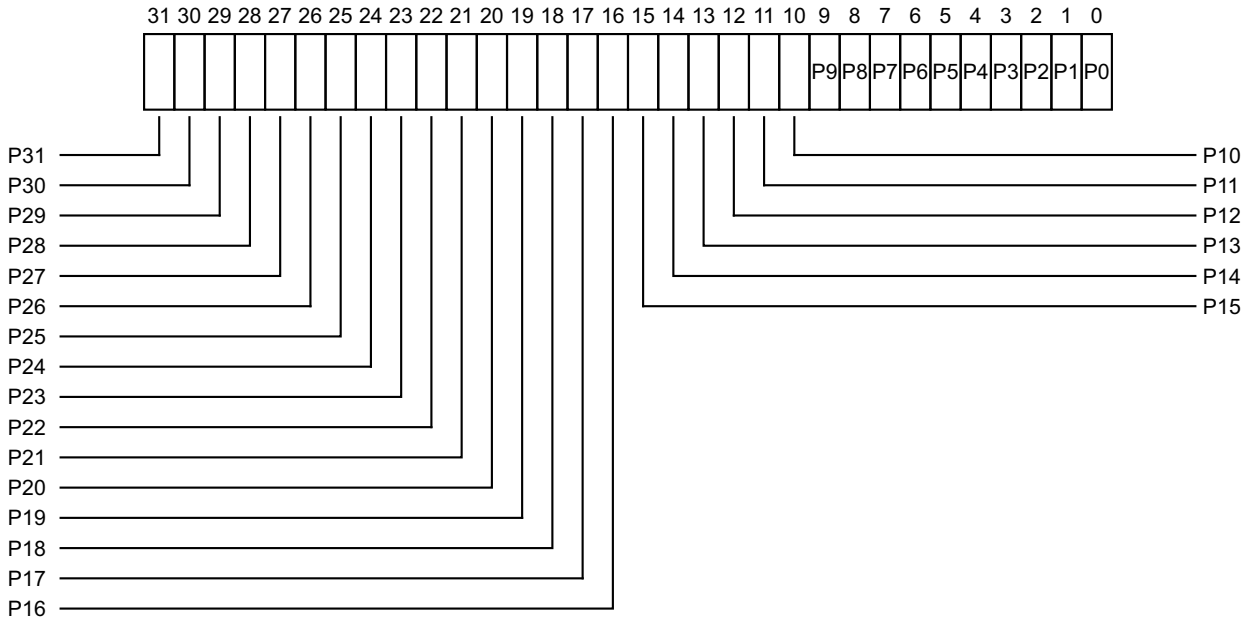

P<x>, bit [x], for x = 0 to 31

Group 1 interrupt active priorities. Possible values of each bit are:

`0` There is no Group 1 interrupt active with this priority level, or all active Group 1 interrupts with this priority level have undergone priority-drop.

`1` There is a Group 1 interrupt active with this priority level which has not undergone priority drop.

The correspondence between priority levels and bits depends on the number of bits of priority that are implemented.

If 5 bits of priority are implemented (bits [7:3] of priority), then there are 32 priority levels, and the active state of these priority levels are held in ICH_AP1R0_EL2 in the bits corresponding to Priority[7:3].

If 6 bits of priority are implemented (bits [7:2] of priority), then there are 64 priority levels, and:

- The active state of priority levels 0 - 124 are held in ICH_AP1R0_EL2 in the bits corresponding to 0:Priority[6:2].
- The active state of priority levels 128 - 252 are held in ICH_AP1R1_EL2 in the bits corresponding to 1:Priority[6:2].

If 7 bits of priority are implemented (bits [7:1] of priority), then there are 128 priority levels, and:

- The active state of priority levels 0 - 62 are held in ICH_AP1R0_EL2 in the bits corresponding to 00:Priority[5:1].
- The active state of priority levels 64 - 126 are held in ICH_AP1R1_EL2 in the bits corresponding to 01:Priority[5:1].
- The active state of priority levels 128 - 190 are held in ICH_AP1R2_EL2 in the bits corresponding to 10:Priority[5:1].
- The active state of priority levels 192 - 254 are held in ICH_AP1R3_EL2 in the bits corresponding to 11:Priority[5:1].

Note

Having the bit corresponding to a priority set to 1 in both ICH_AP0R<n>_EL2 and ICH_AP1R<n>_EL2 might result in UNPREDICTABLE behavior of the interrupt prioritization system for virtual interrupts.

When this register has an architecturally-defined reset value, this field resets to 0.

Always used for memory mapped guests (i.e. when non-secure EL1 is not using system registers) regardless of the group of the virtual interrupt. Virtual non-secure EL1 (i.e. guest) reads and writes to GICV_APR<n> access ICH_AP1R<n>_EL2.

Accessing the ICH_AP1R<n>_EL2:

To access the ICH_AP1R<n>_EL2:

```
MRS <Xt>, ICH_AP1R<n>_EL2 ; Read ICH_AP1R<n>_EL2 into Xt, where n is in the range 0 to 3
MSR ICH_AP1R<n>_EL2, <Xt> ; Write Xt to ICH_AP1R<n>_EL2, where n is in the range 0 to 3
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	100	1100	1001	0:n<1:0>

8.4.3 ICH_EISR_EL2, Interrupt Controller End of Interrupt Status Register

The ICH_EISR_EL2 characteristics are:

Purpose

Indicates which List registers have outstanding EOI maintenance interrupts.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RO	RO	RO

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL2.SRE==0, read accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, read accesses to this register from EL3 are trapped to EL3.

Configurations

AArch64 System register ICH_EISR_EL2 is architecturally mapped to AArch32 System register ICH_EISR.

Attributes

ICH_EISR_EL2 is a 32-bit register.

Field descriptions

The ICH_EISR_EL2 bit assignments are:

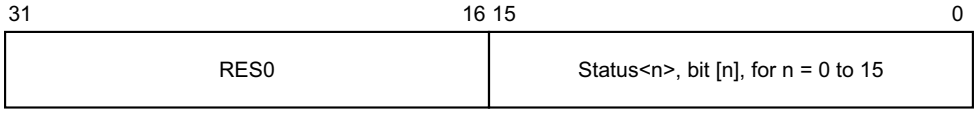

Bits [31:16]

Reserved, RES0.

Status<n>, bit [n], for n = 0 to 15

EOI maintenance interrupt status bit for List register <n>:

`0` List register <n>, ICH_LR<n>_EL2, does not have an EOI maintenance interrupt.

`1` List register <n>, ICH_LR<n>_EL2, has an EOI maintenance interrupt that has not been handled.

For any ICH_LR<n>_EL2, the corresponding status bit is set to 1 if all of the following are true:

- ICH_LR<n>_EL2.State is `0b00`.
- ICH_LR<n>_EL2.HW is 0.
- ICH_LR<n>_EL2.EOI (bit [41]) is 1, indicating that when the interrupt corresponding to that List register is deactivated, a maintenance interrupt is asserted.

Otherwise the status bit takes the value 0.

Accessing the ICH_EISR_EL2:

To access the ICH_EISR_EL2:

```
MRS <Xt>, ICH_EISR_EL2 ; Read ICH_EISR_EL2 into Xt
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	100	1100	1011	011

8.4.4 ICH_ELRSR_EL2, Interrupt Controller Empty List Register Status Register

The ICH_ELRSR_EL2 characteristics are:

Purpose

These registers can be used to locate a usable List register when the hypervisor is delivering an interrupt to a VM.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RO	RO	RO

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL2.SRE==0, read accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, read accesses to this register from EL3 are trapped to EL3.

Configurations

AArch64 System register ICH_ELRSR_EL2 is architecturally mapped to AArch32 System register ICH_ELRSR.

Attributes

ICH_ELRSR_EL2 is a 32-bit register.

Field descriptions

The ICH_ELRSR_EL2 bit assignments are:

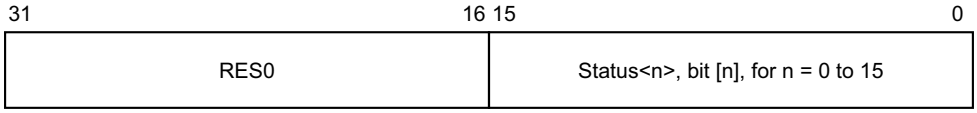

Bits [31:16]

Reserved, RES0.

Status<n>, bit [n], for n = 0 to 15

Status bit for List register <n>, ICH_LR<n>_EL2:

`0` List register ICH_LR<n>_EL2, if implemented, contains a valid interrupt. Using this List register can result in overwriting a valid interrupt.

`1` List register ICH_LR<n>_EL2 does not contain a valid interrupt. The List register is empty and can be used without overwriting a valid interrupt or losing an EOI maintenance interrupt.

For any List register <n>, the corresponding status bit is set to 1 if ICH_LR<n>_EL2.State is `0b00` and either ICH_LR<n>_EL2.HW is 1 or ICH_LR<n>_EL2.EOI (bit [41]) is 0.

Otherwise the status bit takes the value 0.

Accessing the ICH_ELRSR_EL2:

To access the ICH_ELRSR_EL2:

```
MRS <Xt>, ICH_ELRSR_EL2 ; Read ICH_ELRSR_EL2 into Xt
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	100	1100	1011	101

8.4.5 ICH_HCR_EL2, Interrupt Controller Hyp Control Register

The ICH_HCR_EL2 characteristics are:

Purpose

Controls the environment for VMs.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RW	RW	RW

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL2.SRE==0, accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

Configurations

AArch64 System register ICH_HCR_EL2 is architecturally mapped to AArch32 System register ICH_HCR.

Attributes

ICH_HCR_EL2 is a 32-bit register.

Field descriptions

The ICH_HCR_EL2 bit assignments are:

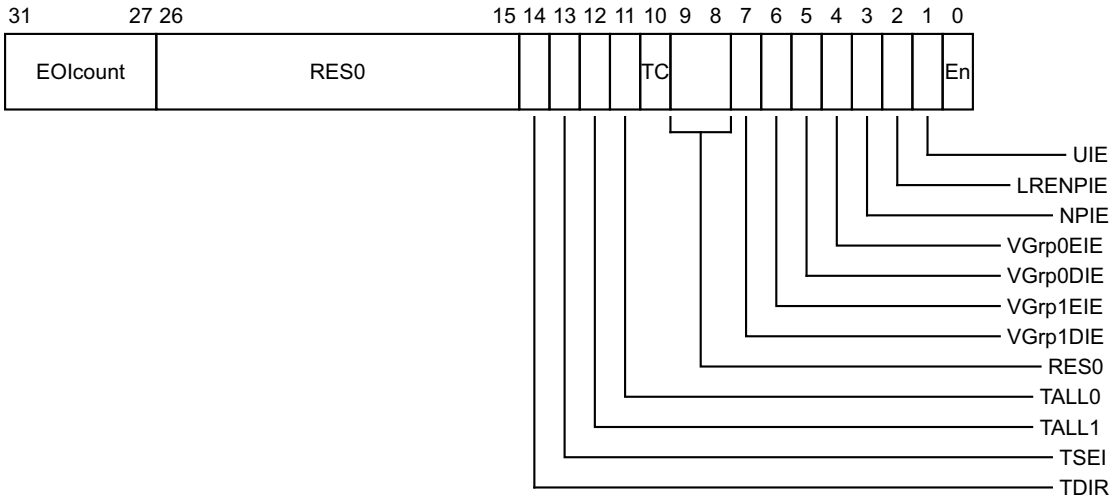

EOIcount, bits [31:27]

This field is incremented whenever a successful write to a virtual EOIR or DIR register would have resulted in a virtual interrupt deactivation. That is:

- A virtual write to EOIR with a valid interrupt identifier that is not in the LPI range (i.e. < 8192) when EOI mode is zero and no List Register was found, or
- A virtual write to DIR with a valid interrupt identifier that is not in the LPI range (i.e. < 8192) when EOI mode is one and no List Register was found

This allows software to manage more active interrupts than there are implemented List Registers.

It is CONSTRAINED UNPREDICTABLE whether a virtual write to EOIR that does not clear a bit in the Active Priorities registers (ICH_AP0R<n>_EL2/ICH_AP1R<n>_EL2) increments EOIcount. Permitted behaviors are:

- Increment EOIcount.
- Leave EOIcount unchanged.

When this register has an architecturally-defined reset value, this field resets to `0`.

Bits [26:15]

Reserved, RES0.

TDIR, bit [14]

Trap Non-secure EL1 writes to ICC_DIR_EL1 and ICV_DIR_EL1.

`0` Non-secure EL1 writes of ICC_DIR_EL1 and ICV_DIR_EL1 are not trapped to EL2, unless trapped by other mechanisms.

`1` Non-secure EL1 writes of ICC_DIR_EL1 and ICV_DIR_EL1 are trapped to EL2.

Support for this bit is OPTIONAL, with support indicated by ICH_VTR_EL2.

If the implementation does not support this trap, this bit is RES0.

ARM deprecates not including this trap bit.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

TSEI, bit [13]

Trap all locally generated SEIs. This bit allows the hypervisor to intercept locally generated SEIs that would otherwise be taken at Non-secure EL1.

`0` Locally generated SEIs do not cause a trap to EL2.

`1` Locally generated SEIs trap to EL2.

If ICH_VTR_EL2.SEIS is 0, this bit is RES0.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

TALL1, bit [12]

Trap all Non-secure EL1 accesses to ICC_* and ICV_* System registers for Group 1 interrupts to EL2.

`0` Non-Secure EL1 accesses to ICC_* and ICV_* registers for Group 1 interrupts proceed as normal.

`1` Non-secure EL1 accesses to ICC_* and ICV_* registers for Group 1 interrupts trap to EL2.

When this register has an architecturally-defined reset value, this field resets to `0`.

TALL0, bit [11]

Trap all Non-secure EL1 accesses to ICC_* and ICV_* System registers for Group 0 interrupts to EL2.

0 Non-Secure EL1 accesses to ICC_* and ICV_* registers for Group 0 interrupts proceed as normal.

1 Non-secure EL1 accesses to ICC_* and ICV_* registers for Group 0 interrupts trap to EL2.

When this register has an architecturally-defined reset value, this field resets to 0.

TC, bit [10]

Trap all Non-secure EL1 accesses to System registers that are common to Group 0 and Group 1 to EL2.

0 Non-secure EL1 accesses to common registers proceed as normal.

1 Non-secure EL1 accesses to common registers trap to EL2.

This affects accesses to ICC_SGI0R_EL1, ICC_SGI1R_EL1, ICC_ASGI1R_EL1, ICC_CTLR_EL1, ICC_DIR_EL1, ICC_PMR_EL1, ICC_RPR_EL1, ICV_CTLR_EL1, ICV_DIR_EL1, ICV_PMR_EL1, and ICV_RPR_EL1.

When this register has an architecturally-defined reset value, this field resets to 0.

Bits [9:8]

Reserved, RES0.

VGrp1DIE, bit [7]

VM Group 1 Disabled Interrupt Enable. Enables the signaling of a maintenance interrupt while signaling of Group 1 interrupts from the virtual CPU interface to the connected vPE is disabled:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled when ICH_VMCR_EL2.VENG1 is 0.

When this register has an architecturally-defined reset value, this field resets to 0.

VGrp1EIE, bit [6]

VM Group 1 Enabled Interrupt Enable. Enables the signaling of a maintenance interrupt while signaling of Group 1 interrupts from the virtual CPU interface to the connected vPE is enabled:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled when ICH_VMCR_EL2.VENG1 is 1.

When this register has an architecturally-defined reset value, this field resets to 0.

VGrp0DIE, bit [5]

VM Group 0 Disabled Interrupt Enable. Enables the signaling of a maintenance interrupt while signaling of Group 0 interrupts from the virtual CPU interface to the connected vPE is disabled:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled when ICH_VMCR_EL2.VENG0 is 0.

When this register has an architecturally-defined reset value, this field resets to 0.

VGrp0EIE, bit [4]

VM Group 0 Enabled Interrupt Enable. Enables the signaling of a maintenance interrupt while signaling of Group 0 interrupts from the virtual CPU interface to the connected vPE is enabled:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled when ICH_VMCR_EL2.VENG0 is 1.

When this register has an architecturally-defined reset value, this field resets to 0.

NPIE, bit [3]

No Pending Interrupt Enable. Enables the signaling of a maintenance interrupt while no pending interrupts are present in the List registers:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled while the List registers contain no interrupts in the pending state.

When this register has an architecturally-defined reset value, this field resets to 0.

LRENPIE, bit [2]

List Register Entry Not Present Interrupt Enable. Enables the signaling of a maintenance interrupt while the virtual CPU interface does not have a corresponding valid List register entry for an EOI request:

0 Maintenance interrupt disabled.

1 Maintenance interrupt is asserted while the EOIcount field is not 0.

When this register has an architecturally-defined reset value, this field resets to 0.

UIE, bit [1]

Underflow Interrupt Enable. Enables the signaling of a maintenance interrupt when the List registers are empty, or hold only one valid entry:

0 Maintenance interrupt disabled.

1 Maintenance interrupt is asserted if none, or only one, of the List register entries is marked as a valid interrupt.

When this register has an architecturally-defined reset value, this field resets to 0.

En, bit [0]

Enable. Global enable bit for the virtual CPU interface:

0 Virtual CPU interface operation disabled.

1 Virtual CPU interface operation enabled.

When this field is set to 0:

- The virtual CPU interface does not signal any maintenance interrupts.
- The virtual CPU interface does not signal any virtual interrupts.
- A read of ICV_IAR0_EL1, ICV_IAR1_EL1, GICV_IAR or GICV_AIAR returns a spurious interrupt ID.

When this register has an architecturally-defined reset value, this field resets to 0.

Accessing the ICH_HCR_EL2:

To access the ICH_HCR_EL2:

```
MRS <Xt>, ICH_HCR_EL2 ; Read ICH_HCR_EL2 into Xt
MSR ICH_HCR_EL2, <Xt> ; Write Xt to ICH_HCR_EL2
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	100	1100	1011	000

8.4.6 ICH_LR<n>_EL2, Interrupt Controller List Registers, n = 0 - 15

The ICH_LR<n>_EL2 characteristics are:

Purpose

Provides interrupt context information for the virtual CPU interface.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RW	RW	RW

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL2.SRE==0, accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

Configurations

AArch64 System register ICH_LR<n>_EL2[31:0] is architecturally mapped to AArch32 System register ICH_LR<n>.

AArch64 System register ICH_LR<n>_EL2[63:32] is architecturally mapped to AArch32 System register ICH_LRC<n>.

Attributes

ICH_LR<n>_EL2 is a 64-bit register.

Field descriptions

The ICH_LR<n>_EL2 bit assignments are:

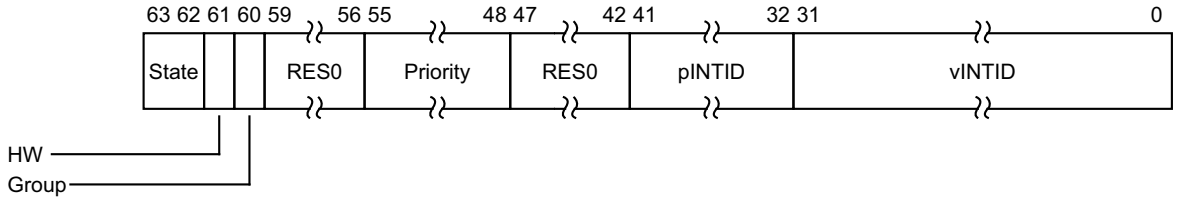

State, bits [63:62]

The state of the interrupt:

`00` Inactive

`01` Pending

`10` Active

`11` Pending and active.

The GIC updates these state bits as virtual interrupts proceed through the interrupt life cycle. Entries in the inactive state are ignored, except for the purpose of generating virtual maintenance interrupts.

For hardware interrupts, the pending and active state is held in the physical Distributor rather than the virtual CPU interface. A hypervisor must only use the pending and active state for software originated interrupts, which are typically associated with virtual devices, or SGIs.

When this register has an architecturally-defined reset value, this field resets to `0`.

HW, bit [61]

Indicates whether this virtual interrupt maps directly to a hardware interrupt, meaning that it corresponds to a physical interrupt. Deactivation of the virtual interrupt also causes the deactivation of the physical interrupt with the ID that the pINTID field indicates.

`0` The interrupt is triggered entirely by software. No notification is sent to the Distributor when the virtual interrupt is deactivated.

`1` The interrupt maps directly to a hardware interrupt. A deactivate interrupt request is sent to the Distributor when the virtual interrupt is deactivated, using the pINTID field from this register to indicate the physical interrupt ID.

If ICH_VMCR_EL2.VEOIM is 0, this request corresponds to a write to ICC_EOIR0_EL1 or ICC_EOIR1_EL1. Otherwise, it corresponds to a write to ICC_DIR_EL1.

When this register has an architecturally-defined reset value, this field resets to `0`.

Group, bit [60]

Indicates the group for this virtual interrupt.

`0` This is a Group 0 virtual interrupt. ICH_VMCR_EL2.VFIQEn determines whether it is signaled as a virtual IRQ or as a virtual FIQ, and ICH_VMCR_EL2.VENG0 enables signaling of this interrupt to the virtual machine.

`1` This is a Group 1 virtual interrupt, signaled as a virtual IRQ. ICH_VMCR_EL2.VENG1 enables the signaling of this interrupt to the virtual machine.

If ICH_VMCR_EL2.VCBPR is 0, then ICC_BPR1_EL1 determines if a pending Group 1 interrupt has sufficient priority to preempt current execution. Otherwise, ICH_LR<n>_EL2 determines preemption.

When this register has an architecturally-defined reset value, this field resets to `0`.

Bits [59:56]

Reserved, RES0.

Priority, bits [55:48]

The priority of this interrupt.

It is IMPLEMENTATION DEFINED how many bits of priority are implemented, though at least five bits must be implemented. Unimplemented bits are RES0 and start from bit [48] up to bit [50]. The number of implemented bits can be discovered from ICH_VTR_EL2.PRIbits.

When this register has an architecturally-defined reset value, this field resets to `0`.

Bits [47:42]

Reserved, RES0.

pINTID, bits [41:32]

Physical INTID, for hardware interrupts.

When the HW bit is 0 (there is no corresponding physical interrupt), this field has the following meaning:

Bit[41] EOI. If this bit is 1, then when the interrupt identified by vINTID is deactivated, a maintenance interrupt is asserted.

Bits[40:32] Reserved, RES0.

When the HW bit is 1 (there is a corresponding physical interrupt):

- This field indicates the physical INTID. This field is only required to implement enough bits to hold a valid value for the implemented INTID size. Any unused higher order bits are RES0.
- If the value of pINTID is 0-15 or 1020-1023, behavior is UNPREDICTABLE. If the value of pINTID is 16-31, this field applies to the PPI associated with this same physical PE ID as the virtual CPU interface requesting the deactivation.

A hardware physical identifier is only required in List Registers for interrupts that require deactivation. This means only 10 bits of Physical INTID are required, regardless of the number specified by ICC_CTLR_EL1.IDbits.

When this register has an architecturally-defined reset value, this field resets to `0`.

vINTID, bits [31:0]

Virtual INTID of the interrupt.

Behavior is UNPREDICTABLE if two or more List Registers specify the same vINTID when:

- ICH_LR<n>_EL2.State == 01.
- ICH_LR<n>_EL2.State == 10.
- ICH_LR<n>_EL2.State == 11.

It is IMPLEMENTATION DEFINED how many bits are implemented, though at least 16 bits must be implemented. Unimplemented bits are RES0. The number of implemented bits can be discovered from ICH_VTR_EL2.IDbits.

Note

When a VM is using memory-mapped access to the GIC, software must ensure that the correct source CPU ID is provided in bits[12:10].

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the ICH_LR<n>_EL2:

To access the ICH_LR<n>_EL2:

```
MRS <Xt>, ICH_LR<n>_EL2 ; Read ICH_LR<n>_EL2 into Xt, where n is in the range 0 to 15
MSR ICH_LR<n>_EL2, <Xt> ; Write Xt to ICH_LR<n>_EL2, where n is in the range 0 to 15
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	100	1100	110:n<3>	n<2:0>

8.4.7 ICH_MISR_EL2, Interrupt Controller Maintenance Interrupt State Register

The ICH_MISR_EL2 characteristics are:

Purpose

Indicates which maintenance interrupts are asserted.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RO	RO	RO

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL2.SRE==0, read accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, read accesses to this register from EL3 are trapped to EL3.

Configurations

AArch64 System register ICH_MISR_EL2 is architecturally mapped to AArch32 System register ICH_MISR.

Attributes

ICH_MISR_EL2 is a 32-bit register.

Field descriptions

The ICH_MISR_EL2 bit assignments are:

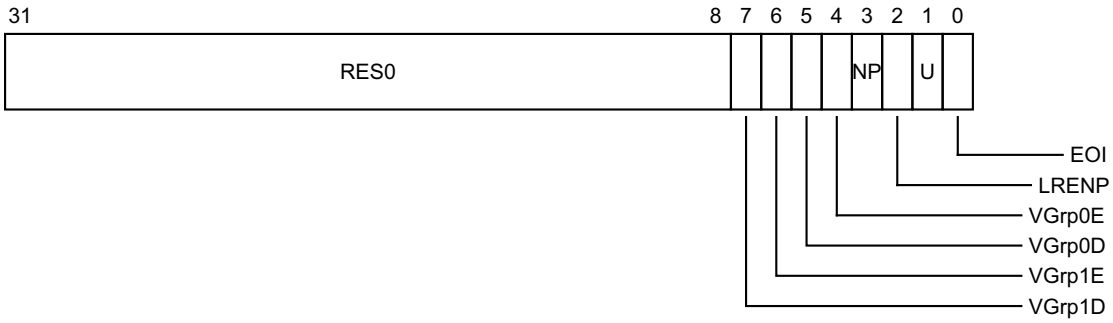

Bits [31:8]

Reserved, RES0.

VGrp1D, bit [7]

vPE Group 1 Disabled.

`0` vPE Group 1 Disabled maintenance interrupt not asserted.

`1` vPE Group 1 Disabled maintenance interrupt asserted.

This maintenance interrupt is asserted when ICH_HCR_EL2.VENG1 is 1 and ICH_VMCR_EL2.VMGrp1En is 0.

When this register has an architecturally-defined reset value, this field resets to 0.

VGrp1E, bit [6]

vPE Group 1 Enabled.

0 vPE Group 1 Enabled maintenance interrupt not asserted.

1 vPE Group 1 Enabled maintenance interrupt asserted.

This maintenance interrupt is asserted when ICH_HCR_EL2.VENG1 is 1 and ICH_VMCR_EL2.VMGrp1En is 1.

When this register has an architecturally-defined reset value, this field resets to 0.

VGrp0D, bit [5]

vPE Group 0 Disabled.

0 vPE Group 0 Disabled maintenance interrupt not asserted.

1 vPE Group 0 Disabled maintenance interrupt asserted.

This maintenance interrupt is asserted when ICH_HCR_EL2.VENG0 is 1 and ICH_VMCR_EL2.VMGrp0En is 0.

When this register has an architecturally-defined reset value, this field resets to 0.

VGrp0E, bit [4]

vPE Group 0 Enabled.

0 vPE Group 0 Enabled maintenance interrupt not asserted.

1 vPE Group 0 Enabled maintenance interrupt asserted.

This maintenance interrupt is asserted when ICH_HCR_EL2.VENG0 is 1 and ICH_VMCR_EL2.VMGrp0En is 1.

When this register has an architecturally-defined reset value, this field resets to 0.

NP, bit [3]

No Pending.

0 No Pending maintenance interrupt not asserted.

1 No Pending maintenance interrupt asserted.

This maintenance interrupt is asserted when ICH_HCR_EL2.NPIE is 1 and no List register is in pending state.

When this register has an architecturally-defined reset value, this field resets to 0.

LRENP, bit [2]

List Register Entry Not Present.

0 List Register Entry Not Present maintenance interrupt not asserted.

1 List Register Entry Not Present maintenance interrupt asserted.

This maintenance interrupt is asserted when ICH_HCR_EL2.LRENPIE is 1 and ICH_HCR_EL2.EOIcount is non-zero.

When this register has an architecturally-defined reset value, this field resets to 0.

U, bit [1]

Underflow.

0 Underflow maintenance interrupt not asserted.

1 Underflow maintenance interrupt asserted.

This maintenance interrupt is asserted when ICH_HCR_EL2.UIE is 1 and zero or one of the List register entries are marked as a valid interrupt, that is, if the corresponding ICH_LR<n>_EL2.State bits do not equal 0x0.

When this register has an architecturally-defined reset value, this field resets to 0.

EOI, bit [0]

End Of Interrupt.

0 End Of Interrupt maintenance interrupt not asserted.

1 End Of Interrupt maintenance interrupt asserted.

This maintenance interrupt is asserted when at least one bit in ICH_EISR_EL2 is 1.

When this register has an architecturally-defined reset value, this field resets to 0.

The U and NP bits do not include the status of any pending/active VSet packets because these bits control generation of interrupts that allow software management of the contents of the List Registers (which are not affected by VSet packets).

Accessing the ICH_MISR_EL2:

To access the ICH_MISR_EL2:

```
MRS <Xt>, ICH_MISR_EL2 ; Read ICH_MISR_EL2 into Xt
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	100	1100	1011	010

8.4.8 ICH_VMCR_EL2, Interrupt Controller Virtual Machine Control Register

The ICH_VMCR_EL2 characteristics are:

Purpose

Enables the hypervisor to save and restore the virtual machine view of the GIC state.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RW	RW	RW

When EL2 is using System register access, EL1 using either System register or memory-mapped access must be supported.

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL2.SRE==0, accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, accesses to this register from EL3 are trapped to EL3.

Configurations

AArch64 System register ICH_VMCR_EL2 is architecturally mapped to AArch32 System register ICH_VMCR.

Attributes

ICH_VMCR_EL2 is a 32-bit register.

Field descriptions

The ICH_VMCR_EL2 bit assignments are:

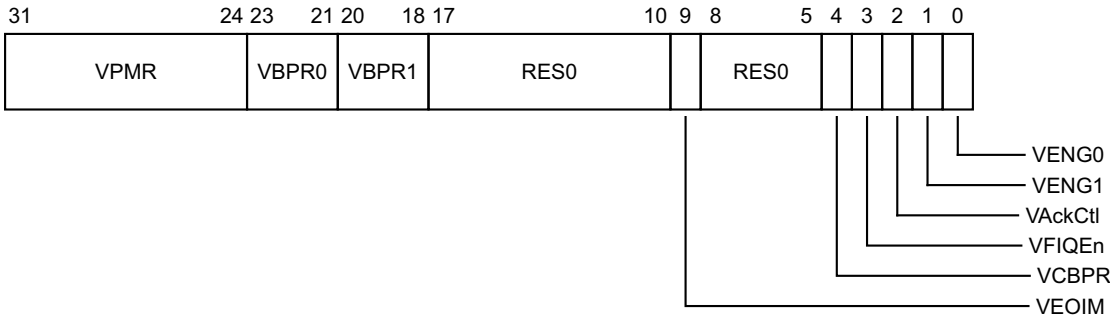

VPMR, bits [31:24]

Virtual Priority Mask. The priority mask level for the virtual CPU interface. If the priority of a pending virtual interrupt is higher than the value indicated by this field, the interface signals the virtual interrupt to the PE.

This field is an alias of ICV_PMR_EL1.Priority.

VBPR0, bits [23:21]

Virtual Binary Point Register, Group 0. Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines Group 0 interrupt preemption, and also determines Group 1 interrupt preemption if ICH_VMCR_EL2.VCBPR == 1.

This field is an alias of ICV_BPR0_EL1.BinaryPoint.

VBPR1, bits [20:18]

Virtual Binary Point Register, Group 1. Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines Group 1 interrupt preemption if ICH_VMCR_EL2.VCBPR == 0.

This field is an alias of ICV_BPR1_EL1.BinaryPoint.

Bits [17:10]

Reserved, RES0.

VEOIM, bit [9]

Virtual EOI mode. Controls whether a write to an End of Interrupt register also deactivates the virtual interrupt:

`0` ICV_EOIR0_EL1 and ICV_EOIR1_EL1 provide both priority drop and interrupt deactivation functionality. Accesses to ICV_DIR_EL1 are UNPREDICTABLE.

`1` ICV_EOIR0_EL1 and ICV_EOIR1_EL1 provide priority drop functionality only. ICV_DIR_EL1 provides interrupt deactivation functionality.

This bit is an alias of ICV_CTLR_EL1.EOImode.

Bits [8:5]

Reserved, RES0.

VCBPR, bit [4]

Virtual Common Binary Point Register. Possible values of this bit are:

`0` ICV_BPR0_EL1 determines the preemption group for virtual Group 0 interrupts only.
ICV_BPR1_EL1 determines the preemption group for virtual Group 1 interrupts.

`1` ICV_BPR0_EL1 determines the preemption group for both virtual Group 0 and virtual Group 1 interrupts.
Reads of ICV_BPR1_EL1 return ICV_BPR0_EL1 plus one, saturated to `0b111`. Writes to ICV_BPR1_EL1 are ignored.

This field is an alias of ICV_CTLR_EL1.CBPR.

VFIQEn, bit [3]

Virtual FIQ enable. Possible values of this bit are:

`0` Group 0 virtual interrupts are presented as virtual IRQs.

`1` Group 0 virtual interrupts are presented as virtual FIQs.

This bit is an alias of GICV_CTLR.FIQEn.

In implementations where the Non-secure copy of ICC_SRE_EL1.SRE is always one, this bit is RES1.

VAckCtl, bit [2]

Virtual AckCtl. Possible values of this bit are:

`0` If the highest priority pending interrupt is Group 1, a read of GICV_IAR or GICV_HPPIR returns an INTID of 1022.

`1` If the highest priority pending interrupt is Group 1, a read of GICV_IAR or GICV_HPPIR returns the INTID of the corresponding interrupt.

This bit is an alias of GICV_CTLR.AckCtl.

This field is supported for backwards compatibility with GICv2. ARM deprecates the use of this field.

In implementations where the Non-secure copy of ICC_SRE_EL1.SRE is always 1, this bit is RES0.

VENG1, bit [1]

Virtual Group 1 interrupt enable. Possible values of this bit are:

0 Virtual Group 1 interrupts are disabled.

1 Virtual Group 1 interrupts are enabled.

This bit is an alias of ICV_IGRPEN1_EL1.Enable.

VENG0, bit [0]

Virtual Group 0 interrupt enable. Possible values of this bit are:

0 Virtual Group 0 interrupts are disabled.

1 Virtual Group 0 interrupts are enabled.

This bit is an alias of ICV_IGRPEN0_EL1.Enable.

Accessing the ICH_VMCR_EL2:

To access the ICH_VMCR_EL2:

```
MRS <Xt>, ICH_VMCR_EL2 ; Read ICH_VMCR_EL2 into Xt
MSR ICH_VMCR_EL2, <Xt> ; Write Xt to ICH_VMCR_EL2
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	100	1100	1011	111

8.4.9 ICH_VTR_EL2, Interrupt Controller VGIC Type Register

The ICH_VTR_EL2 characteristics are:

Purpose

Reports supported GIC virtualisartion features.

Usage constraints

This register is accessible as follows:

EL0	EL1 (NS)	EL1 (S)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RO	RO	RO

Traps and Enables

For a description of the prioritization of any generated exceptions, see section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

Subject to the prioritization rules:

If ICC_SRE_EL2.SRE==0, read accesses to this register from EL2 are trapped to EL2.

If ICC_SRE_EL3.SRE==0, read accesses to this register from EL3 are trapped to EL3.

Configurations

AArch64 System register ICH_VTR_EL2 is architecturally mapped to AArch32 System register ICH_VTR.

Attributes

ICH_VTR_EL2 is a 32-bit register.

Field descriptions

The ICH_VTR_EL2 bit assignments are:

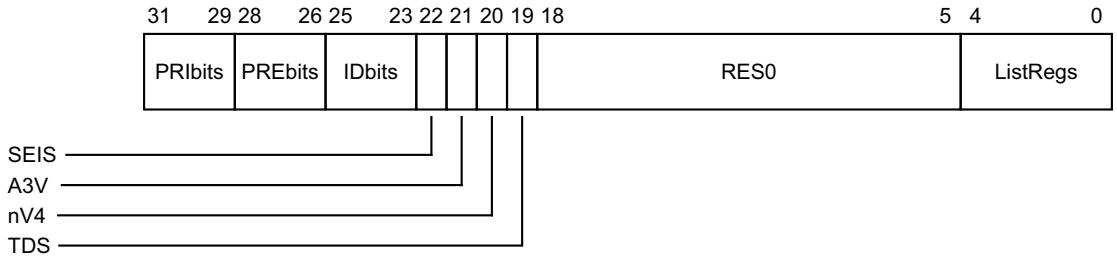

PRIbits, bits [31:29]

Priority bits. The number of virtual priority bits implemented, minus one.

An implementation must implement at least 32 levels of virtual priority (5 priority bits).

This field is an alias of ICV_CTLR_EL1.PRIbits.

PREbits, bits [28:26]

The number of virtual preemption bits implemented, minus one.

An implementation must implement at least 32 levels of virtual preemption priority (5 preemption bits).

The value of this field must be less than or equal to the value of ICH_VTR_EL2.PRIbits.

IDbits, bits [25:23]

The number of virtual interrupt identifier bits supported:

`000` 16 bits.

`001` 24 bits.

All other values are reserved.

This field is an alias of ICV_CTLR_EL1.IDbits.

SEIS, bit [22]

SEI Support. Indicates whether the virtual CPU interface supports generation of SEIs:

`0` The virtual CPU interface logic does not support generation of SEIs.

`1` The virtual CPU interface logic supports generation of SEIs.

This bit is an alias of ICV_CTLR_EL1.SEIS.

A3V, bit [21]

Affinity 3 Valid. Possible values are:

`0` The virtual CPU interface logic only supports zero values of Affinity 3 in SGI generation System registers.

`1` The virtual CPU interface logic supports non-zero values of Affinity 3 in SGI generation System registers.

This bit is an alias of ICV_CTLR_EL1.A3V.

nV4, bit [20]

GICv4 direct injection of virtual interrupts not supported. Possible values are:

`0` The CPU interface logic supports direct injection of virtual interrupts.

`1` The CPU interface logic does not support direct injection of virtual interrupts.

TDS, bit [19]

Separate trapping of Non-secure EL1 writes to ICV_DIR_EL1 supported.

`0` Implementation does not support ICH_HCR_EL2.TDIR.

`1` Implementation supports ICH_HCR_EL2.TDIR.

Bits [18:5]

Reserved, RES0.

ListRegs, bits [4:0]

The number of implemented List registers, minus one. For example, a value of `0b01111` indicates that the maximum of 16 List registers are implemented.

Accessing the ICH_VTR_EL2:

To access the ICH_VTR_EL2:

```
MRS <Xt>, ICH_VTR_EL2 ; Read ICH_VTR_EL2 into Xt
```

Register access is encoded as follows:

op0	op1	CRn	CRm	op2
11	100	1100	1011	001

8.5 AArch32 System register descriptions

This section describes each of the physical AArch32 GIC System registers in register name order. The ICC prefix indicates a GIC CPU interface System register. Each AArch32 System register description contains a reference to the AArch64 register that provides the same functionality.

Unless otherwise stated, the bit assignments for the GIC System registers are the same as those for the equivalent GICC_* and GICV_* memory-mapped registers.

The ICC prefix is used by the System register access mechanism to select the physical or virtual interface System registers according to the setting of HCR. The equivalent memory-mapped physical registers are described in *The GIC CPU interface register descriptions* on page 8-548. The equivalent virtual interface memory-mapped registers are described in *The GIC virtual CPU interface register descriptions* on page 8-587.

Table 8-23 shows the encodings for the AArch32 System registers.

Table 8-23 Encodings for the AArch32 System registers

Register	Width (bits)	opc1	CRn	CRm	opc2	Notes
ICC_PMR	32	0	4	6	0	RW
ICC_SGI1R	64		-	12	-	WO
ICC_IAR0	32		12	8	0	RO
ICC_EOIR0	32				1	WO
ICC_HPPIR0	32				2	RO
ICC_BPR0	32				3	RW
ICC_AP0R<n>	32				4	RW
ICC_AP0R<n>	32				5	RW
ICC_AP0R<n>	32				6	RW
ICC_AP0R<n>	32				7	RW
ICC_AP1R<n>	32			9	0	RW
ICC_AP1R<n>	32				1	RW
ICC_AP1R<n>	32				2	RW
ICC_AP1R<n>	32				3	RW
ICC_DIR	32			11	1	WO
ICC_RPR	32				3	RO

Table 8-23 Encodings for the AArch32 System registers (continued)

Register	Width (bits)	opc1	CRn	CRm	opc2	Notes
ICC_IAR1	32	0	12	12	0	RO
ICC_EOIR1	32				1	WO
ICC_HPPIR1	32				2	RO
ICC_BPR1	32				3	RW
ICC_CTLR	32				4	RW
ICC_SRE	32				5	RW
ICC_IGRPEN0	32				6	RW
ICC_IGRPEN1	32				7	RW
ICC_ASGI1R	64	1	-		-	WO
ICC_SGI0R	64	2	-		-	WO
ICC_MCTLR	32	6	12	12	4	RW
ICC_MSRE	32				5	RW
ICC_MGRPEN1	32				7	RW
ICC_HSRE	32	12	12	9	5	RW

The following access encodings are IMPLEMENTATION DEFINED.

op1	CRn	CRm	op2
000	1100	1101	000

8.5.1 ICC_AP0R<n>, Interrupt Controller Active Priorities Group 0 Registers, n = 0 - 3

The ICC_AP0R<n> characteristics are:

Purpose

Provides information about Group 0 active priorities.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	RW	RW	RW

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	RW	RW

The ICC_AP0R<n> registers are only accessible at Non-secure EL1 when HCR.FMO is set to 0.

Note

When HCR.FMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_AP0R<n> results in an access to ICV_AP0R<n>.

Writing to these registers with any value other than the last read value of the register (or `0x00000000` when there are no Group 0 active priorities) might result in UNPREDICTABLE behavior of the interrupt prioritization system, causing:

- Interrupts that should preempt execution to not preempt execution.
- Interrupts that should not preempt execution to preempt execution.

ICC_AP0R1 is only implemented in implementations that support 6 or more bits of priority. ICC_AP0R2 and ICC_AP0R3 are only implemented in implementations that support 7 bits of priority. Unimplemented registers are UNDEFINED.

Writing to the active priority registers in any order other than the following order will result in UNPREDICTABLE behavior:

- ICC_AP0R<n>.
- Secure ICC_AP1R<n>.
- Non-secure ICC_AP1R<n>.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR.FIQ==1, and EL3 is implemented and configured to use AArch32, accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.FIQ==1, and HCR.FMO==0, and EL2 is implemented and configured to use AArch32, Non-secure accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.FIQ==1, and EL3 is implemented and configured to use AArch64, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and EL3 is implemented and configured to use AArch64, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR.FMO==0, and EL3 is implemented and configured to use AArch64 and EL2 is implemented and configured to use AArch32, Non-secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR_EL2.FMO==0, and EL2 is implemented and configured to use AArch64, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_AP0R<n> is architecturally mapped to AArch64 System register ICC_AP0R<n>_EL1.

Attributes

ICC_AP0R<n> is a 32-bit register.

Field descriptions

The ICC_AP0R<n> bit assignments are:

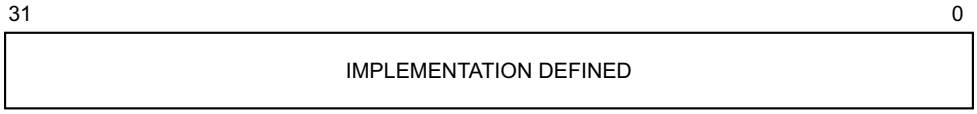

IMPLEMENTATION DEFINED, bits [31:0]

IMPLEMENTATION DEFINED.

When this register has an architecturally-defined reset value, this field resets to `0`.

The contents of these registers are IMPLEMENTATION DEFINED with the one architectural requirement that the value `0x00000000` is consistent with no interrupts being active.

Accessing the ICC_AP0R<n>:

To access the ICC_AP0R<n>:

```
MRC p15,0,<Rt>,c12,c8,<opc2> ; Read ICC_AP0R<n> into Rt, where n is in the range 0 to 3
MCR p15,0,<Rt>,c12,c8,<opc2> ; Write Rt to ICC_AP0R<n>, where n is in the range 0 to 3
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1000	1:n<1:0>

When HCR.FMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_AP0R<n>.

8.5.2 ICC_AP1R<n>, Interrupt Controller Active Priorities Group 1 Registers, n = 0 - 3

The ICC_AP1R<n> characteristics are:

Purpose

Provides information about Group 1 active priorities.

Usage constraints

ICC_AP1R<n>(S) is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	-	-	RW

ICC_AP1R<n>(NS) is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	RW	RW	-

The ICC_AP1R<n> registers are only accessible at Non-secure EL1 when HCR.IMO is set to 0.

Note

When HCR.IMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_AP1R<n> results in an access to ICV_AP1R<n>.

Writing to these registers with any value other than the last read value of the register (or `0x00000000` when there are no Group 1 active priorities) might result in UNPREDICTABLE behavior of the interrupt prioritization system, causing:

- Interrupts that should preempt execution to not preempt execution.
- Interrupts that should not preempt execution to preempt execution.

ICC_AP1R1 is only implemented in implementations that support 6 or more bits of priority. ICC_AP1R2 and ICC_AP1R3 are only implemented in implementations that support 7 bits of priority. Unimplemented registers are UNDEFINED.

Writing to the active priority registers in any order other than the following order will result in UNPREDICTABLE behavior:

- ICC_AP0R<n>.
- Secure ICC_AP1R<n>.
- Non-secure ICC_AP1R<n>.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR.IRQ==1, accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.IRQ==1, and HCR.IMO==0, Non-secure accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.IRQ==1, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR.IMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR_EL2.IMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

AArch32 System register ICC_AP1R<n>(S) is architecturally mapped to AArch64 System register ICC_AP1R<n>_EL1 (S).

AArch32 System register ICC_AP1R<n>(NS) is architecturally mapped to AArch64 System register ICC_AP1R<n>_EL1 (NS).

Attributes

ICC_AP1R<n> is a 32-bit register.

Field descriptions

The ICC_AP1R<n> bit assignments are:

IMPLEMENTATION DEFINED, bits [31:0]

IMPLEMENTATION DEFINED.

When this register has an architecturally-defined reset value, this field resets to `0`.

The contents of these registers are IMPLEMENTATION DEFINED with the one architectural requirement that the value `0x00000000` is consistent with no interrupts being active.

Accessing the ICC_AP1R<n>:

To access the ICC_AP1R<n>:

```
MRC p15,0,<Rt>,c12,c9,<opc2> ; Read ICC_AP1R<n> into Rt, where n is in the range 0 to 3
MCR p15,0,<Rt>,c12,c9,<opc2> ; Write Rt to ICC_AP1R<n>, where n is in the range 0 to 3
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1001	0:n<1:0>

When HCR.IMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_AP1R<n>.

8.5.3 ICC_ASGI1R, Interrupt Controller Alias Software Generated Interrupt Group 1 Register

The ICC_ASGI1R characteristics are:

Purpose

Generates Group 1 SGIs for the Security state that is not the current Security state.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	WO	WO	WO	WO

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	WO	WO

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HCR.FMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HCR_EL2.FMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HCR.IMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HCR_EL2.IMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HSTR.T12==1, Non-secure write accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, write accesses to this register from EL1 are UNDEFINED.

If ICH_HCR.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If SCR.FIQ==1, and SCR.IRQ==1, write accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, Secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, write accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR.IMO==0, and HCR.FMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR_EL2.IMO==0, and HCR_EL2.FMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_ASGI1R performs the same function as AArch64 System operation ICC_ASGI1R_EL1.

Under certain conditions a write to ICC_ASGI1R can generate Group 0 interrupts, see Table 8-14 on page 8-171.

Attributes

ICC_ASGI1R is a 64-bit register.

Field descriptions

The ICC_ASGI1R bit assignments are:

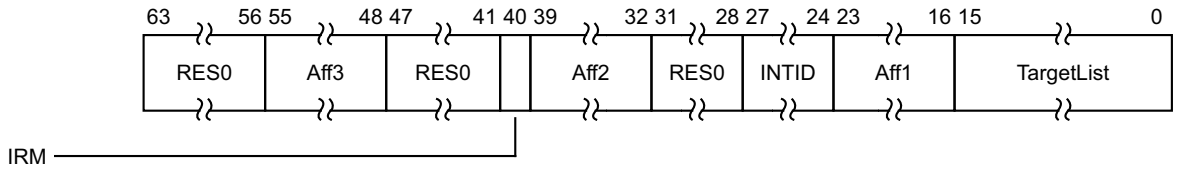

Bits [63:56]

Reserved, RES0.

Aff3, bits [55:48]

The affinity 3 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

Bits [47:41]

Reserved, RES0.

IRM, bit [40]

Interrupt Routing Mode. Determines how the generated interrupts should be distributed to PEs. Possible values are:

`0` Interrupts routed to the PEs specified by Aff3.Aff2.Aff1.<target list>.

`1` Interrupts routed to all PEs in the system, excluding "self".

Aff2, bits [39:32]

The affinity 2 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

Bits [31:28]

Reserved, RES0.

INTID, bits [27:24]

The INTID of the SGI.

Aff1, bits [23:16]

The affinity 1 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

TargetList, bits [15:0]

Target List. The set of PEs for which SGI interrupts will be generated. Each bit corresponds to the PE within a cluster with an Affinity 0 value equal to the bit number.

If a bit is 1 and the bit does not correspond to a valid target PE, the bit must be ignored by the Distributor. It is IMPLEMENTATION DEFINED whether, in such cases, a Distributor can signal a system error.

Note

This restricts a system to sending targeted SGIs to PEs with an affinity 0 number that is less than 16.

If SRE is set only for secure EL3, software executing at EL3 might use the System register interface to generate SGIs. Hence, the Distributor must always be able to receive and acknowledge Generate SGI packets received from CPU interface regardless of the ARE settings for a Security state. However, the Distributor might discard such packets.

If the IRM bit is 1, this field is RES0.

Accessing the ICC_ASGI1R:

To access the ICC_ASGI1R:

```
MCRR p15,1,<Rt>,<Rt2>,c12 ; Write Rt to ICC_ASGI1R[31:0] and Rt2 to ICC_ASGI1R[63:32]
```

Register access is encoded as follows:

coproc	opc1	CRm
1111	0001	1100

8.5.4 ICC_BPR0, Interrupt Controller Binary Point Register 0

The ICC_BPR0 characteristics are:

Purpose

Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines Group 0 interrupt preemption.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	RW	RW	RW

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	RW	RW

ICC_BPR0 is only accessible at Non-secure EL1 when HCR.FMO is set to 0.

Note

When HCR.FMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_BPR0 results in an access to ICV_BPR0.

The minimum binary point value is derived from the number of implemented priority bits. The number of priority bits is IMPLEMENTATION DEFINED, and reported by ICC_CTLR.PRIbits and ICC_MCTLR.PRIbits.

An attempt to program the binary point field to a value less than the minimum value sets the field to the minimum value. On a reset, the binary point field is set to the minimum supported value.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR.FIQ==1, and EL3 is implemented and configured to use AArch32, accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.FIQ==1, and HCR.FMO==0, and EL2 is implemented and configured to use AArch32, Non-secure accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.FIQ==1, and EL3 is implemented and configured to use AArch64, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and EL3 is implemented and configured to use AArch64, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR.FMO==0, and EL3 is implemented and configured to use AArch64 and EL2 is implemented and configured to use AArch32, Non-secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR_EL2.FMO==0, and EL2 is implemented and configured to use AArch64, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_BPR0 is architecturally mapped to AArch64 System register ICC_BPR0_EL1.

Attributes

ICC_BPR0 is a 32-bit register.

Field descriptions

The ICC_BPR0 bit assignments are:

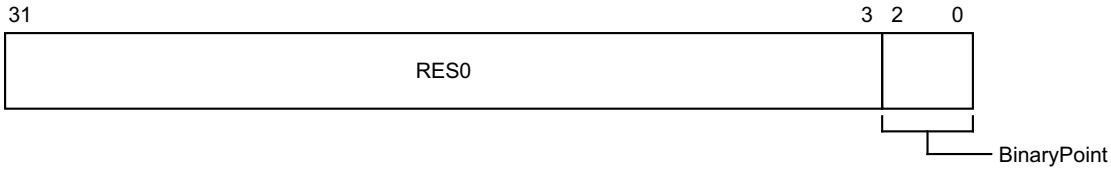

Bits [31:3]

Reserved, RES0.

BinaryPoint, bits [2:0]

The value of this field controls how the 8-bit interrupt priority field is split into a group priority field, that determines interrupt preemption, and a subpriority field. This is done as follows:

Binary point value	Group priority field	Subpriority field	Field with binary point
0	[7:1]	[0]	ggggggg.s
1	[7:2]	[1:0]	gggggg.ss
2	[7:3]	[2:0]	ggggg.sss
3	[7:4]	[3:0]	gggg.ssss
4	[7:5]	[4:0]	ggg.sssss
5	[7:6]	[5:0]	gg.ssssss
6	[7]	[6:0]	g.sssssss
7	No preemption	[7:0]	.ssssssss

Accessing the ICC_BPR0:

To access the ICC_BPR0:

```
MRC p15,0,<Rt>,c12,c8,3 ; Read ICC_BPR0 into Rt
MCR p15,0,<Rt>,c12,c8,3 ; Write Rt to ICC_BPR0
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1000	011

When HCR.FMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_BPR0.

8.5.5 ICC_BPR1, Interrupt Controller Binary Point Register 1

The ICC_BPR1 characteristics are:

Purpose

Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines Group 1 interrupt preemption.

Usage constraints

ICC_BPR1(S) is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	-	-	RW

ICC_BPR1(NS) is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	RW	RW	-

ICC_BPR1 is only accessible at Non-secure EL1 when HCR.IMO is set to 0.

Note

When HCR.IMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_BPR1 results in an access to ICV_BPR1.

When the PE resets into an Exception level that is using AArch32, the reset value is equal to:

- For the Secure copy of the register, the minimum value of ICC_BPR0 plus one.
- For the Non-secure copy of the register, the minimum value of ICC_BPR0.

Where the minimum value of ICC_BPR0 is IMPLEMENTATION DEFINED.

If EL3 is not implemented:

- If the PE is Secure this reset value is (minimum value of ICC_BPR0 plus one).
- If the PE is Non-secure this reset value is (minimum value of ICC_BPR0).

An attempt to program the binary point field to a value less than the reset value sets the field to the reset value.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR.IRQ==1, accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.IRQ==1, and HCR.IMO==0, Non-secure accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.IRQ==1, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR.IMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR_EL2.IMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

AArch32 System register ICC_BPR1(S) is architecturally mapped to AArch64 System register ICC_BPR1_EL1 (S).

AArch32 System register ICC_BPR1(NS) is architecturally mapped to AArch64 System register ICC_BPR1_EL1 (NS).

In GIC implementations supporting two Security states, this register is Banked.

Attributes

ICC_BPR1 is a 32-bit register.

Field descriptions

The ICC_BPR1 bit assignments are:

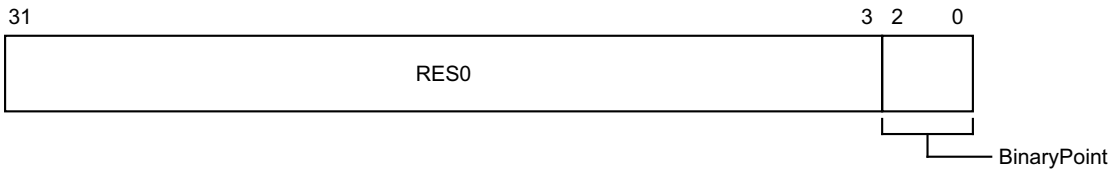

Bits [31:3]

Reserved, RES0.

BinaryPoint, bits [2:0]

If the GIC is configured to use separate binary point fields for Group 0 and Group 1 interrupts, the value of this field controls how the 8-bit interrupt priority field is split into a group priority field, that determines interrupt preemption, and a subpriority field. This is done as follows:

Binary point value	Group priority field	Subpriority field	Field with binary point
0	-	-	-
1	[7:1]	[0]	ggggggg.s
2	[7:2]	[1:0]	gggggg.ss
3	[7:3]	[2:0]	ggggg.sss
4	[7:4]	[3:0]	gggg.ssss
5	[7:5]	[4:0]	ggg.sssss
6	[7:6]	[5:0]	gg.ssssss
7	[7]	[6:0]	g.sssssss

Writing 0 to this field will set this field to its reset value, which is IMPLEMENTATION DEFINED and non-zero.

If EL3 is implemented and ICC_MCTLR.CBPR_EL1S is 1:

- Writing to this register at Secure EL1, or at EL3 not in Monitor mode, modifies ICC_BPR0.
- Reading this register at Secure EL1, or at EL3 not in Monitor mode, returns the value of ICC_BPR0.

If EL3 is implemented and ICC_MCTLR.CBPR_EL1NS is 1, Non-secure accesses to this register at EL1 or EL2 behave as follows, depending on the values of HCR.IMO and SCR.IRQ:

HCR.IMO	SCR.IRQ	Behavior
0	0	Non-secure EL1 and EL2 reads return ICC_BPR0 + 1 saturated to `0b111`. Non-secure EL1 and EL2 writes are ignored.
0	1	Non-secure EL1 and EL2 accesses trap to EL3.
1	0	Non-secure EL1 accesses affect virtual interrupts. Non-secure EL2 reads return ICC_BPR0 + 1 saturated to `0b111`. Non-secure EL2 writes are ignored.
1	1	Non-secure EL1 accesses affect virtual interrupts. Non-secure EL2 accesses trap to EL3.

If EL3 is not implemented and ICC_CTLR.CBPR is 1, Non-secure accesses to this register at EL1 or EL2 behave as follows, depending on the values of HCR.IMO:

HCR.IMO	Behavior
0	Non-secure EL1 and EL2 reads return ICC_BPR0 + 1 saturated to `0b111`. Non-secure EL1 and EL2 writes are ignored.
1	Non-secure EL1 accesses affect virtual interrupts. Non-secure EL2 reads return ICC_BPR0 + 1 saturated to `0b111`. Non-secure EL2 writes are ignored.

Accessing the ICC_BPR1:

To access the ICC_BPR1:

```
MRC p15,0,<Rt>,c12,c12,3 ; Read ICC_BPR1 into Rt
MCR p15,0,<Rt>,c12,c12,3 ; Write Rt to ICC_BPR1
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	011

When HCR.IMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_BPR1.

8.5.6 ICC_CTLR, Interrupt Controller Control Register

The ICC_CTLR characteristics are:

Purpose

Controls aspects of the behavior of the GIC CPU interface and provides information about the features implemented.

Usage constraints

ICC_CTLR(S) is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	-	-	RW

ICC_CTLR(NS) is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	RW	RW	-

ICC_CTLR is only accessible at Non-secure EL1 when HCR.{FMO, IMO} == {0, 0}.

Note

When HCR.{FMO, IMO} != {0, 0}, at Non-secure EL1, the instruction encoding used to access ICC_CTLR results in an access to ICV_CTLR.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TC==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TC==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR.IRQ==1, and SCR.FIQ==1, accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.IRQ==1, SCR.FIQ==1, HCR.IMO==0, and HCR.FMO==0, Non-secure accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR.IMO==0, and HCR.FMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR_EL2.IMO==0, and HCR_EL2.FMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

AArch32 System register ICC_CTLR(S) is architecturally mapped to AArch64 System register ICC_CTLR_EL1 (S).

AArch32 System register ICC_CTLR(NS) is architecturally mapped to AArch64 System register ICC_CTLR_EL1 (NS).

Attributes

ICC_CTLR is a 32-bit register.

Field descriptions

The ICC_CTLR bit assignments are:

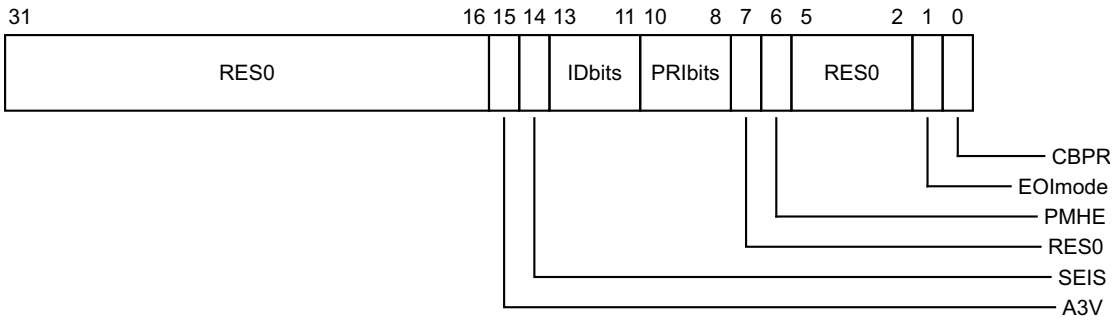

Bits [31:16]

Reserved, RES0.

A3V, bit [15]

Affinity 3 Valid. Read-only and writes are ignored. Possible values are:

`0` The CPU interface logic only supports zero values of Affinity 3 in SGI generation System registers.

`1` The CPU interface logic supports non-zero values of Affinity 3 in SGI generation System registers.

If EL3 is implemented and using AArch32, this bit is an alias of ICC_MCTLR.A3V.

If EL3 is implemented and using AArch64, this bit is an alias of ICC_CTLR_EL3.A3V.

SEIS, bit [14]

SEI Support. Read-only and writes are ignored. Indicates whether the CPU interface supports local generation of SEIs:

`0` The CPU interface logic does not support local generation of SEIs.

`1` The CPU interface logic supports local generation of SEIs.

If EL3 is implemented and using AArch32, this bit is an alias of ICC_MCTLR.SEIS.

If EL3 is implemented and using AArch64, this bit is an alias of ICC_CTLR_EL3.SEIS.

IDbits, bits [13:11]

Identifier bits. Read-only and writes are ignored. The number of physical interrupt identifier bits supported:

`000` 16 bits.

`001` 24 bits.

All other values are reserved.

If EL3 is implemented and using AArch32, this field is an alias of ICC_MCTLR.IDbits.

If EL3 is implemented and using AArch64, this field is an alias of ICC_CTLR_EL3.IDbits.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

PRIbits, bits [10:8]

Priority bits. Read-only and writes are ignored. The number of priority bits implemented, minus one.

An implementation that supports two Security states must implement at least 32 levels of physical priority (5 priority bits).

An implementation that supports only a single Security state must implement at least 16 levels of physical priority (4 priority bits).

Note

This field always returns the number of priority bits implemented, regardless of the Security state of the access or the value of GICD_CTLR.DS.

The division between group priority and subpriority is defined in the binary point registers ICC_BPR0 and ICC_BPR1.

If EL3 is implemented and using AArch32, physical accesses return the value from ICC_MCTLR.PRIbits.

If EL3 is implemented and using AArch64, physical accesses return the value from ICC_CTLR_EL3.PRIbits.

If EL3 is not implemented, physical accesses return the value from this field.

Bit [7]

Reserved, RES0.

PMHE, bit [6]

Priority Mask Hint Enable. Controls whether the priority mask register is used as a hint for interrupt distribution:

`0` Disables use of ICC_PMR as a hint for interrupt distribution.

`1` Enables use of ICC_PMR as a hint for interrupt distribution.

If EL3 is implemented:

- If EL3 is using AArch32, this bit is an alias of ICC_MCTLR.PMHE.
- If EL3 is using AArch64, this bit is an alias of ICC_CTLR_EL3.PMHE.
- If GICD_CTLR.DS == 0, this bit is read-only.
- If GICD_CTLR.DS == 1, this bit is read/write.

If EL3 is not implemented, it is IMPLEMENTATION DEFINED whether this bit is read-only or read-write:

- If this bit is read-only, an implementation can choose to make this field RAZ/WI or RAO/WI.
- If this bit is read/write, it resets to zero.

Bits [5:2]

Reserved, RES0.

EOImode, bit [1]

EOI mode for the current Security state. Controls whether a write to an End of Interrupt register also deactivates the interrupt:

`0` ICC_EOIR0 and ICC_EOIR1 provide both priority drop and interrupt deactivation functionality. Accesses to ICC_DIR are UNPREDICTABLE.

1 ICC_EOIR0 and ICC_EOIR1 provide priority drop functionality only. ICC_DIR provides interrupt deactivation functionality.

If EL3 is implemented:

- If EL3 is using AArch32, this bit is an alias of ICC_MCTLR.EOImode_EL1{S, NS} where S or NS corresponds to the current Security state.
- If EL3 is using AArch64, this bit is an alias of ICC_CTLR_EL3.EOImode_EL1{S, NS} where S or NS corresponds to the current Security state.

If EL3 is not implemented, it is IMPLEMENTATION DEFINED whether this bit is read-only or read-write:

- If this bit is read-only, an implementation can choose to make this field RAZ/WI or RAO/WI.
- If this bit is read/write, it resets to zero.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

CBPR, bit [0]

Common Binary Point Register. Controls whether the same register is used for interrupt preemption of both Group 0 and Group 1 interrupts:

0 ICC_BPR0 determines the preemption group for Group 0 interrupts only.
ICC_BPR1 determines the preemption group for Group 1 interrupts.

1 ICC_BPR0 determines the preemption group for both Group 0 and Group 1 interrupts.

If EL3 is implemented:

- If EL3 is using AArch32, this bit is an alias of ICC_MCTLR.CBPR_EL1{S,NS} where S or NS corresponds to the current Security state.
- If EL3 is using AArch64, this bit is an alias of ICC_CTLR_EL3.CBPR_EL1{S,NS} where S or NS corresponds to the current Security state.
- If GICD_CTLR.DS == 0, this bit is read-only.
- If GICD_CTLR.DS == 1, this bit is read/write.

If EL3 is not implemented, it is IMPLEMENTATION DEFINED whether this bit is read-only or read-write:

- If this bit is read-only, an implementation can choose to make this field RAZ/WI or RAO/WI.
- If this bit is read/write, it resets to zero.

Accessing the ICC_CTLR:

To access the ICC_CTLR:

```
MRC p15,0,<Rt>,c12,c12,4 ; Read ICC_CTLR into Rt
MCR p15,0,<Rt>,c12,c12,4 ; Write Rt to ICC_CTLR
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	100

When HCR.{FMO, IMO} != {0, 0}, execution of this encoding at Non-secure EL1 results in an access to ICV_CTLR.

8.5.7 ICC_DIR, Interrupt Controller Deactivate Interrupt Register

The ICC_DIR characteristics are:

Purpose

When interrupt priority drop is separated from interrupt deactivation, a write to this register deactivates the specified interrupt.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	WO	WO	WO	WO

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	WO	WO

The ICC_DIR register is only accessible at Non-secure EL1 when HCR.{FMO, IMO} == {0, 0}.

Note

At Non-secure EL1, the instruction encoding used to access ICC_DIR results in an access to ICV_DIR in the following cases:

- When HCR.FMO is set to 1.
- When HCR.IMO is set to 1.

There are two cases when writing to ICC_DIR_EL1 that were UNPREDICTABLE for a corresponding GICv2 write to GICC_DIR:

- When EOImode == '0'. GICv3 implementations must ignore such writes. In systems supporting system error generation, an implementation might generate an SEI.
- When EOImode == '1' but no EOI has been issued. The interrupt will be de-activated by the Distributor, however the active priority in the CPU interface for the interrupt will remain set (because no EOI was issued).

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure write accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, write accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, write accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, write accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If SCR.IRQ==1, and SCR.FIQ==1, write accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.IRQ==1, SCR.FIQ==1, HCR.IMO==0, and HCR.FMO==0, Non-secure write accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, Secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, write accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR.IMO==0, and HCR.FMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR_EL2.IMO==0, and HCR_EL2.FMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_DIR performs the same function as AArch64 System operation ICC_DIR_EL1.

Attributes

ICC_DIR is a 32-bit register.

Field descriptions

The ICC_DIR bit assignments are:

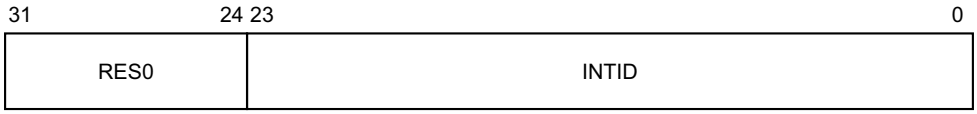

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the interrupt to be deactivated.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICC_CTLR.IDbits and ICC_MCTLR.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICC_DIR:

To access the ICC_DIR:

```
MCR p15,0,<Rt>,c12,c11,1 ; Write Rt to ICC_DIR
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1011	001

This encoding results in an access to ICV_DIR at Non-secure EL1 in the following cases:

- When HCR.FMO is set to 1, and the INTID field refers to a Group 0 interrupt.

- When HCR.IMO is set to 1, and the INTID field refers to a Group 1 interrupt.

8.5.8 ICC_EOIR0, Interrupt Controller End Of Interrupt Register 0

The ICC_EOIR0 characteristics are:

Purpose

A PE writes to this register to inform the CPU interface that it has completed the processing of the specified Group 0 interrupt.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	WO	WO	WO	WO

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	WO	WO

ICC_EOIR0 is only accessible at Non-secure EL1 when HCR.FMO is set to 0.

Note

When HCR.FMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_EOIR0 results in an access to ICV_EOIR0.

A write to this register must correspond to the most recent valid read from an Interrupt Acknowledge Register. A valid read is a read that returns a valid interrupt ID, that is not a special INTID.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure write accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, write accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, write accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, write accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TALL0==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL0==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If SCR.FIQ==1, and EL3 is implemented and configured to use AArch32, write accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.FIQ==1, and HCR.FMO==0, and EL2 is implemented and configured to use AArch32, Non-secure write accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.FIQ==1, and EL3 is implemented and configured to use AArch64, Secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and EL3 is implemented and configured to use AArch64, write accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR.FMO==0, and EL3 is implemented and configured to use AArch64 and EL2 is implemented and configured to use AArch32, Non-secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR_EL2.FMO==0, and EL2 is implemented and configured to use AArch64, Non-secure write accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_EOIR0 performs the same function as AArch64 System operation ICC_EOIR0_EL1.

Attributes

ICC_EOIR0 is a 32-bit register.

Field descriptions

The ICC_EOIR0 bit assignments are:

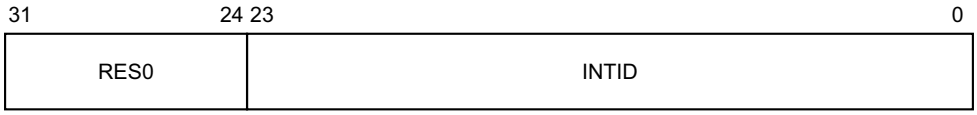

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID from the corresponding ICC_IAR0 access.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICC_CTLR.IDbits and ICC_MCTLR.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

If the EOImode bit for the current Exception level and Security state is 0, a write to this register drops the priority for the interrupt, and also deactivates the interrupt.

If the EOImode bit for the current Exception level and Security state is 1, a write to this register only drops the priority for the interrupt. Software must write to ICC_DIR to deactivate the interrupt.

The appropriate EOImode bit varies as follows:

- If EL3 is not implemented, the appropriate bit is ICC_CTLR.EOImode.
- If EL3 is implemented and the software is executing in Monitor mode, the appropriate bit is ICC_MCTLR.EOImode_EL3.
- If EL3 is implemented and the software is not executing in Monitor mode, the bit depends on the current Security state:
  - — If the software is executing in Secure state, the bit is ICC_CTLR.EOImode in the Secure instance of ICC_CTLR. This is an alias of ICC_MCTLR.EOImode_EL1S.
  - — If the software is executing in Non-secure state, the bit is ICC_CTLR.EOImode in the Non-secure instance of ICC_CTLR. This is an alias of ICC_MCTLR.EOImode_EL1NS.

Accessing the ICC_EOIR0:

To access the ICC_EOIR0:

```
MCR p15,0,<Rt>,c12,c8,1 ; Write Rt to ICC_EOIR0
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1000	001

When HCR.FMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_EOIR0.

8.5.9 ICC_EOIR1, Interrupt Controller End Of Interrupt Register 1

The ICC_EOIR1 characteristics are:

Purpose

A PE writes to this register to inform the CPU interface that it has completed the processing of the specified Group 1 interrupt.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	WO	WO	WO	WO

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	WO	WO

ICC_EOIR1 is only accessible at Non-secure EL1 when HCR.IMO is set to 0.

Note

When HCR.IMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_EOIR1 results in an access to ICV_EOIR1.

A write to this register must correspond to the most recent valid read from an Interrupt Acknowledge Register. A valid read is a read that returns a valid interrupt ID, that is not a special INTID.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure write accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, write accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, write accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, write accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TALL1==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL1==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If SCR.IRQ==1, write accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.IRQ==1, and HCR.IMO==0, Non-secure write accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.IRQ==1, Secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, write accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR.IMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR_EL2.IMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_EOIR1 performs the same function as AArch64 System operation ICC_EOIR1_EL1.

Attributes

ICC_EOIR1 is a 32-bit register.

Field descriptions

The ICC_EOIR1 bit assignments are:

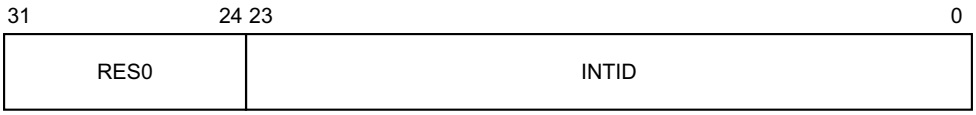

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID from the corresponding ICC_IAR1 access.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICC_CTLR.IDbits and ICC_MCTLR.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

If the EOImode bit for the current Exception level and Security state is 0, a write to this register drops the priority for the interrupt, and also deactivates the interrupt.

If the EOImode bit for the current Exception level and Security state is 1, a write to this register only drops the priority for the interrupt. Software must write to ICC_DIR to deactivate the interrupt.

The appropriate EOImode bit varies as follows:

- If EL3 is not implemented, the appropriate bit is ICC_CTLR.EOImode.
- If EL3 is implemented and the software is executing in Monitor mode, the appropriate bit is ICC_MCTLR.EOImode_EL3.
- If EL3 is implemented and the software is not executing in Monitor mode, the bit depends on the current Security state:
  - If the software is executing in Secure state, the bit is ICC_CTLR.EOImode in the Secure instance of ICC_CTLR. This is an alias of ICC_MCTLR.EOImode_EL1S.
  - If the software is executing in Non-secure state, the bit is ICC_CTLR.EOImode in the Non-secure instance of ICC_CTLR. This is an alias of ICC_MCTLR.EOImode_EL1NS.

Accessing the ICC_EOIR1:

To access the ICC_EOIR1:

```
MCR p15,0,<Rt>,c12,c12,1 ; Write Rt to ICC_EOIR1
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	001

When HCR.IMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_EOIR1.

8.5.10 ICC_HPPIR0, Interrupt Controller Highest Priority Pending Interrupt Register 0

The ICC_HPPIR0 characteristics are:

Purpose

Indicates the highest priority pending Group 0 interrupt on the CPU interface.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RO	RO	RO	RO

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	RO	RO

ICC_HPPIR0 is only accessible at Non-secure EL1 when HCR.FMO is set to 0.

Note

When HCR.FMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_HPPIR0 results in an access to ICV_HPPIR0.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure read accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, read accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, read accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, read accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TALL0==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL0==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If SCR.FIQ==1, and EL3 is implemented and configured to use AArch32, read accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.FIQ==1, and HCR.FMO==0, and EL2 is implemented and configured to use AArch32, Non-secure read accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.FIQ==1, and EL3 is implemented and configured to use AArch64, Secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and EL3 is implemented and configured to use AArch64, read accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR.FMO==0, and EL3 is implemented and configured to use AArch64 and EL2 is implemented and configured to use AArch32, Non-secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR_EL2.FMO==0, and EL2 is implemented and configured to use AArch64, Non-secure read accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_HPPIR0 performs the same function as AArch64 System operation ICC_HPPIR0_EL1.

Attributes

ICC_HPPIR0 is a 32-bit register.

Field descriptions

The ICC_HPPIR0 bit assignments are:

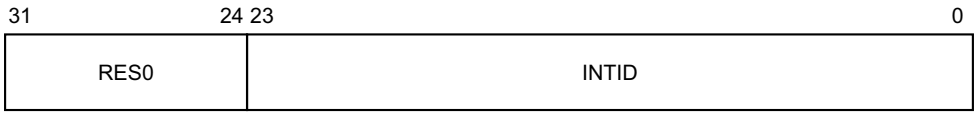

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the highest priority pending interrupt, if that interrupt is observable at the current Security state and Exception level.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. These special INTIDs can be one of: 1020, 1021, or 1023. See *Special INTIDs* on page 3-40, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICC_CTLR.IDbits and ICC_MCTLR.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICC_HPPIR0:

To access the ICC_HPPIR0:

```
MRC p15,0,<Rt>,c12,c8,2 ; Read ICC_HPPIR0 into Rt
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1000	010

When HCR.FMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_HPPIR0.

8.5.11 ICC_HPPIR1, Interrupt Controller Highest Priority Pending Interrupt Register 1

The ICC_HPPIR1 characteristics are:

Purpose

Indicates the highest priority pending Group 1 interrupt on the CPU interface.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RO	RO	RO	RO

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	RO	RO

ICC_HPPIR1 is only accessible at Non-secure EL1 when HCR.IMO is set to 0.

Note

When HCR.IMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_HPPIR1 results in an access to ICV_HPPIR1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure read accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, read accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, read accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, read accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TALL1==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL1==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If SCR.IRQ==1, read accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.IRQ==1, and HCR.IMO==0, Non-secure read accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.IRQ==1, Secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, read accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR.IMO==0, Non-secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR_EL2.IMO==0, Non-secure read accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_HPPIR1 performs the same function as AArch64 System operation ICC_HPPIR1_EL1.

Attributes

ICC_HPPIR1 is a 32-bit register.

Field descriptions

The ICC_HPPIR1 bit assignments are:

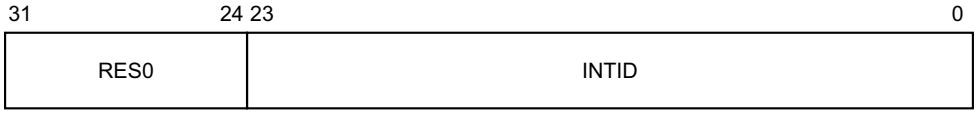

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the highest priority pending interrupt, if that interrupt is observable at the current Security state and Exception level.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. This special INTID can take the value 1023 only. See *Special INTIDs* on page 3-40, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICC_CTLR.IDbits and ICC_MCTLR.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICC_HPPIR1:

To access the ICC_HPPIR1:

```
MRC p15,0,<Rt>,c12,c12,2 ; Read ICC_HPPIR1 into Rt
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	010

When HCR.IMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_HPPIR1.

8.5.12 ICC_HSRE, Interrupt Controller Hyp System Register Enable register

The ICC_HSRE characteristics are:

Purpose

Controls whether the System register interface or the memory-mapped interface to the GIC CPU interface is used for EL2.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RW	RW	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	-	RW

The GIC architecture permits, but does not require, that registers can be shared between memory-mapped registers and the equivalent System registers. This means that if the memory-mapped registers have been accessed while ICC_HSRE.SRE==0, then the System registers might be modified. Therefore, software must only rely on the reset values of the System registers if there has been no use of the GIC functionality while the memory-mapped registers are in use. Otherwise, the System register values must be treated as UNKNOWN.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_MSRE.Enable==0, accesses to this register from EL2 are UNDEFINED.

If ICC_SRE_EL3.Enable==0, accesses to this register from EL2 are trapped to EL3.

Configurations

AArch32 System register ICC_HSRE is architecturally mapped to AArch64 System register ICC_SRE_EL2.

Attributes

ICC_HSRE is a 32-bit register.

Field descriptions

The ICC_HSRE bit assignments are:

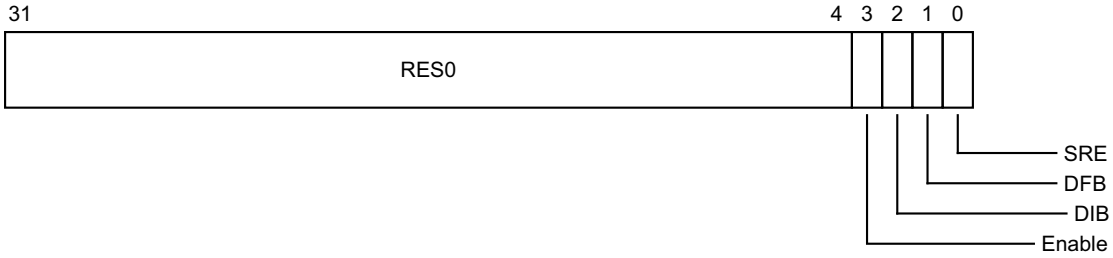

Bits [31:4]

Reserved, RES0.

Enable, bit [3]

Enable. Enables lower Exception level access to ICC_SRE.

0 Non-secure EL1 accesses to ICC_SRE trap to EL2.

1 Non-secure EL1 accesses to ICC_SRE do not trap to EL2.

If ICC_HSRE.SRE is RAO/WI, an implementation is permitted to make the Enable bit RAO/WI.

If ICC_HSRE.SRE is 0, the Enable bit behaves as 1 for all purposes other than reading the value of the bit.

DIB, bit [2]

Disable IRQ bypass.

0 IRQ bypass enabled.

1 IRQ bypass disabled.

If EL3 is implemented and GICD_CTLR.DS is 0, this field is a read-only alias of ICC_MSRE.DIB.

If EL3 is implemented and GICD_CTLR.DS is 1, this field is a read-write alias of ICC_MSRE.DIB.

In systems that do not support IRQ bypass, this bit is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

DFB, bit [1]

Disable FIQ bypass.

0 FIQ bypass enabled.

1 FIQ bypass disabled.

If EL3 is implemented and GICD_CTLR.DS is 0, this field is a read-only alias of ICC_MSRE.DFB.

If EL3 is implemented and GICD_CTLR.DS is 1, this field is a read-write alias of ICC_MSRE.DFB.

In systems that do not support FIQ bypass, this bit is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

SRE, bit [0]

System Register Enable.

0 The memory-mapped interface must be used. Accesses at EL2 or below to any ICH_* System register, or any EL1 or EL2 ICC_* register other than ICC_SRE or ICC_HSRE, are UNDEFINED.

1 The System register interface to the ICH_* registers and the EL1 and EL2 ICC_* registers is enabled for EL2.

If software changes this bit from 1 to 0, the results are UNPREDICTABLE.

If an implementation supports only a System register interface to the GIC CPU interface, this bit is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

Accessing the ICC_HSRE:

To access the ICC_HSRE:

```
MRC p15,4,<Rt>,c12,c9,5 ; Read ICC_HSRE into Rt
MCR p15,4,<Rt>,c12,c9,5 ; Write Rt to ICC_HSRE
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	100	1100	1001	101

8.5.13 ICC_IAR0, Interrupt Controller Interrupt Acknowledge Register 0

The ICC_IAR0 characteristics are:

Purpose

The PE reads this register to obtain the INTID of the signaled Group 0 interrupt. This read acts as an acknowledge for the interrupt.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RO	RO	RO	RO

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	RO	RO

ICC_IAR0 is only accessible at Non-secure EL1 when HCR.FMO is set to 0.

Note

When HCR.FMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_IAR0 results in an access to ICV_IAR0.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure read accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, read accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, read accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, read accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TALL0==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL0==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If SCR.FIQ==1, and EL3 is implemented and configured to use AArch32, read accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.FIQ==1, and HCR.FMO==0, and EL2 is implemented and configured to use AArch32, Non-secure read accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.FIQ==1, and EL3 is implemented and configured to use AArch64, Secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and EL3 is implemented and configured to use AArch64, read accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR.FMO==0, and EL3 is implemented and configured to use AArch64 and EL2 is implemented and configured to use AArch32, Non-secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR_EL2.FMO==0, and EL2 is implemented and configured to use AArch64, Non-secure read accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_IAR0 performs the same function as AArch64 System operation ICC_IAR0_EL1.

To allow software to ensure appropriate observability of actions initiated by GIC register accesses, the PE and CPU interface logic must ensure that reads of this register are self-synchronising when interrupts are masked by the PE (that is when PSTATE.{I,F} == {0,0}). This ensures that the effect of activating an interrupt on the signaling of interrupt exceptions is observed when a read of this register is architecturally executed so that no spurious interrupt exception occurs if interrupts are unmasked by an instruction immediately following the read. See *Observability of the effects of accesses to the GIC registers* on page 8-159, for more information.

Attributes

ICC_IAR0 is a 32-bit register.

Field descriptions

The ICC_IAR0 bit assignments are:

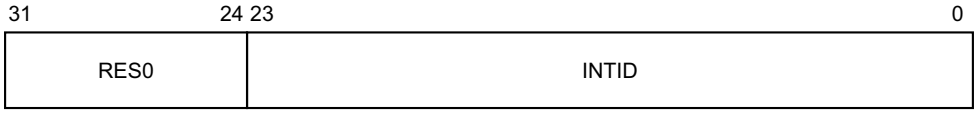

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled interrupt.

This is the INTID of the highest priority pending interrupt, if that interrupt is of sufficient priority for it to be signaled to the PE, and if it can be acknowledged at the current Security state and Exception level.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. These special INTIDs can be one of: 1020, 1021, or 1023. See *Special INTIDs on page 3-40*, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICC_CTLR.IDbits and ICC_MCTLR.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICC_IAR0:

To access the ICC_IAR0:

```
MRC p15,0,<Rt>,c12,c8,0 ; Read ICC_IAR0 into Rt
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1000	000

When HCR.FMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_IAR0.

8.5.14 ICC_IAR1, Interrupt Controller Interrupt Acknowledge Register 1

The ICC_IAR1 characteristics are:

Purpose

The PE reads this register to obtain the INTID of the signaled Group 1 interrupt. This read acts as an acknowledge for the interrupt.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RO	RO	RO	RO

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	RO	RO

ICC_IAR1 is only accessible at Non-secure EL1 when HCR.IMO is set to 0.

Note

When HCR.IMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_IAR1 results in an access to ICV_IAR1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure read accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, read accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, read accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, read accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TALL1==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL1==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If SCR.IRQ==1, read accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.IRQ==1, and HCR.IMO==0, Non-secure read accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.IRQ==1, Secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, read accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR.IMO==0, Non-secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR_EL2.IMO==0, Non-secure read accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_IAR1 performs the same function as AArch64 System operation ICC_IAR1_EL1.

To allow software to ensure appropriate observability of actions initiated by GIC register accesses, the PE and CPU interface logic must ensure that reads of this register are self-synchronising when interrupts are masked by the PE (that is when PSTATE.{I,F} == {0,0}). This ensures that the effect of activating an interrupt on the signaling of interrupt exceptions is observed when a read of this register is architecturally executed so that no spurious interrupt exception occurs if interrupts are unmasked by an instruction immediately following the read. See *Observability of the effects of accesses to the GIC registers* on page 8-159, for more information.

Attributes

ICC_IAR1 is a 32-bit register.

Field descriptions

The ICC_IAR1 bit assignments are:

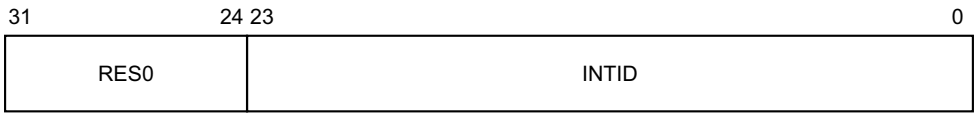

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled interrupt.

This is the INTID of the highest priority pending interrupt, if that interrupt is of sufficient priority for it to be signaled to the PE, and if it can be acknowledged at the current Security state and Exception level.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. This special INTID can take the value 1023 only. See *Special INTIDs* on page 3-40, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICC_CTLR.IDbits and ICC_MCTLR.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICC_IAR1:

To access the ICC_IAR1:

```
MRC p15,0,<Rt>,c12,c12,0 ; Read ICC_IAR1 into Rt
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	000

When HCR.IMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_IAR1.

8.5.15 ICC_IGRPEN0, Interrupt Controller Interrupt Group 0 Enable register

The ICC_IGRPEN0 characteristics are:

Purpose

Controls whether Group 0 interrupts are enabled or not.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	RW	RW	RW

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	RW	RW

ICC_IGRPEN0 is only accessible at Non-secure EL1 when HCR.FMO is set to 0.

Note

When HCR.FMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_IGRPEN0 results in an access to ICV_IGRPEN0.

The lowest Exception level at which this register can be accessed is governed by the Exception level to which FIQ is routed. This routing depends on SCR.FIQ, SCR.NS and HCR.FMO.

If an interrupt is pending within the CPU interface when Enable becomes 0, the interrupt must be released to allow the Distributor to forward the interrupt to a different PE.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR.FIQ==1, and EL3 is implemented and configured to use AArch32, accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.FIQ==1, and HCR.FMO==0, and EL2 is implemented and configured to use AArch32, Non-secure accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.FIQ==1, and EL3 is implemented and configured to use AArch64, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and EL3 is implemented and configured to use AArch64, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR.FMO==0, and EL3 is implemented and configured to use AArch64 and EL2 is implemented and configured to use AArch32, Non-secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and HCR_EL2.FMO==0, and EL2 is implemented and configured to use AArch64, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_IGRPEN0 is architecturally mapped to AArch64 System register ICC_IGRPEN0_EL1.

Attributes

ICC_IGRPEN0 is a 32-bit register.

Field descriptions

The ICC_IGRPEN0 bit assignments are:

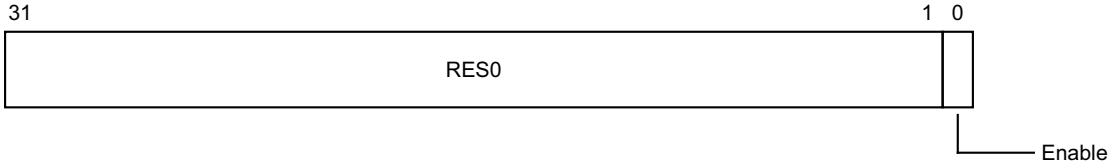

Bits [31:1]

Reserved, RES0.

Enable, bit [0]

Enables Group 0 interrupts.

`0` Group 0 interrupts are disabled.

`1` Group 0 interrupts are enabled.

Virtual accesses to this register update ICH_VMCR.VENG0.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the ICC_IGRPEN0:

To access the ICC_IGRPEN0:

```
MRC p15,0,<Rt>,c12,c12,6 ; Read ICC_IGRPEN0 into Rt
MCR p15,0,<Rt>,c12,c12,6 ; Write Rt to ICC_IGRPEN0
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	110

When HCR.FMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_IGRPEN0.

8.5.16 ICC_IGRPEN1, Interrupt Controller Interrupt Group 1 Enable register

The ICC_IGRPEN1 characteristics are:

Purpose

Controls whether Group 1 interrupts are enabled for the current Security state.

Usage constraints

ICC_IGRPEN1(S) is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	-	-	RW

ICC_IGRPEN1(NS) is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	RW	RW	-

ICC_IGRPEN1 is only accessible at Non-secure EL1 when HCR.IMO is set to 0.

Note

When HCR.IMO is set to 1, at Non-secure EL1, the instruction encoding used to access ICC_IGRPEN1 results in an access to ICV_IGRPEN1.

The lowest Exception level at which this register can be accessed is governed by the Exception level to which IRQ is routed. This routing depends on SCR.IRQ, SCR.NS and HCR.IMO.

If an interrupt is pending within the CPU interface when Enable becomes 0, the interrupt must be released to allow the Distributor to forward the interrupt to a different PE.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR.IRQ==1, accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.IRQ==1, and HCR.IMO==0, Non-secure accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.IRQ==1, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR.IMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.IRQ==1, and HCR_EL2.IMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

AArch32 System register ICC_IGRPEN1(S) is architecturally mapped to AArch64 System register ICC_IGRPEN1_EL1 (S).

AArch32 System register ICC_IGRPEN1(NS) is architecturally mapped to AArch64 System register ICC_IGRPEN1_EL1 (NS).

Attributes

ICC_IGRPEN1 is a 32-bit register.

Field descriptions

The ICC_IGRPEN1 bit assignments are:

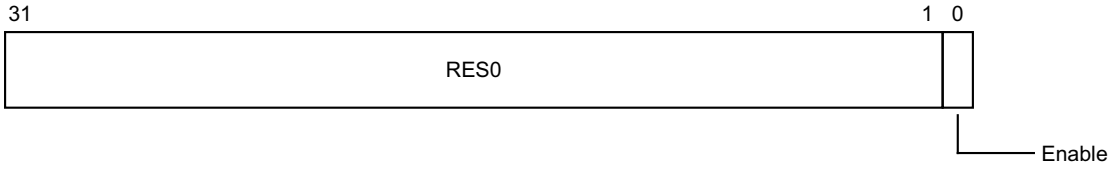

Bits [31:1]

Reserved, RES0.

Enable, bit [0]

Enables Group 1 interrupts for the current Security state.

0 Group 1 interrupts are disabled for the current Security state.

1 Group 1 interrupts are enabled for the current Security state.

Virtual accesses to this register update ICH_VMCR.VENG1.

If EL3 is present:

- This bit is a read/write alias of ICC_MGRPEN1.EnableGrp1{S, NS} as appropriate if EL3 is using AArch32, or ICC_IGRPEN1_EL3.EnableGrp1{S, NS} as appropriate if EL3 is using AArch64.
- When this register is accessed at EL3, the copy of this register appropriate to the current setting of SCR.NS is accessed.

When this register has an architecturally-defined reset value, this field resets to 0.

Accessing the ICC_IGRPEN1:

To access the ICC_IGRPEN1:

```
MRC p15,0,<Rt>,c12,c12,7 ; Read ICC_IGRPEN1 into Rt
MCR p15,0,<Rt>,c12,c12,7 ; Write Rt to ICC_IGRPEN1
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	111

When HCR.IMO is set to 1, execution of this encoding at Non-secure EL1 results in an access to ICV_IGRPEN1.

8.5.17 ICC_MCTLR, Interrupt Controller Monitor Control Register

The ICC_MCTLR characteristics are:

Purpose

Controls aspects of the behavior of the GIC CPU interface and provides information about the features implemented.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1, Monitor mode only)	EL3 (SCR.NS=0, Monitor mode only)
-	-	-	-	RW	RW

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	-	-

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_MSRE.SRE==0, accesses to this register from EL3 are UNDEFINED.

Configurations

This register is only accessible in Secure state.

AArch32 System register ICC_MCTLR can be mapped to AArch64 System register ICC_CTLR_EL3, but this is not architecturally mandated.

Attributes

ICC_MCTLR is a 32-bit register.

Field descriptions

The ICC_MCTLR bit assignments are:

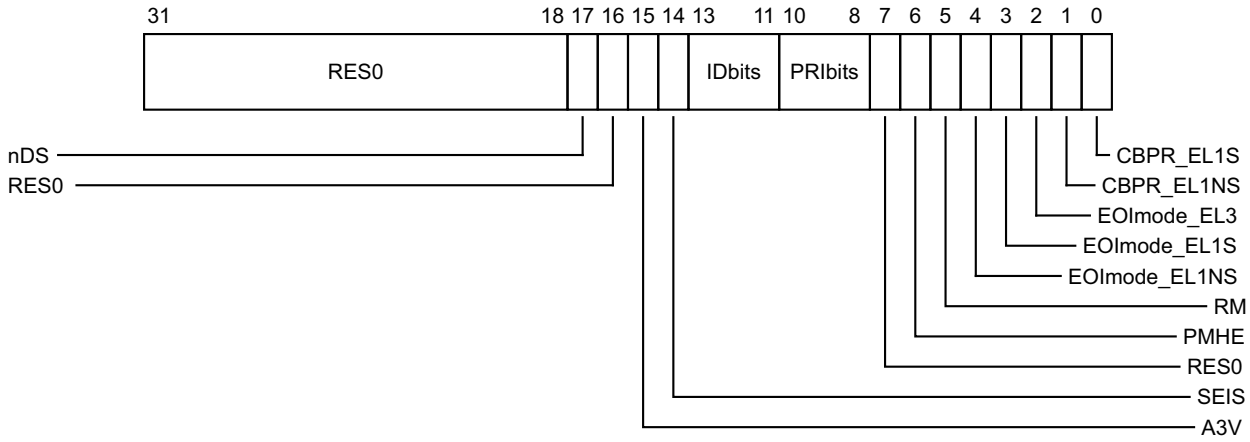

Bits [31:18]

Reserved, RES0.

nDS, bit [17]

Disable Security not supported. Read-only and writes are ignored. Possible values are:

`0` The CPU interface logic supports disabling of security.

`1` The CPU interface logic does not support disabling of security, and requires that security is not disabled.

Bit [16]

Reserved, RES0.

A3V, bit [15]

Affinity 3 Valid. Read-only and writes are ignored. Possible values are:

`0` The CPU interface logic does not support non-zero values of the Aff3 field in SGI generation System registers.

`1` The CPU interface logic supports non-zero values of the Aff3 field in SGI generation System registers.

SEIS, bit [14]

SEI Support. Read-only and writes are ignored. Indicates whether the CPU interface supports generation of SEIs:

`0` The CPU interface logic does not support generation of SEIs.

`1` The CPU interface logic supports generation of SEIs.

IDbits, bits [13:11]

Identifier bits. Read-only and writes are ignored. The number of physical interrupt identifier bits supported:

`000` 16 bits.

`001` 24 bits.

All other values are reserved.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

PRIbits, bits [10:8]

Priority bits. Read-only and writes are ignored. The number of priority bits implemented, minus one.

An implementation that supports two Security states must implement at least 32 levels of physical priority (5 priority bits).

An implementation that supports only a single Security state must implement at least 16 levels of physical priority (4 priority bits).

Note

This field always returns the number of priority bits implemented, regardless of the value of SCR.NS or the value of GICD_CTLR.DS.

The division between group priority and subpriority is defined in the binary point registers ICC_BPR0 and ICC_BPR1.

This field determines the minimum value of ICC_BPR0.

Bit [7]

Reserved, RES0.

PMHE, bit [6]

Priority Mask Hint Enable.

`0` Disables use of the priority mask register as a hint for interrupt distribution.

`1` Enables use of the priority mask register as a hint for interrupt distribution.

Software must write ICC_PMR to `0xFF` before clearing this field to 0.

An implementation might choose to make this field RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

RM, bit [5]

SBZ.

The equivalent bit in AArch64 is the Routing Modifier bit. This feature is not supported when EL3 is using AArch32.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

EOImode_EL1NS, bit [4]

EOI mode for interrupts handled at Non-secure EL1 and EL2. Controls whether a write to an End of Interrupt register also deactivates the interrupt:

`0` ICC_EOIR0 and ICC_EOIR1 provide both priority drop and interrupt deactivation functionality. Accesses to ICC_DIR are UNPREDICTABLE.

`1` ICC_EOIR0 and ICC_EOIR1 provide priority drop functionality only. ICC_DIR provides interrupt deactivation functionality.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

EOImode_EL1S, bit [3]

EOI mode for interrupts handled at Secure EL1. Controls whether a write to an End of Interrupt register also deactivates the interrupt:

`0` ICC_EOIR0 and ICC_EOIR1 provide both priority drop and interrupt deactivation functionality. Accesses to ICC_DIR are UNPREDICTABLE.

`1` ICC_EOIR0 and ICC_EOIR1 provide priority drop functionality only. ICC_DIR provides interrupt deactivation functionality.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

EOImode_EL3, bit [2]

EOI mode for interrupts handled at EL3. Controls whether a write to an End of Interrupt register also deactivates the interrupt:

0 ICC_EOIR0 and ICC_EOIR1 provide both priority drop and interrupt deactivation functionality. Accesses to ICC_DIR are UNPREDICTABLE.

1 ICC_EOIR0 and ICC_EOIR1 provide priority drop functionality only. ICC_DIR provides interrupt deactivation functionality.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

CBPR_EL1NS, bit [1]

Common Binary Point Register, EL1 Non-secure. Controls whether the same register is used for interrupt preemption of both Group 0 and Group 1 Non-secure interrupts at EL1 and EL2:

0 ICC_BPR0 determines the preemption group for Group 0 interrupts only.
ICC_BPR1 determines the preemption group for Non-secure Group 1 interrupts.

1 ICC_BPR0 determines the preemption group for Group 0 interrupts and Non-secure Group 1 interrupts. Non-secure accesses to GICC_BPR and ICC_BPR1 access the state of ICC_BPR0.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

CBPR_EL1S, bit [0]

Common Binary Point Register, EL1 Secure. Controls whether the same register is used for interrupt preemption of both Group 0 and Group 1 Secure interrupts in Secure non-Monitor modes:

0 ICC_BPR0 determines the preemption group for Group 0 interrupts only.
ICC_BPR1 determines the preemption group for Secure Group 1 interrupts.

1 ICC_BPR0 determines the preemption group for Group 0 interrupts and Secure Group 1 interrupts. Secure accesses to ICC_BPR1 access the state of ICC_BPR0.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the ICC_MCTLR:

To access the ICC_MCTLR:

```
MRC p15,6,<Rt>,c12,c12,4 ; Read ICC_MCTLR into Rt
MCR p15,6,<Rt>,c12,c12,4 ; Write Rt to ICC_MCTLR
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	110	1100	1100	100

8.5.18 ICC_MGRPEN1, Interrupt Controller Monitor Interrupt Group 1 Enable register

The ICC_MGRPEN1 characteristics are:

Purpose

Controls whether Group 1 interrupts are enabled or not.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1, Monitor mode only)	EL3 (SCR.NS=0, Monitor mode only)
-	-	-	-	RW	RW

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	-	-

If an interrupt is pending within the CPU interface when an Enable bit becomes 0, the interrupt must be released to allow the Distributor to forward the interrupt to a different PE.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_MSRE.SRE==0, accesses to this register from EL3 are UNDEFINED.

Configurations

This register is only accessible in Secure state.

AArch32 System register ICC_MGRPEN1 can be mapped to AArch64 System register ICC_IGRPEN1_EL3, but this is not architecturally mandated.

Attributes

ICC_MGRPEN1 is a 32-bit register.

Field descriptions

The ICC_MGRPEN1 bit assignments are:

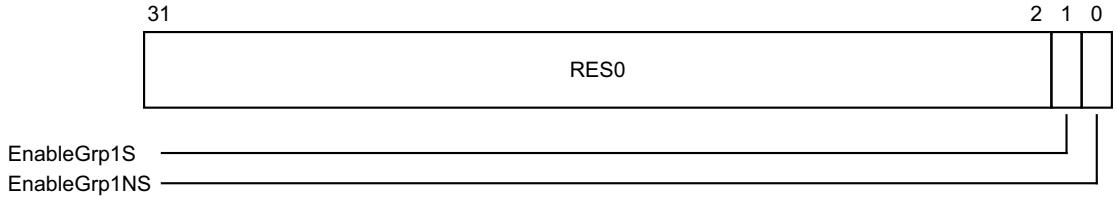

Bits [31:2]

Reserved, RES0.

EnableGrp1S, bit [1]

Enables Group 1 interrupts for the Secure state.

0 Secure Group 1 interrupts are disabled.

1 Secure Group 1 interrupts are enabled.

If EL3 is present, the Secure ICC_IGRPEN1.Enable bit is a read/write alias of the ICC_MGRPEN1.EnableGrp1S bit.

If the highest priority pending interrupt for that PE is a Group 1 interrupt using 1 of N model, then the interrupt will target another PE as a result of the Enable bit changing from 1 to 0.

When this register has an architecturally-defined reset value, this field resets to 0.

EnableGrp1NS, bit [0]

Enables Group 1 interrupts for the Non-secure state.

0 Non-secure Group 1 interrupts are disabled.

1 Non-secure Group 1 interrupts are enabled.

If EL3 is present, the Non-secure ICC_IGRPEN1.Enable bit is a read/write alias of the ICC_MGRPEN1.EnableGrp1NS bit.

If the highest priority pending interrupt for that PE is a Group 1 interrupt using 1 of N model, then the interrupt will target another PE as a result of the Enable bit changing from 1 to 0.

When this register has an architecturally-defined reset value, this field resets to 0.

Accessing the ICC_MGRPEN1:

To access the ICC_MGRPEN1:

```
MRC p15,6,<Rt>,c12,c12,7 ; Read ICC_MGRPEN1 into Rt
MCR p15,6,<Rt>,c12,c12,7 ; Write Rt to ICC_MGRPEN1
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	110	1100	1100	111

8.5.19 ICC_MSRE, Interrupt Controller Monitor System Register Enable register

The ICC_MSRE characteristics are:

Purpose

Controls whether the System register interface or the memory-mapped interface to the GIC CPU interface is used for EL3.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1, Monitor mode only)	EL3 (SCR.NS=0, Monitor mode only)
-	-	-	-	RW	RW

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	-	-

This register is always System register accessible.

The GIC architecture permits, but does not require, that registers can be shared between memory-mapped registers and the equivalent System registers. This means that if the memory-mapped registers have been accessed while ICC_MSRE.SRE==0, then the System registers might be modified. Therefore, software must only rely on the reset values of the System registers if there has been no use of the GIC functionality while the memory-mapped registers are in use. Otherwise, the System register values must be treated as UNKNOWN.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

This register is only accessible in Secure state.

AArch32 System register ICC_MSRE can be mapped to AArch64 System register ICC_SRE_EL3, but this is not architecturally mandated.

Attributes

ICC_MSRE is a 32-bit register.

Field descriptions

The ICC_MSRE bit assignments are:

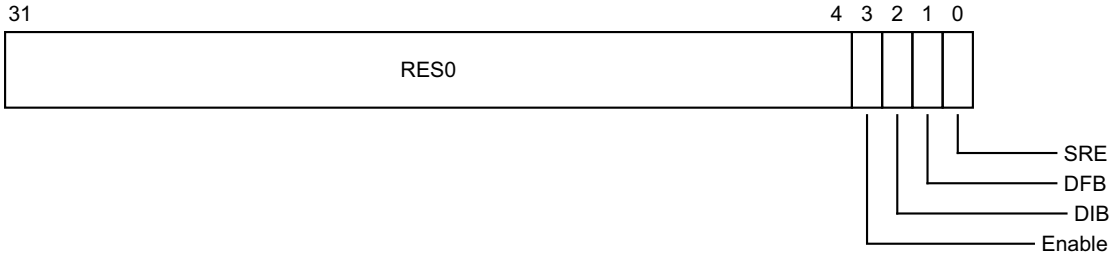

Bits [31:4]

Reserved, RES0.

Enable, bit [3]

Enable. Enables lower Exception level access to ICC_SRE and ICC_HSRE.

`0` Secure EL1 accesses to Secure ICC_SRE trap to EL3.
EL2 accesses to Non-secure ICC_SRE and ICC_HSRE trap to EL3.
Non-secure EL1 accesses to ICC_SRE trap to EL3, unless these accesses are trapped to EL2 as a result of ICC_MSRE.Enable == 0.

`1` Secure EL1 accesses to Secure ICC_SRE do not trap to EL3.
EL2 accesses to Non-secure ICC_SRE and ICC_HSRE do not trap to EL3.
Non-secure EL1 accesses to ICC_SRE do not trap to EL3.

If ICC_MSRE.SRE is RAO/WI, an implementation is permitted to make the Enable bit RAO/WI.

If ICC_MSRE.SRE is 0, the Enable bit behaves as 1 for all purposes other than reading the value of the bit.

DIB, bit [2]

Disable IRQ bypass.

`0` IRQ bypass enabled.

`1` IRQ bypass disabled.

In systems that do not support IRQ bypass, this bit is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

DFB, bit [1]

Disable FIQ bypass.

`0` FIQ bypass enabled.

`1` FIQ bypass disabled.

In systems that do not support FIQ bypass, this bit is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

SRE, bit [0]

System Register Enable.

`0` The memory-mapped interface must be used. Accesses at EL3 or below to any ICH_* System register, or any EL1, EL2, or EL3 ICC_* register other than ICC_SRE, ICC_HSRE, or ICC_MSRE, are UNDEFINED.

`1` The System register interface to the ICH_* registers and the EL1, EL2, and EL3 ICC_* registers is enabled for EL3.

If software changes this bit from 1 to 0, the results are UNPREDICTABLE.

If an implementation supports only a System register interface to the GIC CPU interface, this bit is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

Accessing the ICC_MSRE:

To access the ICC_MSRE:

```
MRC p15,6,<Rt>,c12,c12,5 ; Read ICC_MSRE into Rt
MCR p15,6,<Rt>,c12,c12,5 ; Write Rt to ICC_MSRE
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	110	1100	1100	101

8.5.20 ICC_PMR, Interrupt Controller Interrupt Priority Mask Register

The ICC_PMR characteristics are:

Purpose

Provides an interrupt priority filter. Only interrupts with a higher priority than the value in this register are signaled to the PE.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	RW	RW	RW

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	RW	RW

ICC_PMR is only accessible at Non-secure EL1 when HCR.{FMO, IMO} == {0, 0}.

Note

When HCR.{FMO, IMO} != {0, 0}, at Non-secure EL1, the instruction encoding used to access ICC_PMR results in an access to ICV_PMR.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If ICC_SRE.SRE==0, accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TC==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TC==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If SCR.IRQ==1, and SCR.FIQ==1, accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.IRQ==1, SCR.FIQ==1, HCR.IMO==0, and HCR.FMO==0, Non-secure accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, Secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR.IMO==0, and HCR.FMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR_EL2.IMO==0, and HCR_EL2.FMO==0, Non-secure accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_PMR is architecturally mapped to AArch64 System register ICC_PMR_EL1.

To allow software to ensure appropriate observability of actions initiated by GIC register accesses, the PE and CPU interface logic must ensure that writes to this register are self-synchronising. This ensures that no interrupts below the written PMR value will be taken after a write to this register is architecturally executed. See *Observability of the effects of accesses to the GIC registers* on page 8-159, for more information.

Attributes

ICC_PMR is a 32-bit register.

Field descriptions

The ICC_PMR bit assignments are:

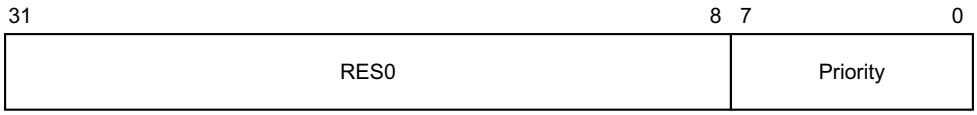

Bits [31:8]

Reserved, RES0.

Priority, bits [7:0]

The priority mask level for the CPU interface. If the priority of an interrupt is higher than the value indicated by this field, the interface signals the interrupt to the PE.

The possible priority field values are as follows:

Implemented priority bits	Possible priority field values	Number of priority levels
[7:0]	0x00-0xFF (0-255), all values	256
[7:1]	0x00-0xFE (0-254), even values only	128
[7:2]	0x00-0xFC (0-252), in steps of 4	64
[7:3]	0x00-0xF8 (0-248), in steps of 8	32
[7:4]	0x00-0xF0 (0-240), in steps of 16	16

Unimplemented priority bits are RAZ/WI.

When this register has an architecturally-defined reset value, this field resets to 0.

Accessing the ICC_PMR:

To access the ICC_PMR:

```
MRC p15,0,<Rt>,c4,c6,0 ; Read ICC_PMR into Rt
MCR p15,0,<Rt>,c4,c6,0 ; Write Rt to ICC_PMR
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	0100	0110	000

When HCR.{FMO, IMO} != {0, 0}, execution of this encoding at Non-secure EL1 results in an access to ICV_PMR.

8.5.21 ICC_RPR, Interrupt Controller Running Priority Register

The ICC_RPR characteristics are:

Purpose

Indicates the Running priority of the CPU interface.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RO	RO	RO	RO

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	RO	RO

ICC_RPR is only accessible at Non-secure EL1 when HCR.{FMO, IMO} == {0, 0}.

Note

When HCR.{FMO, IMO} != {0, 0}, at Non-secure EL1, the instruction encoding used to access ICC_RPR results in an access to ICV_RPR.

If there are no active interrupts on the CPU interface, or all active interrupts have undergone a priority drop, the value returned is the Idle priority.

Software cannot determine the number of implemented priority bits from a read of this register.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure read accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, read accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, read accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, read accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TC==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TC==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If SCR.IRQ==1, and SCR.FIQ==1, read accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR.IRQ==1, SCR.FIQ==1, HCR.IMO==0, and HCR.FMO==0, Non-secure read accesses to this register from EL1 are UNDEFINED.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, Secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, read accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR.IMO==0, and HCR.FMO==0, Non-secure read accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR_EL2.IMO==0, and HCR_EL2.FMO==0, Non-secure read accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_RPR performs the same function as AArch64 System operation ICC_RPR_EL1.

Attributes

ICC_RPR is a 32-bit register.

Field descriptions

The ICC_RPR bit assignments are:

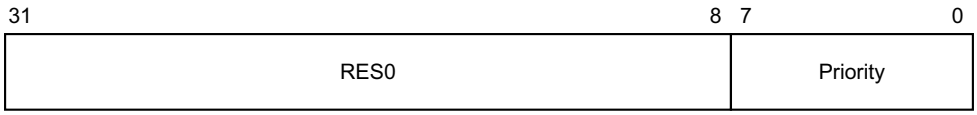

Bits [31:8]

Reserved, RES0.

Priority, bits [7:0]

The current running priority on the CPU interface. This is the group priority of the current active interrupt.

The priority returned is the group priority as if the BPR for the current Exception level and Security state was set to the minimum value of BPR for the number of implemented priority bits.

Note

If 8 bits of priority are implemented the group priority is bits[7:1] of the priority.

Accessing the ICC_RPR:

To access the ICC_RPR:

```
MRC p15,0,<Rt>,c12,c11,3 ; Read ICC_RPR into Rt
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1011	011

When HCR.{FMO, IMO} != {0, 0}, execution of this encoding at Non-secure EL1 results in an access to ICV_RPR.

8.5.22 ICC_SGI0R, Interrupt Controller Software Generated Interrupt Group 0 Register

The ICC_SGI0R characteristics are:

Purpose

Generates Secure Group 0 SGIs, including from the Non-secure state when permitted by GICR_NSACR.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	WO	WO	WO	WO

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	WO	WO

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HCR.FMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HCR_EL2.FMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HCR.IMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HCR_EL2.IMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HSTR.T12==1, Non-secure write accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, write accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, write accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, write accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If SCR.FIQ==1, and SCR.IRQ==1, write accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, Secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, write accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR.IMO==0, and HCR.FMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR_EL2.IMO==0, and HCR_EL2.FMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_SGI0R performs the same function as AArch64 System operation ICC_SGI0R_EL1.

Attributes

ICC_SGI0R is a 64-bit register.

Field descriptions

The ICC_SGI0R bit assignments are:

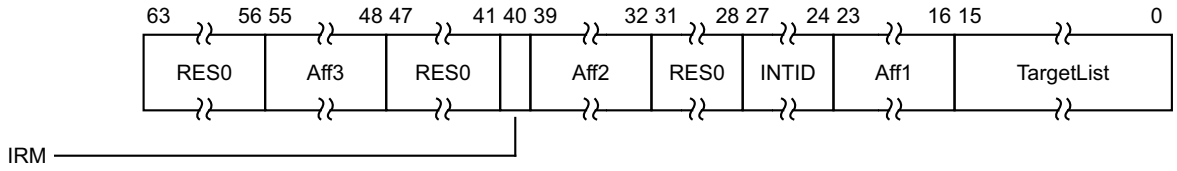

Bits [63:56]

Reserved, RES0.

Aff3, bits [55:48]

The affinity 3 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

Bits [47:41]

Reserved, RES0.

IRM, bit [40]

Interrupt Routing Mode. Determines how the generated interrupts should be distributed to PEs. Possible values are:

`0` Interrupts routed to the PEs specified by Aff3.Aff2.Aff1.<target list>.

`1` Interrupts routed to all PEs in the system, excluding "self".

Aff2, bits [39:32]

The affinity 2 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

Bits [31:28]

Reserved, RES0.

INTID, bits [27:24]

The INTID of the SGI.

Aff1, bits [23:16]

The affinity 1 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

TargetList, bits [15:0]

Target List. The set of PEs for which SGI interrupts will be generated. Each bit corresponds to the PE within a cluster with an Affinity 0 value equal to the bit number.

If a bit is 1 and the bit does not correspond to a valid target PE, the bit must be ignored by the Distributor. It is IMPLEMENTATION DEFINED whether, in such cases, a Distributor can signal a system error.

Note

This restricts a system to sending targeted SGIs to PE with an affinity 0 number of less than 16.

If SRE is set only for secure EL3, software executing at EL3 might use the system register interface to generate SGIs. Hence, the Distributor must always be able to receive and acknowledge Generate SGI packets received from CPU interface regardless of the ARE settings for a security state. However, the Distributor may discard such packets.

If the IRM bit is 1, this field is RES0.

Accessing the ICC_SGI0R:

To access the ICC_SGI0R:

```
MCRR p15,2,<Rt>,<Rt2>,c12 ; Write Rt to ICC_SGI0R[31:0] and Rt2 to ICC_SGI0R[63:32]
```

Register access is encoded as follows:

coproc	opc1	CRm
1111	0010	1100

8.5.23 ICC_SGI1R, Interrupt Controller Software Generated Interrupt Group 1 Register

The ICC_SGI1R characteristics are:

Purpose

Generates Group 1 SGIs for the current Security state.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	WO	WO	WO	WO

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	WO	WO

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HCR.FMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HCR_EL2.FMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HCR.IMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HCR_EL2.IMO==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If HSTR.T12==1, Non-secure write accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, write accesses to this register from EL1 are UNDEFINED.

If ICC_HSRE.SRE==0, write accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, write accesses to this register from EL3 are UNDEFINED.

If ICH_HCR.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If SCR.FIQ==1, and SCR.IRQ==1, write accesses to this register from EL2 and EL3 modes other than Monitor mode are UNDEFINED.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, Secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, and SCR_EL3.IRQ==1, write accesses to this register from EL2 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR.IMO==0, and HCR.FMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

If SCR_EL3.FIQ==1, SCR_EL3.IRQ==1, HCR_EL2.IMO==0, and HCR_EL2.FMO==0, Non-secure write accesses to this register from EL1 are trapped to EL3.

Configurations

There is one instance of this register that is used in both Secure and Non-secure states.

AArch32 System register ICC_SGI1R performs the same function as AArch64 System operation ICC_SGI1R_EL1.

Under certain conditions a write to ICC_SGI1R can generate Group 0 interrupts, see Table 8-14 on page 8-171.

Attributes

ICC_SGI1R is a 64-bit register.

Field descriptions

The ICC_SGI1R bit assignments are:

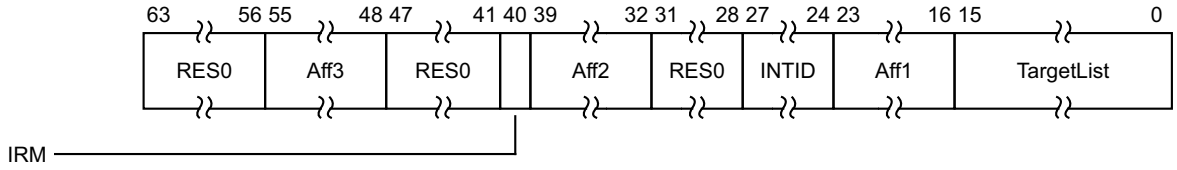

Bits [63:56]

Reserved, RES0.

Aff3, bits [55:48]

The affinity 3 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

Bits [47:41]

Reserved, RES0.

IRM, bit [40]

Interrupt Routing Mode. Determines how the generated interrupts should be distributed to PEs. Possible values are:

`0` Interrupts routed to the PEs specified by Aff3.Aff2.Aff1.<target list>.

`1` Interrupts routed to all PEs in the system, excluding "self".

Aff2, bits [39:32]

The affinity 2 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

Bits [31:28]

Reserved, RES0.

INTID, bits [27:24]

The INTID of the SGI.

Aff1, bits [23:16]

The affinity 1 value of the affinity path of the cluster for which SGI interrupts will be generated.

If the IRM bit is 1, this field is RES0.

TargetList, bits [15:0]

Target List. The set of PEs for which SGI interrupts will be generated. Each bit corresponds to the PE within a cluster with an Affinity 0 value equal to the bit number.

If a bit is 1 and the bit does not correspond to a valid target PE, the bit must be ignored by the Distributor. It is IMPLEMENTATION DEFINED whether, in such cases, a Distributor can signal a system error.

Note

This restricts a system to sending targeted SGIs to PE with an affinity 0 number of less than 16.

If SRE is set only for secure EL3, software executing at EL3 might use the system register interface to generate SGIs. Hence, the Distributor must always be able to receive and acknowledge Generate SGI packets received from CPU interface regardless of the ARE settings for a security state. However, the Distributor may discard such packets.

If the IRM bit is 1, this field is RES0.

Accessing the ICC_SGI1R:

To access the ICC_SGI1R:

```
MCRR p15,0,<Rt>,<Rt2>,c12 ; Write Rt to ICC_SGI1R[31:0] and Rt2 to ICC_SGI1R[63:32]
```

Register access is encoded as follows:

coproc	opc1	CRm
1111	0000	1100

8.5.24 ICC_SRE, Interrupt Controller System Register Enable register

The ICC_SRE characteristics are:

Purpose

Controls whether the System register interface or the memory-mapped interface to the GIC CPU interface is used for EL0 and EL1.

Usage constraints

ICC_SRE(S) is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	-	-	RW

ICC_SRE(NS) is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	RW	RW	-

The GIC architecture permits, but does not require, that registers can be shared between memory-mapped registers and the equivalent System registers. This means that if the memory-mapped registers have been accessed while ICC_SRE.SRE==0, then the System registers might be modified. Therefore, software must only rely on the reset values of the System registers if there has been no use of the GIC functionality while the memory-mapped registers are in use. Otherwise, the System register values must be treated as UNKNOWN.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_HSRE.Enable==0, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE_EL2.Enable==0, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_MSRE.Enable==0, Secure accesses to this register from EL1 are trapped to EL3.

If ICC_MSRE.Enable==0, accesses to this register from EL2 are trapped to EL3.

If ICC_MSRE.Enable==0, Non-secure accesses to this register from EL1 are trapped to EL3 unless already trapped to EL2.

If ICC_SRE_EL3.Enable==0, Secure accesses to this register from EL1 are trapped to EL3.

If ICC_SRE_EL3.Enable==0, accesses to this register from EL2 are trapped to EL3.

If ICC_SRE_EL3.Enable==0, Non-secure accesses to this register from EL1 are trapped to EL3 unless already trapped to EL2.

Configurations

AArch32 System register ICC_SRE(S) is architecturally mapped to AArch64 System register ICC_SRE_EL1 (S).

AArch32 System register ICC_SRE(NS) is architecturally mapped to AArch64 System register ICC_SRE_EL1 (NS).

Attributes

ICC_SRE is a 32-bit register.

Field descriptions

The ICC_SRE bit assignments are:

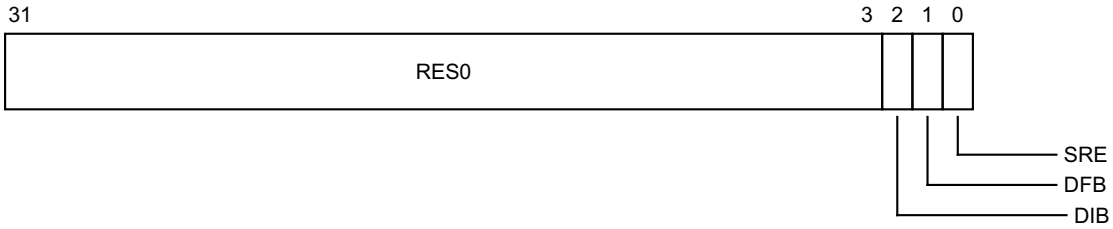

Bits [31:3]

Reserved, RES0.

DIB, bit [2]

Disable IRQ bypass.

0 IRQ bypass enabled.

1 IRQ bypass disabled.

If EL3 is implemented and GICD_CTLR.DS == 0, this field is a read-only alias of ICC_MSRE.DIB.

If EL3 is implemented and GICD_CTLR.DS == 1, and EL2 is not implemented, this field is a read-write alias of ICC_MSRE.DIB.

If EL3 is not implemented or GICD_CTLR.DS == 1, and EL2 is implemented, this field is a read-only alias of ICC_HSRE.DIB.

In systems that do not support IRQ bypass, this field is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

DFB, bit [1]

Disable FIQ bypass.

0 FIQ bypass enabled.

1 FIQ bypass disabled.

If EL3 is implemented and GICD_CTLR.DS == 0, this field is a read-only alias of ICC_MSRE.DFB.

If EL3 is implemented and GICD_CTLR.DS == 1, and EL2 is not implemented, this field is a read-write alias of ICC_MSRE.DFB.

If EL3 is not implemented or GICD_CTLR.DS == 1, and EL2 is implemented, this field is a read-only alias of ICC_HSRE.DFB.

In systems that do not support FIQ bypass, this field is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

SRE, bit [0]

System Register Enable.

0 The memory-mapped interface must be used. Accesses at EL1 to any ICC_* System register other than ICC_SRE are UNDEFINED.

1 The System register interface for the current Security state is enabled.

If software changes this bit from 1 to 0 in the Secure instance of this register, the results are UNPREDICTABLE.

If an implementation supports only a System register interface to the GIC CPU interface, this bit is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

Accessing the ICC_SRE:

To access the ICC_SRE:

```
MRC p15,0,<Rt>,c12,c12,5 ; Read ICC_SRE into Rt
MCR p15,0,<Rt>,c12,c12,5 ; Write Rt to ICC_SRE
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	101

8.6 AArch32 System register descriptions of the virtual registers

This section describes each of the virtual AArch32 GIC System registers in register name order. The ICV prefix indicates a virtual GIC CPU interface System register. Each AArch32 System register description contains a reference to the AArch64 register that provides the same functionality.

Unless otherwise stated, the bit assignments for the GIC System registers are the same as those for the equivalent GICC_* and GICV_* memory-mapped registers.

The ICV_* registers are only accessible at Non-secure EL1. Whether an access encoding maps to an ICC_* register or the equivalent ICV_* register is determined by HCR, see Chapter 5 *Virtual Interrupt Handling and Prioritization*. The equivalent virtual interface memory-mapped registers are described in *The GIC virtual CPU interface register descriptions* on page 8-587.

The encodings for the virtual registers are the same as for the physical registers, see Table 8-23 on page 8-298.

8.6.1 ICV_AP0R<n>, Interrupt Controller Virtual Active Priorities Group 0 Registers, n = 0 - 3

The ICV_AP0R<n> characteristics are:

Purpose

Provides information about virtual Group 0 active priorities.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	RW	-	-

The ICV_AP0R<n> registers are only accessible at Non-secure EL1 when *HCR* on page 1-25.FMO is set to 1.

Note

When *HCR* on page 1-25.FMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_AP0R<n> results in an access to ICC_AP0R<n>.

Writing to these registers with any value other than the last read value of the register (or `0x00000000` when there are no Group 0 active priorities) might result in UNPREDICTABLE behavior of the virtual interrupt prioritization system, causing:

- Interrupts that should preempt execution to not preempt execution.
- Interrupts that should not preempt execution to preempt execution.

ICV_AP0R1 is only implemented in implementations that support 6 or more bits of priority. ICV_AP0R2 and ICV_AP0R3 are only implemented in implementations that support 7 bits of priority. Unimplemented registers are UNDEFINED.

Writing to the active priority registers in any order other than the following order might result in UNPREDICTABLE behavior of the interrupt prioritization system:

- ICV_AP0R<n>.
- ICV_AP1R<n>.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_AP0R<n> is architecturally mapped to AArch64 System register ICV_AP0R<n>_EL1.

Attributes

ICV_AP0R<n> is a 32-bit register.

Field descriptions

The ICV_AP0R<n> bit assignments are:

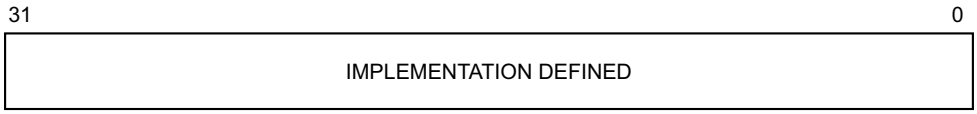

IMPLEMENTATION DEFINED, bits [31:0]

IMPLEMENTATION DEFINED.

When this register has an architecturally-defined reset value, this field resets to `0`.

The contents of these registers are IMPLEMENTATION DEFINED with the one architectural requirement that the value `0x00000000` is consistent with no interrupts being active.

Accessing the ICV_AP0R<n>:

To access the ICV_AP0R<n> when HCR.FMO is set to 1, and executing at Non-secure EL1:

```
MRC p15,0,<Rt>,c12,c8,<opc2> ; Read ICV_AP0R<n> into Rt, where n is in the range 0 to 3
MCR p15,0,<Rt>,c12,c8,<opc2> ; Write Rt to ICV_AP0R<n>, where n is in the range 0 to 3
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1000	1:n<1:0>

When *HCR* on page 1-25.FMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_AP0R<n>.

8.6.2 ICV_AP1R<n>, Interrupt Controller Virtual Active Priorities Group 1 Registers, n = 0 - 3

The ICV_AP1R<n> characteristics are:

Purpose

Provides information about virtual Group 1 active priorities.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	RW	-	-

The ICV_AP1R<n> registers are only accessible at Non-secure EL1 when *HCR* on page 1-25.IMO == 1.

Note

When *HCR* on page 1-25.IMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_AP1R<n> results in an access to ICC_AP1R<n>.

Writing to these registers with any value other than the last read value of the register (or `0x00000000` when there are no Group 1 active priorities) might result in UNPREDICTABLE behavior of the virtual interrupt prioritization system, causing:

- Interrupts that should preempt execution to not preempt execution.
- Interrupts that should not preempt execution to preempt execution.

ICV_AP1R1 is only implemented in implementations that support 6 or more bits of priority. ICV_AP1R2 and ICV_AP1R3 are only implemented in implementations that support 7 bits of priority. Unimplemented registers are UNDEFINED.

Writing to the active priority registers in any order other than the following order might result in UNPREDICTABLE behavior of the interrupt prioritization system:

- ICV_AP0R<n>.
- ICV_AP1R<n>.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_AP1R<n> is architecturally mapped to AArch64 System register ICV_AP1R<n>_EL1.

Attributes

ICV_AP1R<n> is a 32-bit register.

Field descriptions

The ICV_AP1R<n> bit assignments are:

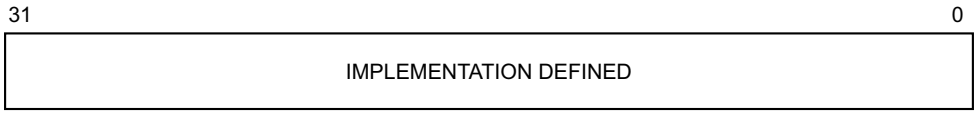

IMPLEMENTATION DEFINED, bits [31:0]

IMPLEMENTATION DEFINED.

When this register has an architecturally-defined reset value, this field resets to `0`.

The contents of these registers are IMPLEMENTATION DEFINED with the one architectural requirement that the value `0x00000000` is consistent with no interrupts being active.

Accessing the ICV_AP1R<n>:

To access the ICV_AP1R<n> when HCR.IMO is set to 1, and executing at Non-secure EL1:

```
MRC p15,0,<Rt>,c12,c9,<opc2> ; Read ICV_AP1R<n> into Rt, where n is in the range 0 to 3
MCR p15,0,<Rt>,c12,c9,<opc2> ; Write Rt to ICV_AP1R<n>, where n is in the range 0 to 3
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1001	0:n<1:0>

When *HCR* on page 1-25.IMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_AP1R<n>.

8.6.3 ICV_BPR0, Interrupt Controller Virtual Binary Point Register 0

The ICV_BPR0 characteristics are:

Purpose

Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines virtual Group 0 interrupt preemption.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	RW	-	-

ICV_BPR0 is only accessible at Non-secure EL1 when *HCR* on page 1-25.FMO is set to 1.

Note

When *HCR* on page 1-25.FMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_BPR0 results in an access to ICC_BPR0.

The minimum binary point value is derived from the number of implemented priority bits. The number of priority bits is IMPLEMENTATION DEFINED, and reported by ICV_CTLR.PRIbits.

An attempt to program the binary point field to a value less than the minimum value sets the field to the minimum value. On a reset, the binary point field is set to the minimum supported value.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_BPR0 is architecturally mapped to AArch64 System register ICV_BPR0_EL1.

Attributes

ICV_BPR0 is a 32-bit register.

Field descriptions

The ICV_BPR0 bit assignments are:

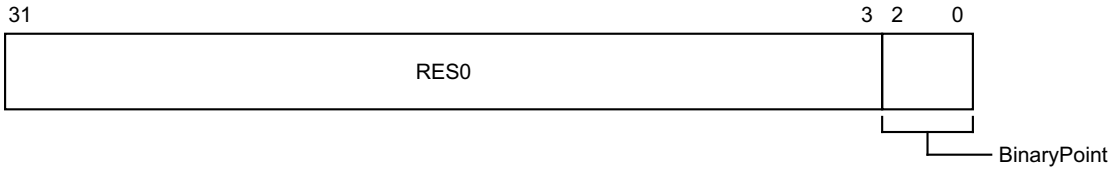

Bits [31:3]

Reserved, RES0.

BinaryPoint, bits [2:0]

The value of this field controls how the 8-bit interrupt priority field is split into a group priority field, that determines interrupt preemption, and a subpriority field. This is done as follows:

Binary point value	Group priority field	Subpriority field	Field with binary point
0	[7:1]	[0]	ggggggg.s
1	[7:2]	[1:0]	gggggg.ss
2	[7:3]	[2:0]	ggggg.sss
3	[7:4]	[3:0]	gggg.ssss
4	[7:5]	[4:0]	ggg.sssss
5	[7:6]	[5:0]	gg.ssssss
6	[7]	[6:0]	g.sssssss
7	No preemption	[7:0]	.ssssssss

Accessing the ICV_BPR0:

To access the ICV_BPR0 when HCR.FMO is set to 1, and executing at Non-secure EL1:

```
MRC p15,0,<Rt>,c12,c8,3 ; Read ICV_BPR0 into Rt
MCR p15,0,<Rt>,c12,c8,3 ; Write Rt to ICV_BPR0
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1000	011

When *HCR* on page 1-25.FMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_BPR0.

8.6.4 ICV_BPR1, Interrupt Controller Virtual Binary Point Register 1

The ICV_BPR1 characteristics are:

Purpose

Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines virtual Group 1 interrupt preemption.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	RW	-	-

ICV_BPR1 is only accessible at Non-secure EL1 when *HCR* on page 1-25.IMO is set to 1.

Note

When *HCR* on page 1-25.IMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_BPR1 results in an access to ICC_BPR1.

The reset value is IMPLEMENTATION DEFINED, but is equal to the minimum value of ICV_BPR0 plus one.

An attempt to program the binary point field to a value less than the reset value sets the field to the reset value.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_BPR1 is architecturally mapped to AArch64 System register ICV_BPR1_EL1.

Attributes

ICV_BPR1 is a 32-bit register.

Field descriptions

The ICV_BPR1 bit assignments are:

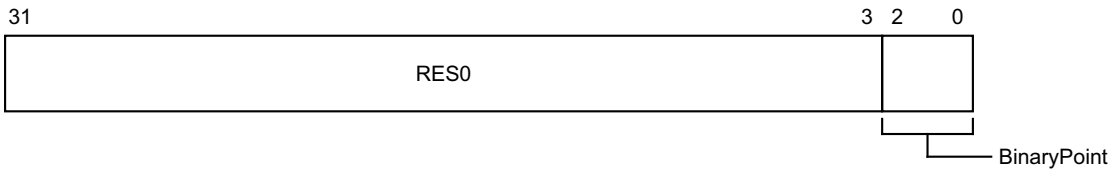

Bits [31:3]

Reserved, RES0.

BinaryPoint, bits [2:0]

If the GIC is configured to use separate binary point fields for virtual Group 0 and virtual Group 1 interrupts, the value of this field controls how the 8-bit interrupt priority field is split into a group priority field, that determines interrupt preemption, and a subpriority field. This is done as follows:

Binary point value	Group priority field	Subpriority field	Field with binary point
0	-	-	-
1	[7:1]	[0]	ggggggg.s
2	[7:2]	[1:0]	gggggg.ss
3	[7:3]	[2:0]	ggggg.sss
4	[7:4]	[3:0]	gggg.ssss
5	[7:5]	[4:0]	ggg.sssss
6	[7:6]	[5:0]	gg.ssssss
7	[7]	[6:0]	g.sssssss

Writing 0 to this field will set this field to its reset value, which is IMPLEMENTATION DEFINED and non-zero.

If ICV_CTLR.CBPR is set to 1, Non-secure EL1 reads return ICV_BPR0 + 1 saturated to `0b111`. Non-secure EL1 writes are ignored.

Accessing the ICV_BPR1:

To access the ICV_BPR1 when HCR.IMO is set to 1, and executing at Non-secure EL1:

```
MRC p15,0,<Rt>,c12,c12,3 ; Read ICV_BPR1 into Rt
MCR p15,0,<Rt>,c12,c12,3 ; Write Rt to ICV_BPR1
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	011

When *HCR* on page 1-25.IMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_BPR1.

8.6.5 ICV_CTLR, Interrupt Controller Virtual Control Register

The ICV_CTLR characteristics are:

Purpose

Controls aspects of the behavior of the GIC virtual CPU interface and provides information about the features implemented.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	RW	-	-

ICV_CTLR is only accessible at Non-secure EL1 when *HCR* on page 1-25.{FMO, IMO} != {0, 0}.

Note

When *HCR* on page 1-25.{FMO, IMO} == {0, 0}, at Non-secure EL1, the instruction encoding used to access ICV_CTLR results in an access to ICC_CTLR.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TC==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TC==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_CTLR is architecturally mapped to AArch64 System register ICV_CTLR_EL1.

Attributes

ICV_CTLR is a 32-bit register.

Field descriptions

The ICV_CTLR bit assignments are:

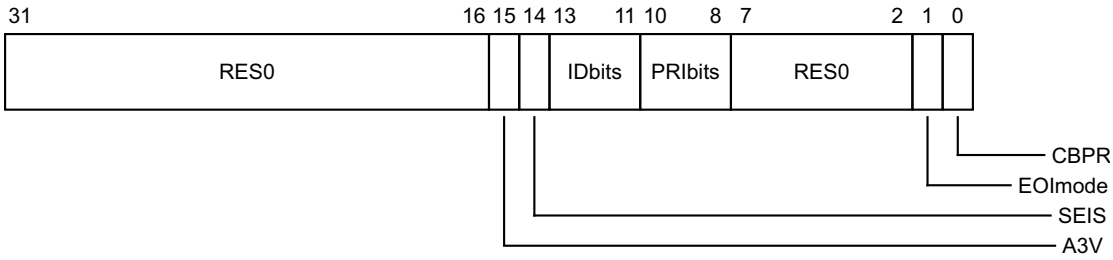

Bits [31:16]

Reserved, RES0.

A3V, bit [15]

Affinity 3 Valid. Read-only and writes are ignored. Possible values are:

`0` The virtual CPU interface logic only supports zero values of Affinity 3 in SGI generation System registers.

`1` The virtual CPU interface logic supports non-zero values of Affinity 3 in SGI generation System registers.

SEIS, bit [14]

SEI Support. Read-only and writes are ignored. Indicates whether the virtual CPU interface supports local generation of SEIs:

`0` The virtual CPU interface logic does not support local generation of SEIs.

`1` The virtual CPU interface logic supports local generation of SEIs.

IDbits, bits [13:11]

Identifier bits. Read-only and writes are ignored. The number of virtual interrupt identifier bits supported:

`000` 16 bits.

`001` 24 bits.

All other values are reserved.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

PRIbits, bits [10:8]

Priority bits. Read-only and writes are ignored. The number of priority bits implemented, minus one.

An implementation must implement at least 32 levels of physical priority (5 priority bits).

Note

This field always returns the number of priority bits implemented.

The division between group priority and subpriority is defined in the binary point registers ICV_BPR0 and ICV_BPR1.

Bits [7:2]

Reserved, RES0.

EOImode, bit [1]

Virtual EOI mode. Controls whether a write to an End of Interrupt register also deactivates the virtual interrupt:

0 ICV_EOIR0 and ICV_EOIR1 provide both priority drop and interrupt deactivation functionality. Accesses to ICV_DIR are UNPREDICTABLE.

1 ICV_EOIR0 and ICV_EOIR1 provide priority drop functionality only. ICV_DIR provides interrupt deactivation functionality.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

CBPR, bit [0]

Common Binary Point Register. Controls whether the same register is used for interrupt preemption of both virtual Group 0 and virtual Group 1 interrupts:

0 ICV_BPR0 determines the preemption group for virtual Group 0 interrupts only.
ICV_BPR1 determines the preemption group for virtual Group 1 interrupts.

1 ICV_BPR0 determines the preemption group for both virtual Group 0 and virtual Group 1 interrupts.
Reads of ICV_BPR1 return ICV_BPR0 plus one, saturated to 0b111. Writes to ICV_BPR1 are ignored.

Accessing the ICV_CTLR:

To access the ICV_CTLR when HCR.{FMO, IMO} != {0, 0}, and executing at Non-secure EL1:

```
MRC p15,0,<Rt>,c12,c12,4 ; Read ICV_CTLR into Rt
MCR p15,0,<Rt>,c12,c12,4 ; Write Rt to ICV_CTLR
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	100

When *HCR* on page 1-25.{FMO, IMO} == {0, 0}, execution of this encoding at Non-secure EL1 results in an access to ICC_CTLR.

8.6.6 ICV_DIR, Interrupt Controller Deactivate Virtual Interrupt Register

The ICV_DIR characteristics are:

Purpose

When interrupt priority drop is separated from interrupt deactivation, a write to this register deactivates the specified virtual interrupt.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	WO	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	WO	-	-

The ICV_DIR register is only accessible at Non-secure EL1 in the following cases:

- When *HCR* on page 1-25.FMO is set to 1.
- When *HCR* on page 1-25.IMO is set to 1.

Note

At Non-secure EL1, the instruction encoding used to access ICV_DIR results in an access to ICC_DIR when *HCR* on page 1-25.{FMO, IMO} == {0, 0}.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure write accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure write accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure write accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TC==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_DIR performs the same function as AArch64 System operation ICV_DIR_EL1.

Attributes

ICV_DIR is a 32-bit register.

Field descriptions

The ICV_DIR bit assignments are:

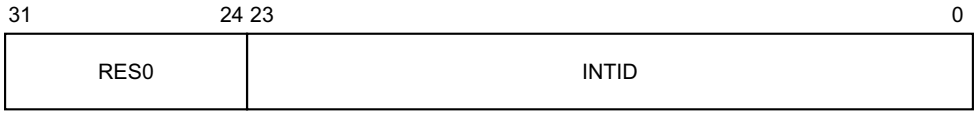

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the virtual interrupt to be deactivated.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICV_CTLR.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICV_DIR:

To access the ICV_DIR:

```
MCR p15,0,<Rt>,c12,c11,1 ; Write Rt to ICV_DIR
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1011	001

This encoding results in an access to ICV_DIR at Non-secure EL1 in the following cases:

- When *HCR* on page 1-25.FMO is set to 1, and the INTID field refers to a Group 0 interrupt.
- When *HCR* on page 1-25.IMO is set to 1, and the INTID field refers to a Group 1 interrupt.

This encoding results in an access to ICC_DIR at Non-secure EL1 in the following cases:

- When HCR2.{FMO, IMO} == {0, 0}.
- When *HCR* on page 1-25.FMO is set to 1, and the INTID field does not refer to a Group 0 interrupt.
- When *HCR* on page 1-25.IMO is set to 1, and the INTID field does not refer to a Group 1 interrupt.

8.6.7 ICV_EOIR0, Interrupt Controller Virtual End Of Interrupt Register 0

The ICV_EOIR0 characteristics are:

Purpose

A PE writes to this register to inform the CPU interface that it has completed the processing of the specified virtual Group 0 interrupt.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	WO	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	WO	-	-

ICV_EOIR0 is only accessible at Non-secure EL1 when *HCR* on page 1-25.FMO is set to 1.

Note

When *HCR* on page 1-25.FMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_EOIR0 results in an access to ICC_EOIR0.

A write to this register must correspond to the most recent valid read from an Interrupt Acknowledge Register. A valid read is a read that returns a valid interrupt ID, that is not a special INTID.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure write accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure write accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure write accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TALL0==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL0==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_EOIR0 performs the same function as AArch64 System operation ICV_EOIR0_EL1.

Attributes

ICV_EOIR0 is a 32-bit register.

Field descriptions

The ICV_EOIR0 bit assignments are:

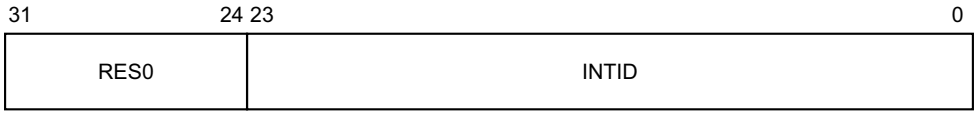

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID from the corresponding ICV_IAR0 access.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICV_CTLR.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

If the ICV_CTLR.EOImode bit is 0, a write to this register drops the priority for the virtual interrupt, and also deactivates the virtual interrupt.

If the ICV_CTLR.EOImode bit is 1, a write to this register only drops the priority for the virtual interrupt. Software must write to ICV_DIR to deactivate the virtual interrupt.

Accessing the ICV_EOIR0:

To access the ICV_EOIR0 when HCR.FMO is set to 1, and executing at Non-secure EL1:

```
MCR p15,0,<Rt>,c12,c8,1 ; Write Rt to ICV_EOIR0
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1000	001

When *HCR* on page 1-25.FMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_EOIR0.

8.6.8 ICV_EOIR1, Interrupt Controller Virtual End Of Interrupt Register 1

The ICV_EOIR1 characteristics are:

Purpose

A PE writes to this register to inform the CPU interface that it has completed the processing of the specified virtual Group 1 interrupt.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	WO	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	WO	-	-

ICV_EOIR1 is only accessible at Non-secure EL1 when *HCR* on page 1-25.IMO is set to 1.

Note

When *HCR* on page 1-25.IMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_EOIR1 results in an access to ICC_EOIR1.

A write to this register must correspond to the most recent valid read from an Interrupt Acknowledge Register. A valid read is a read that returns a valid interrupt ID, that is not a special INTID.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure write accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure write accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure write accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TALL1==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL1==1, Non-secure write accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_EOIR1 performs the same function as AArch64 System operation ICV_EOIR1_EL1.

Attributes

ICV_EOIR1 is a 32-bit register.

Field descriptions

The ICV_EOIR1 bit assignments are:

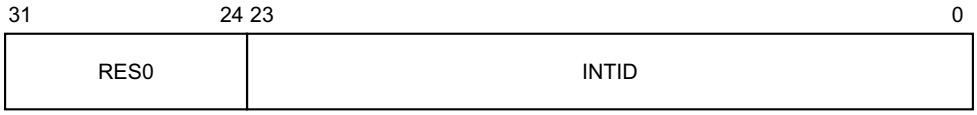

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID from the corresponding ICV_IAR1 access.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICV_CTLR.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

If the ICV_CTLR.EOImode bit is 0, a write to this register drops the priority for the virtual interrupt, and also deactivates the virtual interrupt.

If the ICV_CTLR.EOImode bit is 1, a write to this register only drops the priority for the virtual interrupt. Software must write to ICV_DIR to deactivate the virtual interrupt.

Accessing the ICV_EOIR1:

To access the ICV_EOIR1 when HCR.IMO is set to 1, and executing at Non-secure EL1:

```
MCR p15,0,<Rt>,c12,c12,1 ; Write Rt to ICV_EOIR1
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	001

When *HCR on page 1-25*.IMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_EOIR1.

8.6.9 ICV_HPPIR0, Interrupt Controller Virtual Highest Priority Pending Interrupt Register 0

The ICV_HPPIR0 characteristics are:

Purpose

Indicates the highest priority pending virtual Group 0 interrupt on the virtual CPU interface.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RO	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	RO	-	-

ICV_HPPIR0 is only accessible at Non-secure EL1 when *HCR* on page 1-25.FMO is set to 1.

Note

When *HCR* on page 1-25.FMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_HPPIR0 results in an access to ICC_HPPIR0.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure read accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure read accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure read accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TALL0==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL0==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_HPPIR0 performs the same function as AArch64 System operation ICV_HPPIR0_EL1.

Attributes

ICV_HPPIR0 is a 32-bit register.

Field descriptions

The ICV_HPPIR0 bit assignments are:

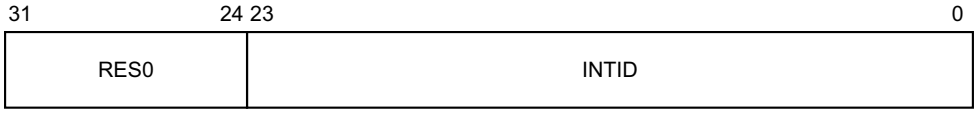

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the highest priority pending virtual interrupt.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. This special INTID can take the value 1023 only. See special interrupt, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICV_CTLR.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICV_HPPIR0:

To access the ICV_HPPIR0 when HCR.FMO is set to 1, and executing at Non-secure EL1:

```
MRC p15,0,<Rt>,c12,c8,2 ; Read ICV_HPPIR0 into Rt
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1000	010

When *HCR* on page 1-25.FMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_HPPIR0.

8.6.10 ICV_HPPIR1, Interrupt Controller Virtual Highest Priority Pending Interrupt Register 1

The ICV_HPPIR1 characteristics are:

Purpose

Indicates the highest priority pending virtual Group 1 interrupt on the virtual CPU interface.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RO	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	RO	-	-

ICV_HPPIR1 is only accessible at Non-secure EL1 when *HCR* on page 1-25.IMO is set to 1.

Note

When *HCR* on page 1-25.IMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_HPPIR1 results in an access to ICC_HPPIR1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure read accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure read accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure read accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TALL1==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL1==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_HPPIR1 performs the same function as AArch64 System operation ICV_HPPIR1_EL1.

Attributes

ICV_HPPIR1 is a 32-bit register.

Field descriptions

The ICV_HPPIR1 bit assignments are:

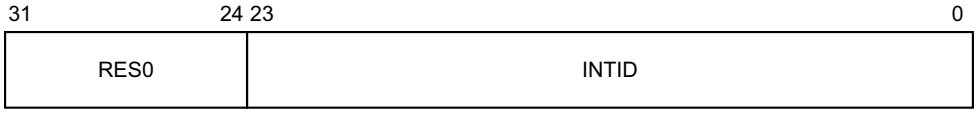

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the highest priority pending virtual interrupt.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. This special INTID can take the value 1023 only. See special interrupt, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICV_CTLR.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICV_HPPIR1:

To access the ICV_HPPIR1 when HCR.IMO is set to 1, and executing at Non-secure EL1:

```
MRC p15,0,<Rt>,c12,c12,2 ; Read ICV_HPPIR1 into Rt
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	010

When *HCR* on page 1-25.IMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_HPPIR1.

8.6.11 ICV_IAR0, Interrupt Controller Virtual Interrupt Acknowledge Register 0

The ICV_IAR0 characteristics are:

Purpose

The PE reads this register to obtain the INTID of the signaled virtual Group 0 interrupt. This read acts as an acknowledge for the interrupt.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RO	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	RO	-	-

ICV_IAR0 is only accessible at Non-secure EL1 when *HCR* on page 1-25.FMO is set to 1.

Note

When *HCR* on page 1-25.FMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_IAR0 results in an access to ICC_IAR0.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure read accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure read accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure read accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TALL0==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL0==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_IAR0 performs the same function as AArch64 System operation ICV_IAR0_EL1.

To allow software to ensure appropriate observability of actions initiated by GIC register accesses, the PE and CPU interface logic must ensure that reads of this register are self-synchronising when interrupts are masked by the PE (that is when PSTATE.{I,F} == {0,0}). This ensures that the effect of activating an interrupt on the signaling of interrupt exceptions is observed when a read of this register is architecturally executed so that no spurious interrupt exception occurs if interrupts are unmasked by an instruction immediately following the read. See *Observability of the effects of accesses to the GIC registers* on page 8-159, for more information.

Attributes

ICV_IAR0 is a 32-bit register.

Field descriptions

The ICV_IAR0 bit assignments are:

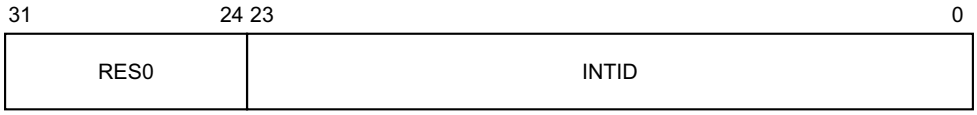

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled virtual interrupt.

This is the INTID of the highest priority pending virtual interrupt, if that interrupt is of sufficient priority for it to be signaled to the PE, and if it can be acknowledged.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. This special INTID can take the value 1023 only. See special interrupt, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICV_CTLR.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICV_IAR0:

To access the ICV_IAR0 when HCR.FMO is set to 1, and executing at Non-secure EL1:

```
MRC p15,0,<Rt>,c12,c8,0 ; Read ICV_IAR0 into Rt
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1000	000

When *HCR on page 1-25*.FMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_IAR0.

8.6.12 ICV_IAR1, Interrupt Controller Virtual Interrupt Acknowledge Register 1

The ICV_IAR1 characteristics are:

Purpose

The PE reads this register to obtain the INTID of the signaled virtual Group 1 interrupt. This read acts as an acknowledge for the interrupt.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RO	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	RO	-	-

ICV_IAR1 is only accessible at Non-secure EL1 when *HCR* on page 1-25.IMO is set to 1.

Note

When *HCR* on page 1-25.IMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_IAR1 results in an access to ICC_IAR1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure read accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure read accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure read accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TALL1==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL1==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_IAR1 performs the same function as AArch64 System operation ICV_IAR1_EL1.

To allow software to ensure appropriate observability of actions initiated by GIC register accesses, the PE and CPU interface logic must ensure that reads of this register are self-synchronising when interrupts are masked by the PE (that is when PSTATE.{I,F} == {0,0}). This ensures that the effect of activating an interrupt on the signaling of interrupt exceptions is observed when a read of this register is architecturally executed so that no spurious interrupt exception occurs if interrupts are unmasked by an instruction immediately following the read. See *Observability of the effects of accesses to the GIC registers* on page 8-159, for more information.

Attributes

ICV_IAR1 is a 32-bit register.

Field descriptions

The ICV_IAR1 bit assignments are:

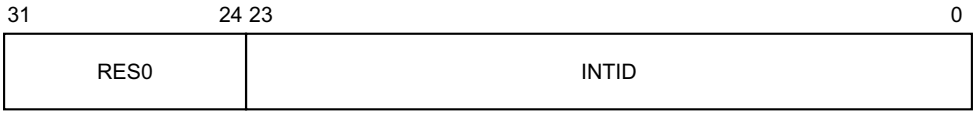

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled virtual interrupt.

This is the INTID of the highest priority pending virtual interrupt, if that interrupt is of sufficient priority for it to be signaled to the PE, and if it can be acknowledged.

If the highest priority pending interrupt is not observable, this field contains a special INTID to indicate the reason. This special INTID can take the value 1023 only. See special interrupt, for more information.

This field has either 16 or 24 bits implemented. The number of implemented bits can be found in ICV_CTLR.IDbits. If only 16 bits are implemented, bits [23:16] of this register are RES0.

Accessing the ICV_IAR1:

To access the ICV_IAR1 when HCR.IMO is set to 1, and executing at Non-secure EL1:

```
MRC p15,0,<Rt>,c12,c12,0 ; Read ICV_IAR1 into Rt
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	000

When *HCR* on page 1-25.IMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_IAR1.

8.6.13 ICV_IGRPEN0, Interrupt Controller Virtual Interrupt Group 0 Enable register

The ICV_IGRPEN0 characteristics are:

Purpose

Controls whether virtual Group 0 interrupts are enabled or not.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	RW	-	-

ICV_IGRPEN0 is only accessible at Non-secure EL1 when *HCR* on page 1-25.FMO is set to 1.

Note

When *HCR* on page 1-25.FMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_IGRPEN0 results in an access to ICC_IGRPEN0.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL0==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_IGRPEN0 is architecturally mapped to AArch64 System register ICV_IGRPEN0_EL1.

Attributes

ICV_IGRPEN0 is a 32-bit register.

Field descriptions

The ICV_IGRPEN0 bit assignments are:

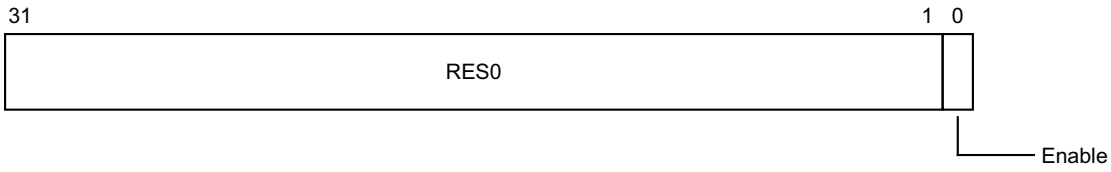

Bits [31:1]

Reserved, RES0.

Enable, bit [0]

Enables virtual Group 0 interrupts.

0 Virtual Group 0 interrupts are disabled.

1 Virtual Group 0 interrupts are enabled.

When this register has an architecturally-defined reset value, this field resets to 0.

Accessing the ICV_IGRPEN0:

To access the ICV_IGRPEN0 when HCR.FMO is set to 1, and executing at Non-secure EL1:

```
MRC p15,0,<Rt>,c12,c12,6 ; Read ICV_IGRPEN0 into Rt
MCR p15,0,<Rt>,c12,c12,6 ; Write Rt to ICV_IGRPEN0
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	110

When *HCR* on page 1-25.FMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_IGRPEN0.

8.6.14 ICV_IGRPEN1, Interrupt Controller Virtual Interrupt Group 1 Enable register

The ICV_IGRPEN1 characteristics are:

Purpose

Controls whether virtual Group 1 interrupts are enabled for the current Security state.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	RW	-	-

ICV_IGRPEN1 is only accessible at Non-secure EL1 when *HCR* on page 1-25.IMO is set to 1.

Note

When *HCR* on page 1-25.IMO is set to 0, at Non-secure EL1, the instruction encoding used to access ICV_IGRPEN1 results in an access to ICC_IGRPEN1.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TALL1==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_IGRPEN1 is architecturally mapped to AArch64 System register ICV_IGRPEN1_EL1.

Attributes

ICV_IGRPEN1 is a 32-bit register.

Field descriptions

The ICV_IGRPEN1 bit assignments are:

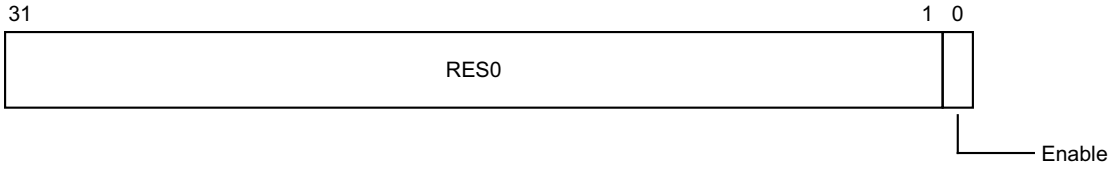

Bits [31:1]

Reserved, RES0.

Enable, bit [0]

Enables virtual Group 1 interrupts.

0 Virtual Group 1 interrupts are disabled.

1 Virtual Group 1 interrupts are enabled.

When this register has an architecturally-defined reset value, this field resets to 0.

Accessing the ICV_IGRPEN1:

To access the ICV_IGRPEN1 when HCR.IMO is set to 1, and executing at Non-secure EL1:

```
MRC p15,0,<Rt>,c12,c12,7 ; Read ICV_IGRPEN1 into Rt
MCR p15,0,<Rt>,c12,c12,7 ; Write Rt to ICV_IGRPEN1
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1100	111

When *HCR* on page 1-25.IMO is set to 0, execution of this encoding at Non-secure EL1 results in an access to ICC_IGRPEN1.

8.6.15 ICV_PMR, Interrupt Controller Virtual Interrupt Priority Mask Register

The ICV_PMR characteristics are:

Purpose

Provides a virtual interrupt priority filter. Only virtual interrupts with a higher priority than the value in this register are signaled to the PE.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RW	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	RW	-	-

ICV_PMR is only accessible at Non-secure EL1 when *HCR* on page 1-25.{FMO, IMO} != {0, 0}.

Note

When *HCR* on page 1-25.{FMO, IMO} == {0, 0}, at Non-secure EL1, the instruction encoding used to access ICV_PMR results in an access to ICC_PMR.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If ICC_SRE.SRE==0, Non-secure accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TC==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TC==1, Non-secure accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_PMR is architecturally mapped to AArch64 System register ICV_PMR_EL1.

To allow software to ensure appropriate observability of actions initiated by GIC register accesses, the PE and CPU interface logic must ensure that writes to this register are self-synchronising. This ensures that no interrupts below the written PMR value will be taken after a write to this register is architecturally executed. See *Observability of the effects of accesses to the GIC registers* on page 8-159, for more information.

Attributes

ICV_PMR is a 32-bit register.

Field descriptions

The ICV_PMR bit assignments are:

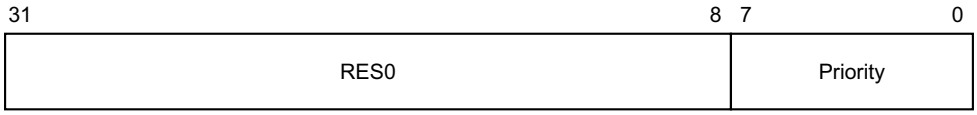

Bits [31:8]

Reserved, RES0.

Priority, bits [7:0]

The priority mask level for the virtual CPU interface. If the priority of a virtual interrupt is higher than the value indicated by this field, the interface signals the virtual interrupt to the PE.

The possible priority field values are as follows:

Implemented priority bits	Possible priority field values	Number of priority levels
[7:0]	0x00-0xFF (0-255), all values	256
[7:1]	0x00-0xFE (0-254), even values only	128
[7:2]	0x00-0xFC (0-252), in steps of 4	64
[7:3]	0x00-0xF8 (0-248), in steps of 8	32
[7:4]	0x00-0xF0 (0-240), in steps of 16	16

Unimplemented priority bits are RAZ/WI.

When this register has an architecturally-defined reset value, this field resets to 0.

Accessing the ICV_PMR:

To access the ICV_PMR when HCR.{FMO, IMO} != {0, 0}, and executing at Non-secure EL1:

```
MRC p15,0,<Rt>,c4,c6,0 ; Read ICV_PMR into Rt
MCR p15,0,<Rt>,c4,c6,0 ; Write Rt to ICV_PMR
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	0100	0110	000

When *HCR* on page 1-25.{FMO, IMO} == {0, 0}, execution of this encoding at Non-secure EL1 results in an access to ICC_PMR.

8.6.16 ICV_RPR, Interrupt Controller Virtual Running Priority Register

The ICV_RPR characteristics are:

Purpose

Indicates the Running priority of the virtual CPU interface.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	RO	-	-	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL1 (S)	EL2 (NS)
-	-	RO	-	-

ICV_RPR is only accessible at Non-secure EL1 when *HCR* on page 1-25.{FMO, IMO} != {0, 0}.

Note

When *HCR* on page 1-25.{FMO, IMO} == {0, 0}, at Non-secure EL1, the instruction encoding used to access ICV_RPR results in an access to ICC_RPR.

If there are no active interrupts on the virtual CPU interface, or all active interrupts have undergone a priority drop, the value returned is the Idle priority.

Software cannot determine the number of implemented priority bits from a read of this register.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure read accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICC_SRE.SRE==0, Non-secure read accesses to this register from EL1 are UNDEFINED.

If ICC_SRE_EL1.SRE==0, Non-secure read accesses to this register from EL1 are trapped to EL1.

If ICH_HCR.TC==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

If ICH_HCR_EL2.TC==1, Non-secure read accesses to this register from EL1 are trapped to EL2.

Configurations

AArch32 System register ICV_RPR performs the same function as AArch64 System operation ICV_RPR_EL1.

Attributes

ICV_RPR is a 32-bit register.

Field descriptions

The ICV_RPR bit assignments are:

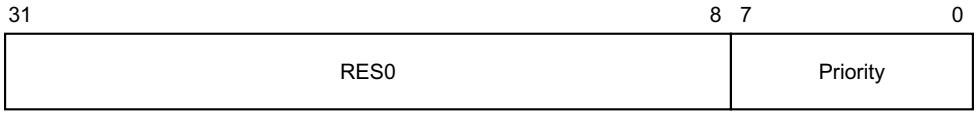

Bits [31:8]

Reserved, RES0.

Priority, bits [7:0]

The current running priority on the virtual CPU interface. This is the group priority of the current active virtual interrupt.

The priority returned is the group priority as if the BPR for the current Exception level and Security state was set to the minimum value of BPR for the number of implemented priority bits.

Note

If 8 bits of priority are implemented the group priority is bits[7:1] of the priority.

Accessing the ICV_RPR:

To access the ICV_RPR when HCR.{FMO, IMO} != {0, 0}, and executing at Non-secure EL1:

```
MRC p15,0,<Rt>,c12,c11,3 ; Read ICV_RPR into Rt
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	000	1100	1011	011

When *HCR* on page 1-25.{FMO, IMO} == {0, 0}, execution of this encoding at Non-secure EL1 results in an access to ICC_RPR.

8.7 AArch32 virtualization control System registers

This section describes each of the virtualization control AArch32 GIC System registers in register name order. The ICH prefix indicates a virtual interface control System register. Each AArch32 System register description contains a reference to the AArch64 register that provides the same functionality.

Unless otherwise stated, the bit assignments for the GIC System registers are the same as those for the equivalent GICH_* memory-mapped registers, see *The GIC virtual interface control register descriptions* on page 8-619.

Table 8-24 shows the encodings for the AArch 32 virtualization control System registers.

Table 8-24 Encodings for the AArch32 virtualization control System registers

Register	Width (bits)	opc1	CRn	CRm	opc2	Notes
ICH_AP0R<n>	32	4	12	8	0	RW
ICH_AP0R<n>	32				1	RW
ICH_AP0R<n>	32				2	RW
ICH_AP0R<n>	32				3	RW
ICH_AP1R<n>	32			9	0	RW
ICH_AP1R<n>	32				1	RW
ICH_AP1R<n>	32				2	RW
ICH_AP1R<n>	32				3	RW
ICH_HCR	32			11	0	RW
ICH_VTR	32				1	RO
ICH_MISR	32				2	RO
ICH_EISR	32				3	RO
ICH_ELRSR	32				5	RO
ICH_VMCR	32				7	RW
ICH_LR<n>, for n=0 - 7	32			12	0-7	RW
ICH_LR<n>, for n=8 - 15	32			13	0-7	RW
ICH_LRC<n>, for n=0 - 7	32			14	0-7	RW
ICH_LRC<n>, for n=8 - 15	32			15	0-7	RW

8.7.1 ICH_AP0R<n>, Interrupt Controller Hyp Active Priorities Group 0 Registers, n = 0 - 3

The ICH_AP0R<n> characteristics are:

Purpose

Provides information about Group 0 active priorities for EL2.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RW	RW	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	-	RW

ICH_AP0R1 is only implemented in implementations that support 6 or more bits of priority. ICH_AP0R2 and ICH_AP0R3 are only implemented in implementations that support 7 bits of priority. Unimplemented registers are UNDEFINED.

Writing to the active priority registers in any order other than the following order will result in UNPREDICTABLE behavior:

- ICH_AP0R<n>.
- ICH_AP1R<n>.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_HSRE.SRE==0, accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, Non-secure accesses to this register from EL3 are UNDEFINED.

Configurations

AArch32 System register ICH_AP0R<n> is architecturally mapped to AArch64 System register ICH_AP0R<n>_EL2.

Attributes

ICH_AP0R<n> is a 32-bit register.

Field descriptions

The ICH_AP0R<n> bit assignments are:

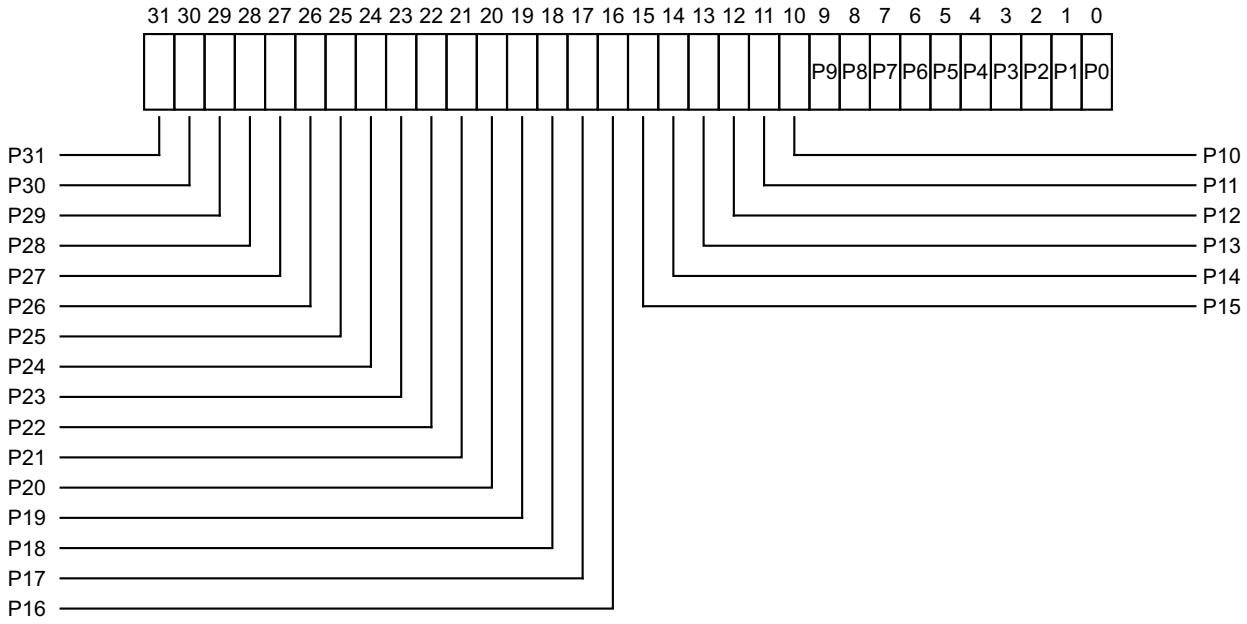

P<x>, bit [x], for x = 0 to 31

Provides the access to the virtual active priorities for Group 0 interrupts. Possible values of each bit are:

`0` There is no Group 0 interrupt active at the priority corresponding to that bit.

`1` There is a Group 0 interrupt active at the priority corresponding to that bit.

The correspondence between priority levels and bits depends on the number of bits of priority that are implemented.

If 5 bits of priority are implemented (bits [7:3] of priority), then there are 32 priority levels, and the active state of these priority levels are held in ICH_AP0R0 in the bits corresponding to Priority[7:3].

If 6 bits of priority are implemented (bits [7:2] of priority), then there are 64 priority levels, and:

- The active state of priority levels 0 - 124 are held in ICH_AP0R0 in the bits corresponding to 0:Priority[6:2].
- The active state of priority levels 128 - 252 are held in ICH_AP0R1 in the bits corresponding to 1:Priority[6:2].

If 7 bits of priority are implemented (bits [7:1] of priority), then there are 128 priority levels, and:

- The active state of priority levels 0 - 62 are held in ICH_AP0R0 in the bits corresponding to 00:Priority[5:1].
- The active state of priority levels 64 - 126 are held in ICH_AP0R1 in the bits corresponding to 01:Priority[5:1].
- The active state of priority levels 128 - 190 are held in ICH_AP0R2 in the bits corresponding to 10:Priority[5:1].
- The active state of priority levels 192 - 254 are held in ICH_AP0R3 in the bits corresponding to 11:Priority[5:1].

Note

Having the bit corresponding to a priority set to 1 in both ICH_AP0R<n> and ICH_AP1R<n> might result in UNPREDICTABLE behavior of the interrupt prioritization system for virtual interrupts.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the ICH_AP0R<n>:

To access the ICH_AP0R<n>:

```
MRC p15,4,<Rt>,c12,c8,<opc2> ; Read ICH_AP0R<n> into Rt, where n is in the range 0 to 3
MCR p15,4,<Rt>,c12,c8,<opc2> ; Write Rt to ICH_AP0R<n>, where n is in the range 0 to 3
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	100	1100	1000	0:n<1:0>

8.7.2 ICH_AP1R<n>, Interrupt Controller Hyp Active Priorities Group 1 Registers, n = 0 - 3

The ICH_AP1R<n> characteristics are:

Purpose

Provides information about Group 1 active priorities for EL2.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RW	RW	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	-	RW

ICH_AP1R1 is only implemented in implementations that support 6 or more bits of priority. ICH_AP1R2 and ICH_AP1R3 are only implemented in implementations that support 7 bits of priority. Unimplemented registers are UNDEFINED.

Writing to the active priority registers in any order other than the following order will result in UNPREDICTABLE behavior:

- ICH_AP0R<n>.
- ICH_AP1R<n>.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_HSRE.SRE==0, accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, Non-secure accesses to this register from EL3 are UNDEFINED.

Configurations

AArch32 System register ICH_AP1R<n> is architecturally mapped to AArch64 System register ICH_AP1R<n>_EL2.

Attributes

ICH_AP1R<n> is a 32-bit register.

Field descriptions

The ICH_AP1R<n> bit assignments are:

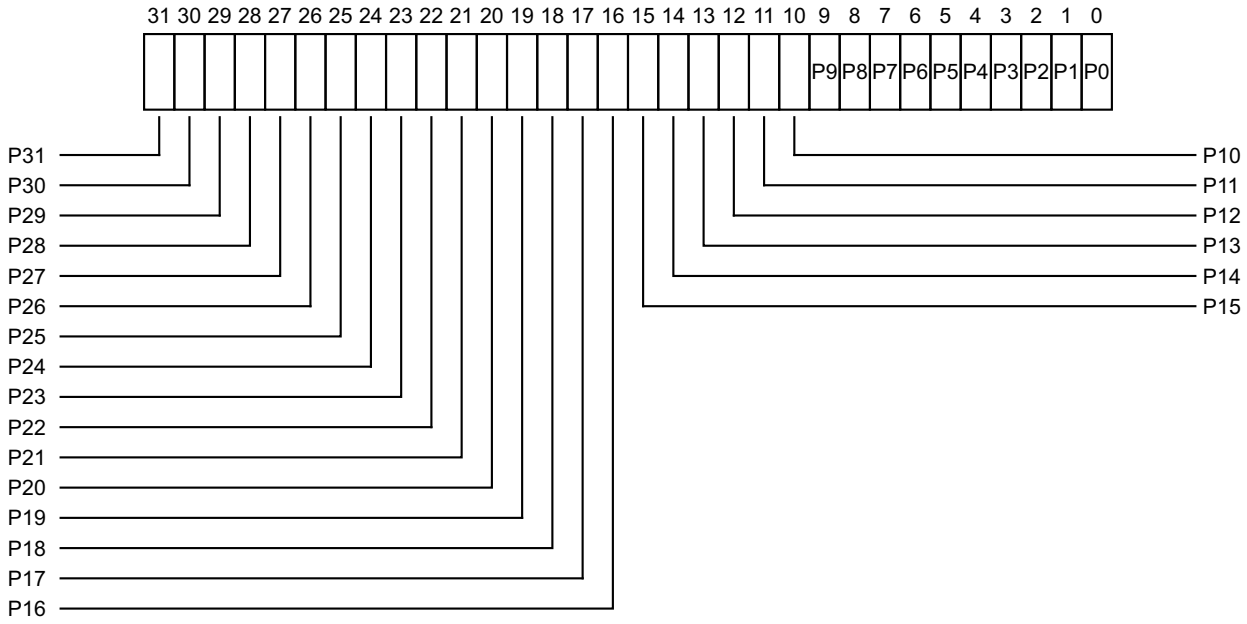

P<x>, bit [x], for x = 0 to 31

Group 1 interrupt active priorities. Possible values of each bit are:

`0` There is no Group 1 interrupt active at the priority corresponding to that bit.

`1` There is a Group 1 interrupt active at the priority corresponding to that bit.

The correspondence between priority levels and bits depends on the number of bits of priority that are implemented.

If 5 bits of priority are implemented (bits [7:3] of priority), then there are 32 priority levels, and the active state of these priority levels are held in ICH_AP1R0 in the bits corresponding to Priority[7:3].

If 6 bits of priority are implemented (bits [7:2] of priority), then there are 64 priority levels, and:

- The active state of priority levels 0 - 124 are held in ICH_AP1R0 in the bits corresponding to 0:Priority[6:2].
- The active state of priority levels 128 - 252 are held in ICH_AP1R1 in the bits corresponding to 1:Priority[6:2].

If 7 bits of priority are implemented (bits [7:1] of priority), then there are 128 priority levels, and:

- The active state of priority levels 0 - 62 are held in ICH_AP1R0 in the bits corresponding to 00:Priority[5:1].
- The active state of priority levels 64 - 126 are held in ICH_AP1R1 in the bits corresponding to 01:Priority[5:1].
- The active state of priority levels 128 - 190 are held in ICH_AP1R2 in the bits corresponding to 10:Priority[5:1].
- The active state of priority levels 192 - 254 are held in ICH_AP1R3 in the bits corresponding to 11:Priority[5:1].

Note

Having the bit corresponding to a priority set to 1 in both ICH_AP0R<n> and ICH_AP1R<n> might result in UNPREDICTABLE behavior of the interrupt prioritization system for virtual interrupts.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the ICH_AP1R<n>:

To access the ICH_AP1R<n>:

```
MRC p15,4,<Rt>,c12,c9,<opc2> ; Read ICH_AP1R<n> into Rt, where n is in the range 0 to 3
MCR p15,4,<Rt>,c12,c9,<opc2> ; Write Rt to ICH_AP1R<n>, where n is in the range 0 to 3
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	100	1100	1001	0:n<1:0>

8.7.3 ICH_EISR, Interrupt Controller End of Interrupt Status Register

The ICH_EISR characteristics are:

Purpose

Indicates which List registers have outstanding EOI maintenance interrupts.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RO	RO	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	-	RO

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_HSRE.SRE==0, read accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, Non-secure read accesses to this register from EL3 are UNDEFINED.

Configurations

AArch32 System register ICH_EISR is architecturally mapped to AArch64 System register ICH_EISR_EL2.

Attributes

ICH_EISR is a 32-bit register.

Field descriptions

The ICH_EISR bit assignments are:

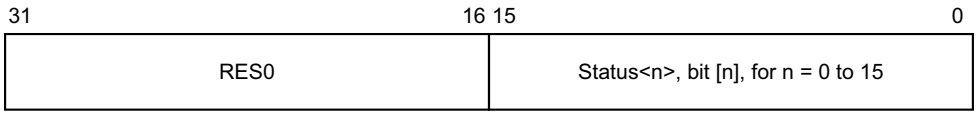

Bits [31:16]

Reserved, RES0.

Status<n>, bit [n], for n = 0 to 15

EOI maintenance interrupt status bit for List register <n>:

`0` List register <n>, ICH_LR<n>, does not have an EOI maintenance interrupt.

`1` List register <n>, ICH_LR<n>, has an EOI maintenance interrupt that has not been handled.

For any ICH_LR<n>, the corresponding status bit is set to 1 if all of the following are true:

- ICH_LRC<n>.State is `0b00`.
- ICH_LRC<n>.HW is 0.
- ICH_LRC<n>.EOI (bit [9]) is 1, indicating that when the interrupt corresponding to that List register is deactivated, a maintenance interrupt is asserted.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the ICH_EISR:

To access the ICH_EISR:

```
MRC p15,4,<Rt>,c12,c11,3 ; Read ICH_EISR into Rt
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	100	1100	1011	011

8.7.4 ICH_ELRSR, Interrupt Controller Empty List Register Status Register

The ICH_ELRSR characteristics are:

Purpose

Indicates which List registers contain valid interrupts.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RO	RO	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	-	RO

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_HSRE.SRE==0, read accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, Non-secure read accesses to this register from EL3 are UNDEFINED.

Configurations

AArch32 System register ICH_ELRSR is architecturally mapped to AArch64 System register ICH_ELRSR_EL2.

Attributes

ICH_ELRSR is a 32-bit register.

Field descriptions

The ICH_ELRSR bit assignments are:

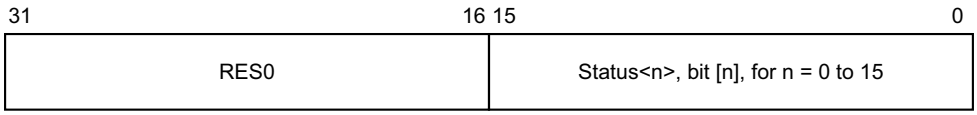

Bits [31:16]

Reserved, RES0.

Status<n>, bit [n], for n = 0 to 15

Status bit for List register <n>, ICH_LR<n>:

`0` List register ICH_LR<n>, if implemented, contains a valid interrupt. Using this List register can result in overwriting a valid interrupt.

`1` List register ICH_LR<n> does not contain a valid interrupt. The List register is empty and can be used without overwriting a valid interrupt or losing an EOI maintenance interrupt.

For any List register <n>, the corresponding status bit is set to 1 if ICH_LRC<n>.State is `0b00` and either ICH_LRC<n>.HW is 1 or ICH_LRC<n>.EOI (bit [9]) is 0.

Accessing the ICH_ELRSR:

To access the ICH_ELRSR:

```
MRC p15,4,<Rt>,c12,c11,5 ; Read ICH_ELRSR into Rt
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	100	1100	1011	101

8.7.5 ICH_HCR, Interrupt Controller Hyp Control Register

The ICH_HCR characteristics are:

Purpose

Controls the environment for VMs.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RW	RW	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	-	RW

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_HSRE.SRE==0, accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, Non-secure accesses to this register from EL3 are UNDEFINED.

Configurations

AArch32 System register ICH_HCR is architecturally mapped to AArch64 System register ICH_HCR_EL2.

Attributes

ICH_HCR is a 32-bit register.

Field descriptions

The ICH_HCR bit assignments are:

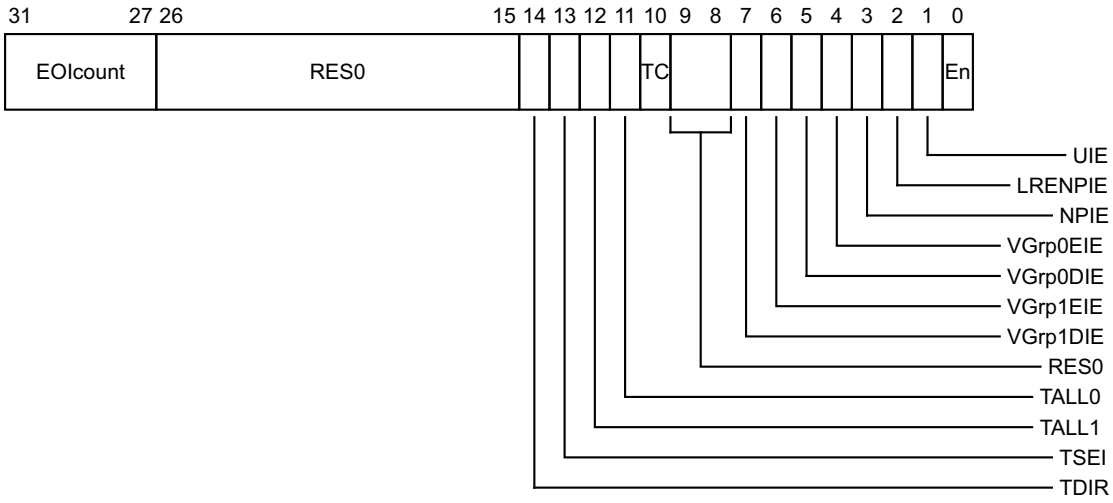

EOIcount, bits [31:27]

This field is incremented whenever a successful write to a virtual EOIR or DIR register would have resulted in a virtual interrupt deactivation. That is:

- A virtual write to EOIR with a valid interrupt identifier that is not in the LPI range (i.e. < 8192) when EOI mode is zero and no List Register was found, or
- A virtual write to DIR with a valid interrupt identifier that is not in the LPI range (i.e. < 8192) when EOI mode is one and no List Register was found

This allows software to manage more active interrupts than there are implemented List Registers.

It is CONSTRAINED UNPREDICTABLE whether a virtual write to EOIR that does not clear a bit in the Active Priorities registers (ICH_AP0R<n>/ICH_AP1R<n>) increments EOIcount. Permitted behaviors are:

- Increment EOIcount.
- Leave EOIcount unchanged.

When this register has an architecturally-defined reset value, this field resets to `0`.

Bits [26:15]

Reserved, RES0.

TDIR, bit [14]

Trap Non-secure EL1 writes to ICC_DIR and ICV_DIR.

`0` Non-secure EL1 writes of ICC_DIR and ICV_DIR are not trapped to EL2, unless trapped by other mechanisms.

`1` Non-secure EL1 writes of ICC_DIR and ICV_DIR are trapped to EL2.

Support for this bit is OPTIONAL, with support indicated by ICH_VTR.

If the implementation does not support this trap, this bit is RES0.

ARM deprecates not including this trap bit.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

TSEI, bit [13]

Trap all locally generated SEIs. This bit allows the hypervisor to intercept locally generated SEIs that would otherwise be taken at Non-secure EL1.

`0` Locally generated SEIs do not cause a trap to EL2.

1 Locally generated SEIs trap to EL2.

If ICH_VTR.SEIS is 0, this bit is RES0.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

TALL1, bit [12]

Trap all Non-secure EL1 accesses to ICC_* and ICV_* System registers for Group 1 interrupts to EL2.

0 Non-Secure EL1 accesses to ICC_* and ICV_* registers for Group 1 interrupts proceed as normal.

1 Non-secure EL1 accesses to ICC_* and ICV_* registers for Group 1 interrupts trap to EL2.

When this register has an architecturally-defined reset value, this field resets to 0.

TALL0, bit [11]

Trap all Non-secure EL1 accesses to ICC_* and ICV_* System registers for Group 0 interrupts to EL2.

0 Non-Secure EL1 accesses to ICC_* and ICV_* registers for Group 0 interrupts proceed as normal.

1 Non-secure EL1 accesses to ICC_* and ICV_* registers for Group 0 interrupts trap to EL2.

When this register has an architecturally-defined reset value, this field resets to 0.

TC, bit [10]

Trap all Non-secure EL1 accesses to System registers that are common to Group 0 and Group 1 to EL2.

0 Non-secure EL1 accesses to common registers proceed as normal.

1 Non-secure EL1 accesses to common registers trap to EL2.

This affects accesses to ICC_SGI0R, ICC_SGI1R, ICC_ASGI1R, ICC_CTLR, ICC_DIR, ICC_PMR, ICC_RPR, ICV_CTLR, ICV_DIR, ICV_PMR, and ICV_RPR.

When this register has an architecturally-defined reset value, this field resets to 0.

Bits [9:8]

Reserved, RES0.

VGrp1DIE, bit [7]

VM Group 1 Disabled Interrupt Enable. Enables the signaling of a maintenance interrupt while signaling of Group 1 interrupts from the virtual CPU interface to the connected vPE is disabled:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled when ICH_VMCR.VENG1 is 0.

When this register has an architecturally-defined reset value, this field resets to 0.

VGrp1EIE, bit [6]

VM Group 1 Enabled Interrupt Enable. Enables the signaling of a maintenance interrupt while signaling of Group 1 interrupts from the virtual CPU interface to the connected vPE is enabled:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled when ICH_VMCR.VENG1 is 1.

When this register has an architecturally-defined reset value, this field resets to 0.

VGrp0DIE, bit [5]

VM Group 0 Disabled Interrupt Enable. Enables the signaling of a maintenance interrupt while signaling of Group 0 interrupts from the virtual CPU interface to the connected vPE is disabled:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled when ICH_VMCR.VENG0 is 0.

When this register has an architecturally-defined reset value, this field resets to 0.

VGrp0EIE, bit [4]

VM Group 0 Enabled Interrupt Enable. Enables the signaling of a maintenance interrupt while signaling of Group 0 interrupts from the virtual CPU interface to the connected vPE is enabled:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled when ICH_VMCR.VENG0 is 1.

When this register has an architecturally-defined reset value, this field resets to 0.

NPIE, bit [3]

No Pending Interrupt Enable. Enables the signaling of a maintenance interrupt while no pending interrupts are present in the List registers:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled while the List registers contain no interrupts in the pending state.

When this register has an architecturally-defined reset value, this field resets to 0.

LRENPIE, bit [2]

List Register Entry Not Present Interrupt Enable. Enables the signaling of a maintenance interrupt while the virtual CPU interface does not have a corresponding valid List register entry for an EOI request:

0 Maintenance interrupt disabled.

1 Maintenance interrupt is asserted while the EOIcount field is not 0.

When this register has an architecturally-defined reset value, this field resets to 0.

UIE, bit [1]

Underflow Interrupt Enable. Enables the signaling of a maintenance interrupt when the List registers are empty, or hold only one valid entry:

0 Maintenance interrupt disabled.

1 Maintenance interrupt is asserted if none, or only one, of the List register entries is marked as a valid interrupt.

When this register has an architecturally-defined reset value, this field resets to 0.

En, bit [0]

Enable. Global enable bit for the virtual CPU interface:

0 Virtual CPU interface operation disabled.

1 Virtual CPU interface operation enabled.

When this field is set to 0:

- The virtual CPU interface does not signal any maintenance interrupts.
- The virtual CPU interface does not signal any virtual interrupts.
- A read of ICV_IAR0, ICV_IAR1, GICV_IAR or GICV_AIAR returns a spurious interrupt ID.

When this register has an architecturally-defined reset value, this field resets to 0.

Accessing the ICH_HCR:

To access the ICH_HCR:

```
MRC p15,4,<Rt>,c12,c11,0 ; Read ICH_HCR into Rt
MCR p15,4,<Rt>,c12,c11,0 ; Write Rt to ICH_HCR
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	100	1100	1011	000

8.7.6 ICH_LR<n>, Interrupt Controller List Registers, n = 0 - 15

The ICH_LR<n> characteristics are:

Purpose

Provides interrupt context information for the virtual CPU interface.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RW	RW	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	-	RW

ICH_LR<n> and ICH_LRC<n> can be updated independently.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_HSRE.SRE==0, accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, Non-secure accesses to this register from EL3 are UNDEFINED.

Configurations

AArch32 System register ICH_LR<n> is architecturally mapped to AArch64 System register ICH_LR<n>_EL2[31:0].

Attributes

ICH_LR<n> is a 32-bit register.

Field descriptions

The ICH_LR<n> bit assignments are:

vINTID, bits [31:0]

Virtual INTID of the interrupt.

Behavior is UNPREDICTABLE if two or more List Registers specify the same vINTID when:

- ICH_LR<n>.State == 01.
- ICH_LR<n>.State == 10.
- ICH_LR<n>.State == 11.

It is IMPLEMENTATION DEFINED how many bits are implemented, though at least 16 bits must be implemented. Unimplemented bits are RES0. The number of implemented bits can be discovered from ICH_VTR.IDbits.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the ICH_LR<n>:

To access the ICH_LR<n>:

```
MRC p15,4,<Rt>,c12,<CRm>,<opc2> ; Read ICH_LR<n> into Rt, where n is in the range 0 to 15
MCR p15,4,<Rt>,c12,<CRm>,<opc2> ; Write Rt to ICH_LR<n>, where n is in the range 0 to 15
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	100	1100	110:n<3>	n<2:0>

8.7.7 ICH_LRC<n>, Interrupt Controller List Registers, n = 0 - 15

The ICH_LRC<n> characteristics are:

Purpose

Provides interrupt context information for the virtual CPU interface.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RW	RW	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	-	RW

ICH_LR<n> and ICH_LRC<n> can be updated independently.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_HSRE.SRE==0, accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, Non-secure accesses to this register from EL3 are UNDEFINED.

Configurations

AArch32 System register ICH_LRC<n> is architecturally mapped to AArch64 System register ICH_LR<n>_EL2[63:32].

Attributes

ICH_LRC<n> is a 32-bit register.

Field descriptions

The ICH_LRC<n> bit assignments are:

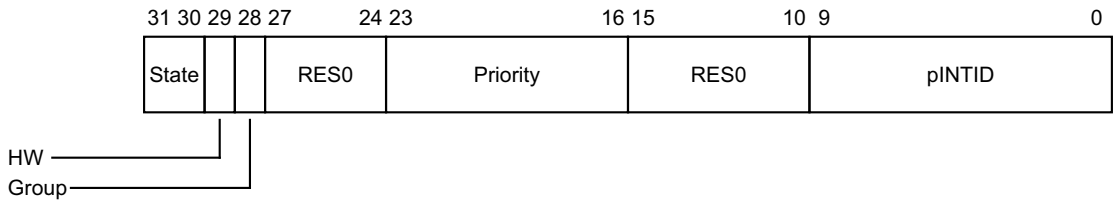

State, bits [31:30]

The state of the interrupt:

00 Inactive

01 Pending

10 Active

11 Pending and active.

The GIC updates these state bits as virtual interrupts proceed through the interrupt life cycle. Entries in the inactive state are ignored, except for the purpose of generating virtual maintenance interrupts.

For hardware interrupts, the pending and active state is held in the physical Distributor rather than the virtual CPU interface. A hypervisor must only use the pending and active state for software originated interrupts, which are typically associated with virtual devices, or SGIs.

When this register has an architecturally-defined reset value, this field resets to 0.

HW, bit [29]

Indicates whether this virtual interrupt maps directly to a hardware interrupt, meaning that it corresponds to a physical interrupt. Deactivation of the virtual interrupt also causes the deactivation of the physical interrupt with the INTID that the pINTID field indicates.

0 The interrupt is triggered entirely by software. No notification is sent to the Distributor when the virtual interrupt is deactivated.

1 The interrupt maps directly to a hardware interrupt. A deactivate interrupt request is sent to the Distributor when the virtual interrupt is deactivated, using the pINTID field from this register to indicate the physical INTID.

If ICH_VMCR.VEOIM is 0, this request corresponds to a write to ICC_EOIR0 or ICC_EOIR1. Otherwise, it corresponds to a write to ICC_DIR.

When this register has an architecturally-defined reset value, this field resets to 0.

Group, bit [28]

Indicates the group for this virtual interrupt.

0 This is a Group 0 virtual interrupt. ICH_VMCR.VFIQEn determines whether it is signaled as a virtual IRQ or as a virtual FIQ, and ICH_VMCR.VENG0 enables signaling of this interrupt to the virtual machine.

1 This is a Group 1 virtual interrupt, signaled as a virtual IRQ. ICH_VMCR.VENG1 enables the signaling of this interrupt to the virtual machine.

If ICH_VMCR.VCBPR is 0, then ICC_BPR1 determines if a pending Group 1 interrupt has sufficient priority to preempt current execution. Otherwise, ICH_LR<n> determines preemption.

When this register has an architecturally-defined reset value, this field resets to 0.

Bits [27:24]

Reserved, RES0.

Priority, bits [23:16]

The priority of this interrupt.

It is IMPLEMENTATION DEFINED how many bits of priority are implemented, though at least five bits must be implemented. Unimplemented bits are RES0 and start from bit [16] up to bit [18]. The number of implemented bits can be discovered from ICH_VTR.PRIbits.

When this register has an architecturally-defined reset value, this field resets to 0.

Bits [15:10]

Reserved, RES0.

pINTID, bits [9:0]

Physical INTID, for hardware interrupts.

When the HW bit is 0 (there is no corresponding physical interrupt), this field has the following meaning:

Bit[9] EOI. If this bit is 1, then when the interrupt identified by vINTID is deactivated, an EOI maintenance interrupt is asserted.

Bits[8:0] Reserved, RES0.

When the HW bit is 1 (there is a corresponding physical interrupt):

- This field indicates the physical INTID. This field is only required to implement enough bits to hold a valid value for the implemented INTID size. Any unused higher order bits are RES0.
- If the value of pINTID is 0-15 or 1020-1023, behavior is UNPREDICTABLE. If the value of pINTID is 16-31, this field applies to the PPI associated with this same physical PE ID as the virtual CPU interface requesting the deactivation.

A hardware physical identifier is only required in List Registers for interrupts that require deactivation. This means only 10 bits of Physical INTID are required, regardless of the number specified by ICC_CTLR.IDbits.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the ICH_LRC<n>:

To access the ICH_LRC<n>:

```
MRC p15,4,<Rt>,c12,<CRm>,<opc2> ; Read ICH_LRC<n> into Rt, where n is in the range 0 to 15
MCR p15,4,<Rt>,c12,<CRm>,<opc2> ; Write Rt to ICH_LRC<n>, where n is in the range 0 to 15
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	100	1100	111:n<3>	n<2:0>

8.7.8 ICH_MISR, Interrupt Controller Maintenance Interrupt State Register

The ICH_MISR characteristics are:

Purpose

Indicates which maintenance interrupts are asserted.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RO	RO	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	-	RO

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_HSRE.SRE==0, read accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, Non-secure read accesses to this register from EL3 are UNDEFINED.

Configurations

AArch32 System register ICH_MISR is architecturally mapped to AArch64 System register ICH_MISR_EL2.

Attributes

ICH_MISR is a 32-bit register.

Field descriptions

The ICH_MISR bit assignments are:

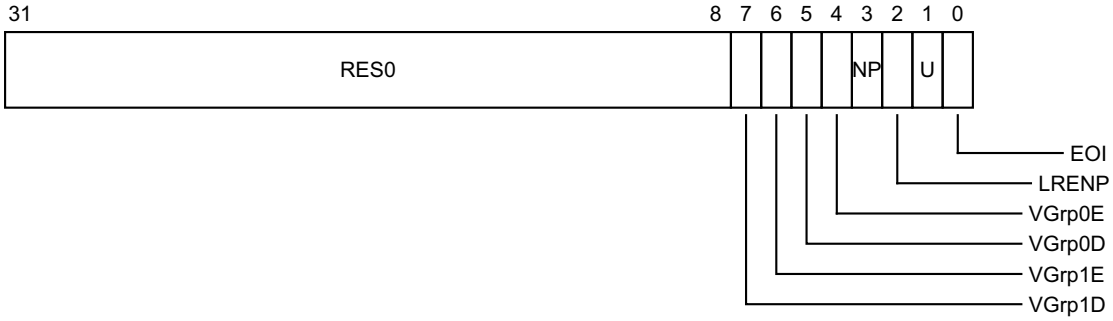

Bits [31:8]

Reserved, RES0.

VGrp1D, bit [7]

vPE Group 1 Disabled.

0 vPE Group 1 Disabled maintenance interrupt not asserted.

1 vPE Group 1 Disabled maintenance interrupt asserted.

This maintenance interrupt is asserted when ICH_HCR.VGrp1DIE is 1 and ICH_VMCR.VMGrp1En is 0.

When this register has an architecturally-defined reset value, this field resets to 0.

VGrp1E, bit [6]

vPE Group 1 Enabled.

0 vPE Group 1 Enabled maintenance interrupt not asserted.

1 vPE Group 1 Enabled maintenance interrupt asserted.

This maintenance interrupt is asserted when ICH_HCR.VGrp1EIE is 1 and ICH_VMCR.VMGrp1En is 1.

When this register has an architecturally-defined reset value, this field resets to 0.

VGrp0D, bit [5]

vPE Group 0 Disabled.

0 vPE Group 0 Disabled maintenance interrupt not asserted.

1 vPE Group 0 Disabled maintenance interrupt asserted.

This maintenance interrupt is asserted when ICH_HCR.VGrp0DIE is 1 and ICH_VMCR.VMGrp0En is 0.

When this register has an architecturally-defined reset value, this field resets to 0.

VGrp0E, bit [4]

vPE Group 0 Enabled.

0 vPE Group 0 Enabled maintenance interrupt not asserted.

1 vPE Group 0 Enabled maintenance interrupt asserted.

This maintenance interrupt is asserted when ICH_HCR.VGrp0EIE is 1 and ICH_VMCR.VMGrp0En is 1.

When this register has an architecturally-defined reset value, this field resets to 0.

NP, bit [3]

No Pending.

0 No Pending maintenance interrupt not asserted.

1 No Pending maintenance interrupt asserted.

This maintenance interrupt is asserted when ICH_HCR.NPIE is 1 and no List register is in pending state.

When this register has an architecturally-defined reset value, this field resets to 0.

LRENP, bit [2]

List Register Entry Not Present.

0 List Register Entry Not Present maintenance interrupt not asserted.

1 List Register Entry Not Present maintenance interrupt asserted.

This maintenance interrupt is asserted when ICH_HCR.LRENPIE is 1 and ICH_HCR.EOIcount is non-zero.

When this register has an architecturally-defined reset value, this field resets to 0.

U, bit [1]

Underflow.

0 Underflow maintenance interrupt not asserted.

1 Underflow maintenance interrupt asserted.

This maintenance interrupt is asserted when ICH_HCR.UIE is 1 and zero or one of the List register entries are marked as a valid interrupt, that is, if the corresponding ICH_LRC<n>.State bits do not equal 0x0.

When this register has an architecturally-defined reset value, this field resets to 0.

EOI, bit [0]

End Of Interrupt.

0 End Of Interrupt maintenance interrupt not asserted.

1 End Of Interrupt maintenance interrupt asserted.

This maintenance interrupt is asserted when at least one bit in ICH_EISR is 1.

When this register has an architecturally-defined reset value, this field resets to 0.

Accessing the ICH_MISR:

To access the ICH_MISR:

```
MRC p15,4,<Rt>,c12,c11,2 ; Read ICH_MISR into Rt
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	100	1100	1011	010

8.7.9 ICH_VMCR, Interrupt Controller Virtual Machine Control Register

The ICH_VMCR characteristics are:

Purpose

Enables the hypervisor to save and restore the virtual machine view of the GIC state.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RW	RW	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	-	RW

When EL2 is using System register access, EL1 using either System register or memory-mapped access must be supported.

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_HSRE.SRE==0, accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, Non-secure accesses to this register from EL3 are UNDEFINED.

Configurations

AArch32 System register ICH_VMCR is architecturally mapped to AArch64 System register ICH_VMCR_EL2.

Attributes

ICH_VMCR is a 32-bit register.

Field descriptions

The ICH_VMCR bit assignments are:

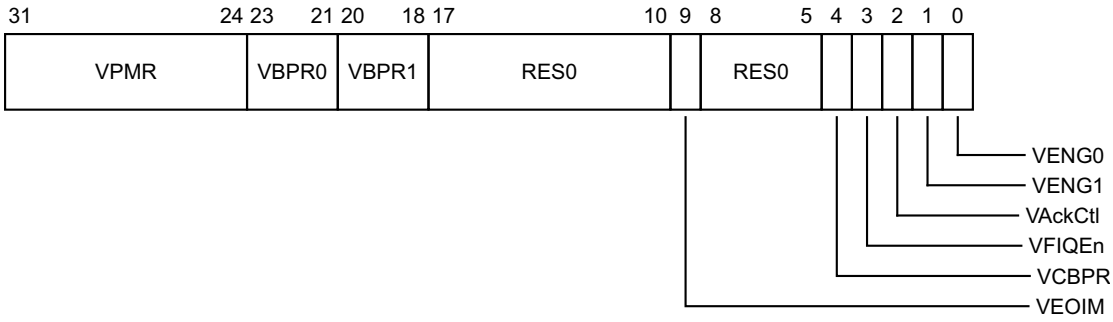

VPMR, bits [31:24]

Virtual Priority Mask. The priority mask level for the virtual CPU interface. If the priority of a pending virtual interrupt is higher than the value indicated by this field, the interface signals the virtual interrupt to the PE.

This field is an alias of ICV_PMR.Priority.

VBPR0, bits [23:21]

Virtual Binary Point Register, Group 0. Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines Group 0 interrupt preemption, and also determines Group 1 interrupt preemption if ICH_VMCR.VCBPR == 1.

This field is an alias of ICV_BPR0.BinaryPoint.

VBPR1, bits [20:18]

Virtual Binary Point Register, Group 1. Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines Group 1 interrupt preemption if ICH_VMCR.VCBPR == 0.

This field is an alias of ICV_BPR1.BinaryPoint.

Bits [17:10]

Reserved, RES0.

VEOIM, bit [9]

Virtual EOI mode. Controls whether a write to an End of Interrupt register also deactivates the virtual interrupt:

`0` ICV_EOIR0 and ICV_EOIR1 provide both priority drop and interrupt deactivation functionality. Accesses to ICV_DIR are UNPREDICTABLE.

`1` ICV_EOIR0 and ICV_EOIR1 provide priority drop functionality only. ICV_DIR provides interrupt deactivation functionality.

This bit is an alias of ICV_CTLR.EOImode.

Bits [8:5]

Reserved, RES0.

VCBPR, bit [4]

Virtual Common Binary Point Register. Possible values of this bit are:

`0` ICV_BPR0 determines the preemption group for virtual Group 0 interrupts only.
ICV_BPR1 determines the preemption group for virtual Group 1 interrupts.

`1` ICV_BPR0 determines the preemption group for both virtual Group 0 and virtual Group 1 interrupts.

Reads of ICV_BPR1 return ICV_BPR0 plus one, saturated to 0b111. Writes to ICV_BPR1 are ignored.

This field is an alias of ICV_CTLR.CBPR.

VFIQEn, bit [3]

Virtual FIQ enable. Possible values of this bit are:

0 Group 0 virtual interrupts are presented as virtual IRQs.

1 Group 0 virtual interrupts are presented as virtual FIQs.

This bit is an alias of GICV_CTLR.FIQEn.

In implementations where the Non-secure copy of ICC_SRE.SRE is always one, this bit is RES1.

VAckCtl, bit [2]

Virtual AckCtl. Possible values of this bit are:

0 If the highest priority pending interrupt is Group 1, a read of GICV_IAR or GICV_HPPIR returns an INTID of 1022.

1 If the highest priority pending interrupt is Group 1, a read of GICV_IAR or GICV_HPPIR returns the INTID of the corresponding interrupt.

This bit is an alias of GICV_CTLR.AckCtl.

This field is supported for backwards compatibility with GICv2. ARM deprecates the use of this field.

In implementations where the Non-secure copy of ICC_SRE.SRE is always one, this bit is RES0.

VENG1, bit [1]

Virtual Group 1 interrupt enable. Possible values of this bit are:

0 Virtual Group 1 interrupts are disabled.

1 Virtual Group 1 interrupts are enabled.

This bit is an alias of ICV_IGRPEN1.Enable.

VENG0, bit [0]

Virtual Group 0 interrupt enable. Possible values of this bit are:

0 Virtual Group 0 interrupts are disabled.

1 Virtual Group 0 interrupts are enabled.

This bit is an alias of ICV_IGRPEN0.Enable.

Accessing the ICH_VMCR:

To access the ICH_VMCR:

```
MRC p15,4,<Rt>,c12,c11,7 ; Read ICH_VMCR into Rt
MCR p15,4,<Rt>,c12,c11,7 ; Write Rt to ICH_VMCR
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	100	1100	1011	111

8.7.10 ICH_VTR, Interrupt Controller VGIC Type Register

The ICH_VTR characteristics are:

Purpose

Reports supported GIC virtualisartion features.

Usage constraints

If EL3 is implemented and is using AArch32, this register is accessible as follows:

EL0 (NS)	EL0 (S)	EL1 (NS)	EL2 (NS)	EL3 (SCR.NS=1)	EL3 (SCR.NS=0)
-	-	-	RO	RO	-

If EL3 is not implemented or EL3 is implemented and is using AArch64, this register is accessible as follows:

EL0	EL1	EL2 (NS)
-	-	RO

Traps and Enables

For a description of the prioritization of any generated exceptions, see G1.11.2 (Exception priority order) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch32 state, and section D1.13.2 (Synchronous exception prioritization for exceptions taken to AArch64) in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for exceptions taken to AArch64 state.

Subject to the prioritization rules:

If HSTR.T12==1, Non-secure accesses to this register from EL1 are trapped to Hyp mode.

If HSTR_EL2.T12==1, Non-secure accesses to this register from EL1 are trapped to EL2.

If ICC_HSRE.SRE==0, read accesses to this register from EL2 are UNDEFINED.

If ICC_MSRE.SRE==0, Non-secure read accesses to this register from EL3 are UNDEFINED.

Configurations

AArch32 System register ICH_VTR is architecturally mapped to AArch64 System register ICH_VTR_EL2.

Attributes

ICH_VTR is a 32-bit register.

Field descriptions

The ICH_VTR bit assignments are:

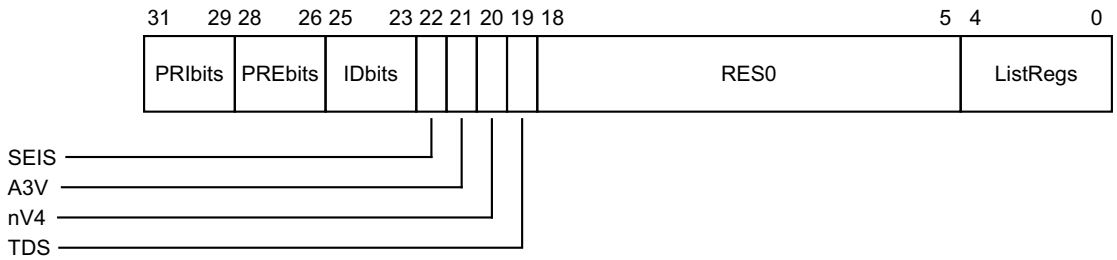

PRIbits, bits [31:29]

Priority bits. The number of virtual priority bits implemented, minus one.

An implementation must implement at least 32 levels of virtual priority (5 priority bits).

This field is an alias of ICV_CTLR.PRIbits.

PREbits, bits [28:26]

The number of virtual preemption bits implemented, minus one.

An implementation must implement at least 32 levels of virtual preemption priority (5 preemption bits).

The value of this field must be less than or equal to the value of ICH_VTR.PRIbits.

IDbits, bits [25:23]

The number of virtual interrupt identifier bits supported:

`000` 16 bits.

`001` 24 bits.

All other values are reserved.

This field is an alias of ICV_CTLR.IDbits.

SEIS, bit [22]

SEI Support. Indicates whether the virtual CPU interface supports generation of SEIs:

`0` The virtual CPU interface logic does not support generation of SEIs.

`1` The virtual CPU interface logic supports generation of SEIs.

This bit is an alias of ICV_CTLR.SEIS.

A3V, bit [21]

Affinity 3 Valid. Possible values are:

`0` The virtual CPU interface logic only supports zero values of Affinity 3 in SGI generation System registers.

`1` The virtual CPU interface logic supports non-zero values of Affinity 3 in SGI generation System registers.

This bit is an alias of ICV_CTLR.A3V.

nV4, bit [20]

GICv4 direct injection of virtual interrupts not supported. Possible values are:

`0` The CPU interface logic supports direct injection of virtual interrupts.

`1` The CPU interface logic does not support direct injection of virtual interrupts.

TDS, bit [19]

Separate trapping of Non-secure EL1 writes to ICV_DIR supported.

`0` Implementation does not support ICH_HCR.TDIR.

`1` Implementation supports ICH_HCR.TDIR.

Bits [18:5]

Reserved, RES0.

ListRegs, bits [4:0]

The number of implemented List registers, minus one. For example, a value of `0b01111` indicates that the maximum of 16 List registers are implemented.

Accessing the ICH_VTR:

To access the ICH_VTR:

```
MRC p15,4,<Rt>,c12,c11,1 ; Read ICH_VTR into Rt
```

Register access is encoded as follows:

coproc	opc1	CRn	CRm	opc2
1111	100	1100	1011	001

8.8 The GIC Distributor register map

Table 8-25 shows the Distributor register map. Address offsets are relative to the *Distributor base address* defined by the system memory map. Unless otherwise stated in the register description, all GIC registers are 32-bits wide. Reserved register addresses are RAZ/WI.

Table 8-25 Distributor register map

Offset	Name	Type	Reset[a]	Description
0x0000	GICD_CTLR	RW	See the register description	Distributor Control Register
0x0004	GICD_TYPER	RO	IMPLEMENTATION DEFINED	Interrupt Controller Type Register
0x0008	GICD_IIDR	RO	IMPLEMENTATION DEFINED	Distributor Implementer Identification Register
0x000C	-	-	-	Reserved
0x0010	GICD_STATUSR	RW	0x000 00000	Error Reporting Status Register, optional
0x0014-0x001C	-	-	-	Reserved
0x0020-0x003C	-	-	-	IMPLEMENTATION DEFINED registers
0x0040	GICD_SETSPI_NSR	WO	-	Set SPI Register
0x0044	-	-	-	Reserved
0x0048	GICD_CLRSPI_NSR	WO	-	Clear SPI Register
0x004C	-	-	-	Reserved
0x0050	GICD_SETSPI_SR	WO	-	Set SPI, Secure Register
0x0054	-	-	-	Reserved
0x0058	GICD_CLRSPI_SR	WO	-	Clear SPI, Secure Register
0x005C-0x007C	-	-	-	Reserved
0x0080	GICD_IGROUPR<n>	RW	IMPLEMENTATION DEFINED	Interrupt Group Registers
0x0084-0x00FC			0x0000 0000	
0x0100-0x017C	GICD_ISENABLER<n>	RW	IMPLEMENTATION DEFINED	Interrupt Set-Enable Registers
0x0180-0x01FC	GICD_ICENABLER<n>	RW	IMPLEMENTATION DEFINED	Interrupt Clear-Enable Registers
0x0200-0x027C	GICD_ISPENDR<n>	RW	0x0000 0000	Interrupt Set-Pending Registers
0x0280-0x02FC	GICD_ICPENDR<n>	RW	0x0000 0000	Interrupt Clear-Pending Registers
0x0300-0x037C	GICD_ISACTIVER<n>	RW	0x0000 0000	Interrupt Set-Active Registers
0x0380-0x03FC	GICD_ICACTIVER<n>	RW	0x0000 0000	Interrupt Clear-Active Registers
0x0400-0x07F8	GICD_IPRIORITYR<n>	RW	IMPLEMENTATION DEFINED	Interrupt Priority Registers
0x0800-0x081C	GICD_ITARGETSR<n>[bc]	RO	IMPLEMENTATION DEFINED	Interrupt Processor Targets Registers
0x0820-0xBF8		RW	0x0000 0000	
0x0C00-0x0CFC	GICD_ICFGR<n>	RW	IMPLEMENTATION DEFINED	Interrupt Configuration Registers
0x0D00-0x0D7C	GICD_IGRPMODR<n>[d]	-	0x0000 0000	Interrupt Group Modifier Registers

Table 8-25 Distributor register map (continued)

Offset	Name	Type	Reset[a]	Description
0x0E00-0x0EFC	GICD_NSACR<n>	RW	0x0000 0000	Non-secure Access Control Registers
0x0F00	GICD_SGIR	WO	-	Software Generated Interrupt Register
0x0F10-0x0F1C	GICD_CPENDSGIR<n>[e]	RW	0x0000 0000	SGI Clear-Pending Registers
0x0F20-0x0F2C	GICD_SPENDSGIR<n>[e]	RW	0x0000 0000	SGI Set-Pending Registers
0x0F30-0x60FC	-	-	-	Reserved
0x6100-0x7FD8	GICD_IROUTER<n>	RW	-	Interrupt Routing Registers
0x7FE0-0xBFFC	-	-	-	Reserved
0xC000-0xFFCC	-	-	-	IMPLEMENTATION DEFINED registers
0xFFD0-0xFFFC	-	RO	IMPLEMENTATION DEFINED	Reserved for ID registers, see *Identification registers* on page 8-173

a. For details of any restrictions that apply to the reset values that are IMPLEMENTATION DEFINED, see the appropriate register description.
b. When affinity routing is enabled, GICD_IROUTER<n> are used instead of these registers.
c. In an implementation with a single connected PE, these registers are RAZ/WI.
d. These registers are RES0 when affinity routing is not enabled for the Secure state.
e. Used only when affinity routing is not enabled.

A Distributor might optionally provide an IMPLEMENTATION DEFINED set of aliases for message-based interrupt requests.

Table 8-26 shows the Distributor message-based interrupt register map.

Table 8-26 Distributor message-based interrupt register map

Offset	Name	Type	Reset	Description
0x0000-0x003C	-	-	-	Reserved
0x0040	GICD_SETSPI_NSR	WO	-	Set SPI Register
0x044	-	-	-	Reserved
0x0048	GICD_CLRSPI_NSR	WO	-	Clear SPI Register
0x004C	-	-	-	Reserved
0x0050	GICD_SETSPI_SR	WO	-	Set SPI, Secure Register
0x0054	-	-	-	Reserved
0x0058	GICD_CLRSPI_NSR	WO	-	Clear SPI, Secure Register
0x005C	-	-	-	Reserved
0x0060-0xFFFC	-	-	-	Reserved

8.9 The GIC Distributor register descriptions

This section describes each of the GIC Distributor registers in register name order.

8.9.1 GICD_CLRSPI_NSR, Clear Non-secure SPI Pending Register

The GICD_CLRSPI_NSR characteristics are:

Purpose

Removes the pending state from a valid SPI if permitted by the Security state of the access and the GICD_NSACR<n> value for that SPI.

A write to this register changes the state of a pending SPI to inactive, and the state of an active and pending SPI to active.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WO	WO	WO

Writes to this register have no effect if:

- The value written specifies a Secure SPI, the value is written by a Non-secure access, and the value of the corresponding GICD_NSACR<n> register is less than `0b10`.
- The value written specifies an invalid SPI.
- The SPI is not pending.

16-bit accesses to bits [15:0] of this register must be supported.

Note

A Secure access to this register can clear the pending state of any valid SPI.

Configurations

If GICD_TYPER.MBIS == 0, this register is reserved.

In implementations that support only a single Security state, this register provides functionality for all SPIs.

Attributes

GICD_CLRSPI_NSR is a 32-bit register.

Field descriptions

The GICD_CLRSPI_NSR bit assignments are:

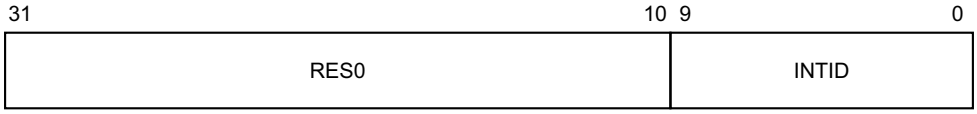

Bits [31:10]

Reserved, RES0.

INTID, bits [9:0]

The INTID of the SPI.

The function of this register depends on whether the targeted SPI is configured to be an edge-triggered or level-sensitive interrupt:

- For an edge-triggered interrupt, a write to GICD_SETSPI_NSR or GICD_SETSPI_SR adds the pending state to the targeted interrupt. It will stop being pending on activation, or if the pending state is removed by a write to GICD_CLRSPI_NSR, GICD_CLRSPI_SR, or GICD_ICPENDR<n>.
- For a level-sensitive interrupt, a write to GICD_SETSPI_NSR or GICD_SETSPI_SR adds the pending state to the targeted interrupt. It will remain pending until it is deasserted by a write to GICD_CLRSPI_NSR or GICD_CLRSPI_SR. If the interrupt is activated between having the pending state added and being deactivated, then the interrupt will be active and pending.

Accessing the GICD_CLRSPI_NSR:

GICD_CLRSPI_NSR can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	0x0048

8.9.2 GICD_CLRSPI_SR, Clear Secure SPI Pending Register

The GICD_CLRSPI_SR characteristics are:

Purpose

Removes the pending state from a valid SPI.

A write to this register changes the state of a pending SPI to inactive, and the state of an active and pending SPI to active.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WI	WO	WI

Writes to this register have no effect if:

- The value is written by a Non-secure access.
- The value written specifies an invalid SPI.
- The SPI is not pending.

16-bit accesses to bits [15:0] of this register must be supported.

Configurations

If GICD_TYPER.MBIS == 0, this register is reserved.

In implementations that support only a single Security state, this register is WI.

Attributes

GICD_CLRSPI_SR is a 32-bit register.

Field descriptions

The GICD_CLRSPI_SR bit assignments are:

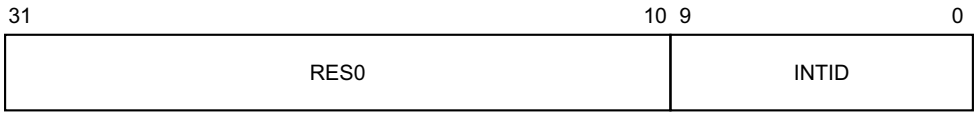

Bits [31:10]

Reserved, RES0.

INTID, bits [9:0]

The INTID of the SPI.

The function of this register depends on whether the targeted SPI is configured to be an edge-triggered or level-sensitive interrupt:

- For an edge-triggered interrupt, a write to GICD_SETSPI_NSR or GICD_SETSPI_SR adds the pending state to the targeted interrupt. It will stop being pending on activation, or if the pending state is removed by a write to GICD_CLRSPI_NSR, GICD_CLRSPI_SR, or GICD_ICPENDR<n>.
- For a level-sensitive interrupt, a write to GICD_SETSPI_NSR or GICD_SETSPI_SR adds the pending state to the targeted interrupt. It will remain pending until it is deasserted by a write to GICD_CLRSPI_NSR or GICD_CLRSPI_SR. If the interrupt is activated between having the pending state added and being deactivated, then the interrupt will be active and pending.

Accessing the GICD_CLRSPI_SR:

GICD_CLRSPI_SR can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	0x0058

8.9.3 GICD_CPENDSGIR<n>, SGI Clear-Pending Registers, n = 0 - 3

The GICD_CPENDSGIR<n> characteristics are:

Purpose

Removes the pending state from an SGI.

A write to this register changes the state of a pending SGI to inactive, and the state of an active and pending SGI to active.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

These registers are used only when affinity routing is not enabled. When affinity routing is enabled, this register is RES0. An implementation is permitted to make the register RAZ/WI in this case.

A register bit that corresponds to an unimplemented SGI is RAZ/WI.

These registers are byte-accessible.

If the GIC implementation supports two Security states:

- A register bit that corresponds to a Group 0 interrupt is RAZ/WI to Non-secure accesses.
- Register bits corresponding to unimplemented PEs are RAZ/WI.

Configurations

Four SGI clear-pending registers are implemented. Each register contains eight clear-pending bits for each of four SGIs, for a total of 16 possible SGIs.

In multiprocessor implementations, each PE has a copy of these registers.

Attributes

GICD_CPENDSGIR<n> is a 32-bit register.

Field descriptions

The GICD_CPENDSGIR<n> bit assignments are:

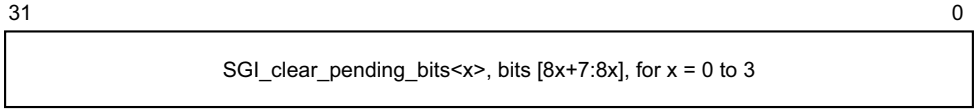

SGI_clear_pending_bits<x>, bits [8x+7:8x], for x = 0 to 3

Removes the pending state from SGI number 4n + x for the PE corresponding to the bit number written to.

Reads and writes have the following behavior:

`0` If read, indicates that the SGI from the corresponding PE is not pending and is not active and pending.
If written, has no effect.

`1` If read, indicates that the SGI from the corresponding PE is pending or is active and pending.
If written, removes the pending state from the SGI for the corresponding PE.

When this register has an architecturally-defined reset value, this field resets to `0`.

For SGI ID m, generated by processing element C writing to the corresponding GICD_SGIR field, where DIV and MOD are the integer division and modulo operations:

- The corresponding GICD_CPENDSGIR<n> number is given by n = m DIV 4.
- The offset of the required register is (0xF10 + (4n)).
- The offset of the required field within the register GICD_CPENDSGIR<n> is given by m MOD 4.
- The required bit in the 8-bit SGI clear-pending field m is bit C.

Accessing the GICD_CPENDSGIR<n>:

GICD_CPENDSGIR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	0x0F10 + 4n

8.9.4 GICD_CTLR, Distributor Control Register

The GICD_CTLR characteristics are:

Purpose

Enables interrupts and affinity routing.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

If an interrupt is pending within a CPU interface when the corresponding GICD_CTLR.EnableGrpX bit is written from 1 to 0 the interrupt must be retrieved from the CPU interface.

Note

This might have no effect on the forwarded interrupt if it has already been activated.

Configurations

The format of this register depends on the Security state of the access and the number of Security states supported, which is specified by GICD_CTLR.DS.

Attributes

GICD_CTLR is a 32-bit register.

Field descriptions

The GICD_CTLR bit assignments are:

When access is Secure, in a system that supports two Security states:

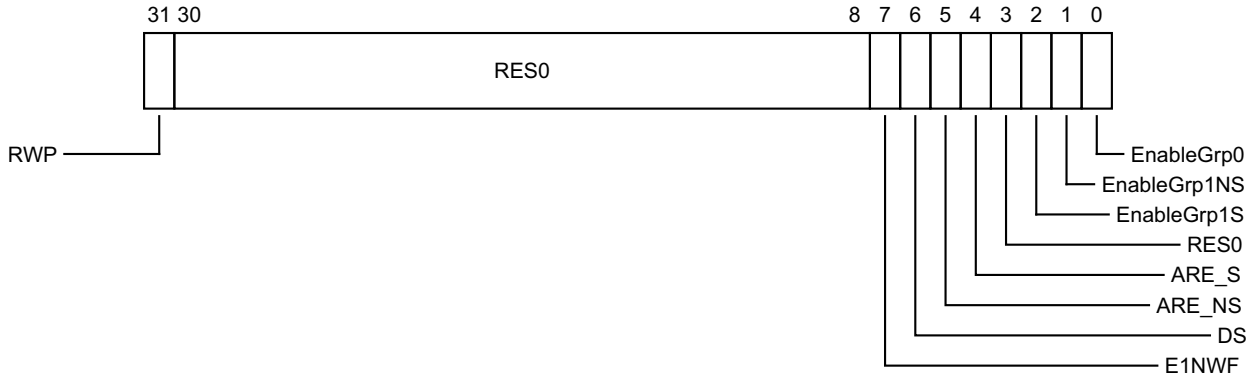

RWP, bit [31]

Register Write Pending. Read only. Iindicates whether a register write is in progress or not:

`0` No register write in progress. The effects of previous register writes to the affected register fields are visible to all logical components of the GIC architecture, including the CPU interfaces.

1 Register write in progress. The effects of previous register writes to the affected register fields are not guaranteed to be visible to all logical components of the GIC architecture, including the CPU interfaces, as the effects of the changes are still being propagated.

This field tracks updates to:

- GICD_CTLR[2:0], the Group Enables, for transitions from 1 to 0 only.
- GICD_CTLR[7:4], the ARE bits, E1NWF bit and DS bit.
- GICD_ICENABLER<n>, the bits that allow disabling of SPIs.

Updates to other register fields are not tracked by this field.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Bits [30:8]

Reserved, RES0.

E1NWF, bit [7]

Enable 1 of N Wakeup Functionality.

This bit is OPTIONAL in the architecture. If it is not implemented, it is RAZ/WI.

If it is implemented, then it has the following behavior:

0 A PE that is asleep cannot be picked for 1 of N interrupts.

1 A PE that is asleep can be picked for 1 of N interrupts as determined by IMPLEMENTATION DEFINED controls.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

DS, bit [6]

Disable Security.

0 Non-secure accesses are not permitted to access and modify registers that control Group 0 interrupts.

1 Non-secure accesses are permitted to access and modify registers that control Group 0 interrupts.

If DS is written from 0 to 1 when GICD_CTLR.ARE_S == 1, then GICD_CTLR.ARE for the single security state is RAO/WI.

If the Distributor only supports a single Security state, this bit is RAO/WI.

If the Distributor supports two Security states, it IMPLEMENTATION DEFINED whether this bit is programmable or implemented as RAZ/WI.

When this field is set to 1, all accesses to GICD_CTLR access the single Security state view, and all bits are accessible.

When set to 1, this field can only be cleared by a hardware reset.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

ARE_NS, bit [5]

Affinity Routing Enable, Non-secure state.

0 Affinity routing disabled for Non-secure state.

1 Affinity routing enabled for Non-secure state.

When affinity routing is enabled for the Secure state, this field is RAO/WI.

Changing the ARE_NS settings from 0 to 1 is UNPREDICTABLE except when GICD_CTLR.EnableGrp1 Non-Secure == 0.

Changing the ARE_NS settings from 1 to 0 is UNPREDICTABLE.

In systems where GICv2 backwards compatibility for the Non-secure state is implemented, this bit resets to 0. Otherwise this bit is RAO/WI.

ARE_S, bit [4]

Affinity Routing Enable, Secure state.

0 Affinity routing disabled for Secure state.

1 Affinity routing enabled for Secure state.

Changing the ARE_S setting from 0 to 1 is UNPREDICTABLE except when all of the following apply:

- GICD_CTLR.EnableGrp0==0.
- GICD_CTLR.EnableGrp1S==0.
- GICD_CTLR.EnableGrp1NS==0.

Changing the ARE_S settings from 1 to 0 is UNPREDICTABLE.

When this register has an architecturally-defined reset value, this field resets to 0.

Bit [3]

Reserved, RES0.

EnableGrp1S, bit [2]

Enable Secure Group 1 interrupts.

0 Secure Group 1 interrupts are disabled.

1 Secure Group 1 interrupts are enabled.

If GICD_CTLR.ARE_S == 0, this field is RES0.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

EnableGrp1NS, bit [1]

Enable Non-secure Group 1 interrupts.

0 Non-secure Group 1 interrupts are disabled.

1 Non-secure Group 1 interrupts are enabled.

Note

This field also controls whether LPIs are forwarded to the PE.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

EnableGrp0, bit [0]

Enable Group 0 interrupts.

0 Group 0 interrupts are disabled.

1 Group 0 interrupts are enabled.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

When access is Non-secure, in a system that supports two Security states:

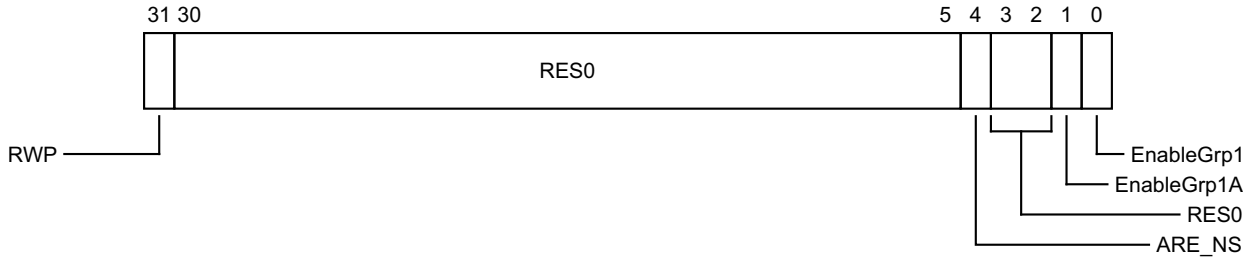

RWP, bit [31]

This bit is a read-only alias of the Secure GICD_CTLR.RWP bit.

Bits [30:5]

Reserved, RES0.

ARE_NS, bit [4]

This bit is a read-write alias of the Secure GICD_CTLR.ARE_NS bit.

Bits [3:2]

Reserved, RES0.

EnableGrp1A, bit [1]

If ARE_NS == 1, then this bit is a read-write alias of the Secure GICD_CTLR.EnableGrp1NS bit.

If ARE_NS == 0, then this bit is RES0.

EnableGrp1, bit [0]

If ARE_NS == 0, then this bit is a read-write alias of the Secure GICD_CTLR.EnableGrp1NS bit.

If ARE_NS == 1, then this bit is RES0.

When in a system that supports only a single Security state:

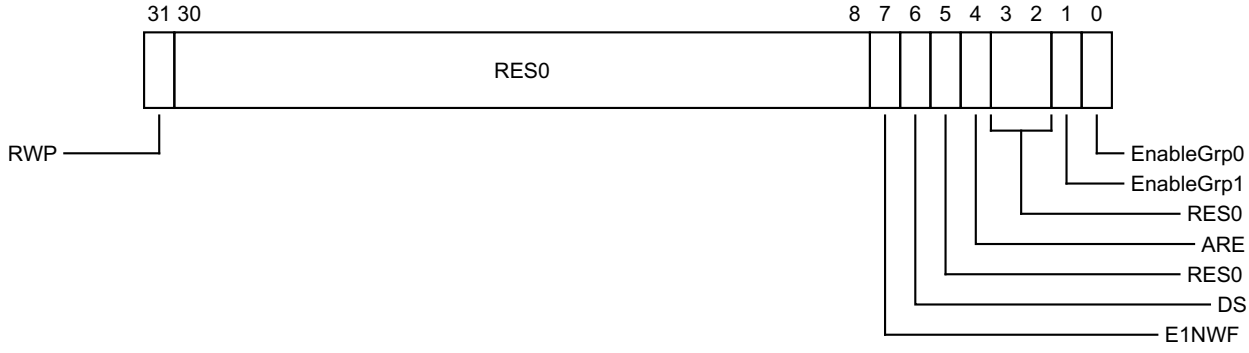

RWP, bit [31]

Register Write Pending. Read only. Indicates whether a register write is in progress or not:

`0` No register write in progress. The effects of previous register writes to the affected register fields are visible to all logical components of the GIC architecture, including the CPU interfaces.

1 Register write in progress. The effects of previous register writes to the affected register fields are not guaranteed to be visible to all logical components of the GIC architecture, including the CPU interfaces, as the effects of the changes are still being propagated.

This field tracks updates to:

- GICD_CTLR[2:0], the Group Enables, for transitions from 1 to 0 only.
- GICD_CTLR[7:4], the ARE bits, E1NWF bit and DS bit.
- GICD_ICENABLER<n>, the bits that allow disabling of SPIs.

Updates to other register fields are not tracked by this field.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Bits [30:8]

Reserved, RES0.

E1NWF, bit [7]

Enable 1 of N Wakeup Functionality.

This bit is OPTIONAL in the architecture. If it is not implemented, it is RAZ/WI.

If it is implemented, then it has the following behavior:

0 A PE that is asleep cannot be picked for 1 of N interrupts.

1 A PE that is asleep can be picked for 1 of N interrupts as determined by IMPLEMENTATION DEFINED controls.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

DS, bit [6]

Disable Security. This field is RAO/WI.

Bit [5]

Reserved, RES0.

ARE, bit [4]

Affinity Routing Enable.

0 Affinity routing disabled.

1 Affinity routing enabled.

Changing the ARE settings from 0 to 1 is UNPREDICTABLE except when all of the following apply:

- GICD_CTLR.EnableGrp1==0.
- GICD_CTLR.EnableGrp0==0.

Changing ARE from 1 to 0 is UNPREDICTABLE.

If GICv2 backwards compatibility is not implemented, this field is RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

Bits [3:2]

Reserved, RES0.

EnableGrp1, bit [1]

Enable Group 1 interrupts.

0 Group 1 interrupts disabled.

1 Group 1 interrupts enabled.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

EnableGrp0, bit [0]

Enable Group 0 interrupts.

`0` Group 0 interrupts are disabled.

`1` Group 0 interrupts are enabled.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICD_CTLR:

GICD_CTLR can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	`0x0000`

8.9.5 GICD_ICACTIVER<n>, Interrupt Clear-Active Registers, n = 0 - 31

The GICD_ICACTIVER<n> characteristics are:

Purpose

Deactivates the corresponding interrupt. These registers are used when saving and restoring GIC state.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

When affinity routing is enabled for the security state of an interrupt, the bits corresponding to SGIs and PPIs in that security state are RAZ/WI, and equivalent functionality for SGIs and PPIs is provided by GICR_ICACTIVER0.

Bits corresponding to unimplemented interrupts are RAZ/WI.

If GICD_CTLR.DS==0, unless the GICD_NSACR<n> registers permit Non-secure software to control Group 0 and Secure Group 1 interrupts, any bits that correspond to Group 0 or Secure Group 1 interrupts are accessible only by Secure accesses and are RAZ/WI to Non-secure accesses.

Configurations

These registers are available in all GIC configurations. If the GIC implementation supports two Security states, these registers are Common.

The number of implemented GICD_ICACTIVER<n> registers is (GICD_TYPER.ITLinesNumber+1). Registers are numbered from 0.

In a multiprocessor implementation, GICD_ICACTIVER0 is Banked for each connected PE. This register provides the Clear-active bits for interrupts 0-31.

Attributes

GICD_ICACTIVER<n> is a 32-bit register.

Field descriptions

The GICD_ICACTIVER<n> bit assignments are:

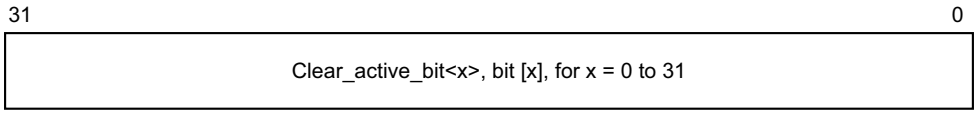

Clear_active_bit<x>, bit [x], for x = 0 to 31

Removes the active state from interrupt number 32n + x. Reads and writes have the following behavior:

`0` If read, indicates that the corresponding interrupt is not active, and is not active and pending.

If written, has no effect.

`1` If read, indicates that the corresponding interrupt is active, or is active and pending.

If written, deactivates the corresponding interrupt, if the interrupt is active. If the interrupt is already deactivated, the write has no effect.

After a write of 1 to this bit, a subsequent read of this bit returns 0.

When this register has an architecturally-defined reset value, this field resets to `0`.

For INTID m, when DIV and MOD are the integer division and modulo operations:

- The corresponding GICD_ICACTIVER<n> number, n, is given by n = m DIV 32.
- The offset of the required GICD_ICACTIVER is (0x380 + (4*n)).
- The bit number of the required group modifier bit in this register is m MOD 32.

Accessing the GICD_ICACTIVER<n>:

GICD_ICACTIVER<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	0x0380 + 4n

8.9.6 GICD_ICENABLER<n>, Interrupt Clear-Enable Registers, n = 0 - 31

The GICD_ICENABLER<n> characteristics are:

Purpose

Disables forwarding of the corresponding interrupt to the CPU interfaces.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

For SGIs and PPIs:

- When ARE is 1 for the Security state of an interrupt, the field for that interrupt is RES0 and an implementation is permitted to make the field RAZ/WI in this case.
- Equivalent functionality is provided by GICR_ICENABLER0.

Bits corresponding to unimplemented interrupts are RAZ/WI.

When GICD_CTLR.DS==0, bits corresponding to Group 0 and Secure Group 1 interrupts are RAZ/WI to Non-secure accesses.

It is IMPLEMENTATION DEFINED whether implemented SGIs are permanently enabled, or can be enabled and disabled by writes to GICD_ISENABLER<n> and GICD_ICENABLER<n> where n=0.

Completion of a write to this register does not guarantee that the effects of the write are visible throughout the affinity hierarchy. To ensure an enable has been cleared, software must write to the register with bits set to 1 to clear the required enables. Software must then poll GICD_CTLR.RWP until it has the value zero.

Configurations

These registers are available in all GIC configurations. If the GIC implementation supports two Security states, these registers are Common.

The number of implemented GICD_ICENABLER<n> registers is (GICD_TYPER.ITLinesNumber+1). Registers are numbered from 0.

In a multiprocessor implementation, GICD_ICENABLER0 is Banked for each connected PE. This register provides the Clear-enable bits for interrupts 0-31.

Attributes

GICD_ICENABLER<n> is a 32-bit register.

Field descriptions

The GICD_ICENABLER<n> bit assignments are:

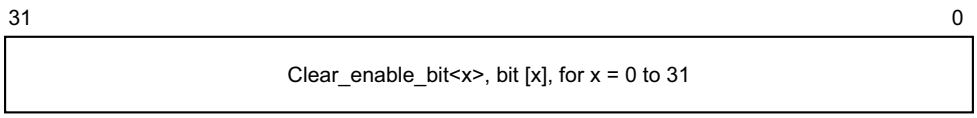

Clear_enable_bit<x>, bit [x], for x = 0 to 31

For SPIs and PPIs, controls the forwarding of interrupt number 32n + x to the CPU interfaces. Reads and writes have the following behavior:

`0` If read, indicates that forwarding of the corresponding interrupt is disabled.

If written, has no effect.

1 If read, indicates that forwarding of the corresponding interrupt is enabled.
If written, disables forwarding of the corresponding interrupt.
After a write of 1 to this bit, a subsequent read of this bit returns 0.

For SGIs, the behavior of this bit is IMPLEMENTATION DEFINED.

When this register has an architecturally-defined reset value, this field resets to an IMPLEMENTATION DEFINED value, that might be UNKNOWN.

For INTID m, when DIV and MOD are the integer division and modulo operations:

- The corresponding GICD_ICENABLER<n> number, n, is given by n = m DIV 32.
- The offset of the required GICD_ICENABLER is (0x180 + (4*n)).
- The bit number of the required group modifier bit in this register is m MOD 32.

Note

Writing a 1 to a GICD_ICENABLER<n> bit only disables the forwarding of the corresponding interrupt from the Distributor to any CPU interface. It does not prevent the interrupt from changing state, for example becoming pending or active and pending if it is already active.

Accessing the GICD_ICENABLER<n>:

GICD_ICENABLER<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	0x0180 + 4n

8.9.7 GICD_ICFGR<n>, Interrupt Configuration Registers, n = 0 - 63

The GICD_ICFGR<n> characteristics are:

Purpose

Determines whether the corresponding interrupt is edge-triggered or level-sensitive.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

For SGIs and PPIs:

- When ARE is 1 for the Security state of an interrupt, the field for that interrupt is RES0 and an implementation is permitted to make the field RAZ/WI in this case.
- Equivalent functionality is provided by GICR_ICFGR<n>.

For each supported PPI, it is IMPLEMENTATION DEFINED whether software can program the corresponding Int_config field.

For SGIs, Int_config fields are RO, meaning that GICD_ICFGR0 is RO.

Software must disable an interrupt before the value of the corresponding programmable Int_config field is changed. GIC behavior is otherwise UNPREDICTABLE.

Changing the interrupt configuration between level-sensitive and edge-triggered (in either direction) at a time when there is a pending interrupt will leave the interrupt in an UNKNOWN pending state.

Fields corresponding to unimplemented interrupts are RAZ/WI.

Configurations

These registers are available in all GIC configurations. If the GIC implementation supports two Security states, these registers are Common.

In a multiprocessor implementation, if Int_config[1] is programmable for any PPI, then GICD_ICFGR1 is Banked for each connected PE. This register holds the Int_config fields for the PPIs (interrupts 16-31).

Attributes

GICD_ICFGR<n> is a 32-bit register.

Field descriptions

The GICD_ICFGR<n> bit assignments are:

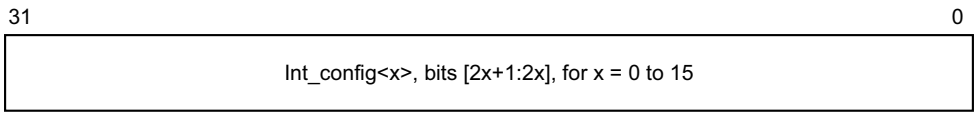

Int_config<x>, bits [2x+1:2x], for x = 0 to 15

Indicates whether the interrupt with ID 16n + x is level-sensitive or edge-triggered.

Int_config[0] (bit [2x]) is RES0.

Possible values of Int_config[1] (bit [2x+1]) are:

`0` Corresponding interrupt is level-sensitive.

1 Corresponding interrupt is edge-triggered.

For SGIs, Int_config[1] is RAO/WI.

For SPIs and PPIs, Int_config[1] is programmable unless the implementation supports two Security states and the bit corresponds to a Group 0 or Secure Group 1 interrupt, in which case the bit is RAZ/WI to Non-secure accesses.

When this register has an architecturally-defined reset value, this field resets to an IMPLEMENTATION DEFINED value, that might be UNKNOWN.

Accessing the GICD_ICFGR<n>:

GICD_ICFGR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	`0x0C00` + 4n

8.9.8 GICD_ICPENDR<n>, Interrupt Clear-Pending Registers, n = 0 - 31

The GICD_ICPENDR<n> characteristics are:

Purpose

Removes the pending state from the corresponding interrupt.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

Clear-pending bits for SGIs are RO/WI.

When affinity routing is enabled for the security state of an interrupt:

- Bits corresponding to SGIs and PPIs are RAZ/WI, and equivalent functionality for SGIs and PPIs is provided by GICR_ICPENDR0.
- Bits corresponding to Group 0 and Group 1 Secure interrupts can only be cleared by secure accesses.

Bits corresponding to unimplemented interrupts are RAZ/WI.

If GICD_CTLR.DS==0, unless the GICD_NSACR<n> registers permit Non-secure software to control Group 0 and Secure Group 1 interrupts, any bits that correspond to Group 0 or Secure Group 1 interrupts are accessible only by Secure accesses and are RAZ/WI to Non-secure accesses.

Configurations

These registers are available in all GIC configurations. If the GIC implementation supports two Security states, these registers are Common.

The number of implemented GICD_ICPENDR<n> registers is (GICD_TYPER.ITLinesNumber+1). Registers are numbered from 0.

In a multiprocessor implementation, GICD_ICPENDR0 is Banked for each connected PE. This register provides the Clear-pending bits for interrupts 0-31.

Attributes

GICD_ICPENDR<n> is a 32-bit register.

Field descriptions

The GICD_ICPENDR<n> bit assignments are:

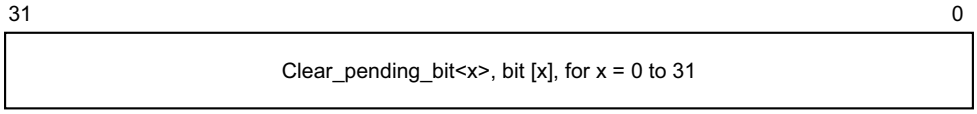

Clear_pending_bit<x>, bit [x], for x = 0 to 31

For SPIs and PPIs, removes the pending state from interrupt number 32n + x. Reads and writes have the following behavior:

`0` If read, indicates that the corresponding interrupt is not pending on any PE.
If written, has no effect.

`1` If read, indicates that the corresponding interrupt is pending, or active and pending:

- On this PE if the interrupt is an SGI or PPI.

- On at least one PE if the interrupt is an SPI.

If written, changes the state of the corresponding interrupt from pending to inactive, or from active and pending to active. This has no effect in the following cases:

- If the interrupt is an SGI. In this case, the write is ignored. The pending state of an SGI can be cleared using GICD_CPENDSGIR<n>.
- If the interrupt is not pending and is not active and pending.
- If the interrupt is a level-sensitive interrupt that is pending or active and pending for a reason other than a write to GICD_ISPENDR<n>. In this case, if the interrupt signal continues to be asserted, the interrupt remains pending or active and pending.

When this register has an architecturally-defined reset value, this field resets to `0`.

For INTID m, when DIV and MOD are the integer division and modulo operations:

- The corresponding GICD_ICPENDR<n> number, n, is given by n = m DIV 32.
- The offset of the required GICD_ICPENDR is (`0x200` + (4*n)).
- The bit number of the required group modifier bit in this register is m MOD 32.

Accessing the GICD_ICPENDR<n>:

GICD_ICPENDR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	`0x0280` + 4n

8.9.9 GICD_IGROUPR<n>, Interrupt Group Registers, n = 0 - 31

The GICD_IGROUPR<n> characteristics are:

Purpose

Controls whether the corresponding interrupt is in Group 0 or Group 1.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RAZ/WI

For SGIs and PPIs:

- When ARE is 1 for the Security state of an interrupt, the field for that interrupt is RES0 and an implementation is permitted to make the field RAZ/WI in this case.
- Equivalent functionality is provided by GICR_IGROUPR0.

When GICD_CTLR.DS==0, the register is RAZ/WI to Non-secure accesses.

Bits corresponding to unimplemented interrupts are RAZ/WI.

Note

Accesses to GICD_IGROUPR0 when affinity routing is not enabled for a Security state access the same state as GICR_IGROUPR0, and must update Redistributor state associated with the PE performing the accesses.

Implementations must ensure that an interrupt that is pending at the time of the write uses either the old value or the new value and must ensure that the interrupt is neither lost nor handled more than once. The effect of the change must be visible in finite time.

Configurations

These registers are available in all GIC configurations. If the GIC implementation supports two Security states, these registers are Secure.

The number of implemented GICD_IGROUPR<n> registers is (GICD_TYPER.ITLinesNumber+1). Registers are numbered from 0.

GICD_IGROUPR0 is Banked for each connected PE. This register provides the Group status bits for interrupts 0-31.

Attributes

GICD_IGROUPR<n> is a 32-bit register.

Field descriptions

The GICD_IGROUPR<n> bit assignments are:

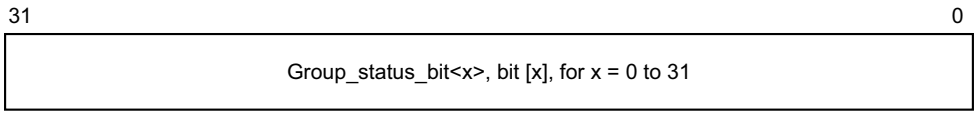

Group_status_bit<x>, bit [x], for x = 0 to 31

Group status bit.

0 When GICD_CTLR.DS==1, the corresponding interrupt is Group 0.
When GICD_CTLR.DS==0, the corresponding interrupt is Secure.

1 When GICD_CTLR.DS==0, the corresponding interrupt is Group 1.
When GICD_CTLR.DS==1, the corresponding interrupt is Non-secure.

If affinity routing is enabled for the Security state of an interrupt, the bit that corresponds to the interrupt is concatenated with the equivalent bit in GICD_IGRPMODR<n> to form a 2-bit field that defines an interrupt group. The encoding of this field is described in GICD_IGRPMODR<n>.

If affinity routing is disabled for the Security state of an interrupt, then:

- The corresponding GICD_IGRPMODR<n> bit is RES0.
- For Secure interrupts, the interrupt is Secure Group 0.
- For Non-secure interrupts, the interrupt is Non-secure Group 1.

For INTID m, when DIV and MOD are the integer division and modulo operations:

- The corresponding GICD_IGROUP<n> number, n, is given by n = m DIV 32.
- The offset of the required GICD_IGROUP is (`0x080` + (4*n)).
- The bit number of the required group modifier bit in this register is m MOD 32.

Typically, the reset value of all GICD_IGROUPR<n> and GICD_IGRPMODR<n> registers is 0, so that all interrupts are Group 0 unless reprogrammed as Group 1 by Secure accesses to the appropriate registers.

Accessing the GICD_IGROUPR<n>:

GICD_IGROUPR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	`0x0080` + 4n

8.9.10 GICD_IGRPMODR<n>, Interrupt Group Modifier Registers, n = 0 - 31

The GICD_IGRPMODR<n> characteristics are:

Purpose

In GIC implementations which support two Security states, along with the GICD_IGROUPR<n> registers, controls whether the corresponding interrupt is in:

- Secure Group 0.
- Non-secure Group 1.
- Secure Group 1.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RAZ/WI

When affinity routing is enabled for Secure state, GICD_IGRPMODR0 is RES0 and equivalent functionality is proved by GICR_IGRPMODR0.

When GICD_CTLR.DS==0, the register is RAZ/WI to Non-secure accesses.

In implementations that support a single Security state, Secure Group 1 interrupts are treated as Group 0 accesses.

Note

Implementations must ensure that an interrupt that is pending at the time of the write uses either the old value or the new value and must ensure that the interrupt is neither lost nor handled more than once. The effect of the change must be visible in finite time.

Configurations

If the GIC implementation supports two Security states, these registers are Secure.

The number of implemented GICD_IGROUPR<n> registers is (GICD_TYPER.ITLinesNumber+1). Registers are numbered from 0.

GICD_IGRPMODR0 is Banked for each connected PE. This register provides the Group modifier bits for interrupts 0-31.

When GICD_CTLR.ARE_S==0 or GICD_CTLR.DS==1, the GICD_IGRPMODR<n> registers are RES0. An implementation can make these registers RAZ/WI in this case.

Attributes

GICD_IGRPMODR<n> is a 32-bit register.

Field descriptions

The GICD_IGRPMODR<n> bit assignments are:

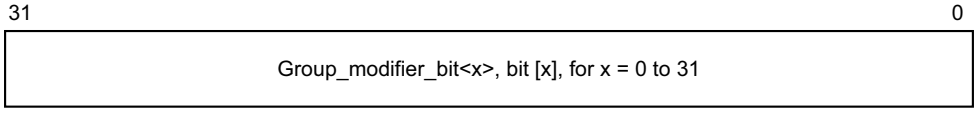

Group_modifier_bit<x>, bit [x], for x = 0 to 31

Group modifier bit. When affinity routing is enabled for the Security state of an interrupt, the bit that corresponds to the interrupt is concatenated with the equivalent bit in GICD_IGROUPR<n> to form a 2-bit field that defines an interrupt group:

Group modifier bit	Group status bit	Definition	Short name
0	0	Secure Group 0	G0S
0	1	Non-secure Group 1	G1NS
1	0	Secure Group 1	G1S
1	1	Reserved, treated as Non-secure Group 1	-

When this register has an architecturally-defined reset value, this field resets to `0`.

For INTID m, when DIV and MOD are the integer division and modulo operations:

- The corresponding GICD_IGRPMODR<n> number, n, is given by n = m DIV 32.
- The offset of the required GICD_IGRPMODR is (`0x080` + (4*n)).
- The bit number of the required group modifier bit in this register is m MOD 32.

See GICD_IGROUPR<n> for information about the GICD_IGRPMODR0 reset value.

Accessing the GICD_IGRPMODR<n>:

GICD_IGRPMODR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	`0x0D00` + 4n

8.9.11 GICD_IIDR, Distributor Implementer Identification Register

The GICD_IIDR characteristics are:

Purpose

Provides information about the implementer and revision of the Distributor.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

Configurations

This register is available in all configurations of the GIC. If the GIC implementation supports two Security states, this register is Common.

Attributes

GICD_IIDR is a 32-bit register.

Field descriptions

The GICD_IIDR bit assignments are:

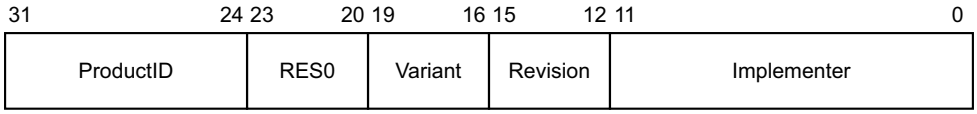

ProductID, bits [31:24]

An IMPLEMENTATION DEFINED product identifier.

Bits [23:20]

Reserved, RES0.

Variant, bits [19:16]

An IMPLEMENTATION DEFINED variant number. Typically, this field is used to distinguish product variants, or major revisions of a product.

Revision, bits [15:12]

An IMPLEMENTATION DEFINED revision number. Typically, this field is used to distinguish minor revisions of a product.

Implementer, bits [11:0]

Contains the JEP106 code of the company that implemented the Distributor:

- Bits [11:8] are the JEP106 continuation code of the implementer. For an ARM implementation, this field is `0x4`.
- Bit [7] is always 0.
- Bits [6:0] are the JEP106 identity code of the implementer. For an ARM implementation, bits [7:0] are therefore `0x3B`.

Accessing the GICD_IIDR:

GICD_IIDR can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	0x0008

8.9.12 GICD_IPRIORITYR<n>, Interrupt Priority Registers, n = 0 - 254

The GICD_IPRIORITYR<n> characteristics are:

Purpose

Holds the priority of the corresponding interrupt.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

These registers are always used when affinity routing is not enabled. When affinity routing is enabled for the security state of an interrupt:

- GICR_IPRIORITYR<n> is used instead of GICD_IPRIORITYR<n> where n = 0 to 7 (that is, for SGIs and PPIs).
- GICD_IPRIORITYR<n> is RAZ/WI where n = 0 to 7.

These registers are byte-accessible.

A register field corresponding to an unimplemented interrupt is RAZ/WI.

A GIC might implement fewer than eight priority bits, but must implement at least bits [7:4] of each field. In each field, unimplemented bits are RAZ/WI (*Interrupt prioritization* on page 4-65).

If the GIC implementation supports only a single Security state, the full priority range is programmable by Non-secure accesses.

If the GIC implementation supports two Security states:

- A register bit that corresponds to a Group 0 or Secure Group 1 interrupt is RAZ/WI to Non-secure accesses.
- A Non-secure access to a field that corresponds to a Non-secure Group 1 interrupt behaves as described in *Software accesses of interrupt priority* on page 4-72.

It is IMPLEMENTATION DEFINED whether changing the value of a priority field changes the priority of an active interrupt.

Note

Implementations must ensure that an interrupt that is pending at the time of the write uses either the old value or the new value and must ensure that the interrupt is neither lost nor handled more than once. The effect of the change must be visible in finite time.

Configurations

These registers are available in all configurations of the GIC. If the GIC implementation supports two Security states, these registers are Common.

The number of implemented GICD_IPRIORITYR<n> registers is 8*(GICD_TYPER.ITLinesNumber+1). Registers are numbered from 0.

In a multiprocessor implementation, GICD_IPRIORITYR0 to GICD_IPRIORITYR7 are Banked for each connected PE. These registers provide the Priority fields for interrupts 0-31.

Attributes

GICD_IPRIORITYR<n> is a 32-bit register.

Field descriptions

The GICD_IPRIORITYR<n> bit assignments are:

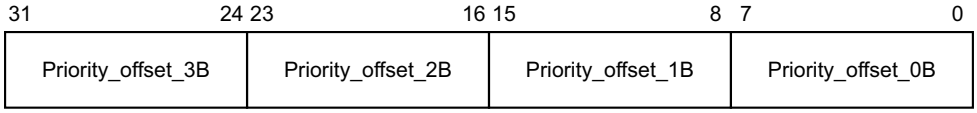

Priority_offset_3B, bits [31:24]

Interrupt priority value from an IMPLEMENTATION DEFINED range, at byte offset 3. Lower priority values correspond to greater priority of the interrupt.

When this register has an architecturally-defined reset value, this field resets to `0`.

Priority_offset_2B, bits [23:16]

Interrupt priority value from an IMPLEMENTATION DEFINED range, at byte offset 2. Lower priority values correspond to greater priority of the interrupt.

When this register has an architecturally-defined reset value, this field resets to `0`.

Priority_offset_1B, bits [15:8]

Interrupt priority value from an IMPLEMENTATION DEFINED range, at byte offset 1. Lower priority values correspond to greater priority of the interrupt.

When this register has an architecturally-defined reset value, this field resets to `0`.

Priority_offset_0B, bits [7:0]

Interrupt priority value from an IMPLEMENTATION DEFINED range, at byte offset 0. Lower priority values correspond to greater priority of the interrupt.

When this register has an architecturally-defined reset value, this field resets to `0`.

For interrupt ID m, when DIV and MOD are the integer division and modulo operations:

- The corresponding GICD_IPRIORITYR<n> number, n, is given by n = m DIV 4.
- The offset of the required GICD_IPRIORITYR<n> register is (`0x400` + (4*n)).
- The byte offset of the required Priority field in this register is m MOD 4, where:
  - Byte offset 0 refers to register bits [7:0].
  - Byte offset 1 refers to register bits [15:8].
  - Byte offset 2 refers to register bits [23:16].
  - Byte offset 3 refers to register bits [31:24].

Accessing the GICD_IPRIORITYR<n>:

GICD_IPRIORITYR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	`0x0400` + 4n

8.9.13 GICD_IROUTER<n>, Interrupt Routing Registers, n = 32 - 1019

The GICD_IROUTER<n> characteristics are:

Purpose

When affinity routing is enabled, provides routing information for the SPI with INTID n.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

These registers are used only when affinity routing is enabled. When affinity routing is not enabled:

- These registers are RES0. An implementation is permitted to make the register RAZ/WI in this case.
- The GICD_ITARGETSR<n> registers provide interrupt routing information.

Note

When affinity routing becomes enabled for a Security state (for example, following a reset or following a write to GICD_CTLR) the value of all writeable fields in this register is UNKNOWN for that Security state. When the group of an interrupt changes so the ARE setting for the interrupt changes to 1, the value of this register is UNKNOWN for that interrupt.

If GICD_CTLR.DS==0, unless the GICD_NSACR<n> registers permit Non-secure software to control Group 0 and Secure Group 1 interrupts, any GICD_IROUTER<n> registers that correspond to Group 0 or Secure Group 1 interrupts are accessible only by Secure accesses and are RAZ/WI to Non-secure accesses.

Note

For each interrupt, a GIC implementation might support fewer than 256 values for an affinity level. In this case, some bits of the corresponding affinity level field might be RO.

Implementations must ensure that an interrupt that is pending at the time of the write uses either the old value or the new value and must ensure that the interrupt is neither lost nor handled more than once. The effect of the change must be visible in finite time.

Configurations

These registers are available in all configurations of the GIC. If the GIC implementation supports two Security states, these registers are Common.

The maximum value of n is given by (32*(GICD_TYPER.ITLinesNumber+1) - 1). GICD_IROUTER<n> registers where n=0 to 31 are reserved.

Attributes

GICD_IROUTER<n> is a 64-bit register.

Field descriptions

The GICD_IROUTER<n> bit assignments are:

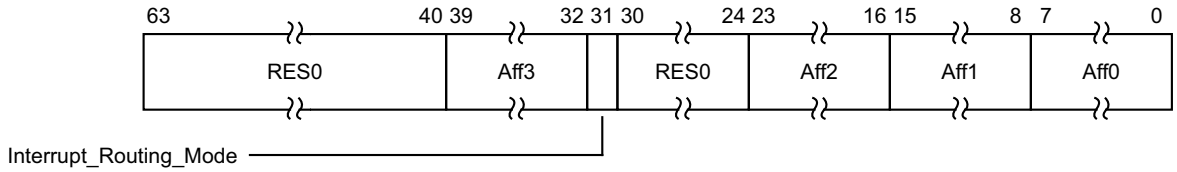

Bits [63:40]

Reserved, RES0.

Aff3, bits [39:32]

Affinity level 3, the least significant affinity level field.

Interrupt_Routing_Mode, bit [31]

Interrupt Routing Mode. Defines how SPIs are routed in an affinity hierarchy:

`0` Interrupts routed to the PE specified by a.b.c.d. In this routing, a, b, c, and d are the values of fields Aff3, Aff2, Aff1, and Aff0 respectively.

`1` Interrupts routed to any PE defined as a participating node.

If GICD_IROUTER<n>.IRM == 0 and the affinity path does not correspond to an implemented PE, then if the corresponding interrupt becomes pending it will not be forwarded to any PE and will remain pending.

In implementations that do not require 1 of N distribution of SPIs, this bit might be RAZ/WI.

When this bit is set to 1, GICD_IROUTER<n>.{Aff3, Aff2, Aff1, Aff0} are UNKNOWN.

Note

An implementation might choose to make the Aff<n> fields RO when this field is 1.

Bits [30:24]

Reserved, RES0.

Aff2, bits [23:16]

Affinity level 2, an intermediate affinity level field.

Aff1, bits [15:8]

Affinity level 1, an intermediate affinity level field.

Aff0, bits [7:0]

Affinity level 0, the most significant affinity level field.

For an SPI with INTID m:

- The corresponding GICD_IROUTER<n> register number, n, is given by n = m.
- The offset of the GICD_IROUTER<n> register is `0x6000` + 8n.

Accessing the GICD_IROUTER<n>:

GICD_IROUTER<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	`0x6000` + 8n

8.9.14 GICD_ISACTIVER<n>, Interrupt Set-Active Registers, n = 0 - 31

The GICD_ISACTIVER<n> characteristics are:

Purpose

Activates the corresponding interrupt. These registers are used when saving and restoring GIC state.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

When affinity routing is enabled for the security state of an interrupt, bits corresponding to SGIs and PPIs are RAZ/WI, and equivalent functionality for SGIs and PPIs is provided by GICR_ISACTIVER0.

Bits corresponding to unimplemented interrupts are RAZ/WI.

If GICD_CTLR.DS==0, unless the GICD_NSACR<n> registers permit Non-secure software to control Group 0 and Secure Group 1 interrupts, any bits that correspond to Group 0 or Secure Group 1 interrupts are accessible only by Secure accesses and are RAZ/WI to Non-secure accesses.

The bit reads as one if the status of the interrupt is active or active and pending. GICD_ISPENDR<n> and GICD_ICPENDR<n> provide the pending status of the interrupt.

Configurations

These registers are available in all GIC configurations. If the GIC implementation supports two Security states, these registers are Common.

The number of implemented GICD_ISACTIVER<n> registers is (GICD_TYPER.ITLinesNumber+1). Registers are numbered from 0.

In a multiprocessor implementation, GICD_ISACTIVER0 is Banked for each connected PE. This register provides the Set-active bits for interrupts 0-31.

Attributes

GICD_ISACTIVER<n> is a 32-bit register.

Field descriptions

The GICD_ISACTIVER<n> bit assignments are:

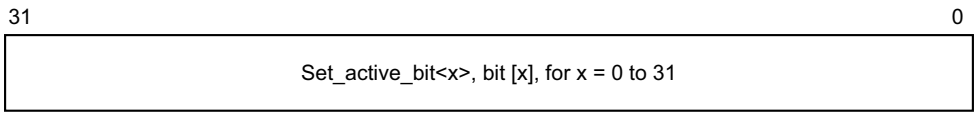

Set_active_bit<x>, bit [x], for x = 0 to 31

Adds the active state to interrupt number 32n + x. Reads and writes have the following behavior:

`0` If read, indicates that the corresponding interrupt is not active, and is not active and pending.
If written, has no effect.

`1` If read, indicates that the corresponding interrupt is active, or is active and pending.
If written, activates the corresponding interrupt, if the interrupt is not already active. If the interrupt is already active, the write has no effect.
After a write of 1 to this bit, a subsequent read of this bit returns 1.

When this register has an architecturally-defined reset value, this field resets to `0`.

For INTID m, when DIV and MOD are the integer division and modulo operations:

- The corresponding GICD_ISACTIVER<n> number, n, is given by n = m DIV 32.
- The offset of the required GICD_ISACTIVER is (`0x300` + (4*n)).
- The bit number of the required group modifier bit in this register is m MOD 32.

Accessing the GICD_ISACTIVER<n>:

GICD_ISACTIVER<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	`0x0300` + 4n

8.9.15 GICD_ISENABLER<n>, Interrupt Set-Enable Registers, n = 0 - 31

The GICD_ISENABLER<n> characteristics are:

Purpose

Enables forwarding of the corresponding interrupt to the CPU interfaces.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

For SGIs and PPIs:

- When ARE is 1 for the Security state of an interrupt, the field for that interrupt is RES0 and an implementation is permitted to make the field RAZ/WI in this case.
- Equivalent functionality is provided by GICR_ISENABLER0.

Bits corresponding to unimplemented interrupts are RAZ/WI.

When GICD_CTLR.DS==0, bits corresponding to Group 0 or Secure Group 1 interrupts are RAZ/WI to Non-secure accesses.

It is IMPLEMENTATION DEFINED whether implemented SGIs are permanently enabled, or can be enabled and disabled by writes to GICD_ISENABLER<n> and GICD_ICENABLER<n> where n=0.

For SPIs and PPIs, each bit controls the forwarding of the corresponding interrupt from the Distributor to the CPU interfaces.

Configurations

These registers are available in all GIC configurations. If the GIC implementation supports two Security states, these registers are Common.

The number of implemented GICD_ISENABLER<n> registers is (GICD_TYPER.ITLinesNumber+1). Registers are numbered from 0.

In a multiprocessor implementation, GICD_ISENABLER0 is Banked for each connected PE. This register provides the Clear-enable bits for interrupts 0-31.

Attributes

GICD_ISENABLER<n> is a 32-bit register.

Field descriptions

The GICD_ISENABLER<n> bit assignments are:

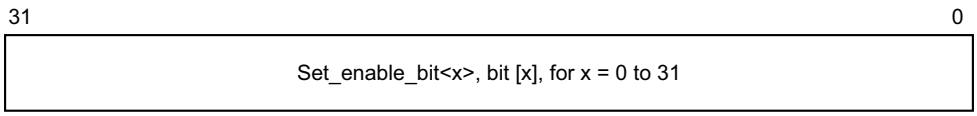

Set_enable_bit<x>, bit [x], for x = 0 to 31

For SPIs and PPIs, controls the forwarding of interrupt number 32n + x to the CPU interfaces. Reads and writes have the following behavior:

`0` If read, indicates that forwarding of the corresponding interrupt is disabled.
If written, has no effect.

`1` If read, indicates that forwarding of the corresponding interrupt is enabled.

If written, enables forwarding of the corresponding interrupt.

After a write of 1 to this bit, a subsequent read of this bit returns 1.

For SGIs, the behavior of this bit is IMPLEMENTATION DEFINED.

When this register has an architecturally-defined reset value, this field resets to an IMPLEMENTATION DEFINED value, that might be UNKNOWN.

For INTID m, when DIV and MOD are the integer division and modulo operations:

- The corresponding GICD_ISENABLER<n> number, n, is given by n = m DIV 32.
- The offset of the required GICD_ISENABLER is (`0x100` + (4*n)).
- The bit number of the required group modifier bit in this register is m MOD 32.

At start-up, and after a reset, a PE can use this register to discover which peripheral INTIDs the GIC supports. If GICD_CTLR.DS==0 in a system that supports EL3, the PE must do this for the Secure view of the available interrupts, and Non-secure software running on the PE must do this discovery after the Secure software has configured interrupts as Group 0/Secure Group 1 and Non-secure Group 1.

Accessing the GICD_ISENABLER<n>:

GICD_ISENABLER<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	`0x0100` + 4n

8.9.16 GICD_ISPENDR<n>, Interrupt Set-Pending Registers, n = 0 - 31

The GICD_ISPENDR<n> characteristics are:

Purpose

Adds the pending state to the corresponding interrupt.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

Set-pending bits for SGIs are read-only and ignore writes. The Set-pending bits for SGIs are provided as GICD_SPENDSGIR<n>.

When affinity routing is enabled for the security state of an interrupt:

- Bits corresponding to SGIs and PPIs are RAZ/WI, and equivalent functionality for SGIs and PPIs is provided by GICR_ISPENDR0.
- Bits corresponding to Group 0 and Group 1 Secure interrupts can only be set by secure accesses.

Bits corresponding to unimplemented interrupts are RAZ/WI.

If GICD_CTLR.DS==0, unless the GICD_NSACR<n> registers permit Non-secure software to control Group 0 and Secure Group 1 interrupts, any bits that correspond to Group 0 or Secure Group 1 interrupts are accessible only by Secure accesses and are RAZ/WI to Non-secure accesses.

Configurations

These registers are available in all GIC configurations. If the GIC implementation supports two Security states, these registers are Common.

The number of implemented GICD_ISPENDR<n> registers is (GICD_TYPER.ITLinesNumber+1). Registers are numbered from 0.

In a multiprocessor implementation, GICD_ISPENDR0 is Banked for each connected PE. This register provides the Set-pending bits for interrupts 0-31.

Attributes

GICD_ISPENDR<n> is a 32-bit register.

Field descriptions

The GICD_ISPENDR<n> bit assignments are:

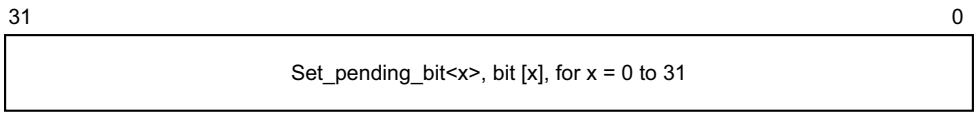

Set_pending_bit<x>, bit [x], for x = 0 to 31

For SPIs and PPIs, adds the pending state to interrupt number 32n + x. Reads and writes have the following behavior:

`0` If read, indicates that the corresponding interrupt is not pending on any PE.
If written, has no effect.

`1` If read, indicates that the corresponding interrupt is pending, or active and pending:

- On this PE if the interrupt is an SGI or PPI.

- On at least one PE if the interrupt is an SPI.

If written, changes the state of the corresponding interrupt from inactive to pending, or from active to active and pending. This has no effect in the following cases:

- If the interrupt is an SGI. The pending state of an SGI can be set using GICD_SPENDSGIR<n>.
- If the interrupt is not inactive and is not active.
- If the interrupt is already pending because of a write to GICD_ISPENDR<n>.
- If the interrupt is already pending because the corresponding interrupt signal is asserted. In this case, the interrupt remains pending if the interrupt signal is deasserted.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the GICD_ISPENDR<n>:

GICD_ISPENDR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	`0x0200` + 4n

8.9.17 GICD_ITARGETSR<n>, Interrupt Processor Targets Registers, n = 0 - 254

The GICD_ITARGETSR<n> characteristics are:

Purpose

When affinity routing is not enabled, holds the list of target PEs for the interrupt. That is, it holds the list of CPU interfaces to which the Distributor forwards the interrupt if it is asserted and has sufficient priority.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

These registers are used when affinity routing is not enabled. When affinity routing is enabled for the security state of an interrupt, the target PEs for an interrupt are defined by GICD_IROUTER<n> and the associated byte in GICD_ITARGETSR<n> is RES0. An implementation is permitted to make the byte RAZ/WI in this case.

- These registers are byte-accessible.
- A register field corresponding to an unimplemented interrupt is RAZ/WI.
- A field bit corresponding to an unimplemented CPU interface is RAZ/WI.
- GICD_ITARGETSR0-GICD_ITARGETSR7 are read-only. Each field returns a value that corresponds only to the PE reading the register.
- It is IMPLEMENTATION DEFINED which, if any, SPIs are statically configured in hardware. The field for such an SPI is read-only, and returns a value that indicates the PE targets for the interrupt.
- If GICD_CTLR.DS==0, unless the GICD_NSACR<n> registers permit Non-secure software to control Group 0 and Secure Group 1 interrupts, any bits that correspond to Group 0 or Secure Group 1 interrupts are accessible only by Secure accesses and are RAZ/WI to Non-secure accesses.

In a single connected PE implementation, all interrupts target one PE, and these registers are RAZ/WI.

Note

Implementations must ensure that an interrupt that is pending at the time of the write uses either the old value or the new value and must ensure that the interrupt is neither lost nor handled more than once. The effect of the change must be visible in finite time.

Configurations

These registers are available in all configurations of the GIC. If the GIC implementation supports two Security states, these registers are Common.

The number of implemented GICD_ITARGETSR<n> registers is 8*(GICD_TYPER.ITLinesNumber+1). Registers are numbered from 0.

In a multiprocessor implementation, GICD_ITARGETSR0 to GICD_ITARGETSR7 are Banked for each connected PE. These registers provide the PE targets fields for interrupts 0-31.

Attributes

GICD_ITARGETSR<n> is a 32-bit register.

Field descriptions

The GICD_ITARGETSR<n> bit assignments are:

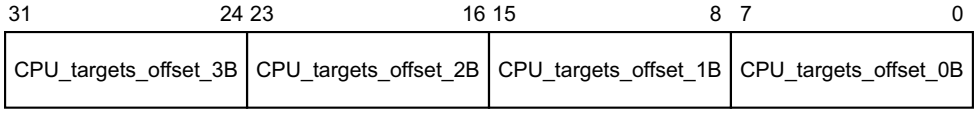

PEs in the system number from 0, and each bit in a PE targets field refers to the corresponding PE. For example, a value of `0x3` means that the Pending interrupt is sent to PEs 0 and 1. For GICD_ITARGETSR0-GICD_ITARGETSR7, a read of any targets field returns the number of the PE performing the read.

CPU_targets_offset_3B, bits [31:24]

PE targets for an interrupt, at byte offset 3.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

CPU_targets_offset_2B, bits [23:16]

PE targets for an interrupt, at byte offset 2.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

CPU_targets_offset_1B, bits [15:8]

PE targets for an interrupt, at byte offset 1.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

CPU_targets_offset_0B, bits [7:0]

PE targets for an interrupt, at byte offset 0.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

The bits that are set to 1 in the PE targets field determine which PEs are targeted:

Value of PE targets field	Interrupt targets
0bxxxxxxx1	CPU interface 0
0bxxxxxx1x	CPU interface 1
0bxxxxx1xx	CPU interface 2
0bxxxx1xxx	CPU interface 3
0bxxx1xxxx	CPU interface 4
0bxx1xxxxx	CPU interface 5
0bx1xxxxxx	CPU interface 6
0b1xxxxxxx	CPU interface 7

For interrupt ID m, when DIV and MOD are the integer division and modulo operations:

- The corresponding GICD_ITARGETSR<n> number, n, is given by n = m DIV 4.
- The offset of the required GICD_ITARGETSR<n> register is (`0x800` + (4*n)).
- The byte offset of the required Priority field in this register is m MOD 4, where:
  - Byte offset 0 refers to register bits [7:0].

— Byte offset 1 refers to register bits [15:8].

— Byte offset 2 refers to register bits [23:16].

— Byte offset 3 refers to register bits [31:24].

Software can write to these registers at any time. Any change to a targets field value:

- Has no effect on any active interrupt. This means that removing a CPU interface from a targets list does not cancel an active state for interrupts on that CPU interface. There is no effect on interrupts that are active and pending until the active status is cleared, at which time it is treated as a pending interrupt.
- Has an effect on any pending interrupts. This means:
  — Adding a CPU interface to the target list of a pending interrupt makes the interrupt pending on that CPU interface.
  — Removing a CPU interface from the target list of a pending interrupt removes the pending state of the interrupt on that CPU interface.

Accessing the GICD_ITARGETSR<n>:

GICD_ITARGETSR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	`0x0800` + 4n

8.9.18 GICD_NSACR<n>, Non-secure Access Control Registers, n = 0 - 63

The GICD_NSACR<n> characteristics are:

Purpose

Enables Secure software to permit Non-secure software on a particular PE to create and control Group 0 interrupts.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RAZ/WI	RW	RAZ/WI

These registers are implemented only in GIC implementations that support two Security states. In implementations that support only a single Security state, the registers are not implemented, and the corresponding address space is reserved.

These registers are Secure, and are RAZ/WI to Non-secure accesses.

These registers are always used when affinity routing is not enabled. When affinity routing is enabled for the Secure state, GICD_NSACR0 is RES0 and GICR_NSACR provides equivalent functionality for SGIs.

These registers do not support PPIs, therefore GICD_NSACR1 is RAZ/WI.

Configurations

The concept of selective enabling of Non-secure access to Group 0 and Secure Group 1 interrupts applies to SGIs and SPIs.

GICD_NSACR0 is a Banked register used for SGIs. A copy is provided for every PE that has a CPU interface and that supports this feature.

Attributes

GICD_NSACR<n> is a 32-bit register.

Field descriptions

The GICD_NSACR<n> bit assignments are:

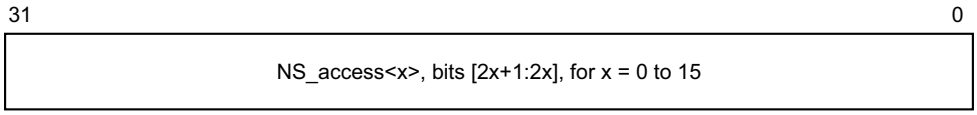

NS_access<x>, bits [2x+1:2x], for x = 0 to 15

Controls Non-secure access of the interrupt with ID 16n + x.

If the corresponding interrupt does not support configurable Non-secure access, the field is RAZ/WI.

Otherwise, the field is RW and determines the level of Non-secure control permitted if the interrupt is a Secure interrupt. If the interrupt is a Non-secure interrupt, this field is ignored.

The possible values of each 2-bit field are:

`00` No Non-secure access is permitted to fields associated with the corresponding interrupt.

`01` Non-secure read and write access is permitted to set-pending bits in GICD_ISPENDR<n> associated with the corresponding interrupt. A Non-secure write access to GICD_SETSPI_NSR is permitted to set the pending state of the corresponding interrupt. A Non-secure write access to GICD_SGIR is permitted to generate a Secure SGI for the corresponding interrupt.

An implementation might also provide read access to clear-pending bits in GICD_ICPENDR<n> associated with the corresponding interrupt.

`10` As `01`, but adds Non-secure read and write access permission to fields associated with the corresponding interrupt in the GICD_ICPENDR<n> registers. A Non-secure write access to GICD_CLRSPI_NSR is permitted to clear the pending state of the corresponding interrupt. Also adds Non-secure read access permission to fields associated with the corresponding interrupt in the GICD_ISACTIVER<n> and GICD_ICACTIVER<n> registers.

`11` For GICD_NSACR0 this encoding is reserved and treated as `10`.

For all other GICD_NSACR<n> registers this encoding is treated as `10`, but adds Non-secure read and write access permission to GICD_ITARGETSR<n> and GICD_IROUTER<n> fields associated with the corresponding interrupt.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

For interrupt ID m, when DIV and MOD are the integer division and modulo operations:

- The corresponding GICD_NSACR<n> number, n, is given by n = m DIV 16.
- The offset of the required GICD_NSACR<n> register is (`0xE00` + (4*n)).

Note

Because each field in this register comprises two bits, GICD_NSACR0 controls access rights to SGI registers, GICD_NSACR1 controls access to PPI registers (and is always RAZ/WI), and all other GICD_NSACR<n> registers control access to SPI registers.

For compatibility with GICv2, writes to GICD_NSACR0 for a particular PE must be coordinated within the Distributor and must update GICR_NSACR for the Redistributor associated with that PE.

Accessing the GICD_NSACR<n>:

GICD_NSACR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	`0x0E00` + 4n

8.9.19 GICD_SETSPI_NSR, Set Non-secure SPI Pending Register

The GICD_SETSPI_NSR characteristics are:

Purpose

Adds the pending state to a valid SPI if permitted by the Security state of the access and the GICD_NSACR<n> value for that SPI.

A write to this register changes the state of an inactive SPI to pending, and the state of an active SPI to active and pending.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WO	WO	WO

Writes to this register have no effect if:

- The value written specifies a Secure SPI, the value is written by a Non-secure access, and the value of the corresponding GICD_NSACR<n> register is 0.
- The value written specifies an invalid SPI.
- The SPI is already pending.

16-bit accesses to bits [15:0] of this register must be supported.

Note

A Secure access to this register can set the pending state of any valid SPI.

Configurations

If GICD_TYPER.MBIS == 0, this register is reserved.

In implementations that support only a single Security state, this register provides functionality for all SPIs.

Attributes

GICD_SETSPI_NSR is a 32-bit register.

Field descriptions

The GICD_SETSPI_NSR bit assignments are:

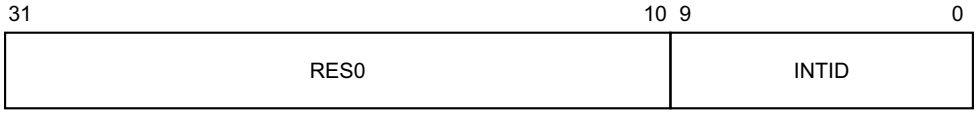

Bits [31:10]

Reserved, RES0.

INTID, bits [9:0]

The INTID of the SPI.

The function of this register depends on whether the targeted SPI is configured to be an edge-triggered or level-sensitive interrupt:

- For an edge-triggered interrupt, a write to GICD_SETSPI_NSR or GICD_SETSPI_SR adds the pending state to the targeted interrupt. It will stop being pending on activation, or if the pending state is removed by a write to GICD_CLRSPI_NSR, GICD_CLRSPI_SR, or GICD_ICPENDR<n>.
- For a level-sensitive interrupt, a write to GICD_SETSPI_NSR or GICD_SETSPI_SR adds the pending state to the targeted interrupt. It will remain pending until it is deasserted by a write to GICD_CLRSPI_NSR or GICD_CLRSPI_SR. If the interrupt is activated between having the pending state added and being deactivated, then the interrupt will be active and pending.

Accessing the GICD_SETSPI_NSR:

GICD_SETSPI_NSR can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	0x0040

8.9.20 GICD_SETSPI_SR, Set Secure SPI Pending Register

The GICD_SETSPI_SR characteristics are:

Purpose

Adds the pending state to a valid SPI.

A write to this register changes the state of an inactive SPI to pending, and the state of an active SPI to active and pending.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WI	WO	WI

Writes to this register have no effect if:

- The value is written by a Non-secure access.
- The value written specifies an invalid SPI.
- The SPI is already pending.

16-bit accesses to bits [15:0] of this register must be supported.

Configurations

If GICD_TYPER.MBIS == 0, this register is reserved.

In implementations that support only a single Security state, this register is WI.

Attributes

GICD_SETSPI_SR is a 32-bit register.

Field descriptions

The GICD_SETSPI_SR bit assignments are:

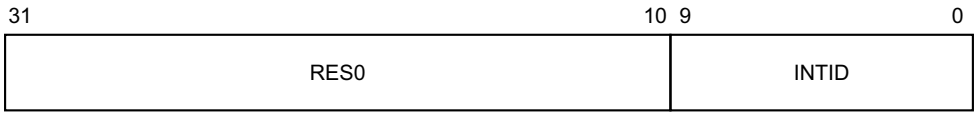

Bits [31:10]

Reserved, RES0.

INTID, bits [9:0]

The INTID of the SPI.

The function of this register depends on whether the targeted SPI is configured to be an edge-triggered or level-sensitive interrupt:

- For an edge-triggered interrupt, a write to GICD_SETSPI_NSR or GICD_SETSPI_SR adds the pending state to the targeted interrupt. It will stop being pending on activation, or if the pending state is removed by a write to GICD_CLRSPI_NSR, GICD_CLRSPI_SR, or GICD_ICPENDR<n>.
- For a level-sensitive interrupt, a write to GICD_SETSPI_NSR or GICD_SETSPI_SR adds the pending state to the targeted interrupt. It will remain pending until it is deasserted by a write to GICD_CLRSPI_NSR or GICD_CLRSPI_SR. If the interrupt is activated between having the pending state added and being deactivated, then the interrupt will be active and pending.

Accessing the GICD_SETSPI_SR:

GICD_SETSPI_SR can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	0x0050

8.9.21 GICD_SGIR, Software Generated Interrupt Register

The GICD_SGIR characteristics are:

Purpose

Controls the generation of SGIs.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WO	WO	WO

This register is used only when affinity routing is not enabled. When affinity routing is enabled, this register is RES0.

It is IMPLEMENTATION DEFINED whether this register has any effect when the forwarding of interrupts by the Distributor is disabled by GICD_CTLR.

Configurations

This register is available in all configurations of the GIC. If the GIC supports two Security states this register is Common.

Attributes

GICD_SGIR is a 32-bit register.

Field descriptions

The GICD_SGIR bit assignments are:

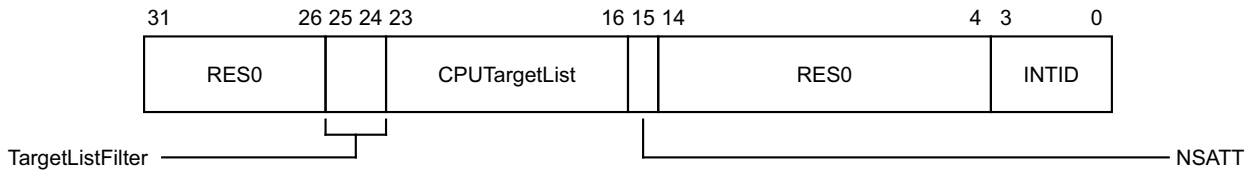

Bits [31:26]

Reserved, RES0.

TargetListFilter, bits [25:24]

Determines how the Distributor processes the requested SGI.

`00` Forward the interrupt to the CPU interfaces specified by GICD_SGIR.CPUTargetList.

`01` Forward the interrupt to all CPU interfaces except that of the PE that requested the interrupt.

`10` Forward the interrupt only to the CPU interface of the PE that requested the interrupt.

`11` Reserved.

When this register has an architecturally-defined reset value, this field resets to `0`.

CPUTargetList, bits [23:16]

When GICD_SGIR.TargetListFilter is `00`, this field defines the CPU interfaces to which the Distributor must forward the interrupt.

Each bit of the field refers to the corresponding CPU interface. For example, CPUTargetList[0] corresponds to interface 0. Setting a bit to 1 indicates that the interrupt must be forwarded to the corresponding interface.

If this field is `00000000` when GICD_SGIR.TargetListFilter is `00`, the Distributor does not forward the interrupt to any CPU interface.

When this register has an architecturally-defined reset value, this field resets to `0b11111111`.

NSATT, bit [15]

Specifies the required group of the SGI.

`0` Forward the SGI specified in the INTID field to a specified CPU interface only if the SGI is configured as Group 0 on that interface.

`1` Forward the SGI specified in the INTID field to a specified CPU interface only if the SGI is configured as Group 1 on that interface.

This field is writable only by a Secure access. Non-secure accesses can also generate Group 0 interrupts, if allowed to do so by GICD_NSACR0. Otherwise, Non-secure writes to GICD_SGIR generate an SGI only if the specified SGI is programmed as Group 1, regardless of the value of bit [15] of the write.

When this register has an architecturally-defined reset value, this field resets to `0`.

Bits [14:4]

Reserved, RES0.

INTID, bits [3:0]

The INTID of the SGI to forward to the specified CPU interfaces.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the GICD_SGIR:

GICD_SGIR can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	`0x0F00`

8.9.22 GICD_SPENDSGIR<n>, SGI Set-Pending Registers, n = 0 - 3

The GICD_SPENDSGIR<n> characteristics are:

Purpose

Adds the pending state to an SGI.

A write to this register changes the state of an inactive SGI to pending, and the state of an active SGI to active and pending.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

These registers are used only when affinity routing is not enabled. When affinity routing is enabled for the Security state of an interrupt then the bit associated with SGI in that security state is RES0. An implementation is permitted to make the register RAZ/WI in this case.

A register bit that corresponds to an unimplemented SGI is RAZ/WI.

These registers are byte-accessible.

If the GIC implementation supports two Security states:

- A register bit that corresponds to a Group 0 interrupt is RAZ/WI to Non-secure accesses.
- Register bits corresponding to unimplemented PEs are RAZ/WI.

Configurations

Four SGI set-pending registers are implemented. Each register contains eight set-pending bits for each of four SGIs, for a total of 16 possible SGIs.

In multiprocessor implementations, each PE has a copy of these registers.

Attributes

GICD_SPENDSGIR<n> is a 32-bit register.

Field descriptions

The GICD_SPENDSGIR<n> bit assignments are:

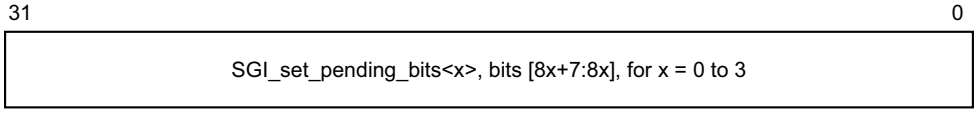

SGI_set_pending_bits<x>, bits [8x+7:8x], for x = 0 to 3

Adds the pending state to SGI number 4n + x for the PE corresponding to the bit number written to.

Reads and writes have the following behavior:

`0` If read, indicates that the SGI from the corresponding PE is not pending and is not active and pending.
If written, has no effect.

`1` If read, indicates that the SGI from the corresponding PE is pending or is active and pending.
If written, adds the pending state to the SGI for the corresponding PE.

When this register has an architecturally-defined reset value, this field resets to `0`.

For SGI ID m, generated by processing element C writing to the corresponding GICD_SGIR field, where DIV and MOD are the integer division and modulo operations:

- The corresponding GICD_SPENDSGIR<n> number is given by n = m DIV 4.
- The offset of the required register is (0xF20 + (4n)).
- The offset of the required field within the register GICD_CPENDSGIR<n> is given by m MOD 4.
- The required bit in the 8-bit SGI set-pending field m is bit C.

Accessing the GICD_SPENDSGIR<n>:

GICD_SPENDSGIR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	0x0F20 + 4n

8.9.23 GICD_STATUSR, Error Reporting Status Register

The GICD_STATUSR characteristics are:

Purpose

Provides software with a mechanism to detect:

- Accesses to reserved locations.
- Writes to read-only locations.
- Reads of write-only locations.

Usage constraints

GICD_STATUSR(S) is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	-

GICD_STATUSR(NS) is accessible as follows:

Security disabled	Secure	Non-secure
RW	-	RW

This is an optional register. If the register is not implemented, the location is RAZ/WI.

Configurations

If the GIC implementation supports two Security states this register is Banked to provide Secure and Non-secure copies.

Attributes

GICD_STATUSR is a 32-bit register.

Field descriptions

The GICD_STATUSR bit assignments are:

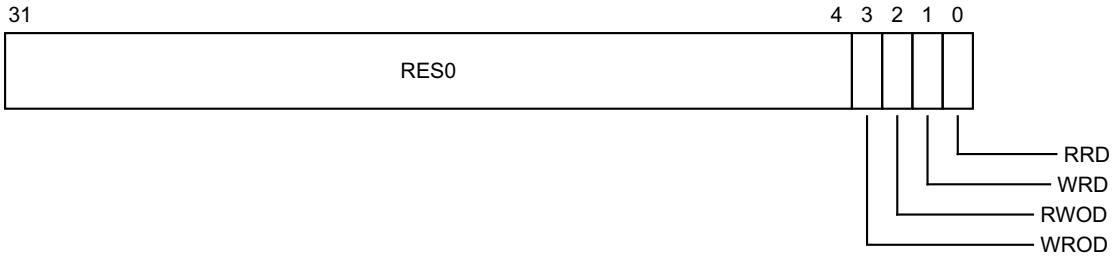

Bits [31:4]

Reserved, RES0.

WROD, bit [3]

Write to an RO location.

`0` Normal operation.

`1` A write to an RO location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

RWOD, bit [2]

Read of a WO location.

0 Normal operation.

1 A read of a WO location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

WRD, bit [1]

Write to a reserved location.

0 Normal operation.

1 A write to a reserved location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

RRD, bit [0]

Read of a reserved location.

0 Normal operation.

1 A read of a reserved location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

Accessing the GICD_STATUSR:

GICD_STATUSR can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	0x0010

8.9.24 GICD_TYPER, Interrupt Controller Type Register

The GICD_TYPER characteristics are:

Purpose

Provides information about what features the GIC implementation supports. It indicates:

- Whether the GIC implementation supports two Security states.
- The maximum number of INTIDs that the GIC implementation supports.
- The number of PEs that can be used as interrupt targets.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

Configurations

This register is available in all configurations of the GIC. If the GIC implementation supports two Security states, this register is Common.

Attributes

GICD_TYPER is a 32-bit register.

Field descriptions

The GICD_TYPER bit assignments are:

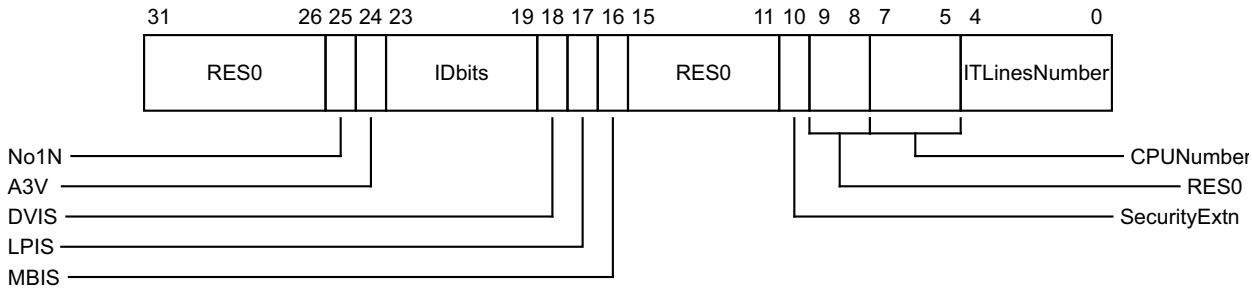

Bits [31:26]

Reserved, RES0.

No1N, bit [25]

Indicates whether 1 of N SPI interrupts are supported.

0 1 of N SPI interrupts are supported.

1 1 of N SPI interrupts are not supported.

A3V, bit [24]

Affinity 3 valid. Indicates whether the Distributor supports nonzero values of Affinity level 3. Possible values are:

0 The Distributor only supports zero values of Affinity level 3.

1 The Distributor supports nonzero values of Affinity level 3.

IDbits, bits [23:19]

The number of interrupt identifier bits supported, minus one.

DVIS, bit [18]

Indicates whether the implementation supports Direct Virtual LPI injection.

0 The implementation does not support Direct Virtual LPI injection.

1 The implementation supports Direct Virtual LPI injection.

For GICv3, this field is RES0.

LPIS, bit [17]

Indicates whether the implementation supports LPIs.

0 The implementation does not support LPIs.

1 The implementation supports LPIs.

MBIS, bit [16]

Indicates whether the implementation supports message-based interrupts by writing to Distributor registers.

0 The implementation does not support message-based interrupts by writing to Distributor registers.

The GICD_CLRSPI_NSR, GICD_SETSPI_NSR, GICD_CLRSPI_SR, and GICD_SETSPI_SR registers are reserved.

1 The implementation supports message-based interrupts by writing to the GICD_CLRSPI_NSR, GICD_SETSPI_NSR, GICD_CLRSPI_SR, or GICD_SETSPI_SR registers.

Bits [15:11]

Reserved, RES0.

SecurityExtn, bit [10]

Indicates whether the GIC implementation supports two Security states:

When GICD_CTLR.DS == 1, this field is RAZ.

0 The GIC implementation supports only a single Security state.

1 The GIC implementation supports two Security states.

Bits [9:8]

Reserved, RES0.

CPUNumber, bits [7:5]

In a GIC implementation where affinity routing is not enabled, this field indicates the number of PEs that can be used as interrupt targets, minus one.

These PEs must be numbered contiguously from zero, but the relationship between this number and the affinity hierarchy from MPIDR is IMPLEMENTATION DEFINED. If the implementation does not support ARE being zero, this field is 000.

ITLinesNumber, bits [4:0]

Indicates the maximum SPI INTID that the GIC implementation supports. If the value of this field is N, the maximum SPI INTID is 32(N+1)-1. For example, 00011 specifies that the maximum SPI INTID is 127.

The maximum SPI INTID an implementation might support is 1019 (field value 11111). Regardless of the range of INTIDs defined by this field, interrupt IDs 1020-1023 are reserved for special purposes.

Note

The value derived from this field specifies the maximum number of SPIs that the GIC implementation might support. An implementation might not implement all SPIs up to this maximum.

The ITLinesNumber field only indicates the maximum number of SPIs that the GIC implementation might support. This value determines the number of instances of the following interrupt registers:

- GICD_IGROUPR<n>.
- GICD_ISENABLER<n>.
- GICD_ICENABLER<n>.
- GICD_ISPENDR<n>.
- GICD_ICPENDR<n>.
- GICD_ISACTIVER<n>.
- GICD_ICACTIVER<n>.
- GICD_IPRIORITYR<n>.
- GICD_ITARGETSR<n>.
- GICD_ICFGR<n>.

The GIC architecture does not require a GIC implementation to support a continuous range of SPI interrupt IDs. Software must check which SPI INTIDs are supported, up to the maximum value indicated by GICD_TYPER.ITLinesNumber.

Accessing the GICD_TYPER:

GICD_TYPER can be accessed through the memory-mapped interface:

Component	Offset
GIC Distributor	0x0004

8.10 The GIC Redistributor register map

This section describes the Redistributor register map.

The mechanism by which an ITS communicates with the Redistributors is IMPLEMENTATION DEFINED. An implementation might perform this communication using memory-mapped functionality, and a portion of the Redistributor memory map is allocated for such communication. The definition of the communication is outside the scope of this GIC architecture specification.

Each Redistributor defines two 64KB frames in the physical address map:

- RD_base for controlling the overall behavior of the Redistributor, for controlling LPIs, and for generating LPIs in a system that does not include at least one ITS..
- SGI_base for controlling and generating PPIs and SGIs.

The frame for each Redistributor must be contiguous and must be ordered as follows:

1. `RD_base`
2. `SGI_base`

In GICv4, there are two additional 64KB frames:

- A frame to control virtual LPIs. The base address of this frame is referred to as `VLPI_base`.
- A frame for a reserved page.

The frames for each Redistributor must be contiguous and must be ordered as follows:

1. `RD_base`
2. `SGI_base`
3. `VLPI_base`
4. Reserved

Table 8-27 shows the GIC Redistributor register map for the physical LPI registers.

Table 8-27 GIC physical LPI Redistributor register map

Offset from `RD_base`	Name	Type	Reset	Description
`0x0000`	GICR_CTLR	RW	See the register description	Redistributor Control Register
`0x0004`	GICR_IIDR	RO	IMPLEMENTATION DEFINED	Implementer Identification Register
`0x0008`	GICR_TYPER	RO	IMPLEMENTATION DEFINED	Redistributor Type Register
`0x0010`	GICR_STATUSR	RW	`0x0000 0000`	Error Reporting Status Register, optional
`0x0014`	GICR_WAKER	RW	See the register description	Redistributor Wake Register
`0x0018`	-	-	-	Reserved
`0x0020`	-	-	IMPLEMENTATION DEFINED	IMPLEMENTATION DEFINED registers
`0x0040`	GICR_SETLPIR[a]	WO	-	Set LPI Pending Register
`0x0048`	GICR_CLRLPIR[a]	WO	-	Clear LPI Pending Register
`0x0050`	-	-	-	Reserved
`0x0070`	GICR_PROPBASER	RW	-	Redistributor Properties Base Address Register
`0x0078`	GICR_PENDBASER	RW	-	Redistributor LPI Pending Table Base Address Register
`0x0080`	-	-	-	Reserved
`0x00A0`	GICR_INVLPIR[a]	WO	-	Redistributor Invalidate LPI Register

Table 8-27 GIC physical LPI Redistributor register map (continued)

Offset from RD_base	Name	Type	Reset	Description
0x00A8	-	-	-	Reserved
0x00B0	GICR_INVALLR[a]	WO	-	Redistributor Invalidate All Register
0x00B8	-	-	-	Reserved
0x00C0	GICR_SYNCR[a]	RO	-	Redistributor Synchronize Register
0x00C8	-	-	-	Reserved
0x0100	-	WO	IMPLEMENTATION DEFINED	IMPLEMENTATION DEFINED registers
0x0108	-	-	-	Reserved
0x0110	-	WO	IMPLEMENTATION DEFINED	IMPLEMENTATION DEFINED registers
0x0118-0xBFFC	-	-	-	Reserved
0xC000-0xFFCC	-	-	IMPLEMENTATION DEFINED	IMPLEMENTATION DEFINED registers
0xFFD0-0xFFFC	-	RO	IMPLEMENTATION DEFINED	Reserved for ID registers, see *Identification registers* on page 8-173

a. This register is IMPLEMENTATION DEFINED in implementations that include an ITS.

Table 8-28 shows the GIC Redistributor register map for the virtual LPI registers.

Table 8-28 GIC virtual LPI Redistributor register map

Offset from VLPI_base	Name	Type	Reset	Description
0x0000	-	-	-	Reserved
0x0040	-	WO	IMPLEMENTATION DEFINED	IMPLEMENTATION DEFINED registers
0x0050	-	-	-	Reserved
0x0070	GICR_VPROPBASER	RW	-	Virtual Redistributor Properties Base Address Register
0x0078	GICR_VPENDBASER	RW	-	Virtual Pending Table Base Address Register
0x0080-0x037C	-	RW	IMPLEMENTATION DEFINED	IMPLEMENTATION DEFINED registers
0x0380-0xBFFC	-	-	-	Reserved
0xC000-0xFFCC	-	-	IMPLEMENTATION DEFINED	IMPLEMENTATION DEFINED registers
0xFFD0-0xFFFC	-	-	-	Reserved

Table 8-29 on page 8-488 shows the GIC Redistributor register map for the SGI and PPI registers.

Table 8-29 GIC SGI and PPI Redistributor register map

Offset from SGI_base	Name	Type	Reset	Description
0x0080	GICR_IGROUPR0	RW	-	Interrupt Group Register 0
0x0100	GICR_ISENABLER0	RW	IMPLEMENTATION DEFINED	Interrupt Set-Enable Register 0
0x0180	GICR_ICENABLER0	RW	IMPLEMENTATION DEFINED	Interrupt Clear-Enable Register 0
0x0200	GICR_ISPENDR0	RW	0xXXXX 0000	Interrupt Set-Pend Register 0
0x0280	GICR_ICPENDR0	RW	0xXXXX 0000	Interrupt Clear-Pend Register 0
0x0300	GICR_ISACTIVER0	RW	0x0000 0000	Interrupt Set-Active Register 0
0x0380	GICR_ICACTIVER0	RW	0x0000 0000	Interrupt Clear-Active Register 0
0x0400-0x041C	GICR_IPRIORITYR<n>	RW	0x0000 0000	Interrupt Priority Registers
0x0C00	GICR_ICFGR0	RW	IMPLEMENTATION DEFINED	SGI Configuration Register
0x0C04	GICR_ICFGR1	RW	IMPLEMENTATION DEFINED	PPI Configuration Register
0x0D00	GICR_IGRPMODR0	RW	-	Interrupt Group Modifier Register 0
0x0E00	GICR_NSACR	RW	0x0000 0000	Non-Secure Access Control Register
0x0E04-0xBFFC	-	-	-	Reserved
0xC000-0xFFCC	-	-	IMPLEMENTATION DEFINED	IMPLEMENTATION DEFINED registers
0xFFD0-0xFFFC	-	-	-	Reserved

8.11 The GIC Redistributor register descriptions

This section describes each of the GIC Redistributor registers in register name order.

8.11.1 GICR_CLRLPIR, Clear LPI Pending Register

The GICR_CLRLPIR characteristics are:

Purpose

Clears the pending state of the specified LPI.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WO	WO	WO

When written with a 32-bit write the data is zero-extended to 64 bits.

This register is mandatory in an implementation that supports LPIs and does not include an ITS. The functionality of this register is IMPLEMENTATION DEFINED in an implementation that does include an ITS.

Writes to this register have no effect if any of the following apply:

- GICR_CTLR.EnableLPIs == 0.
- The pINTID value specifies an unimplemented LPI.
- The pINTID value specifies an LPI that is not pending.

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_CLRLPIR is a 64-bit register.

Field descriptions

The GICR_CLRLPIR bit assignments are:

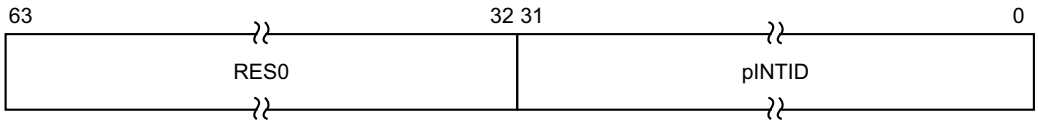

Bits [63:32]

Reserved, RES0.

pINTID, bits [31:0]

The INTID of the physical LPI.

Note

The size of this field is IMPLEMENTATION DEFINED, and is specified by the GICD_TYPER.Idbits field. Unimplemented bits are RES0.

Accessing the GICR_CLRLPIR:

GICR_CLRLPIR can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	RD_base	0x0048-0x004C

8.11.2 GICR_CTLR, Redistributor Control Register

The GICR_CTLR characteristics are:

Purpose

Controls the operation of a Redistributor, and enables the signaling of LPIs by the Redistributor to the connected PE.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_CTLR is a 32-bit register.

Field descriptions

The GICR_CTLR bit assignments are:

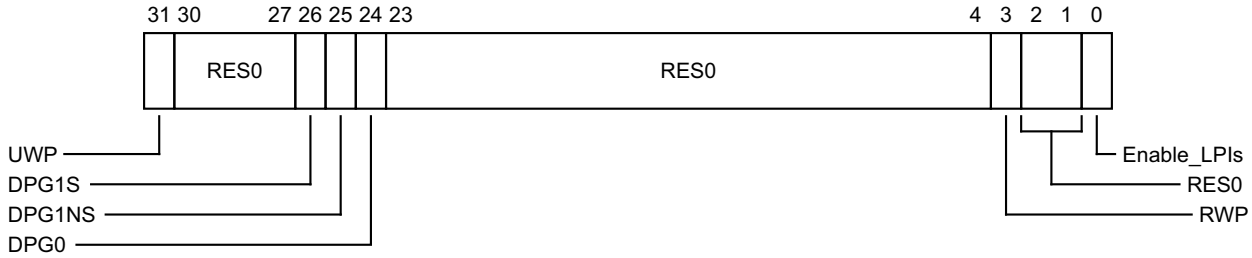

UWP, bit [31]

Upstream Write Pending. Read-only. Indicates whether all upstream writes have been communicated to the Distributor.

`0` The effects of all upstream writes have been communicated to the Distributor, including any Generate SGI packets.

`1` Not all the effects of upstream writes, including any Generate SGI packets, have been communicated to the Distributor.

Bits [30:27]

Reserved, RES0.

DPG1S, bit [26]

Disable Processor selection for Group 1 Secure interrupts. When GICR_TYPER.DPGS == 1:

`0` A Group 1 Secure SPI configured to use the 1 of N distribution model can select this PE, if the PE is not asleep and if Secure Group 1 interrupts are enabled.

`1` A Group 1 Secure SPI configured to use the 1 of N distribution model cannot select this PE.

When GICR_TYPER.DPGS == 0 this bit is RAZ/WI.

When GICD_CTLR.DS==1, this field is RAZ/WI. In GIC implementations that support two Security states, this field is only accessible by Secure accesses, and is RAZ/WI to Non-secure accesses.

It is IMPLEMENTATION DEFINED whether this bit affects the selection of PEs for interrupts using the 1 of N distribution model when GICD_CTLR.ARE_S==0.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

DPG1NS, bit [25]

Disable Processor selection for Group 1 Non-secure interrupts. When GICR_TYPER.DPGS == 1:

`0` A Group 1 Non-secure SPI configured to use the 1 of N distribution model can select this PE, if the PE is not asleep and if Non-secure Group 1 interrupts are enabled.

`1` A Group 1 Non-secure SPI configured to use the 1 of N distribution model cannot select this PE.

When GICR_TYPER.DPGS == 0 this bit is RAZ/WI.

It is IMPLEMENTATION DEFINED whether this bit affects the selection of PEs for interrupts using the 1 of N distribution model when GICD_CTLR.ARE_NS==0.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

DPG0, bit [24]

Disable Processor selection for Group 0 interrupts. When GICR_TYPER.DPGS == 1:

`0` A Group 0 SPI configured to use the 1 of N distribution model can select this PE, if the PE is not asleep and if Group 0 interrupts are enabled.

`1` A Group 0 SPI configured to use the 1 of N distribution model cannot select this PE.

When GICR_TYPER.DPGS == 0 this bit is RAZ/WI.

When GICD_CTLR.DS==1, this field is always accessible. In GIC implementations that support two Security states, this field is RAZ/WI to Non-secure accesses.

It is IMPLEMENTATION DEFINED whether this bit affects the selection of PEs for interrupts using the 1 of N distribution model when GICD_CTLR.ARE_S==0.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

Bits [23:4]

Reserved, RES0.

RWP, bit [3]

Register Write Pending. This bit indicates whether a register write for the current Security state is in progress or not.

`0` The effect of all previous writes to the following registers are visible to all agents in the system:

- GICR_ICENABLER0
- GICR_CTLR.DPG1S
- GICR_CTLR.DPG1NS
- GICR_CTLR.DPG0

`1` The effect of all previous writes to the following registers are not guaranteed by the architecture to be visible yet to the all agents in the system as the changes are still being propagated:

- GICR_ICENABLER0
- GICR_CTLR.DPG1S
- GICR_CTLR.DPG1NS

- GICR_CTLR.DPG0

Bits [2:1]

Reserved, RES0.

Enable_LPIs, bit [0]

In implementations where affinity routing is enabled for the Security state:

0 LPI support is disabled. Any doorbell interrupt generated as a result of a write to a virtual LPI register must be discarded, and any ITS translation requests or commands involving LPIs in this Redistributor are ignored.

1 LPI support is enabled.

Note

If GICR_TYPER.LPIS == 0, this field is RES0.

If GICD_CTLR.ARE_NS is written from 1 to 0 when this bit is 1, behavior is an IMPLEMENTATION DEFINED choice between clearing GICR_CTLR.Enable_LPIs to 0 or maintaining its current value.

When affinity routing is not enabled for the Non-secure state, this bit is RES0. When a write changes this bit from 0 to 1, this bit becomes RES1 and the Redistributor must load the LPI Pending table from memory to check for any pending interrupts.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

The participation of a PE in the 1 of N interrupt distribution model for a given interrupt group is governed by the concatenation of GICR_WAKER.ProcessorSleep, the appropriate GICR_CTLR.DPG{1, 0} bit, and the PE interrupt group enable. The behavior options are:

PS	DPG{1S, 1NS, 0}	Enable	PE behavior
0	0	0	The PE cannot be selected.
0	0	1	The PE can be selected.
0	1	*	The PE cannot be selected.
1	*	*	The PE cannot be selected when GICD_CTLR.E1NWF == 0. When GICD_CTLR.E1NWF == 1, the mechanism by which PEs are selected is IMPLEMENTATION DEFINED.

If an SPI using the 1 of N distribution model has been forwarded to the PE and a write to GICR_CTLR occurs that changes the DPG bit for the interrupt group of the SPI, the IRI must attempt to select a different target PE for the SPI. This might have no effect on the forwarded SPI if it has already been activated.

Accessing the GICR_CTLR:

GICR_CTLR can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	RD_base	0x0000

8.11.3 GICR_ICACTIVER0, Interrupt Clear-Active Register 0

The GICR_ICACTIVER0 characteristics are:

Purpose

Deactivates the corresponding SGI or PPI. These registers are used when saving and restoring GIC state.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

When affinity routing is not enabled for the Security state of an interrupt in GICR_ICACTIVER0, the corresponding bit is RAZ/WI and equivalent functionality is provided by GICD_ICACTIVER<n> with n=0.

This register only applies to SGIs (bits [15:0]) and PPIs (bits [31:16]). For SPIs, this functionality is provided by GICD_ICACTIVER<n>.

When GICD_CTLR.DS == 0, bits corresponding to Secure SGIs and PPIs are RAZ/WI to Non-secure accesses.

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_ICACTIVER0 is a 32-bit register.

Field descriptions

The GICR_ICACTIVER0 bit assignments are:

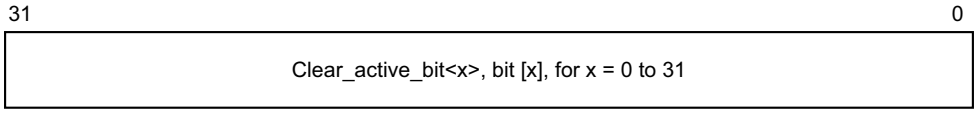

Clear_active_bit<x>, bit [x], for x = 0 to 31

Removes the active state from interrupt number x. Reads and writes have the following behavior:

`0` If read, indicates that the corresponding interrupt is not active, and is not active and pending.

If written, has no effect.

`1` If read, indicates that the corresponding interrupt is active, or is active and pending.

If written, deactivates the corresponding interrupt, if the interrupt is active. If the interrupt is already deactivated, the write has no effect.

After a write of 1 to this bit, a subsequent read of this bit returns 0.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICR_ICACTIVER0:

GICR_ICACTIVER0 can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	SGI_base	0x0380

8.11.4 GICR_ICENABLER0, Interrupt Clear-Enable Register 0

The GICR_ICENABLER0 characteristics are:

Purpose

Disables forwarding of the corresponding SGI or PPI to the CPU interfaces.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

When affinity routing is not enabled for the Security state of an interrupt in GICR_ICENABLER0, the corresponding bit is RAZ/WI and equivalent functionality is provided by GICD_ICENABLER<n> with n=0.

This register only applies to SGIs (bits [15:0]) and PPIs (bits [31:16]). For SPIs, this functionality is provided by GICD_ICENABLER<n>.

When GICD_CTLR.DS == 0, bits corresponding to Secure SGIs and PPIs are RAZ/WI to Non-secure accesses.

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_ICENABLER0 is a 32-bit register.

Field descriptions

The GICR_ICENABLER0 bit assignments are:

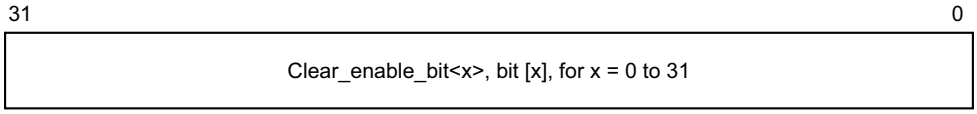

Clear_enable_bit<x>, bit [x], for x = 0 to 31

For PPIs and SGIs, controls the forwarding of interrupt number x to the CPU interfaces. Reads and writes have the following behavior:

`0` If read, indicates that forwarding of the corresponding interrupt is disabled.
If written, has no effect.

`1` If read, indicates that forwarding of the corresponding interrupt is enabled.
If written, disables forwarding of the corresponding interrupt.
After a write of 1 to this bit, a subsequent read of this bit returns 0.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICR_ICENABLER0:

GICR_ICENABLER0 can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	SGI_base	0x0180

8.11.5 GICR_ICFGR0, Interrupt Configuration Register 0

The GICR_ICFGR0 characteristics are:

Purpose

Determines whether the corresponding SGI is edge-triggered or level-sensitive.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

This register is used when affinity routing is enabled.

When affinity routing is disabled for the Security state of an interrupt, the field for that interrupt is RES0 and an implementation is permitted to make the field RAZ/WI in this case. Equivalent functionality is provided by GICD_ICFGR<n> with n=0 .

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_ICFGR0 is a 32-bit register.

Field descriptions

The GICR_ICFGR0 bit assignments are:

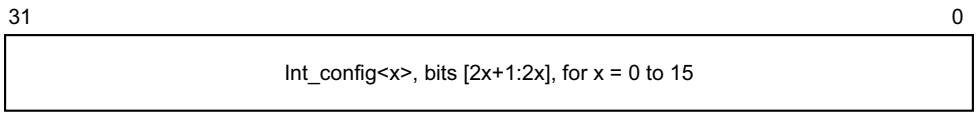

Int_config<x>, bits [2x+1:2x], for x = 0 to 15

Indicates whether the interrupt with ID 16n + x is level-sensitive or edge-triggered.

Int_config[0] (bit [2x]) is RES0.

Possible values of Int_config[1] (bit [2x+1]) are:

`0` Corresponding interrupt is level-sensitive.

`1` Corresponding interrupt is edge-triggered.

For SGIs, Int_config[1] is RAO/WI.

A read of this bit always returns the correct value to indicate the interrupt triggering method.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICR_ICFGR0:

GICR_ICFGR0 can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	SGI_base	`0x0C00`

8.11.6 GICR_ICFGR1, Interrupt Configuration Register 1

The GICR_ICFGR1 characteristics are:

Purpose

Determines whether the corresponding PPI is edge-triggered or level-sensitive.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

This register is used when affinity routing is enabled.

When affinity routing is disabled for the Security state of an interrupt, the field for that interrupt is RES0 and an implementation is permitted to make the field RAZ/WI in this case. Equivalent functionality is provided by GICD_ICFGR<n> with n=1 .

For each supported PPI, it is IMPLEMENTATION DEFINED whether software can program the corresponding Int_config field.

Software must disable an interrupt before the value of the corresponding programmable Int_config field is changed. GIC behavior is otherwise UNPREDICTABLE.

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_ICFGR1 is a 32-bit register.

Field descriptions

The GICR_ICFGR1 bit assignments are:

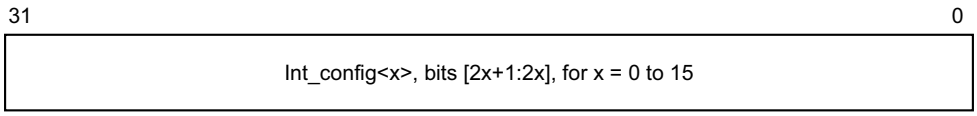

Int_config<x>, bits [2x+1:2x], for x = 0 to 15

Indicates whether the interrupt with ID 16n + x is level-sensitive or edge-triggered.

Int_config[0] (bit [2x]) is RES0.

Possible values of Int_config[1] (bit [2x+1]) are:

`0` Corresponding interrupt is level-sensitive.

`1` Corresponding interrupt is edge-triggered.

A read of this bit always returns the correct value to indicate the interrupt triggering method.

For PPIs, Int_config[1] is programmable unless the implementation supports two Security states and the bit corresponds to a Group 0 or Secure Group 1 interrupt, in which case the bit is RAZ/WI to Non-secure accesses.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICR_ICFGR1:

GICR_ICFGR1 can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	SGI_base	0x0C04

8.11.7 GICR_ICPENDR0, Interrupt Clear-Pending Register 0

The GICR_ICPENDR0 characteristics are:

Purpose

Removes the pending state from the corresponding SGI or PPI.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

When affinity routing is not enabled for the Security state of an interrupt in GICR_ICPENDR0, the corresponding bit is RAZ/WI and equivalent functionality is provided by GICD_ICPENDR<n> with n=0.

This register only applies to SGIs (bits [15:0]) and PPIs (bits [31:16]). For SPIs, this functionality is provided by GICD_ICENABLER<n>.

When GICD_CTLR.DS == 0, bits corresponding to Secure SGIs and PPIs are RAZ/WI to Non-secure accesses.

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_ICPENDR0 is a 32-bit register.

Field descriptions

The GICR_ICPENDR0 bit assignments are:

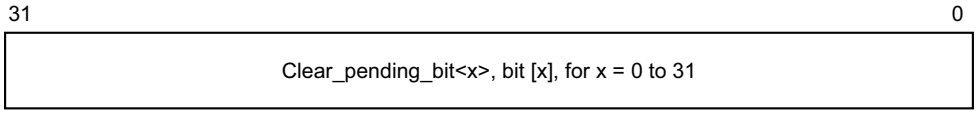

Clear_pending_bit<x>, bit [x], for x = 0 to 31

Removes the pending state from interrupt number x. Reads and writes have the following behavior:

`0` If read, indicates that the corresponding interrupt is not pending.
If written, has no effect.

`1` If read, indicates that the corresponding interrupt is pending, or active and pending.
If written, changes the state of the corresponding interrupt from pending to inactive, or from active and pending to active. This has no effect in the following cases:

- If the interrupt is not pending and is not active and pending.
- If the interrupt is a level-sensitive interrupt that is pending or active and pending for a reason other than a write to GICD_ISPENDR<n>. In this case, if the interrupt signal continues to be asserted, the interrupt remains pending or active and pending.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICR_ICPENDR0:

GICR_ICPENDR0 can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	SGI_base	0x0280

8.11.8 GICR_IGROUPR0, Interrupt Group Register 0

The GICR_IGROUPR0 characteristics are:

Purpose

Controls whether the corresponding SGI or PPI is in Group 0 or Group 1.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

When affinity routing is not enabled for the Security state of an interrupt in GICR_IGROUPR0, an alias is provided for the first 8 PEs by GICD_IGROUPR<n> with n=0.

When GICD_CTLR.DS == 0, the register is RAZ/WI to Non-secure accesses.

Bits corresponding to unimplemented interrupts are RAZ/WI.

Note

Implementations must ensure that an interrupt that is pending at the time of the write uses either the old value or the new value and must ensure that the interrupt is neither lost nor handled more than once. The effect of the change must be visible in finite time.

Configurations

This register is available in all GIC configurations. If the GIC implementation supports two Security states, this register is Secure.

A copy of this register is provided for each Redistributor.

Attributes

GICR_IGROUPR0 is a 32-bit register.

Field descriptions

The GICR_IGROUPR0 bit assignments are:

Redistributor_group_status_bit<x>, bit [x], for x = 0 to 31

Group status bit. In this register:

- Bits [31:16] are group status bits for PPIs.
- Bits [15:0] are group status bits for SGIs.

`0` When GICD_CTLR.DS==0, the corresponding interrupt is Group 0.

`1` When GICD_CTLR.DS==0, the corresponding interrupt is Group 1.

When GICD_CTLR.DS == 1, the bit that corresponds to the interrupt is concatenated with the equivalent bit in GICR_IGRPMODR0 to form a 2-bit field that defines an interrupt group. The encoding of this field is at GICR_IGRPMODR0.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

The considerations for the reset value of this register are the same as those for GICD_IGROUPR<n> with n=0.

Accessing the GICR_IGROUPR0:

GICR_IGROUPR0 can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	SGI_base	0x0080

8.11.9 GICR_IGRPMODR0, Interrupt Group Modifier Register 0

The GICR_IGRPMODR0 characteristics are:

Purpose

In GIC implementations which support two Security states, along with the GICR_IGROUPR0 register, controls whether the corresponding SGI or PPI is in:

- Secure Group 0.
- Non-secure Group 1.
- When System register access is enabled, Secure Group 1.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RES0	RW	RW

When affinity routing is not enabled for the Security state of an interrupt in GICR_IGRPMODR0, the corresponding bit is RES0 and equivalent functionality is provided by GICD_IGRPMODR<n> with n=0.

This register only applies to SGIs (bits [15:0]) and PPIs (bits [31:16]). For SPIs, this functionality is provided by GICD_IGRPMODR<n>.

When GICD_CTLR.ARE_S == 0 or GICD_CTLR.DS == 1, GICR_IGRPMODR0 is RES0. An implementation can make this register RAZ/WI in this case.

When GICD_CTLR.DS==0, the register is RAZ/WI to Non-secure accesses.

Note

Implementations must ensure that an interrupt that is pending at the time of the write uses either the old value or the new value and must ensure that the interrupt is neither lost nor handled more than once. The effect of the change must be visible in finite time.

Configurations

If the GIC implementation supports two Security states, these registers are Secure.

A copy of this register is provided for each Redistributor.

Attributes

GICR_IGRPMODR0 is a 32-bit register.

Field descriptions

The GICR_IGRPMODR0 bit assignments are:

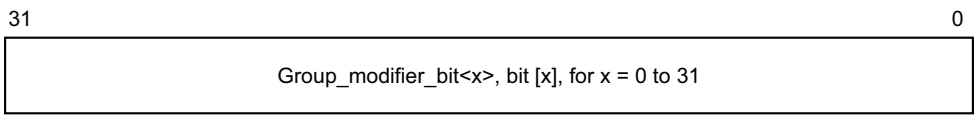

Group_modifier_bit<x>, bit [x], for x = 0 to 31

Group modifier bit. In implementations where affinity routing is enabled for the Security state of an interrupt, the bit that corresponds to the interrupt is concatenated with the equivalent bit in GICR_IGROUPR0 to form a 2-bit field that defines an interrupt group:

Group modifier bit	Group status bit	Definition	Short name
0	0	Secure Group 0	G0S
0	1	Non-secure Group 1	G1NS
1	0	Secure Group 1	G1S
1	1	Reserved, treated as Non-secure Group 1	-

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICR_IGRPMODR0:

GICR_IGRPMODR0 can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	SGI_base	0x0D00

8.11.10 GICR_IIDR, Redistributor Implementer Identification Register

The GICR_IIDR characteristics are:

Purpose

Provides information about the implementer and revision of the Redistributor.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

Configurations

This register is available in all configurations of the GIC. If the GIC implementation supports two Security states, this register is Common.

Attributes

GICR_IIDR is a 32-bit register.

Field descriptions

The GICR_IIDR bit assignments are:

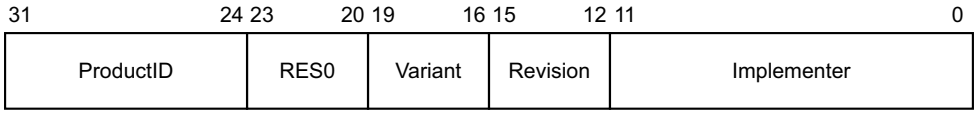

ProductID, bits [31:24]

An IMPLEMENTATION DEFINED product identifier.

Bits [23:20]

Reserved, RES0.

Variant, bits [19:16]

An IMPLEMENTATION DEFINED variant number. Typically, this field is used to distinguish product variants, or major revisions of a product.

Revision, bits [15:12]

An IMPLEMENTATION DEFINED revision number. Typically, this field is used to distinguish minor revisions of a product.

Implementer, bits [11:0]

Contains the JEP106 code of the company that implemented the Redistributor:

- Bits [11:8] are the JEP106 continuation code of the implementer. For an ARM implementation, this field is `0x4`.
- Bit [7] is always 0.
- Bits [6:0] are the JEP106 identity code of the implementer. For an ARM implementation, bits [7:0] are therefore `0x3B`.

Accessing the GICR_IIDR:

GICR_IIDR can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	RD_base	0x0004

8.11.11 GICR_INVALLR, Redistributor Invalidate All Register

The GICR_INVALLR characteristics are:

Purpose

Invalidates any cached configuration data of all physical LPIs, causing the GIC to reload the interrupt configuration from the physical configuration table at the address specified by GICR_PROPBASER.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WO	WO	WO

This register is mandatory in an implementation that supports LPIs and does not include an ITS. The functionality is IMPLEMENTATION DEFINED in an implementation that does include an ITS.

Writes to this register have no effect if no physical LPIs are currently stored in the local Redistributor cache.

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_INVALLR is a 64-bit register.

Field descriptions

The GICR_INVALLR bit assignments are:

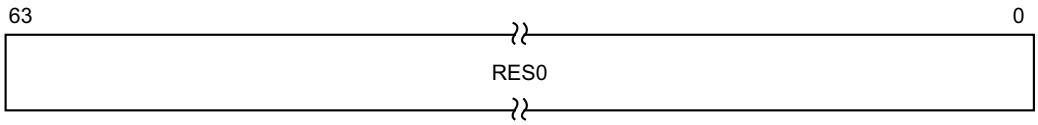

Bits [63:0]

Reserved, RES0.

Note

If any LPI has been forwarded to the PE and a valid write to GICR_INVALLR is received, the Redistributor must ensure it reloads its properties from memory. This has no effect on the forwarded LPI if it has already been activated.

Accessing the GICR_INVALLR:

GICR_INVALLR can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	RD_base	0x00B0-0x00B4

8.11.12 GICR_INVLPIR, Redistributor Invalidate LPI Register

The GICR_INVLPIR characteristics are:

Purpose

Invalidates the cached configuration data of a specified LPI, causing the GIC to reload the interrupt configuration from the physical configuration table at the address specified by GICR_PROPBASER.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WO	WO	WO

When written with a 32-bit write the data is zero-extended to 64 bits.

This register is mandatory in an implementation that supports LPIs and does not include an ITS. The functionality is IMPLEMENTATION DEFINED in an implementation that does include an ITS.

Writes to this register have no effect if either:

- The specified LPI is not currently stored in the local Redistributor.
- The pINTID field corresponds to an unimplemented LPI.

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_INVLPIR is a 64-bit register.

Field descriptions

The GICR_INVLPIR bit assignments are:

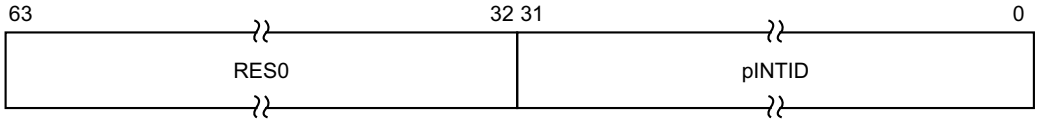

Bits [63:32]

Reserved, RES0.

pINTID, bits [31:0]

The INTID of the physical LPI to be cleaned.

Note

The size of this field is IMPLEMENTATION DEFINED, and is specified by the GICD_TYPER.IDbits field. Unimplemented bits are RES0.

Note

If any LPI has been forwarded to the PE and a valid write to GICR_INVLPIR is received, the Redistributor must ensure it reloads its properties from memory and apply any changes by retrieving and reforwarding the LPI as required. This has no effect on the forwarded LPI if it has already been activated.

Accessing the GICR_INVLPIR:

GICR_INVLPIR can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	RD_base	0x00A0-0x00A4

8.11.13 GICR_IPRIORITYR<n>, Interrupt Priority Registers, n = 0 - 7

The GICR_IPRIORITYR<n> characteristics are:

Purpose

Holds the priority of the corresponding interrupt for each SGI and PPI supported by the GIC.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

These registers are used when affinity routing is enabled for the Security state of the interrupt. When affinity routing is not enabled the bits corresponding to the interrupt are RAZ/WI and GICD_IPRIORITYR<n> provides equivalent functionality.

These registers are used for SGIs and PPIs only. Equivalent functionality for SPIs is provided by GICD_IPRIORITYR<n>.

These registers are byte-accessible.

When GICD_CTLR.DS == 0, bits corresponding to Secure SGIs and PPIs are RAZ/WI to Non-secure accesses.

Note

Implementations must ensure that an interrupt that is pending at the time of the write uses either the old value or the new value and must ensure that the interrupt is neither lost nor handled more than once. The effect of the change must be visible in finite time.

Configurations

A copy of these registers is provided for each Redistributor.

These registers are configured as follows:

- GICR_IPRIORITYR0-GICR_IPRIORITYR3 store the priority of SGIs.
- GICR_IPRIORITYR4-GICR_IPRIORITYR7 store the priority of PPIs.

Attributes

GICR_IPRIORITYR<n> is a 32-bit register.

Field descriptions

The GICR_IPRIORITYR<n> bit assignments are:

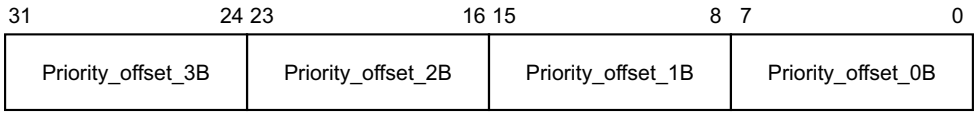

Priority_offset_3B, bits [31:24]

Interrupt priority value from an IMPLEMENTATION DEFINED range, at byte offset 3. Lower priority values correspond to greater priority of the interrupt.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Priority_offset_2B, bits [23:16]

Interrupt priority value from an IMPLEMENTATION DEFINED range, at byte offset 2. Lower priority values correspond to greater priority of the interrupt.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Priority_offset_1B, bits [15:8]

Interrupt priority value from an IMPLEMENTATION DEFINED range, at byte offset 1. Lower priority values correspond to greater priority of the interrupt.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Priority_offset_0B, bits [7:0]

Interrupt priority value from an IMPLEMENTATION DEFINED range, at byte offset 0. Lower priority values correspond to greater priority of the interrupt.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICR_IPRIORITYR<n>:

GICR_IPRIORITYR<n> can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	SGI_base	`0x0400` + 4n

8.11.14 GICR_ISACTIVER0, Interrupt Set-Active Register 0

The GICR_ISACTIVER0 characteristics are:

Purpose

Activates the corresponding SGI or PPI. These registers are used when saving and restoring GIC state.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

When affinity routing is not enabled for the Security state of an interrupt in GICR_ISACTIVER0, the corresponding bit is RAZ/WI and equivalent functionality is provided by GICD_ISACTIVER<n> with n=0.

This register only applies to SGIs (bits [15:0]) and PPIs (bits [31:16]). For SPIs, this functionality is provided by GICD_ISACTIVER<n>.

When GICD_CTLR.DS == 0, bits corresponding to Secure SGIs and PPIs are RAZ/WI to Non-secure accesses.

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_ISACTIVER0 is a 32-bit register.

Field descriptions

The GICR_ISACTIVER0 bit assignments are:

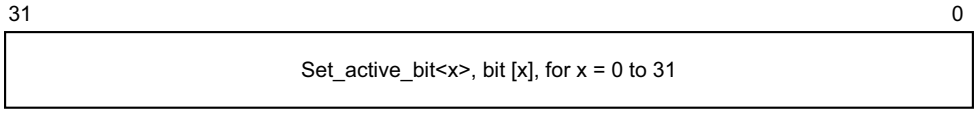

Set_active_bit<x>, bit [x], for x = 0 to 31

Adds the active state to interrupt number x. Reads and writes have the following behavior:

`0` If read, indicates that the corresponding interrupt is not active, and is not active and pending.

If written, has no effect.

`1` If read, indicates that the corresponding interrupt is active, or is active and pending.

If written, activates the corresponding interrupt, if the interrupt is not already active. If the interrupt is already active, the write has no effect.

After a write of 1 to this bit, a subsequent read of this bit returns 1.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICR_ISACTIVER0:

GICR_ISACTIVER0 can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	SGI_base	0x0300

8.11.15 GICR_ISENABLER0, Interrupt Set-Enable Register 0

The GICR_ISENABLER0 characteristics are:

Purpose

Enables forwarding of the corresponding SGI or PPI to the CPU interfaces.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

When affinity routing is not enabled for the Security state of an interrupt in GICR_ISENABLER0, the corresponding bit is RAZ/WI and equivalent functionality is provided by GICD_ISENABLER<n> with n=0.

This register only applies to SGIs (bits [15:0]) and PPIs (bits [31:16]). For SPIs, this functionality is provided by GICD_ISENABLER<n>.

When GICD_CTLR.DS == 0, bits corresponding to Secure SGIs and PPIs are RAZ/WI to Non-secure accesses.

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_ISENABLER0 is a 32-bit register.

Field descriptions

The GICR_ISENABLER0 bit assignments are:

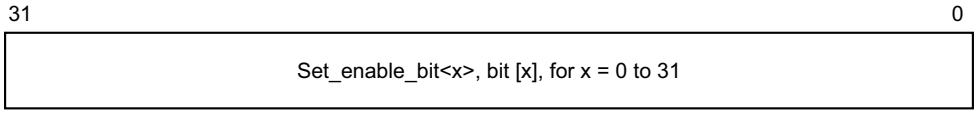

Set_enable_bit<x>, bit [x], for x = 0 to 31

For PPIs and SGIs, controls the forwarding of interrupt number x to the CPU interface. Reads and writes have the following behavior:

`0` If read, indicates that forwarding of the corresponding interrupt is disabled.
If written, has no effect.

`1` If read, indicates that forwarding of the corresponding interrupt is enabled.
If written, enables forwarding of the corresponding interrupt.
After a write of 1 to this bit, a subsequent read of this bit returns 1.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the GICR_ISENABLER0:

GICR_ISENABLER0 can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	SGI_base	0x0100

8.11.16 GICR_ISPENDR0, Interrupt Set-Pending Register 0

The GICR_ISPENDR0 characteristics are:

Purpose

Adds the pending state to the corresponding SGI or PPI.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

When affinity routing is not enabled for the Security state of an interrupt in GICR_ISPENDR0, the corresponding bit is RAZ/WI and equivalent functionality is provided by GICD_ISPENDR<n> with n=0.

This register only applies to SGIs (bits [15:0]) and PPIs (bits [31:16]). For SPIs, this functionality is provided by GICD_ISPENDR<n>.

When GICD_CTLR.DS == 0, bits corresponding to Secure SGIs and PPIs are RAZ/WI to Non-secure accesses.

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_ISPENDR0 is a 32-bit register.

Field descriptions

The GICR_ISPENDR0 bit assignments are:

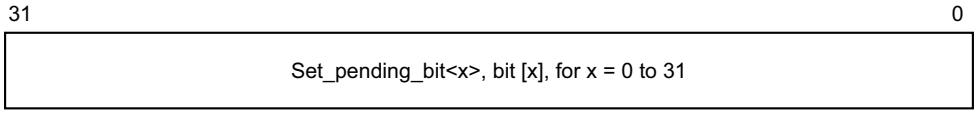

Set_pending_bit<x>, bit [x], for x = 0 to 31

For PPIs and SGIs, adds the pending state to interrupt number x. Reads and writes have the following behavior:

`0` If read, indicates that the corresponding interrupt is not pending on this PE.
If written, has no effect.

`1` If read, indicates that the corresponding interrupt is pending, or active and pending on this PE.
If written, changes the state of the corresponding interrupt from inactive to pending, or from active to active and pending. This has no effect in the following cases:

- If the interrupt is already pending because of a write to GICR_ISPENDR0.
- If the interrupt is already pending because the corresponding interrupt signal is asserted. In this case, the interrupt remains pending if the interrupt signal is deasserted.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICR_ISPENDR0:

GICR_ISPENDR0 can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	SGI_base	0x0200

8.11.17 GICR_NSACR, Non-secure Access Control Register

The GICR_NSACR characteristics are:

Purpose

Enables Secure software to permit Non-secure software on a particular PE to create and control SGIs by writing to ICC_SGI1R_EL1, ICC_ASGI1R_EL1 or ICC_SGI0R_EL1, see Table 8-14 on page 8-171.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

When GICD_CTLR.DS == 1, this register is RAZ/WI.

When GICD_CTLR.DS == 0, this register is Secure, and is RAZ/WI to Non-secure accesses.

This register is used when affinity routing is enabled. When affinity routing is not enabled for the Security state of the interrupt, GICD_NSACR<n> with n=0 provides equivalent functionality.

Configurations

There are no configuration notes.

Attributes

GICR_NSACR is a 32-bit register.

Field descriptions

The GICR_NSACR bit assignments are:

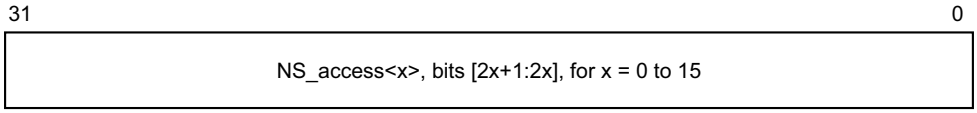

NS_access<x>, bits [2x+1:2x], for x = 0 to 15

Configures the level of Non-secure access permitted when the SGI is in Secure Group 0 or Secure Group 1, as defined from GICR_IGROUPR0 and GICR_IGRPMODR0. A field is provided for each SGI. The possible values of each 2-bit field are:

`00` Non-secure access is not permitted to fields associated with the corresponding SGI.

`01` Non-secure writes are permitted to generate a Secure Group 0 SGI.

`10` As `0b01`, but additionally Non-secure writes to are permitted to generate a Secure Group 1 SGI.

11 Reserved.

Note

It is a programming error if software uses this value. However, ARM strongly recommends that implementations treat this value as `10`.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICR_NSACR:

GICR_NSACR can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	SGI_base	0x0E00

8.11.18 GICR_PENDBASER, Redistributor LPI Pending Table Base Address Register

The GICR_PENDBASER characteristics are:

Purpose

Specifies the base address of the LPI pending table, and the Shareability and Cacheability of accesses to the LPI pending table.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

Having the GICR_PENDBASER OuterCache, Shareability or InnerCache fields programmed to different values on different Redistributors with GICR_CTLR.EnableLPIs == 1 in the system is UNPREDICTABLE.

Changing GICR_PENDBASER with GICR_CTLR.EnableLPIs == 1 is UNPREDICTABLE.

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_PENDBASER is a 64-bit register.

Field descriptions

The GICR_PENDBASER bit assignments are:

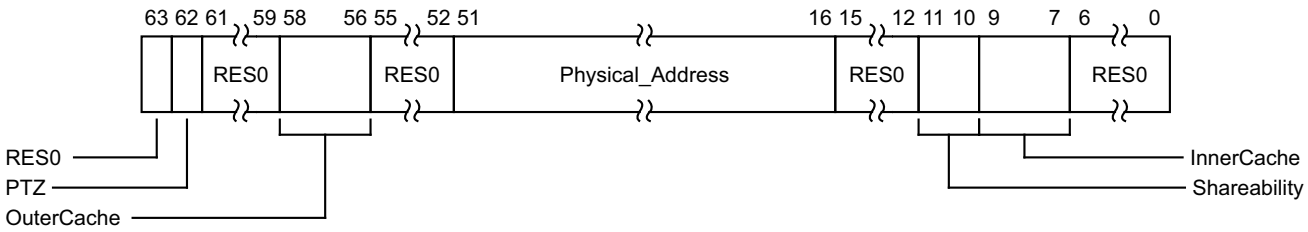

Bit [63]

Reserved, RES0.

PTZ, bit [62]

Pending Table Zero. Indicates to the Redistributor whether the LPI Pending table is zero when GICR_CTLR.EnableLPIs == 1.

This field is WO, and reads as 0.

`0` The LPI Pending table is not zero, and contains live data.

`1` The LPI Pending table is zero. Software must ensure the LPI Pending table is zero before this value is written.

Bits [61:59]

Reserved, RES0.

OuterCache, bits [58:56]

Indicates the Outer Cacheability attributes of accesses to the LPI Pending table. The possible values of this field are:

000 Memory type defined in InnerCache field. For Normal memory, Outer Cacheability is the same as Inner Cacheability.

001 Normal Outer Non-cacheable.

010 Normal Outer Cacheable Read-allocate, Write-through.

011 Normal Outer Cacheable Read-allocate, Write-back.

100 Normal Outer Cacheable Write-allocate, Write-through.

101 Normal Outer Cacheable Write-allocate, Write-back.

110 Normal Outer Cacheable Read-allocate, Write-allocate, Write-through.

111 Normal Outer Cacheable Read-allocate, Write-allocate, Write-back.

It is IMPLEMENTATION DEFINED whether this field has a fixed value or can be programmed by software. Implementing this field with a fixed value is deprecated.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

Bits [55:52]

Reserved, RES0.

Physical_Address, bits [51:16]

Bits [51:16] of the physical address containing the LPI Pending table.

In implementations supporting fewer than 52 bits of physical address, unimplemented upper bits are RAZ/WI.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Bits [15:12]

Reserved, RES0.

Shareability, bits [11:10]

Indicates the Shareability attributes of accesses to the LPI Pending table. The possible values of this field are:

00 Non-shareable.

01 Inner Shareable.

10 Outer Shareable.

11 Reserved. Treated as 00.

It is IMPLEMENTATION DEFINED whether this field has a fixed value or can be programmed by software. Implementing this field with a fixed value is deprecated.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

InnerCache, bits [9:7]

Indicates the Inner Cacheability attributes of accesses to the LPI Pending table. The possible values of this field are:

000 Device-nGnRnE.

001 Normal Inner Non-cacheable.

010 Normal Inner Cacheable Read-allocate, Write-through.

011 Normal Inner Cacheable Read-allocate, Write-back.

100 Normal Inner Cacheable Write-allocate, Write-through.

101 Normal Inner Cacheable Write-allocate, Write-back.

110 Normal Inner Cacheable Read-allocate, Write-allocate, Write-through.

111 Normal Inner Cacheable Read-allocate, Write-allocate, Write-back.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Bits [6:0]

Reserved, RES0.

Accessing the GICR_PENDBASER:

GICR_PENDBASER can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	RD_base	0x0078-0x007C

8.11.19 GICR_PROPBASER, Redistributor Properties Base Address Register

The GICR_PROPBASER characteristics are:

Purpose

Specifies the base address of the LPI Configuration table, and the Shareability and Cacheability of accesses to the LPI Configuration table.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

It is IMPLEMENTATION DEFINED whether GICR_PROPBASER can be set to different values on different Redistributors. GICR_TYPER.CommonLPIAff identifies the Redistributors that must have GICR_PROPBASER set to the same values whenever GICR_CTLR.EnableLPIs == 1.

Setting different values in different copies of GICR_PROPBASER on Redistributors that are required to use a common LPI Configuration table when GICR_CTLR.EnableLPIs == 1 leads to UNPREDICTABLE behavior.

Changing GICR_PROPBASER when GICR_CTLR.EnableLPIs == 1 is UNPREDICTABLE.

Other restrictions apply when a Redistributor caches information from GICR_PROPBASER. See *LPI Configuration tables* on page 6-95 for more information.

Configurations

A copy of this register is provided for each Redistributor.

An implementation might make this register RO, for example to correspond to a configuration table in read-only memory.

Attributes

GICR_PROPBASER is a 64-bit register.

Field descriptions

The GICR_PROPBASER bit assignments are:

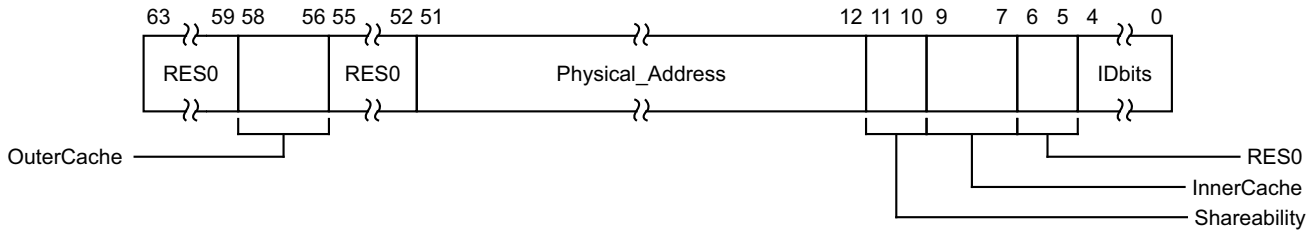

Bits [63:59]

Reserved, RES0.

OuterCache, bits [58:56]

Indicates the Outer Cacheability attributes of accesses to the LPI Configuration table. The possible values of this field are:

`000` Memory type defined in InnerCache field. For Normal memory, Outer Cacheability is the same as Inner Cacheability.

001 Normal Outer Non-cacheable.

010 Normal Outer Cacheable Read-allocate, Write-through.

011 Normal Outer Cacheable Read-allocate, Write-back.

100 Normal Outer Cacheable Write-allocate, Write-through.

101 Normal Outer Cacheable Write-allocate, Write-back.

110 Normal Outer Cacheable Read-allocate, Write-allocate, Write-through.

111 Normal Outer Cacheable Read-allocate, Write-allocate, Write-back.

It is IMPLEMENTATION DEFINED whether this field has a fixed value or can be programmed by software. Implementing this field with a fixed value is deprecated.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

Bits [55:52]

Reserved, RES0.

Physical_Address, bits [51:12]

Bits [51:12] of the physical address containing the LPI Configuration table.

In implementations supporting fewer than 52 bits of physical address, unimplemented upper bits are RAZ/WI.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Shareability, bits [11:10]

Indicates the Shareability attributes of accesses to the LPI Configuration table. The possible values of this field are:

00 Non-shareable.

01 Inner Shareable.

10 Outer Shareable.

11 Reserved. Treated as 00.

It is IMPLEMENTATION DEFINED whether this field has a fixed value or can be programmed by software. Implementing this field with a fixed value is deprecated.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

InnerCache, bits [9:7]

Indicates the Inner Cacheability attributes of accesses to the LPI Configuration table. The possible values of this field are:

000 Device-nGnRnE.

001 Normal Inner Non-cacheable.

010 Normal Inner Cacheable Read-allocate, Write-through.

011 Normal Inner Cacheable Read-allocate, Write-back.

100 Normal Inner Cacheable Write-allocate, Write-through.

101 Normal Inner Cacheable Write-allocate, Write-back.

110 Normal Inner Cacheable Read-allocate, Write-allocate, Write-through.

111 Normal Inner Cacheable Read-allocate, Write-allocate, Write-back.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Bits [6:5]

Reserved, RES0.

IDbits, bits [4:0]

The number of bits of LPI INTID supported, minus one, by the LPI Configuration table starting at Physical_Address.

If the value of this field is larger than the value of GICD_TYPER.IDbits, the GICD_TYPER.IDbits value applies.

If the value of this field is less than `0b1101`, indicating that the largest INTID is less than 8192 (the smallest LPI INTID), the GIC will behave as if all physical LPIs are out of range.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICR_PROPBASER:

GICR_PROPBASER can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	RD_base	`0x0070-0x0074`

8.11.20 GICR_SETLPIR, Set LPI Pending Register

The GICR_SETLPIR characteristics are:

Purpose

Generates an LPI by setting the pending state of the specified LPI.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WO	WO	WO

When written with a 32-bit write the data is zero-extended to 64 bits.

This register is mandatory in an implementation that supports LPIs and does not include an ITS. The functionality is IMPLEMENTATION DEFINED in an implementation that does include an ITS.

Writes to this register have no effect if either:

- The pINTID field corresponds to an LPI that is already pending.
- The pINTID field corresponds to an unimplemented LPI.
- GICR_CTLR.EnableLPIs == 0.

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_SETLPIR is a 64-bit register.

Field descriptions

The GICR_SETLPIR bit assignments are:

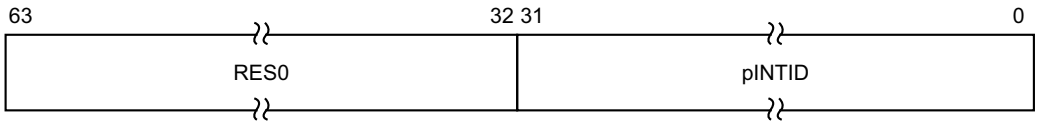

Bits [63:32]

Reserved, RES0.

pINTID, bits [31:0]

The INTID of the physical LPI to be generated.

Note

The size of this field is IMPLEMENTATION DEFINED, and is specified by the GICD_TYPER.IDbits field. Unimplemented bits are RES0.

Accessing the GICR_SETLPIR:

GICR_SETLPIR can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	RD_base	0x0040-0x0044

8.11.21 GICR_STATUSR, Error Reporting Status Register

The GICR_STATUSR characteristics are:

Purpose

Provides software with a mechanism to detect:

- Accesses to reserved locations.
- Writes to read-only locations.
- Reads of write-only locations.

Usage constraints

GICR_STATUSR(S) is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	-

GICR_STATUSR(NS) is accessible as follows:

Security disabled	Secure	Non-secure
RW	-	RW

This is an optional register. If the register is not implemented, the location is RAZ/WI.

Configurations

A copy of this register is provided for each Redistributor.

If the GIC implementation supports two Security states this register is Banked to provide Secure and Non-secure copies.

Attributes

GICR_STATUSR is a 32-bit register.

Field descriptions

The GICR_STATUSR bit assignments are:

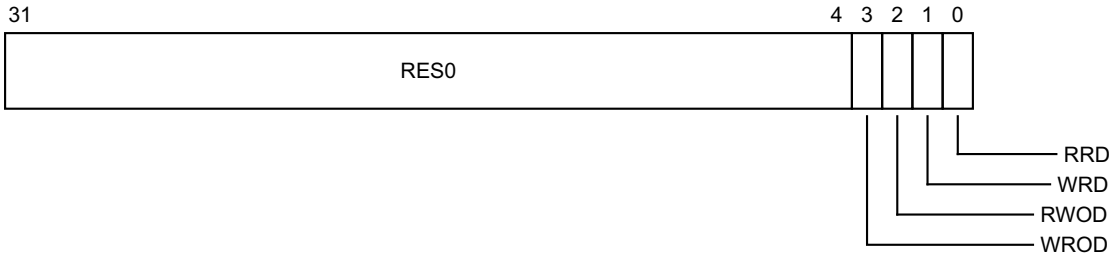

Bits [31:4]

Reserved, RES0.

WROD, bit [3]

Write to an RO location.

0 Normal operation.

1 A write to an RO location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

RWOD, bit [2]

Read of a WO location.

0 Normal operation.

1 A read of a WO location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

WRD, bit [1]

Write to a reserved location.

0 Normal operation.

1 A write to a reserved location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

RRD, bit [0]

Read of a reserved location.

0 Normal operation.

1 A read of a reserved location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

Accessing the GICR_STATUSR:

GICR_STATUSR can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	RD_base	0x0010

8.11.22 GICR_SYNCR, Redistributor Synchronize Register

The GICR_SYNCR characteristics are:

Purpose

Indicates completion of physical Redistributor operations.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

Optionally, when this register is accessed, an implementation might wait until all operations are complete before returning a value, in which case GICR_SYNCR.Busy is always 0.

This register is mandatory in an implementation that supports LPIs and does not include an ITS. The functionality is IMPLEMENTATION DEFINED in an implementation that does include an ITS.

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_SYNCR is a 32-bit register.

Field descriptions

The GICR_SYNCR bit assignments are:

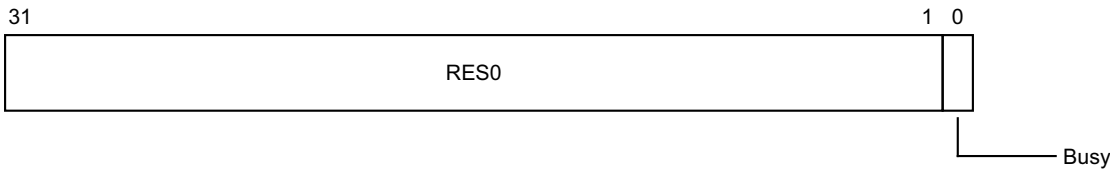

Bits [31:1]

Reserved, RES0.

Busy, bit [0]

Indicates completion of any Redistributor operations as follows:

`0` No operations are in progress.

`1` A write is in progress to one or more of the following registers:

- GICR_CLRLPIR.
- GICR_INVLPIR.
- GICR_INVALLR.

This field also indicates completion of any operations initiated by writes to GICR_PENDBASER or GICR_PROPBASER.

Accessing the GICR_SYNCR:

GICR_SYNCR can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	RD_base	0x00C0-0x00C4

8.11.23 GICR_TYPER, Redistributor Type Register

The GICR_TYPER characteristics are:

Purpose

Provides information about the configuration of this Redistributor.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_TYPER is a 64-bit register.

Field descriptions

The GICR_TYPER bit assignments are:

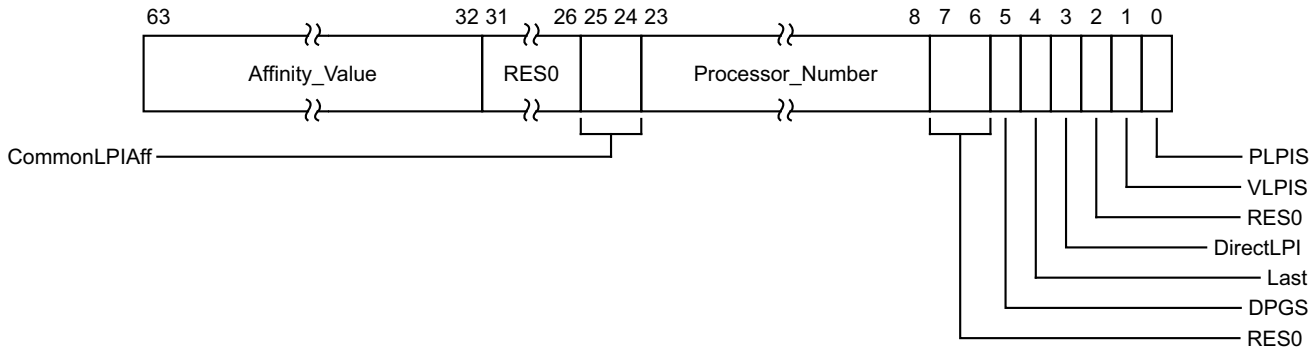

Affinity_Value, bits [63:32]

The identity of the PE associated with this Redistributor.

Bits [63:56] provide Aff3, the Affinity level 3 value for the Redistributor.

Bits [55:48] provide Aff2, the Affinity level 2 value for the Redistributor.

Bits [47:40] provide Aff1, the Affinity level 1 value for the Redistributor.

Bits [39:32] provide Aff0, the Affinity level 0 value for the Redistributor.

Bits [31:26]

Reserved, RES0.

CommonLPIAff, bits [25:24]

The affinity level at which Redistributors share a LPI Configuration table.

`00` All Redistributors must share a LPI Configuration table.

`01` All Redistributors with the same Aff3 value must share an LPI Configurationt table.

`10` All Redistributors with the same Aff3.Aff2 value must share an LPI Configuration table.

11 All Redistributors with the same Aff3.Aff2.Aff1 value must share an LPI Configuration table.

Processor_Number, bits [23:8]

A unique identifier for the PE. When GITS_TYPER.PTA == 0, an ITS uses this field to identify the interrupt target.

When affinity routing is disabled for a Security state, this field indicates which GICD_ITARGETSR<n> corresponds to this Redistributor.

Bits [7:6]

Reserved, RES0.

DPGS, bit [5]

Sets support for GICR_CTLR.DPG* bits.

0 GICR_CTLR.DPG* bits are not supported.

1 GICR_CTLR.DPG* bits are supported.

Last, bit [4]

Indicates whether this Redistributor is the highest-numbered Redistributor in a series of contiguous Redistributor pages.

0 This Redistributor is not the highest-numbered Redistributor in a series of contiguous Redistributor pages.

1 This Redistributor is the highest-numbered Redistributor in a series of contiguous Redistributor pages.

DirectLPI, bit [3]

Indicates whether this Redistributor supports direct injection of LPIs.

0 This Redistributor does not support direct injection of LPIs. The GICR_SETLPIR, GICR_CLRLPIR, GICR_INVLPIR, GICR_INVALLR, and GICR_SYNCR registers are either not implemented, or have an IMPLEMENTATION DEFINED purpose.

1 This Redistributor supports direct injection of LPIs. The GICR_SETLPIR, GICR_CLRLPIR, GICR_INVLPIR, GICR_INVALLR, and GICR_SYNCR registers are implemented.

Bit [2]

Reserved, RES0.

VLPIS, bit [1]

Indicates whether the GIC implementation supports virtual LPIs and the direct injection of virtual LPIs.

0 The implementation does not support virtual LPIs or the direct injection of virtual LPIs.

1 The implementation supports virtual LPIs and the direct injection of virtual LPIs.

Note

In GICv3 implementations this field is RES0.

PLPIS, bit [0]

Indicates whether the GIC implementation supports physical LPIs.

0 The implementation does not support physical LPIs.

1 The implementation supports physical LPIs.

Accessing the GICR_TYPER:

GICR_TYPER can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	RD_base	0x0008-0x000C

8.11.24 GICR_VPENDBASER, Virtual Redistributor LPI Pending Table Base Address Register

The GICR_VPENDBASER characteristics are:

Purpose

Specifies the base address of the memory that holds the virtual LPI Pending table for the currently scheduled virtual machine.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

The effect of a write to this register is not guaranteed to be visible throughout the affinity hierarchy, as indicated by GICR_CTLR.RWP == 0.

Configurations

This register is provided only in GICv4 implementations.

Attributes

GICR_VPENDBASER is a 64-bit register.

Field descriptions

The GICR_VPENDBASER bit assignments are:

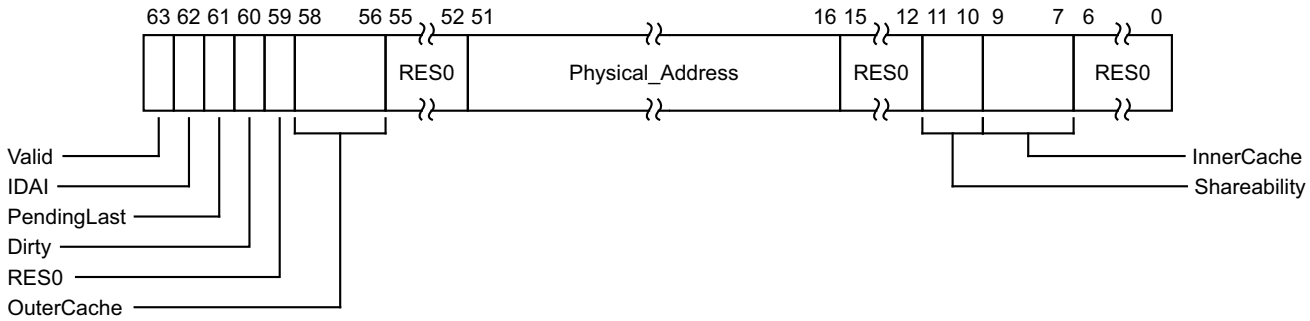

Valid, bit [63]

This bit controls whether the virtual LPI Pending table is valid:

`0` The virtual LPI Pending table is not valid. No vPE is scheduled.

`1` The virtual LPI Pending table is valid. All other fields in this register are RO when this field is 1.

When this register has an architecturally-defined reset value, this field resets to `0`.

IDAI, bit [62]

Implementation Defined Area Invalid. Indicates whether the IMPLEMENTATION DEFINED area in the LPI Pending table is valid:

`0` The IMPLEMENTATION DEFINED area is valid.

`1` The IMPLEMENTATION DEFINED area is invalid and all pending interrupt information is held in the architecturally defined part of the LPI Pending table.

For more information, see *LPI Pending tables* on page 6-97 and *Virtual LPI Configuration tables and virtual LPI Pending tables* on page 6-97.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

PendingLast, bit [61]

Indicates whether there are pending and enabled interrupts for the last scheduled vPE.

This value is set by the implementation when GICR_VPENDBASER.Valid has been written from 1 to 0 and is otherwise UNKNOWN.

`0` There are no pending and enabled interrupts for the last scheduled vPE.

`1` There is at least one pending interrupt for the last scheduled vPE. It is IMPLEMENTATION DEFINED whether this bit is set when the only pending interrupts for the last scheduled vPE are not enabled.

ARM deprecates setting PendingLast to 1 when the only pending interrupts for the last scheduled virtual machine are not enabled.

When the GICR_VPENDBASER.Valid bit is transitioned from 0 to 1, then the state of this bit indicates to the hardware whether the LPI Pending table contains no pending interrupts:

- `0b0`: The LPI Pending table is known to be zero, and so the pending table does not need to be read by hardware to determine which pending interrupts are present.
- `0b1`: The LPI Pending table is not known to be zero, and so the hardware must read the LPI Pending table to determine which pending interrupts are present.

When this register has an architecturally-defined reset value, this field resets to `0`.

Dirty, bit [60]

Read-only. Indicates whether there are any virtual LPIs for the last scheduled vPE that have not completed. This field is used only when GICR_VPENDBASER.Valid has been cleared to 0, and is otherwise UNKNOWN:

`0` There are no uncompleted virtual LPIs for the last scheduled vPE.

`1` There is at least one uncompleted virtual LPI for the last scheduled vPE.

Note

When GICR_VPENDBASER.Valid == 0, the Redistributor must ensure any outstanding pending virtual interrupts are cleared from the CPU interface.

When this register has an architecturally-defined reset value, this field resets to `0`.

Bit [59]

Reserved, RES0.

OuterCache, bits [58:56]

Indicates the Outer Cacheability attributes of accesses to virtual LPI Pending tables of vPEs targeting this Redistributor. The possible values of this field are:

`000` Memory type defined in InnerCache field. For Normal memory, Outer Cacheability is the same as Inner Cacheability.

`001` Normal Outer Non-cacheable.

`010` Normal Outer Cacheable Read-allocate, Write-through.

`011` Normal Outer Cacheable Read-allocate, Write-back.

`100` Normal Outer Cacheable Write-allocate, Write-through.

`101` Normal Outer Cacheable Write-allocate, Write-back.

`110` Normal Outer Cacheable Read-allocate, Write-allocate, Write-through.

`111` Normal Outer Cacheable Read-allocate, Write-allocate, Write-back.

It is IMPLEMENTATION DEFINED whether this field has a fixed value or can be programmed by software. Implementing this field with a fixed value is deprecated.

The Cacheability, Outer Cacheability and Shareability fields are used for accesses to the virtual LPI Pending table of resident and non-resident VPEs.

If the OuterCacheabilty attribute of the virtual LPI Pending tables that are associated with vPEs targeting the same Redistributor are different, behavior is UNPREDICTABLE.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

Bits [55:52]

Reserved, RES0.

Physical_Address, bits [51:16]

Bits [51:16] of the physical address containing the virtual LPI Pending table.

In implementations supporting fewer than 52 bits of physical address, unimplemented upper bits are RAZ/WI.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Bits [15:12]

Reserved, RES0.

Shareability, bits [11:10]

Indicates the Shareability attributes of accesses to the virtual LPI Pending table. The possible values of this field are:

`00` Non-shareable.

`01` Inner Shareable.

`10` Outer Shareable.

`11` Reserved. Treated as `00`.

It is IMPLEMENTATION DEFINED whether this field has a fixed value or can be programmed by software. Implementing this field with a fixed value is deprecated.

The Cacheability, Outer Cacheability and Shareability fields are used for accesses to the virtual LPI Pending table of resident and non-resident vPEs.

If the Shareabilty attribute of the virtual LPI Pending tables that are associated with vPEs targeting the same Redistributor are different, behavior is UNPREDICTABLE.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

InnerCache, bits [9:7]

Indicates the Inner Cacheability attributes of accesses to the virtual LPI Pending table. The possible values of this field are:

`000` Device-nGnRnE.

`001` Normal Inner Non-cacheable.

`010` Normal Inner Cacheable Read-allocate, Write-through.

`011` Normal Inner Cacheable Read-allocate, Write-back.

`100` Normal Inner Cacheable Write-allocate, Write-through.

`101` Normal Inner Cacheable Write-allocate, Write-back.

`110` Normal Inner Cacheable Read-allocate, Write-allocate, Write-through.

`111` Normal Inner Cacheable Read-allocate, Write-allocate, Write-back.

The Cacheability, Outer Cacheability and Shareability fields are used for accesses to the virtual LPI Pending table of resident and non-resident vPEs.

If the Inner Cacheabilty attribute of the virtual LPI Pending tables that are associated with vPEs targeting the same Redistributor are different, behavior is UNPREDICTABLE.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Bits [6:0]

Reserved, RES0.

Accessing the GICR_VPENDBASER:

GICR_VPENDBASER can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	VLPI_base	0x0078-0x007C

8.11.25 GICR_VPROPBASER, Virtual Redistributor Properties Base Address Register

The GICR_VPROPBASER characteristics are:

Purpose

Specifies the base address of the memory that holds the virtual LPI Configuration table for the currently scheduled virtual machine.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

Configurations

This register is provided in GICv4 implementations only.

Attributes

GICR_VPROPBASER is a 64-bit register.

Field descriptions

The GICR_VPROPBASER bit assignments are:

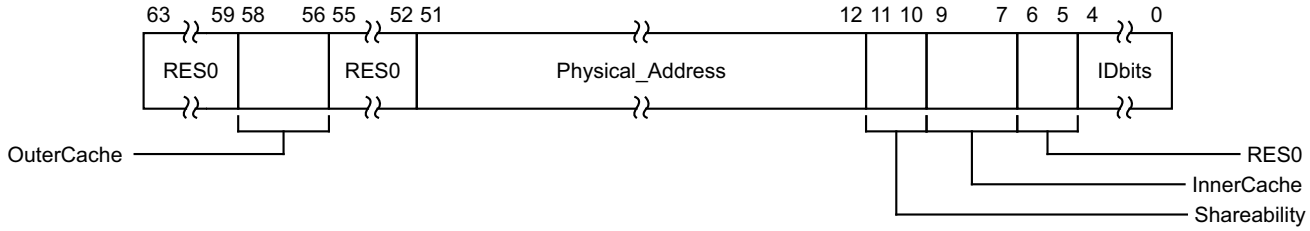

Bits [63:59]

Reserved, RES0.

OuterCache, bits [58:56]

Indicates the Outer Cacheability attributes of accesses to the LPI Configuration table. The possible values of this field are:

`000` Memory type defined in InnerCache field. For Normal memory, Outer Cacheability is the same as Inner Cacheability.

`001` Normal Outer Non-cacheable.

`010` Normal Outer Cacheable Read-allocate, Write-through.

`011` Normal Outer Cacheable Read-allocate, Write-back.

`100` Normal Outer Cacheable Write-allocate, Write-through.

`101` Normal Outer Cacheable Write-allocate, Write-back.

`110` Normal Outer Cacheable Read-allocate, Write-allocate, Write-through.

`111` Normal Outer Cacheable Read-allocate, Write-allocate, Write-back.

It is IMPLEMENTATION DEFINED whether this field has a fixed value or can be programmed by software. Implementing this field with a fixed value is deprecated.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

Bits [55:52]

Reserved, RES0.

Physical_Address, bits [51:12]

Bits [51:12] of the physical address containing the virtual LPI Configuration table.

In implementations supporting fewer than 52 bits of physical address, unimplemented upper bits are RAZ/WI.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Shareability, bits [11:10]

Indicates the Shareability attributes of accesses to the LPI Configuration table. The possible values of this field are:

`00` Non-shareable.

`01` Inner Shareable.

`10` Outer Shareable.

`11` Reserved. Treated as `00`.

It is IMPLEMENTATION DEFINED whether this field has a fixed value or can be programmed by software. Implementing this field with a fixed value is deprecated.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

InnerCache, bits [9:7]

Indicates the Inner Cacheability attributes of accesses to the LPI Configuration table. The possible values of this field are:

`000` Device-nGnRnE.

`001` Normal Inner Non-cacheable.

`010` Normal Inner Cacheable Read-allocate, Write-through.

`011` Normal Inner Cacheable Read-allocate, Write-back.

`100` Normal Inner Cacheable Write-allocate, Write-through.

`101` Normal Inner Cacheable Write-allocate, Write-back.

`110` Normal Inner Cacheable Read-allocate, Write-allocate, Write-through.

`111` Normal Inner Cacheable Read-allocate, Write-allocate, Write-back.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Bits [6:5]

Reserved, RES0.

IDbits, bits [4:0]

The number of bits of virtual LPI INTID supported, minus one.

If the value of this field is less than `0b1101`, indicating that the largest INTID is less than 8192 (the smallest LPI INTID), the GIC will behave as if all virtual LPIs are out of range.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICR_VPROPBASER:

GICR_VPROPBASER can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	VLPI_base	0x0070-0x0074

8.11.26 GICR_WAKER, Redistributor Wake Register

The GICR_WAKER characteristics are:

Purpose

Permits software to control the behavior of the **WakeRequest** power management signal corresponding to the Redistributor. Power management operations follow the rules in *Power management* on page 7-152.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RAZ/WI

In a GIC implementation that supports only a single Security state, this register is always accessible.

In a GIC implementation that supports two Security states, this is a Secure register. This register is RAZ/WI to Non-secure accesses.

To ensure a Redistributor is quiescent, software must write to GICR_WAKER with ProcessorSleep == 1, then poll the register until ChildrenAsleep == 1.

Configurations

A copy of this register is provided for each Redistributor.

Attributes

GICR_WAKER is a 32-bit register.

Field descriptions

The GICR_WAKER bit assignments are:

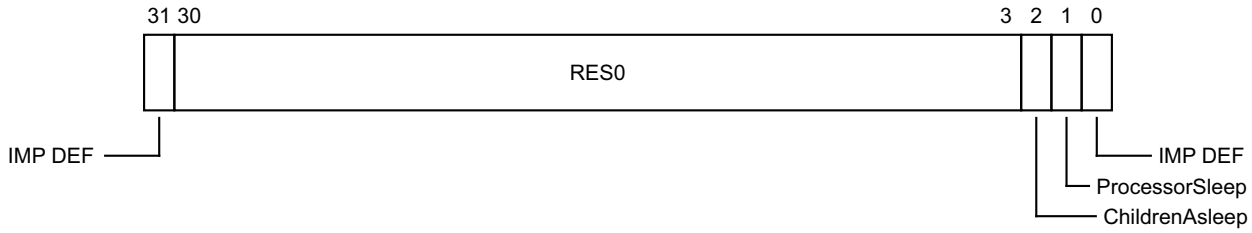

IMPLEMENTATION DEFINED, bit [31]

IMPLEMENTATION DEFINED.

Bits [30:3]

Reserved, RES0.

ChildrenAsleep, bit [2]

Read-only. Indicates whether the connected PE is quiescent:

`0` An interface to the connected PE might be active.

`1` All interfaces to the connected PE are quiescent.

When this register has an architecturally-defined reset value, this field resets to `1`.

ProcessorSleep, bit [1]

Indicates whether the Redistributor can assert the **WakeRequest** signal:

`0` This PE is not in, and is not entering, a low power state.

`1` The PE is either in, or is in the process of entering, a low power state.

All interrupts that arrive at the Redistributor:

- Assert a **WakeRequest** signal.
- Are held in the pending state at the Redistributor, and are not communicated to the CPU interface.

Note

When ProcessorSleep == 1, the Redistributor must ensure that any interrupts that are pending on the CPU interface are released.

For an implementation that is using the GIC Stream Protocol Interface:

- A Quiesce (IRI) command can put the interface between the Redistributor and the CPU interface in a quiescent state.
- A Release (ICC) command can release any interrupts that are pending on the CPU interface.

Note

Before powering down a PE, software must set this bit to 1 and wait until ChildrenAsleep == 1. After powering up a PE, or following a failed powerdown, software must set this bit to 0 and wait until ChildrenAsleep == 0.

Changing ProcessorSleep from 1 to 0 when ChildrenAsleep is not 1 results in UNPREDICTABLE behavior.

Changing ProcessorSleep from 0 to 1 when the Enable for each interrupt group in the associated CPU interface is not 0 results in UNPREDICTABLE behavior.

When this register has an architecturally-defined reset value, this field resets to `1`.

IMPLEMENTATION DEFINED, bit [0]

IMPLEMENTATION DEFINED.

Accessing the GICR_WAKER:

GICR_WAKER can be accessed through the memory-mapped interface:

Component	Frame	Offset
GIC Redistributor	RD_base	`0x0014`

8.12 The GIC CPU interface register map

Table 8-30 shows the CPU interface register maps. In this table, address offsets are relative to the *CPU interface base address* defined by the system memory map. Unless otherwise stated in the register description, all GIC registers are 32-bits wide. Reserved register addresses are RAZ/WI.

For a multiprocessor implementation, the GIC implements a set of CPU interface registers for each CPU interface. ARM strongly recommends that each PE has the same CPU interface base address for the CPU interface that connects it to the GIC. This is the private CPU interface base address for that PE. It is IMPLEMENTATION DEFINED whether a PE can access the CPU interface registers of other PEs in the system.

The CPU interface registers can be accessed using the System register interface. See *GIC System register access* on page 8-160 for more information.

Table 8-30 CPU interface register map

Offset	Name	Type	Reset	Description
0x0000	GICC_CTLR	RW	See the register description	CPU Interface Control Register
0x0004	GICC_PMR	RW	0x0000 0000	Interrupt Priority Mask Register
0x0008	GICC_BPR	RW	0x0000 000x[a]	Binary Point Register
0x000C	GICC_IAR	RO	-	Interrupt Acknowledge Register
0x0010	GICC_EOIR	WO	-	End of Interrupt Register
0x0014	GICC_RPR	RO	-	Running Priority Register
0x0018	GICC_HPPIR	RO	-	Highest Priority Pending Interrupt Register
0x001C	GICC_ABPR	RW	0x0000 000x[a]	Aliased Binary Point Register
0x0020	GICC_AIAR	RO	-	Aliased Interrupt Acknowledge Register
0x0024	GICC_AEOIR	WO	-	Aliased End of Interrupt Register
0x0028	GICC_AHPPIR	RO	-	Aliased Highest Priority Pending Interrupt Register
0x002C	GICC_STATUSR	RW	0x0000 0000	Error Reporting Status Register, optional
0x0030-0x003C	-	-	-	Reserved
0x0040-0x00CF	-	-	-	IMPLEMENTATION DEFINED registers
0x00D0-0x00DC	GICC_APR<n>	RW	0x0000 0000	Active Priorities Registers
0x00E0-0x00EC	GICC_NSAPR<n>	RW	0x0000 0000	Non-secure Active Priorities Registers
0x00ED-0x00F8	-	-	-	Reserved
0x00FC	GICC_IIDR	RO	IMPLEMENTATION DEFINED	CPU Interface Identification Register
0x1000	GICC_DIR	WO	-	Deactivate Interrupt Register

a. See the register description for more information.

8.13 The GIC CPU interface register descriptions

This section describes each of the GIC CPU interface registers in register name order.

8.13.1 GICC_ABPR, CPU Interface Aliased Binary Point Register

The GICC_ABPR characteristics are:

Purpose

Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines Group 1 interrupt preemption.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

This register is used only when System register access is not enabled. When System register access is enabled, the System registers ICC_BPR0_EL1 and ICC_BPR1_EL1 provide equivalent functionality.

Configurations

In systems that support two Security states:

- This register is an alias of the Non-secure copy of GICC_BPR.
- Non-secure accesses to this register return a shifted value of the binary point.
- If ICC_CTLR_EL3.CBPR_EL1NS == 1, Secure accesses to this register access ICC_BPR0_EL1.

Attributes

The reset value of this register is defined as (minimum GICC_BPR.Binary_Point + 1), resulting in a permitted range of `0x1-0x4`.

Field descriptions

The GICC_ABPR bit assignments are:

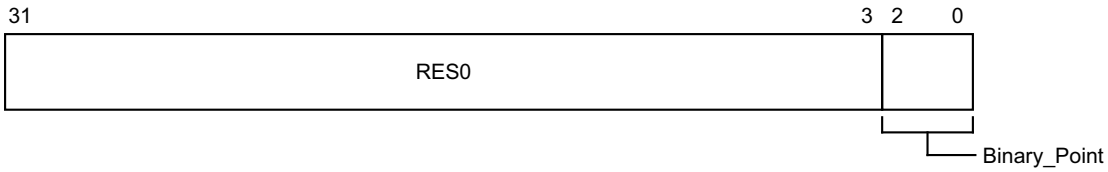

Bits [31:3]

Reserved, RES0.

Binary_Point, bits [2:0]

Controls how the 8-bit interrupt priority field is split into a group priority field, that determines interrupt preemption, and a subpriority field. The following sections describe how this field determines the interrupt priority bits assigned to the group priority field:

- Table 4-8 on page 4-68, for the processing of Group 1 interrupts in a GIC implementation that supports interrupt grouping, when GICC_CTLR.CBPR == 0.
- Table 4-9 on page 4-68, for all other cases.

When this register has an architecturally-defined reset value, this field resets to an IMPLEMENTATION DEFINED value, that might be UNKNOWN.

Accessing the GICC_ABPR:

GICC_ABPR can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	0x001C

8.13.2 GICC_AEOIR, CPU Interface Aliased End Of Interrupt Register

The GICC_AEOIR characteristics are:

Purpose

A write to this register performs priority drop for the specified Group 1 interrupt and, if the appropriate GICC_CTLR.EOImodeS or GICC_CTLR.EOImodeNS field == 0, also deactivates the interrupt.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WO	WO	WO

A write to this register must correspond to the most recently acknowledged Group 1 interrupt. If a value other than the last value read from GICC_AIAR is written to this register, the effect is UNPREDICTABLE.

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_EOIR1 provides equivalent functionality.
- For AArch64 implementations, ICC_EOIR1_EL1 provides equivalent functionality.

When affinity routing is enabled for a Security state, it is a programming error to use memory-mapped registers to access the GIC.

Configurations

In GIC implementations that support two Security states, this register is an alias of the Non-secure view of GICC_EOIR. A Secure access to this register is identical to a Non-secure access to GICC_EOIR.

Attributes

GICC_AEOIR is a 32-bit register.

Field descriptions

The GICC_AEOIR bit assignments are:

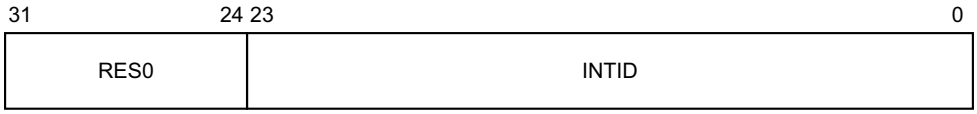

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled interrupt.

Note

INTIDs 1020-1023 are reserved and convey additional information such as spurious interrupts.

When affinity routing is not enabled:

- Bits [23:13] are RES0.

- For SGIs, bits [12:10] identify the CPU interface corresponding to the source PE. For all other interrupts these bits are RES0.

Accessing the GICC_AEOIR:

GICC_AEOIR can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	0x0024

8.13.3 GICC_AHPPIR, CPU Interface Aliased Highest Priority Pending Interrupt Register

The GICC_AHPPIR characteristics are:

Purpose

If the highest priority pending interrupt is in Group 1, this register provides the INTID of the highest priority pending interrupt on the CPU interface.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_HPPIR1 provides equivalent functionality.
- For AArch64 implementations, ICC_HPPIR1_EL1 provides equivalent functionality.

If the highest priority pending interrupt is in Group 0, a read of this register returns the special INTID 1023.

Interrupt identifiers corresponding to an interrupt group that is not enabled are ignored.

If the highest priority pending interrupt is a direct interrupt that is both individually enabled in the Distributor and part of an interrupt group that is enabled in the Distributor, and the interrupt group is disabled in the CPU interface for this PE, this register returns the special INTID 1023.

See *Preemption* on page 4-71 for more information about pending interrupts that are not considered when determining the highest priority pending interrupt.

When affinity routing is enabled for a Security state, it is a programming error to use memory-mapped registers to access the GIC.

Configurations

If the GIC implementation supports two Security states, this register is an alias of the Non-secure view of GICC_HPPIR. A Secure access to this register is identical to a Non-secure access to GICC_HPPIR.

Attributes

GICC_AHPPIR is a 32-bit register.

Field descriptions

The GICC_AHPPIR bit assignments are:

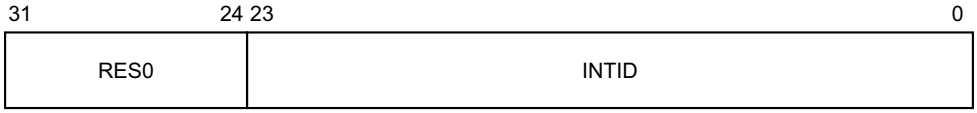

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled interrupt.

Note

INTIDs 1020-1023 are reserved and convey additional information such as spurious interrupts.

When affinity routing is not enabled:

- Bits [23:13] are RES0.
- For SGIs, bits [12:10] identify the CPU interface corresponding to the source PE. For all other interrupts these bits are RES0.

Accessing the GICC_AHPPIR:

GICC_AHPPIR can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	0x0028

8.13.4 GICC_AIAR, CPU Interface Aliased Interrupt Acknowledge Register

The GICC_AIAR characteristics are:

Purpose

Provides the INTID of the signaled Group 1 interrupt. A read of this register by the PE acts as an acknowledge for the interrupt.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

When affinity routing is enabled for a Security state, it is a programming error to use memory-mapped registers to access the GIC.

Configurations

In GIC implementations that support two Security states, this register is an alias of the Non-secure view of GICC_IAR. A Secure access to this register is identical to a Non-secure access to GICC_IAR.

Attributes

GICC_AIAR is a 32-bit register.

Field descriptions

The GICC_AIAR bit assignments are:

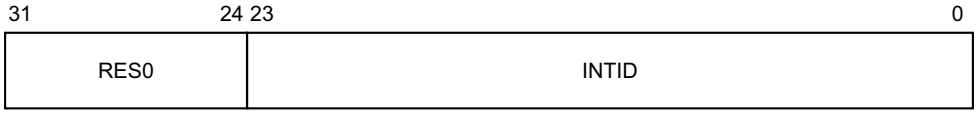

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled interrupt.

Note

INTIDs 1020-1023 are reserved and convey additional information such as spurious interrupts.

When affinity routing is not enabled:

- Bits [23:13] are RES0.
- For SGIs, bits [12:10] identify the CPU interface corresponding to the source PE. For all other interrupts these bits are RES0.

Accessing the GICC_AIAR:

GICC_AIAR can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	0x0020-0x003C

8.13.5 GICC_APR<n>, CPU Interface Active Priorities Registers, n = 0 - 3

The GICC_APR<n> characteristics are:

Purpose

Provides information about interrupt active priorities.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

These registers are used only when System register access is not enabled. When System register access is enabled the following registers provide equivalent functionality:

- In AArch64:
  - For Group 0, ICC_AP0R<n>_EL1.
  - For Group 1, ICC_AP1R<n>_EL1.
- In AArch32:
  - For Group 0, ICC_AP0R<n>.
  - For Group 1, ICC_AP1R<n>.

Configurations

The contents of these registers are IMPLEMENTATION DEFINED with the one architectural requirement that the value `0x00000000` is consistent with no interrupts being active.

When GICD_CTLR.DS == 0, these registers are Banked, and Non-secure accesses do not affect Secure operation. The Secure copies of these registers hold active priorities for Group 0 interrupts, and the Non-secure copies provide a Non-secure view of the active priorities for Group 1 interrupts.

GICC_APR1 is only implemented in implementations that support 6 or more bits of priority. GICC_APR2 and GICC_APR3 are only implemented in implementations that support 7 bits of priority.

When GICD_CTLR.DS==1, these registers hold the active priorities for Group 0 interrupts, and the active priorities for Group 1 interrupts are held by the GICC_NSAPR<n> registers.

Attributes

GICC_APR<n> is a 32-bit register.

Field descriptions

The GICC_APR<n> bit assignments are:

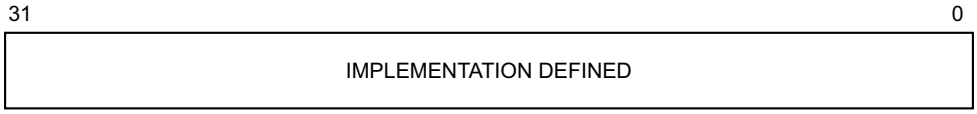

IMPLEMENTATION DEFINED, bits [31:0]

IMPLEMENTATION DEFINED.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the GICC_APR<n>:

GICC_APR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	0x00D0 + 4n

8.13.6 GICC_BPR, CPU Interface Binary Point Register

The GICC_BPR characteristics are:

Purpose

Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

This register is used only when System register access is not enabled. When System register access is enabled this register is RAZ/WI, and the System registers ICC_BPR0_EL1 and ICC_BPR1_EL1 provides equivalent functionality.

Configurations

In systems that support two Security states:

- This register is Banked.
- The Secure instance of this register determines Group 0 interrupt preemption.
- The Non-secure instance of this register determines Group 1 interrupt preemption.

In systems that support only one Security state, when GICC_CTLR.CBPR == 0, this register determines only Group 0 interrupt preemption.

When GICC_CTLR.CBPR == 1, this register determines interrupt preemption for both Group 0 and Group 1 interrupts.

Attributes

GICC_BPR is a 32-bit register.

Field descriptions

The GICC_BPR bit assignments are:

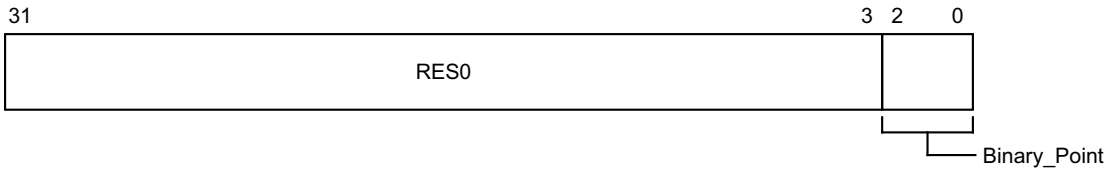

Bits [31:3]

Reserved, RES0.

Binary_Point, bits [2:0]

Controls how the 8-bit interrupt priority field is split into a group priority field, that determines interrupt preemption, and a subpriority field. The following sections describe how this field determines the interrupt priority bits assigned to the group priority field:

- Table 4-8 on page 4-68, for the processing of Group 1 interrupts in a GIC implementation that supports interrupt grouping, when GICC_CTLR.CBPR == 0.
- Table 4-9 on page 4-68, for all other cases.

When this register has an architecturally-defined reset value, this field resets to an IMPLEMENTATION DEFINED value, that might be UNKNOWN.

Note

Aliasing the Non-secure GICC_BPR as GICC_ABPR in a multiprocessor system permits a PE that can make only Secure accesses to configure the preemption setting for Group 1 interrupts by accessing GICC_ABPR.

Accessing the GICC_BPR:

GICC_BPR can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	0x0008

8.13.7 GICC_CTLR, CPU Interface Control Register

The GICC_CTLR characteristics are:

Purpose

Controls the CPU interface, including enabling of interrupt groups, interrupt signal bypass, binary point registers used, and separation of priority drop and interrupt deactivation.

Note

If the GIC implementation supports two Security states, independent EOI controls are provided for accesses from each Security state. Secure accesses handle both Group 0 and Group 1 interrupts, and Non-secure accesses handle Group 1 interrupts only.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_CTLR and ICC_MCTLR provide equivalent functionality.
- For AArch64 implementations, ICC_CTLR_EL1 and ICC_CTLR_EL3 provide equivalent functionality.

Configurations

In a GIC implementation that supports two Security states:

- This register is Banked.
- The register bit assignments are different in the Secure and Non-secure copies.

Attributes

GICC_CTLR is a 32-bit register.

Field descriptions

The GICC_CTLR bit assignments are:

When GICD_CTLR.DS==0, Non-secure access:

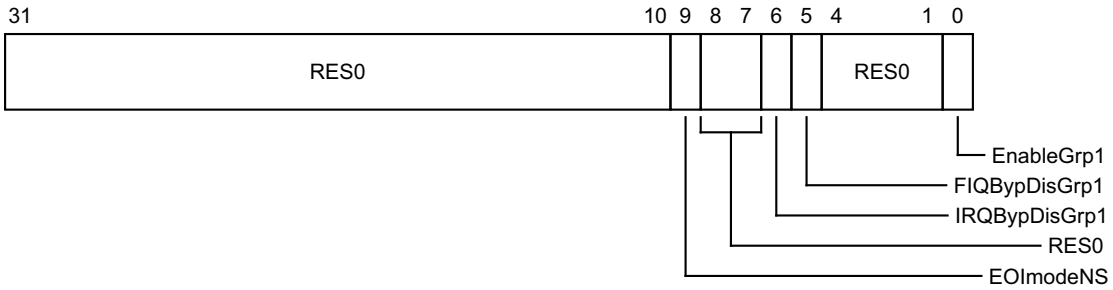

Bits [31:10]

Reserved, RES0.

EOImodeNS, bit [9]

Controls the behavior of Non-secure accesses to GICC_EOIR, GICC_AEOIR, and GICC_DIR.

0 GICC_EOIR and GICC_AEOIR provide both priority drop and interrupt deactivation functionality. Accesses to GICC_DIR are UNPREDICTABLE.

1 GICC_EOIR and GICC_AEOIR provide priority drop functionality only. GICC_DIR provides interrupt deactivation functionality.

Note

An implementation is permitted to make this bit RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

Bits [8:7]

Reserved, RES0.

IRQBypDisGrp1, bit [6]

When the signaling of IRQs by the CPU interface is disabled, this field partly controls whether the bypass IRQ signal is signaled to the PE for Group 1:

0 The bypass IRQ signal is signaled to the PE.

1 The bypass IRQ signal is not signaled to the PE.

If System register access is enabled for EL3 and ICC_SRE_EL3.DIB == 1, this field is RAO/WI.

If System register access is enabled for EL1, this field is ignored.

If an implementation does not support legacy interrupts, this bit is permitted to be RAO/WI.

See *Interrupt bypass support* on page 2-35 for more information.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

FIQBypDisGrp1, bit [5]

When the signaling of FIQs by the CPU interface is disabled, this field partly controls whether the bypass FIQ signal is signaled to the PE for Group 1:

0 The bypass FIQ signal is signaled to the PE.

1 The bypass FIQ signal is not signaled to the PE.

If System register access is enabled for EL3 and ICC_SRE_EL3.DFB == 1, this field is RAO/WI.

If System register access is enabled for EL1, this field is ignored.

If an implementation does not support legacy interrupts, this bit is permitted to be RAO/WI.

See *Interrupt bypass support* on page 2-35 for more information.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

Bits [4:1]

Reserved, RES0.

EnableGrp1, bit [0]

This Non-secure field enables the signaling of Group 1 interrupts by the CPU interface to a target PE:

0 Group 1 interrupt signaling is disabled.

1 Group 1 interrupt signaling is enabled.

When this register has an architecturally-defined reset value, this field resets to 0.

When GICD_CTLR.DS==0, Secure access:

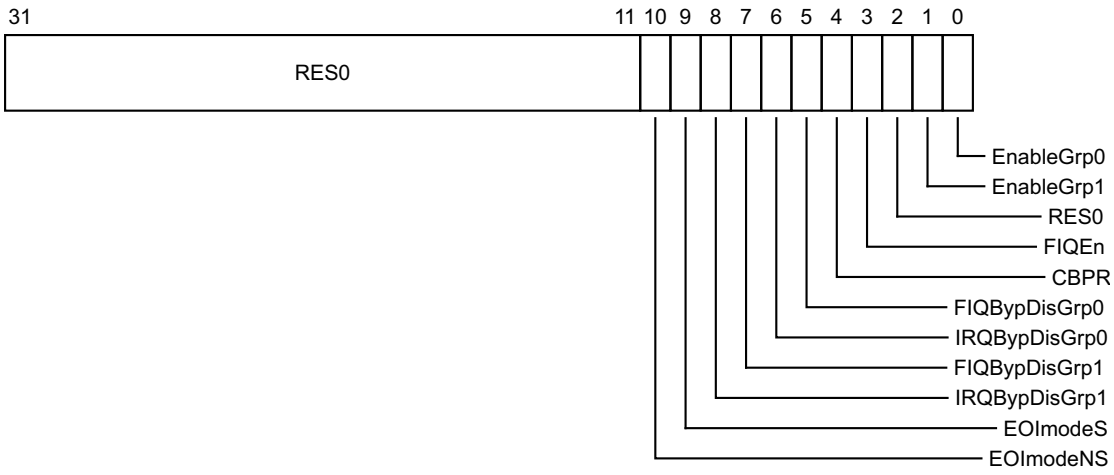

Bits [31:11]

Reserved, RES0.

EOImodeNS, bit [10]

Controls the behavior of Non-secure accesses to GICC_EOIR, GICC_AEOIR, and GICC_DIR.

`0` GICC_EOIR and GICC_AEOIR provide both priority drop and interrupt deactivation functionality. Accesses to GICC_DIR are UNPREDICTABLE.

`1` GICC_EOIR and GICC_AEOIR provide priority drop functionality only. GICC_DIR provides interrupt deactivation functionality.

Note

An implementation is permitted to make this bit RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

EOImodeS, bit [9]

Controls the behavior of Secure accesses to GICC_EOIR, GICC_AEOIR, and GICC_DIR.

`0` GICC_EOIR and GICC_AEOIR provide both priority drop and interrupt deactivation functionality. Accesses to GICC_DIR are UNPREDICTABLE.

`1` GICC_EOIR and GICC_AEOIR provide priority drop functionality only. GICC_DIR provides interrupt deactivation functionality.

Note

An implementation is permitted to make this bit RAO/WI.

This field shares state with GICC_CTLR.EOImode.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

IRQBypDisGrp1, bit [8]

When the signaling of IRQs by the CPU interface is disabled, this field partly controls whether the bypass IRQ signal is signaled to the PE for Group 1:

`0` The bypass IRQ signal is signaled to the PE.

`1` The bypass IRQ signal is not signaled to the PE.

If System register access is enabled for EL3 and ICC_SRE_EL3.DIB == 1, this field is RAO/WI.

If System register access is enabled for EL1, this field is ignored.

If an implementation does not support legacy interrupts, this bit is permitted to be RAO/WI.

See *Interrupt bypass support* on page 2-35 for more information.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

FIQBypDisGrp1, bit [7]

When the signaling of FIQs by the CPU interface is disabled, this field partly controls whether the bypass FIQ signal is signaled to the PE for Group 1:

`0` The bypass FIQ signal is signaled to the PE.

`1` The bypass FIQ signal is not signaled to the PE.

If System register access is enabled for EL3 and ICC_SRE_EL3.DFB == 1, this field is RAO/WI.

If System register access is enabled for EL1, this field is ignored.

If an implementation does not support legacy interrupts, this bit is permitted to be RAO/WI.

See *Interrupt bypass support* on page 2-35 for more information.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

IRQBypDisGrp0, bit [6]

When the signaling of IRQs by the CPU interface is disabled, this field partly controls whether the bypass IRQ signal is signaled to the PE for Group 0:

`0` The bypass IRQ signal is signaled to the PE.

`1` The bypass IRQ signal is not signaled to the PE.

If System register access is enabled for EL3 and ICC_SRE_EL3.DIB == 1, this field is RAO/WI.

If System register access is enabled for EL1, this field is ignored.

If an implementation does not support legacy interrupts, this bit is permitted to be RAO/WI.

See *Interrupt bypass support* on page 2-35 for more information.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

FIQBypDisGrp0, bit [5]

When the signaling of FIQs by the CPU interface is disabled, this field partly controls whether the bypass FIQ signal is signaled to the PE for Group 0:

`0` The bypass FIQ signal is signaled to the PE.

`1` The bypass FIQ signal is not signaled to the PE.

If System register access is enabled for EL3 and ICC_SRE_EL3.DIB == 1, this field is RAO/WI.

If System register access is enabled for EL1, this field is ignored.

If an implementation does not support legacy interrupts, this bit is permitted to be RAO/WI.

See *Interrupt bypass support* on page 2-35 for more information.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

CBPR, bit [4]

Controls whether GICC_BPR provides common control of preemption to Group 0 and Group 1 interrupts:

`0` GICC_BPR determines preemption for Group 0 interrupts only.
GICC_ABPR determines preemption for Group 1 interrupts.

`1` GICC_BPR determines preemption for both Group 0 and Group 1 interrupts.

This field is an alias of ICC_CTLR_EL3.CBPR_EL1NS.

In a GIC that supports two Security states, when CBPR == 1:

- A Non-secure read of GICC_BPR returns the value of Secure GICC_BPR.BinaryPoint, incremented by 1, and saturated to `0b111`.
- Non-secure writes of GICC_BPR are ignored.

When this register has an architecturally-defined reset value, this field resets to `0`.

FIQEn, bit [3]

Controls whether the CPU interface signals Group 0 interrupts to a target PE using the FIQ or IRQ signal:

`0` Group 0 interrupts are signaled using the IRQ signal.

`1` Group 0 interrupts are signaled using the FIQ signal.

Group 1 interrupts are signaled using the IRQ signal only.

If an implementation supports two Security states, this bit is permitted to be RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

Bit [2]

Reserved, RES0.

EnableGrp1, bit [1]

This Non-secure field enables the signaling of Group 1 interrupts by the CPU interface to a target PE:

`0` Group 1 interrupt signaling is disabled.

`1` Group 1 interrupt signaling is enabled.

When this register has an architecturally-defined reset value, this field resets to `0`.

EnableGrp0, bit [0]

Enables the signaling of Group 0 interrupts by the CPU interface to a target PE:

`0` Group 0 interrupt signaling is disabled.

`1` Group 0 interrupt signaling is enabled.

When this register has an architecturally-defined reset value, this field resets to `0`.

When GICD_CTLR.DS==1:

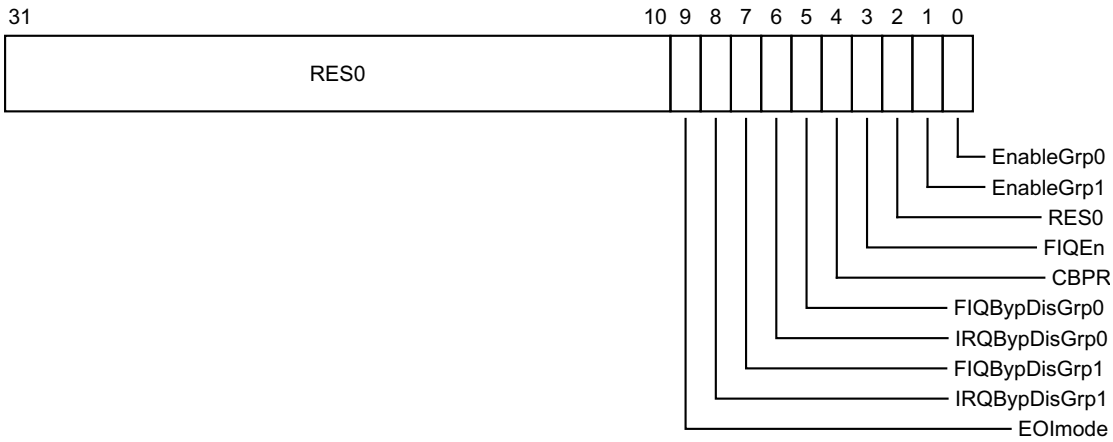

Bits [31:10]

Reserved, RES0.

EOImode, bit [9]

Controls the behavior of accesses to GICC_EOIR, GICC_AEOIR, and GICC_DIR.

0 GICC_EOIR and GICC_AEOIR provide both priority drop and interrupt deactivation functionality. Accesses to GICC_DIR are UNPREDICTABLE.

1 GICC_EOIR and GICC_AEOIR provide priority drop functionality only. GICC_DIR provides interrupt deactivation functionality.

Note

An implementation is permitted to make this bit RAO/WI.

This field shares state with GICC_CTLR.EOImodeS.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

IRQBypDisGrp1, bit [8]

When the signaling of IRQs by the CPU interface is disabled, this field partly controls whether the bypass IRQ signal is signaled to the PE for Group 1:

0 The bypass IRQ signal is signaled to the PE.

1 The bypass IRQ signal is not signaled to the PE.

If System register access is enabled for EL3 and ICC_SRE_EL3.DIB == 1, this field is RAO/WI.

If System register access is enabled for EL1, this field is ignored.

If an implementation does not support legacy interrupts, this bit is permitted to be RAO/WI.

See *Interrupt bypass support* on page 2-35 for more information.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

FIQBypDisGrp1, bit [7]

When the signaling of FIQs by the CPU interface is disabled, this field partly controls whether the bypass FIQ signal is signaled to the PE for Group 1:

0 The bypass FIQ signal is signaled to the PE.

1 The bypass FIQ signal is not signaled to the PE.

If System register access is enabled for EL3 and ICC_SRE_EL3.DFB == 1, this field is RAO/WI.

If System register access is enabled for EL1, this field is ignored.

If an implementation does not support legacy interrupts, this bit is permitted to be RAO/WI.

See *Interrupt bypass support* on page 2-35 for more information.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

IRQBypDisGrp0, bit [6]

When the signaling of IRQs by the CPU interface is disabled, this field partly controls whether the bypass IRQ signal is signaled to the PE for Group 0:

0 The bypass IRQ signal is signaled to the PE.

1 The bypass IRQ signal is not signaled to the PE.

If System register access is enabled for EL3 and ICC_SRE_EL3.DIB == 1, this field is RAO/WI.

If System register access is enabled for EL1, this field is ignored.

If an implementation does not support legacy interrupts, this bit is permitted to be RAO/WI.

See *Interrupt bypass support* on page 2-35 for more information.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

FIQBypDisGrp0, bit [5]

When the signaling of FIQs by the CPU interface is disabled, this field partly controls whether the bypass FIQ signal is signaled to the PE for Group 0:

0 The bypass FIQ signal is signaled to the PE.

1 The bypass FIQ signal is not signaled to the PE.

If System register access is enabled for EL3 and ICC_SRE_EL3.DIB == 1, this field is RAO/WI.

If System register access is enabled for EL1, this field is ignored.

If an implementation does not support legacy interrupts, this bit is permitted to be RAO/WI.

See *Interrupt bypass support* on page 2-35 for more information.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

CBPR, bit [4]

Controls whether GICC_BPR provides common control of preemption to Group 0 and Group 1 interrupts:

0 GICC_BPR determines preemption for Group 0 interrupts only.
GICC_ABPR determines preemption for Group 1 interrupts.

1 GICC_BPR determines preemption for both Group 0 and Group 1 interrupts.

This field is an alias of ICC_CTLR_EL3.CBPR_EL1NS.

In a GIC that supports two Security states, when CBPR == 1:

- A Non-secure read of GICC_BPR returns the value of Secure GICC_BPR.BinaryPoint, incremented by 1, and saturated to 0b111.
- Non-secure writes of GICC_BPR are ignored.

When this register has an architecturally-defined reset value, this field resets to 0.

FIQEn, bit [3]

Controls whether the CPU interface signals Group 0 interrupts to a target PE using the FIQ or IRQ signal:

0 Group 0 interrupts are signaled using the IRQ signal.

1 Group 0 interrupts are signaled using the FIQ signal.

Group 1 interrupts are signaled using the IRQ signal only.

If an implementation supports two Security states, this bit is permitted to be RAO/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to 0.

Bit [2]

Reserved, RES0.

EnableGrp1, bit [1]

This Non-secure field enables the signaling of Group 1 interrupts by the CPU interface to a target PE:

0 Group 1 interrupt signaling is disabled.

1 Group 1 interrupt signaling is enabled.

When this register has an architecturally-defined reset value, this field resets to 0.

EnableGrp0, bit [0]

Enables the signaling of Group 0 interrupts by the CPU interface to a target PE:

`0` Group 0 interrupt signaling is disabled.

`1` Group 0 interrupt signaling is enabled.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the GICC_CTLR:

GICC_CTLR can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	`0x0000`

8.13.8 GICC_DIR, CPU Interface Deactivate Interrupt Register

The GICC_DIR characteristics are:

Purpose

When interrupt priority drop is separated from interrupt deactivation, a write to this register deactivates the specified interrupt.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WO	WO	WO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_DIR provides equivalent functionality.
- For AArch64 implementations, ICC_DIR_EL1 provides equivalent functionality.

Writes to this register have an effect only in the following cases:

- When GICD_CTLR.DS == 1, if GICC_CTLR.EOImode == 1.
- In GIC implementations that support two Security states:
  - If the access is Secure and GICC_CTLR.EOImodeS == 1.
  - If the access is Non-secure and GICC_CTLR.EOImodeNS == 1.

The following writes must be ignored:

- Writes to this register when the corresponding EOImode field in GICC_CTLR == 0. In systems that support system error generation, an implementation might generate a system error.
- Writes to this register when the corresponding EOImode field in GICC_CTLR == 0 and the corresponding interrupt is not active. In systems that support system error generation, an implementation might generate a system error. In implementations using the GIC Stream Protocol interface these writes correspond to a Deactivate for an interrupt that is not active.

If the corresponding EOImode field in GICC_CTLR is 1 and this register is written to without a corresponding write to GICC_EOIR or GICC_AEOIR, the interrupt is deactivated but the bit corresponding to it in the active priorities registers remains set.

When affinity routing is enabled for a Security state, it is a programming error to use memory-mapped registers to access the GIC.

Configurations

There are no configuration notes.

Attributes

GICC_DIR is a 32-bit register.

Field descriptions

The GICC_DIR bit assignments are:

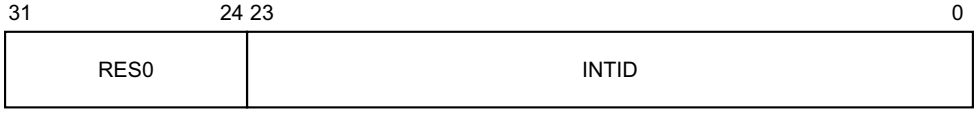

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled interrupt.

Note

INTIDs 1020-1023 are reserved and convey additional information such as spurious interrupts.

When affinity routing is not enabled:

- Bits [23:13] are RES0.
- For SGIs, bits [12:10] identify the CPU interface corresponding to the source PE. For all other interrupts these bits are RES0.

Accessing the GICC_DIR:

GICC_DIR can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	0x1000

8.13.9 GICC_EOIR, CPU Interface End Of Interrupt Register

The GICC_EOIR characteristics are:

Purpose

A write to this register performs priority drop for the specified interrupt and, if the appropriate GICC_CTLR.EOImodeS or GICC_CTLR.EOImodeNS field == 0, also deactivates the interrupt.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WO	WO	WO

The following writes must be ignored:

- Writes of INTIDs 1020-1023.
- Secure writes corresponding to Group 1 interrupts. In systems that support system error generation, an implementation might generate a system error. In this case, GIC behavior is predictable, and the highest Secure active priority (in the Secure copy of GICC_APR<n>) will be reset if the highest active priority is Secure. System behavior is UNPREDICTABLE.
- Non-secure writes corresponding to Group 0 interrupts when GICC_CTLR.EOImodeS == 1. In systems that support system error generation, an implementation might generate a system error. In this case, GIC behavior is predictable, and the highest Non-secure active priority (in the Non-secure copy of GICC_APR<n>) will be reset if the highest active priority is Non-secure. System behavior is UNPREDICTABLE.

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_EOIR0 and ICC_EOIR1 provide equivalent functionality.
- For AArch64 implementations, ICC_EOIR0_EL1 and ICC_EOIR1_EL1 provide equivalent functionality.

When affinity routing is enabled for a Security state, it is a programming error to use memory-mapped registers to access the GIC.

Configurations

If the GIC implementation supports two Security states:

- This register is Common.
- GICC_AEOIR is an alias of the Non-secure view of this register.

In GIC implementations that support only a single Security state, or on a Secure write, the register provides functionality for Group 0 interrupts.

Non-secure writes in a GIC implementation that supports two Security states provide functionality for Group 1 interrupts.

Attributes

GICC_EOIR is a 32-bit register.

Field descriptions

The GICC_EOIR bit assignments are:

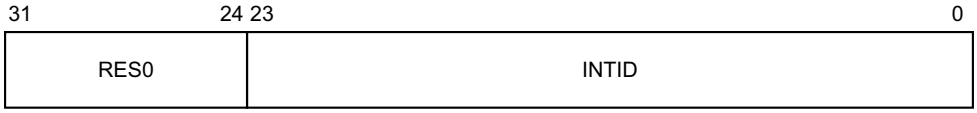

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled interrupt.

Note

INTIDs 1020-1023 are reserved and convey additional information such as spurious interrupts.

When affinity routing is not enabled:

- Bits [23:13] are RES0.
- For SGIs, bits [12:10] identify the CPU interface corresponding to the source PE. For all other interrupts these bits are RES0.

For every read of a valid INTID from GICC_IAR, the connected PE must perform a matching write to GICC_EOIR. The value written to GICC_EOIR must be the INTID from GICC_IAR. Reads of INTIDs 1020-1023 do not require matching writes.

Note

ARM recommends that software preserves the entire register value read from GICC_IAR, and writes that value back to GICC_EOIR on completion of interrupt processing.

For nested interrupts, the order of writes to this register must be the reverse of the order of interrupt acknowledgement. Behavior is UNPREDICTABLE if:

- This ordering constraint is not maintained.
- The value written to this register does not match an active interrupt, or the ID of a spurious interrupt.
- The value written to this register does not match the last valid interrupt value read from GICC_IAR.

See *Interrupt lifecycle* on page 4-46 for general information about the effect of writes to end of interrupt registers, and about the possible separation of the priority drop and interrupt deactivate operations.

If the GIC implementation supports two Security states:

- GICC_CTLR.EOImodeS controls the behavior of Secure accesses to GICC_EOIR and GICC_AEOIR.
- GICC_CTLR.EOImodeNS controls the behavior of Non-secure accesses to GICC_EOIR and GICC_AEOIR.

Accessing the GICC_EOIR:

GICC_EOIR can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	0x0010

8.13.10 GICC_HPPIR, CPU Interface Highest Priority Pending Interrupt Register

The GICC_HPPIR characteristics are:

Purpose

Provides the INTID of the highest priority pending interrupt on the CPU interface.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_HPPIR0 and ICC_HPPIR1 provide equivalent functionality.
- For AArch64 implementations, ICC_HPPIR0_EL1 and ICC_HPPIR1_EL1 provide equivalent functionality.

If the highest priority pending interrupt is in Group 0, a Non-secure read of this register returns the special INTID 1023.

If the highest priority pending interrupt is in Group 1, a Secure read of this register (or a read in a GIC implementation that supports only a single Security state) returns the special INTID 1022.

If no interrupts are in the pending state, a read of this register returns the special INTID 1023.

Interrupt identifiers corresponding to an interrupt group that is not enabled are ignored.

If the highest priority pending interrupt is a direct interrupt that is both individually enabled in the Distributor and part of an interrupt group that is enabled in the Distributor, and the interrupt group is disabled in the CPU interface for this PE, this register returns the special INTID 1023.

See *Preemption* on page 4-71 for more information about pending interrupts that are not considered when determining the highest priority pending interrupt.

When affinity routing is enabled for a Security state, it is a programming error to use memory-mapped registers to access the GIC.

Configurations

If the GIC implementation supports two Security states:

- This register is Common.
- GICC_AHPPIR is an alias of the Non-secure view of this register.

Attributes

GICC_HPPIR is a 32-bit register.

Field descriptions

The GICC_HPPIR bit assignments are:

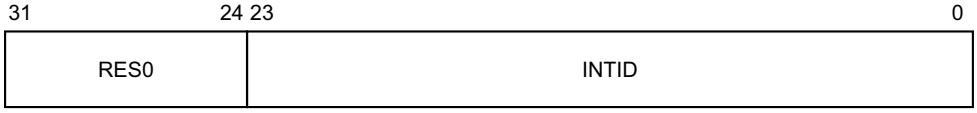

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled interrupt.

Note

INTIDs 1020-1023 are reserved and convey additional information such as spurious interrupts.

When affinity routing is not enabled:

- Bits [23:13] are RES0.
- For SGIs, bits [12:10] identify the CPU interface corresponding to the source PE. For all other interrupts these bits are RES0.

Accessing the GICC_HPPIR:

GICC_HPPIR can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	0x0018

8.13.11 GICC_IAR, CPU Interface Interrupt Acknowledge Register

The GICC_IAR characteristics are:

Purpose

Provides the INTID of the signaled interrupt. A read of this register by the PE acts as an acknowledge for the interrupt.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

When GICD_CTLR.DS==1, if the highest priority pending interrupt is in Group 1, the special INTID 1022 is returned.

In GIC implementations that support two Security states, if the highest priority pending interrupt is in Group 0, Non-secure reads return the special INTID 1023.

In GIC implementations that support two Security states, if the highest priority pending interrupt is in Group 1, Secure reads return the special INTID 1022.

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_IAR0 and ICC_IAR1 provide equivalent functionality.
- For AArch64 implementations, ICC_IAR0_EL1 and ICC_IAR1_EL1 provide equivalent functionality.

When affinity routing is enabled for a Security state, it is a programming error to use memory-mapped registers to access the GIC.

Configurations

This register is available in all configurations of the GIC. If the GIC implementation supports two Security states:

- This register is Common.
- GICC_AIAR is an alias of the Non-secure view of this register.

The format of the INTID is governed by whether affinity routing is enabled for a Security state.

Attributes

GICC_IAR is a 32-bit register.

Field descriptions

The GICC_IAR bit assignments are:

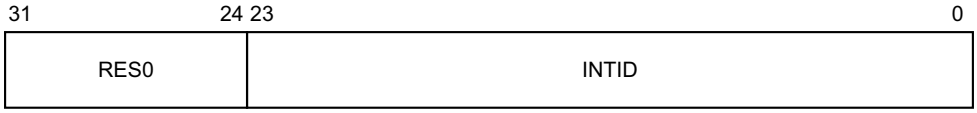

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

The INTID of the signaled interrupt.

Note

INTIDs 1020-1023 are reserved and convey additional information such as spurious interrupts.

When affinity routing is not enabled:

- Bits [23:13] are RES0.
- For SGIs, bits [12:10] identify the CPU interface corresponding to the source PE. For all other interrupts these bits are RES0.

A read of this register returns the INTID of the highest priority pending interrupt for the CPU interface. The read returns a spurious INTID of 1023 if any of the following apply:

- Forwarding of interrupts by the Distributor to the CPU interface is disabled.
- Signaling of interrupts by the CPU interface to the connected PE is disabled.
- There are no pending interrupts on the CPU interface with sufficient priority for the interface to signal it to the PE.

When the GIC returns a valid INTID to a read of this register it treats the read as an acknowledge of that interrupt. In addition, it changes the interrupt status from pending to active, or to active and pending if the pending state of the interrupt persists. Normally, the pending state of an interrupt persists only if the interrupt is level-sensitive and remains asserted.

For every read of a valid INTID from GICC_IAR, the connected PE must perform a matching write to GICC_EOIR.

Note

- ARM recommends that software preserves the entire register value read from this register, and writes that value back to GICC_EOIR on completion of interrupt processing.
- For SPIs, although multiple target PEs might attempt to read this register at any time, only one PE can obtain a valid INTID. See *Activation* on page 4-47 for more information.

Accessing the GICC_IAR:

GICC_IAR can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	0x000C

8.13.12 GICC_IIDR, CPU Interface Identification Register

The GICC_IIDR characteristics are:

Purpose

Provides information about the implementer and revision of the CPU Interface.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

Configurations

There are no configuration notes.

Attributes

GICC_IIDR is a 32-bit register.

Field descriptions

The GICC_IIDR bit assignments are:

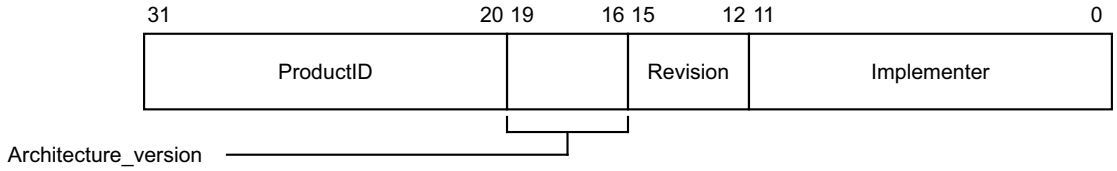

ProductID, bits [31:20]

An IMPLEMENTATION DEFINED product identifier.

Architecture_version, bits [19:16]

The version of the GIC architecture that is implemented.

`0001` GICv1.

`0010` GICv2.

`0011` GICv3.

`0100` GICv4 memory-mapped interface supported.

Other values are reserved.

Revision, bits [15:12]

An IMPLEMENTATION DEFINED revision number for the CPU interface.

Implementer, bits [11:0]

Contains the JEP106 code of the company that implemented the CPU interface.

- Bits [11:8] are the JEP106 continuation code of the implementer. For an ARM implementation, this field is `0x4`.
- Bit [7] is always 0.
- Bits [6:0] are the JEP106 identity code of the implementer. For an ARM implementation, bits [7:0] are therefore `0x3B`.

Accessing the GICC_IIDR:

GICC_IIDR can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	0x00FC

8.13.13 GICC_NSAPR<n>, CPU Interface Non-secure Active Priorities Registers, n = 0 - 3

The GICC_NSAPR<n> characteristics are:

Purpose

Provides information about Group 1 interrupt active priorities.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

Configurations

The contents of these registers are IMPLEMENTATION DEFINED with the one architectural requirement that the value `0x00000000` is consistent with no interrupts being active.

When GICD_CTLR.DS==0, these registers are RAZ/WI to Non-secure accesses.

GICC_NSAPR1 is only implemented in implementations that support 6 or more bits of priority. GICC_NSAPR2 and GICC_NSAPR3 are only implemented in implementations that support 7 bits of priority.

Attributes

GICC_NSAPR<n> is a 32-bit register.

Field descriptions

The GICC_NSAPR<n> bit assignments are:

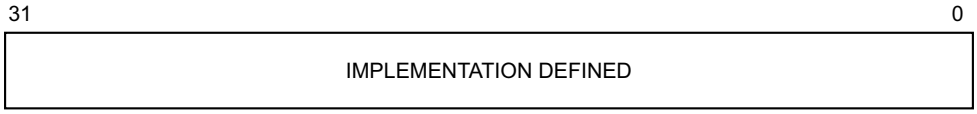

IMPLEMENTATION DEFINED, bits [31:0]

IMPLEMENTATION DEFINED.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the GICC_NSAPR<n>:

GICC_NSAPR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	`0x00E0` + 4n

8.13.14 GICC_PMR, CPU Interface Priority Mask Register

The GICC_PMR characteristics are:

Purpose

This register provides an interrupt priority filter. Only interrupts with higher priority than the value in this register are signaled to the PE.

Note

Higher interrupt priority corresponds to a lower value of the Priority field.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

If the GIC implementation supports two Security states:

- Non-secure accesses to this register can only read or write values corresponding to the lower half of the priority range.
- If a Secure write has programmed the register with a value that corresponds to a value in the upper half of the priority range then:
  - Any Non-secure read of the register returns `0x00`, regardless of the value held in the register.
  - Non-secure writes are ignored.

See *Interrupt prioritization* on page 4-65 for more information.

Configurations

This register is available in all configurations of the GIC. If the GIC implementation supports two Security states this register is Common.

Attributes

GICC_PMR is a 32-bit register.

Field descriptions

The GICC_PMR bit assignments are:

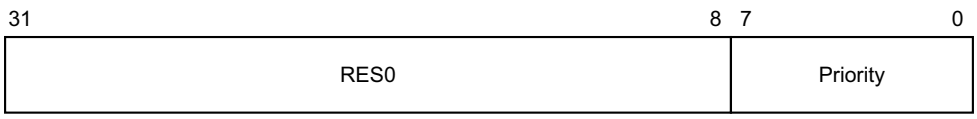

Bits [31:8]

Reserved, RES0.

Priority, bits [7:0]

The priority mask level for the CPU interface. If the priority of the interrupt is higher than the value indicated by this field, the interface signals the interrupt to the PE.

If the GIC implementation supports fewer than 256 priority levels some bits might be RAZ/WI, as follows:

- For 128 supported levels, bit [0] = `0b0`.

- For 64 supported levels, bits [1:0] = `0b00`.
- For 32 supported levels, bits [2:0] = `0b000`.
- For 16 supported levels, bits [3:0] = `0b0000`.

See *Interrupt prioritization* on page 4-65 for more information.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICC_PMR:

GICC_PMR can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	`0x0004`

8.13.15 GICC_RPR, CPU Interface Running Priority Register

The GICC_RPR characteristics are:

Purpose

This register indicates the running priority of the CPU interface.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

If there is no active interrupt on the CPU interface, the idle priority value is returned.

If the GIC implementation supports two Security states, a Non-secure read of the Priority field returns:

- `0x00` if the field value is less than `0x80`.
- The Non-secure view of the Priority value if the field value is `0x80` or more.

See *Interrupt prioritization* on page 4-65 for more information.

Note

Software cannot determine the number of implemented priority bits from this register.

Configurations

This register is available in all configurations of the GIC. If the GIC implementation supports two Security states this register is Common.

Attributes

GICC_RPR is a 32-bit register.

Field descriptions

The GICC_RPR bit assignments are:

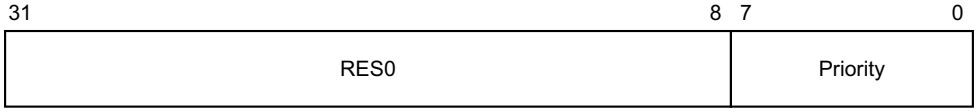

Bits [31:8]

Reserved, RES0.

Priority, bits [7:0]

The current running priority on the CPU interface.

Accessing the GICC_RPR:

GICC_RPR can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	`0x0014`

8.13.16 GICC_STATUSR, CPU Interface Status Register

The GICC_STATUSR characteristics are:

Purpose

Provides software with a mechanism to detect:

- Accesses to reserved locations.
- Writes to read-only locations.
- Reads of write-only locations.

Usage constraints

GICC_STATUSR(S) is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	-

GICC_STATUSR(NS) is accessible as follows:

Security disabled	Secure	Non-secure
RW	-	RW

This is an optional register. If the register is not implemented, the location is RAZ/WI.

If this register is implemented, GICV_STATUSR must also be implemented.

Configurations

If the GIC implementation supports two Security states this register is Banked to provide Secure and Non-secure copies.

This register is used only when System register access is not enabled. If System register access is enabled, this register is not updated. Equivalent functionality might be provided by appropriate traps and exceptions.

Attributes

GICC_STATUSR is a 32-bit register.

Field descriptions

The GICC_STATUSR bit assignments are:

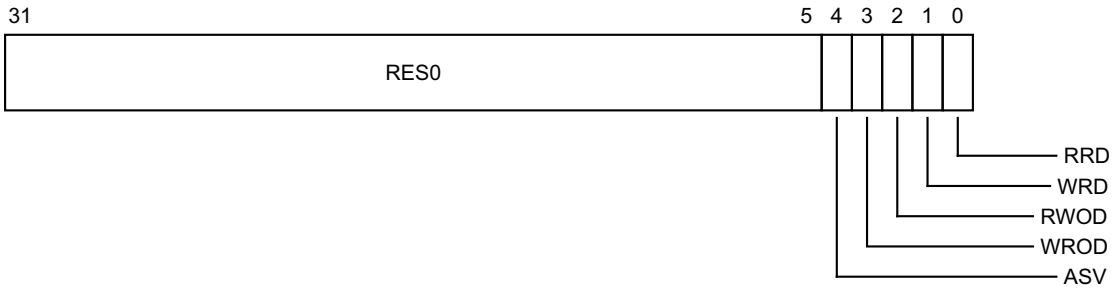

Bits [31:5]

Reserved, RES0.

ASV, bit [4]

Attempted security violation.

0 Normal operation.

1 A Non-secure access to a Secure register has been detected.

Note

This bit is not set to 1 for registers where any of the fields are Non-secure.

WROD, bit [3]

Write to an RO location.

0 Normal operation.

1 A write to an RO location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

RWOD, bit [2]

Read of a WO location.

0 Normal operation.

1 A read of a WO location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

WRD, bit [1]

Write to a reserved location.

0 Normal operation.

1 A write to a reserved location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

RRD, bit [0]

Read of a reserved location.

0 Normal operation.

1 A read of a reserved location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

Accessing the GICC_STATUSR:

GICC_STATUSR can be accessed through the memory-mapped interface:

Component	Offset
GIC CPU interface	0x002C

8.14 The GIC virtual CPU interface register map

Note

Unless explicitly defined otherwise in this section, the GICV_* registers are as defined in the GICv2 specification, see *ARM® Generic Interrupt Controller, Architecture version 2.0, Architecture Specification*.

These registers provide the virtual CPU interface accessed by the virtual machine. Typically, a virtual machine is unaware of any difference between virtual interrupts and physical interrupts. This means the programmers' model for handling virtual interrupts must be identical to that for handling physical interrupts. In general, these registers have the same format as the GIC physical CPU interface registers, but they operate on the interrupt view defined primarily by the List registers.

These registers are memory-mapped, with defined offsets from an IMPLEMENTATION DEFINED GICV_* register base address.

Note

The offset of each GICV_* register is the same as the offset of the corresponding register for the physical CPU interface. For example, GICV_PMR is at offset `0x0004` from the GICV_* register base address, and GICC_PMR is at the same offset from the GICC_* register base address.

This means that:

- The hypervisor can use the stage 2 address translations to map the virtual CPU interface accesses to the correct physical addresses.
- Software, whether accessing the registers of a physical CPU interface or of a virtual CPU interface, uses the same register addresses.

To enable use of 64KB pages, the GICV_* memory map must ensure that:

- The base address of the GICV_* registers is 64KB aligned.
- An alias of the GICV_* registers is provided starting at offset `0xF000` from the start of the page such that a second copy of GICV_DIR exists at the start of the next 64KB page.

This provides support for both 4KB and 64KB pages.

Table 8-31 shows the GIC virtual CPU interface register map.

Table 8-31 GIC virtual CPU interface register map

Offset	Name	Type	Reset	Description
`0x0000`	GICV_CTLR	RW	See the register description	VM Control Register
`0x0004`	GICV_PMR	RW	`0x0000 0000`	VM Priority Mask Register
`0x0008`	GICV_BPR	RW	`0x0000 000x`[a]	VM Binary Point Register
`0x000C`	GICV_IAR	RO	-	VM Interrupt Acknowledge Register
`0x0010`	GICV_EOIR	WO	-	VM End of Interrupt Register
`0x0014`	GICV_RPR	RO	-	VM Running Priority Register
`0x0018`	GICV_HPPIR	RO	-	VM Highest Priority Pending Interrupt Register
`0x001C`	GICV_ABPR	RW	`0x0000 000x`[a]	VM Aliased Binary Point Register
`0x0020`	GICV_AIAR	RO	-	VM Aliased Interrupt Acknowledge Register
`0x0024`	GICV_AEOIR	WO	-	VM Aliased End of Interrupt Register

Table 8-31 GIC virtual CPU interface register map (continued)

Offset	Name	Type	Reset	Description
0x0028	GICV_AHPPIR	RO	-	VM Aliased Highest Priority Pending Interrupt Register
0x002C	GICV_STATUSR	RW	0x0000 0000	VM Error Reporting Status Register, optional
0x0030-0x003C	-	-	-	Reserved
0x0040-0x00CC	-	-	-	IMPLEMENTATION DEFINED
0x00D0-0x00DC	GICV_APR<n>	RW	0x0000 0000	VM Active Priorities Registers
0x00E0-0x00EC	-	-	RAZ/WI	Reserved for second set of Active Priorities Registers, as the Note in the description describes
0x00F0-0x00F8	-	-	-	Reserved
0x00FC	GICV_IIDR	RO	IMPLEMENTATION DEFINED	VM CPU Interface Identification Register
0x0100-0x0FFC	-	-	-	Reserved
0x1000	GICV_DIR	WO	-	VM Deactivate Interrupt Register
0x10000				
0x1004-0x1FFC	-	-	-	Reserved

a. See the register description for more information.

8.15 The GIC virtual CPU interface register descriptions

This section describes each of the GIC virtual CPU interface registers in register name order.

8.15.1 GICV_ABPR, Virtual Machine Aliased Binary Point Register

The GICV_ABPR characteristics are:

Purpose

Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines Group 1 interrupt preemption.

This register corresponds to GICC_ABPR in the physical CPU interface.

Note

GICH_LR<n>.Group determines whether a virtual interrupt is Group 0 or Group 1.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_BPR1 provides equivalent functionality.
- For AArch64 implementations, ICC_BPR1_EL1 provides equivalent functionality.

The value contained in this register is one greater than the actual applied binary point value, as described in *Priority grouping* on page 4-67.

This register is used for Group 1 interrupts when GICV_CTLR.CBPR == 0. GICV_BPR provides equivalent functionality for Group 0 interrupts, and for Group 1 interrupts when GICV_CTLR.CBPR == 1.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICV_ABPR is a 32-bit register.

Field descriptions

The GICV_ABPR bit assignments are:

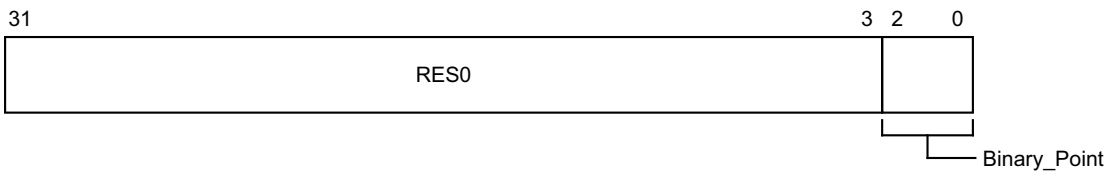

Bits [31:3]

Reserved, RES0.

Binary_Point, bits [2:0]

Controls how the 8-bit interrupt priority field is split into a group priority field, that determines interrupt preemption, and a subpriority field.

For information about how this field determines the interrupt priority bits assigned to the group priority field, see Table 4-8 on page 4-68.

When this register has an architecturally-defined reset value, this field resets to `0`.

The Binary_Point field of this register is aliased to GICH_VMCR.VBPR1.

Accessing the GICV_ABPR:

GICV_ABPR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	`0x001C`

8.15.2 GICV_AEOIR, Virtual Machine Aliased End Of Interrupt Register

The GICV_AEOIR characteristics are:

Purpose

A write to this register performs a priority drop for the specified Group 1 virtual interrupt and, if GICV_CTLR.EOImode == 0, also deactivates the interrupt.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WO	WO	WO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_EOIR1 provides equivalent functionality.
- For AArch64 implementations, ICC_EOIR1_EL1 provides equivalent functionality.

This register is used for Group 1 interrupts only. GICV_EOIR provides equivalent functionality for Group 0 interrupts.

When affinity routing is enabled, it is a programming error to use memory-mapped registers to access the GIC.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICV_AEOIR is a 32-bit register.

Field descriptions

The GICV_AEOIR bit assignments are:

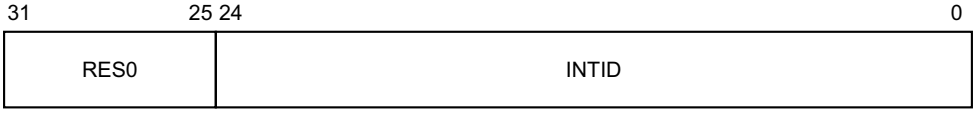

Bits [31:25]

Reserved, RES0.

INTID, bits [24:0]

The INTID of the signaled interrupt.

Note

INTIDs 1020-1023 are reserved and convey additional information such as spurious interrupts.

When affinity routing is not enabled:

- Bits [23:13] are RES0.
- For SGIs, bits [12:10] identify the CPU interface corresponding to the source PE. For all other interrupts these bits are RES0.

A successful EOI request means that:

- The highest priority bit in GICH_APR<n> is cleared, causing the running priority to drop.
- If the appropriate GICV_CTLR.EOImode bit == 0, the interrupt is deactivated in the corresponding List register. If the INTID corresponds to a hardware interrupt, the interrupt is also deactivated in the Distributor.

Note

Only Group 1 interrupts can target the hypervisor, and therefore only Group 1 interrupts are deactivated in the Distributor.

A write to this register is UNPREDICTABLE if INTID corresponds to a Group 0 interrupt. In addition, the following GICv2 UNPREDICTABLE cases require specific actions:

- If highest active priority is Group 0 and the identified interrupt is in the List Registers and it matches the highest active priority. When EL2 is using System registers and ICH_VTR_EL2.SEIS is 1, an IMPLEMENTATION DEFINED SEI might be generated, otherwise GICv3 implementations must ignore such writes.
- If the identified interrupt is in the List Registers, and the HW bit is 1, and the interrupt to be deactivated is an SGI (that is, the value of Physical_ID is between 0 and 15). GICv3 implementations must perform the deactivate operation. This means that a GICv3 implementation in legacy operation must ensure only a single SGI is active for a PE.
- If the identified interrupt is in the List Registers, and the HW bit is 1, and the corresponding pINTID field value is between 1020 and 1023, indicating a special purpose INTID. GICv3 implementations must not perform a deactivate operation but must still change the state of the List register as appropriate. When EL2 is using System registers and ICH_VTR_EL2.SEIS is 1, an implementation might generate a system error.

Accessing the GICV_AEOIR:

GICV_AEOIR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	0x0024

8.15.3 GICV_AHPPIR, Virtual Machine Aliased Highest Priority Pending Interrupt Register

The GICV_AHPPIR characteristics are:

Purpose

Provides the INTID of the highest priority pending Group 1 virtual interrupt in the List registers.

This register corresponds to the physical CPU interface register GICC_AHPPIR.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_HPPIR1 provides equivalent functionality.
- For AArch64 implementations, ICC_HPPIR1_EL1 provides equivalent functionality.

This register is used for Group 1 interrupts only. GICV_HPPIR provides equivalent functionality for Group 0 interrupts.

The register does not return the INTID of an interrupt that is active and pending.

When affinity routing is enabled, it is a programming error to use memory-mapped registers to access the GIC.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICV_AHPPIR is a 32-bit register.

Field descriptions

The GICV_AHPPIR bit assignments are:

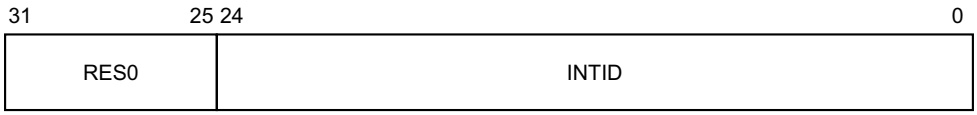

Bits [31:25]

Reserved, RES0.

INTID, bits [24:0]

The INTID of the signaled interrupt.

Note

INTIDs 1020-1023 are reserved and convey additional information such as spurious interrupts.

When affinity routing is not enabled:

- Bits [23:13] are RES0.
- For SGIs, bits [12:10] identify the CPU interface corresponding to the source PE. For all other interrupts these bits are RES0.

A read of this register returns the spurious INTID 1023 if any of the following are true:

- There are no pending interrupts of sufficiently high priority value to be signaled to the PE.
- The highest priority pending interrupt is in Group 0.

Accessing the GICV_AHPPIR:

GICV_AHPPIR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	0x0028

8.15.4 GICV_AIAR, Virtual Machine Aliased Interrupt Acknowledge Register

The GICV_AIAR characteristics are:

Purpose

Provides the INTID of the signaled Group 1 virtual interrupt. A read of this register by the PE acts as an acknowledge for the interrupt.

This register corresponds to the physical CPU interface register GICC_AIAR.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_IAR1 provides equivalent functionality.
- For AArch64 implementations, ICC_IAR1_EL1 provides equivalent functionality.

This register is used for Group 1 interrupts only. GICV_IAR provides equivalent functionality for Group 0 interrupts.

When affinity routing is enabled, it is a programming error to use memory-mapped registers to access the GIC.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICV_AIAR is a 32-bit register.

Field descriptions

The GICV_AIAR bit assignments are:

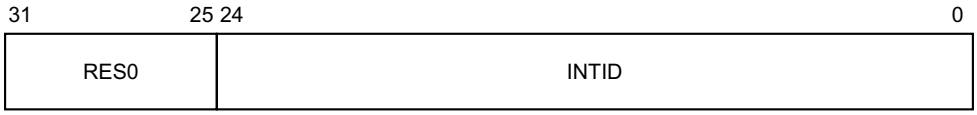

Bits [31:25]

Reserved, RES0.

INTID, bits [24:0]

The INTID of the signaled interrupt.

Note

INTIDs 1020-1023 are reserved and convey additional information such as spurious interrupts.

When affinity routing is not enabled:

- Bits [23:13] are RES0.
- For SGIs, bits [12:10] identify the CPU interface corresponding to the source PE. For all other interrupts these bits are RES0.

The operation of this register is similar to the operation of GICV_IAR. When a vPE reads this register, the corresponding GICH_LR<n>.Group field is checked to determine whether the interrupt is in Group 0 or Group 1:

- If the interrupt is Group 0, the spurious INTID 1023 is returned and the interrupt is not acknowledged.
- If the interrupt is Group 1, the INTID is returned. The List register entry is updated to active state, and the appropriate bit in GICH_APR<n> is set to 1.

A read of this register returns the spurious INTID 1023 if any of the following are true:

- When the virtual CPU interface is enabled and GICH_HCR.En == 1:
  - There are no pending interrupts of sufficiently high priority value to be signaled to the PE.
  - The highest priority pending interrupt is in Group 0.
- Interrupt signaling by the virtual CPU interface is disabled.

Accessing the GICV_AIAR:

GICV_AIAR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	0x0020

8.15.5 GICV_APR<n>, Virtual Machine Active Priorities Registers, n = 0 - 3

The GICV_APR<n> characteristics are:

Purpose

Provides information about interrupt active priorities.

These registers correspond to the physical CPU interface registers GICC_APR<n>.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

If System register access is not enabled for EL2, these registers access GICH_APR<n>. If System register access is enabled for EL2, these registers access ICH_AP1R<n>_EL2. All active priority mapped guests are held in the accessed registers, regardless of interrupt group.

Configurations

When System register access is disabled for EL2, these registers access GICH_APR<n>, and all active priorities for virtual machines are held in GICH_APR<n> regardless of interrupt group.

When System register access is enabled for EL2, these registers access ICH_AP1R<n>_EL2, and all active priorities for virtual machines are held in ICH_AP1R<n>_EL2 regardless of interrupt group.

Attributes

GICV_APR<n> is a 32-bit register.

Field descriptions

The GICV_APR<n> bit assignments are:

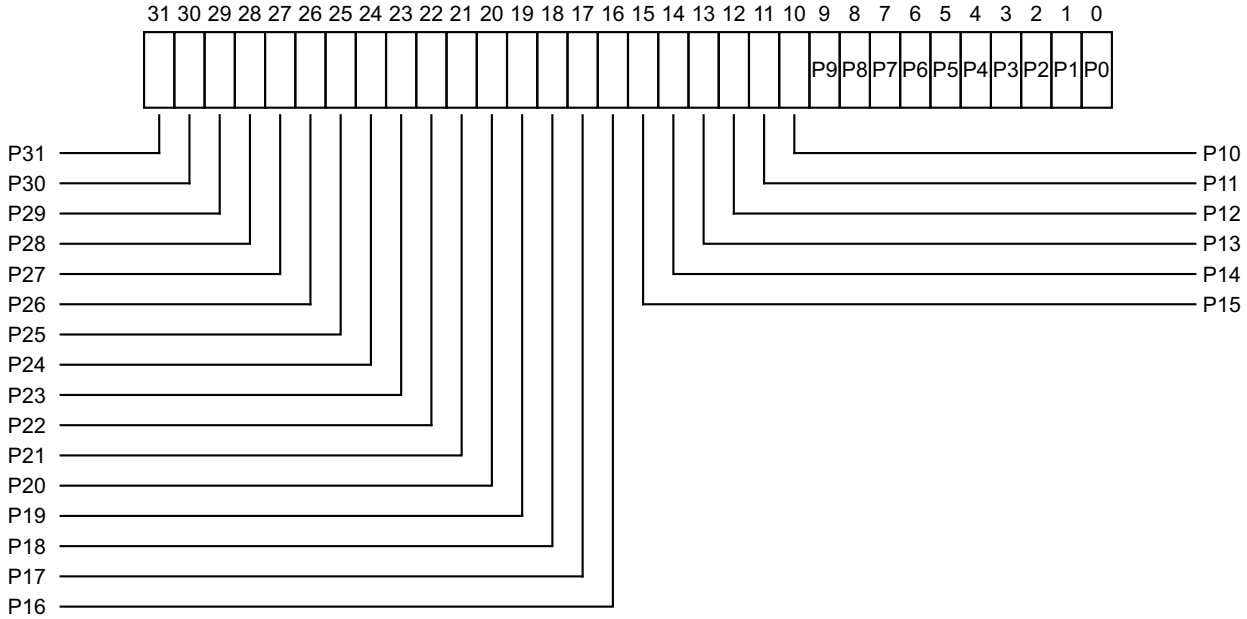

P<x>, bit [x], for x = 0 to 31

Provides information about active priorities for the virtual machine.

See GICH_APR<n> and ICH_AP1R<n>_EL2 for the correspondence between priorities and bits.

Accessing the GICV_APR<n>:

GICV_APR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	0x00D0 + 4n

8.15.6 GICV_BPR, Virtual Machine Binary Point Register

The GICV_BPR characteristics are:

Purpose

Defines the point at which the priority value fields split into two parts, the group priority field and the subpriority field. The group priority field determines Group 0 interrupt preemption.

This register corresponds to GICC_BPR in the physical CPU interface.

Note

GICH_LR<n>.Group determines whether a virtual interrupt is Group 0 or Group 1.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_BPR0 provides equivalent functionality.
- For AArch64 implementations, ICC_BPR0_EL1 provides equivalent functionality.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

When GICV_CTLR.CBPR == 1, this register determines interrupt preemption for both Group 0 and Group 1 interrupts.

Attributes

GICV_BPR is a 32-bit register.

Field descriptions

The GICV_BPR bit assignments are:

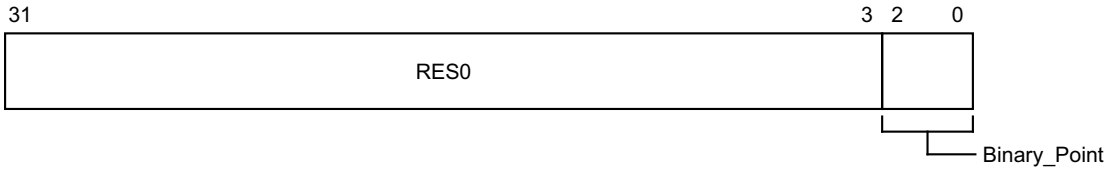

Bits [31:3]

Reserved, RES0.

Binary_Point, bits [2:0]

Controls how the 8-bit interrupt priority field is split into a group priority field, that determines interrupt preemption, and a subpriority field.

For information about how this field determines the interrupt priority bits assigned to the group priority field, see Table 4-9 on page 4-68

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

The Binary_Point field of this register is aliased to GICH_VMCR.VBPR0.

Accessing the GICV_BPR:

GICV_BPR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	0x0008

8.15.7 GICV_CTLR, Virtual Machine Control Register

The GICV_CTLR characteristics are:

Purpose

Controls the behavior of virtual interrupts.

This register corresponds to the physical CPU interface register GICC_CTLR.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_CTLR provides equivalent functionality.
- For AArch64 implementations, ICC_CTLR_EL1 provides equivalent functionality.

Configurations

This register is available when a GIC implementation supports interrupt virtualization.

Attributes

GICV_CTLR is a 32-bit register.

Field descriptions

The GICV_CTLR bit assignments are:

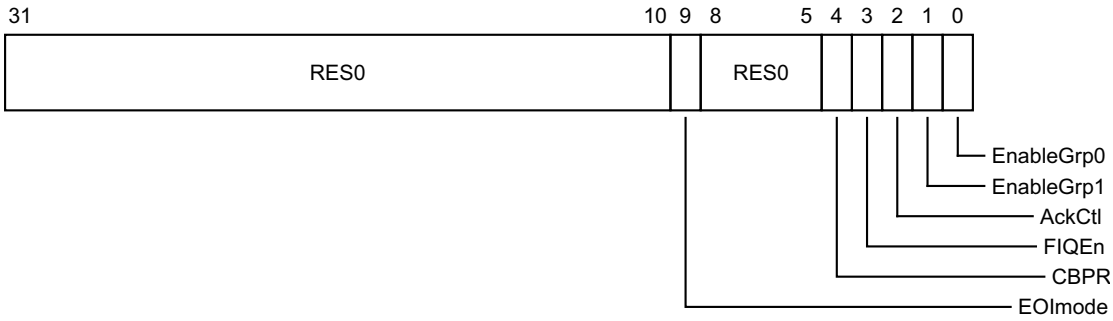

Bits [31:10]

Reserved, RES0.

EOImode, bit [9]

Controls the behavior associated with the GICV_EOIR, GICV_AEOIR, and GICV_DIR registers:

0 Writes to GICV_EOIR and GICV_AEOIR perform priority drop and deactivate interrupt operations simultaneously. Behavior on a write to GICV_DIR is UNPREDICTABLE.

When it has completed processing the interrupt, the virtual machine writes to GICV_EOIR or GICV_AEOIR to deactivate the interrupt. The write updates the List registers and causes the virtual CPU interface to signal the interrupt completion to the physical Distributor.

1 Writes to GICV_EOIR and GICV_AEOIR perform priority drop operation only. Writes to GICV_DIR perform deactivate interrupt operation only.

At some point during interrupt processing, the virtual machine writes to GICV_EOIR or GICV_AEOIR. This write drops the priority of the virtual interrupt by updating its entry in the List registers.

When it has completed processing the interrupt, the virtual machine writes to GICV_DIR to deactivate the interrupt. The write updates the List registers and causes the virtual CPU interface to signal the interrupt completion to the Distributor.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Bits [8:5]

Reserved, RES0.

CBPR, bit [4]

Controls whether GICV_BPR affects both Group 0 and Group 1 interrupts:

0 GICV_BPR affects Group 0 virtual interrupts only. GICV_ABPR affects Group 1 virtual interrupts only.

1 GICV_BPR affects both Group 0 and Group 1 virtual interrupts.

See *Priority grouping* on page 4-67 for more information.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

FIQEn, bit [3]

FIQ Enable. Controls whether Group 0 virtual interrupts are presented as virtual FIQs:

0 Group 0 virtual interrupts are presented as virtual IRQs.

1 Group 0 virtual interrupts are presented as virtual FIQs.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

AckCtl, bit [2]

ARM deprecates use of this bit. ARM strongly recommends that software is written to operate with this bit always cleared to 0.

Acknowledge control. When the highest priority interrupt is Group 1, determines whether GICV_IAR causes the CPU interface to acknowledge the interrupt or returns the spurious identifier 1022, and whether GICV_HPPIR returns the interrupt ID or the special identifier 1022.

0 If the highest priority pending interrupt is Group 1, a read of GICV_IAR or GICV_HPPIR returns an interrupt ID of 1022.

1 If the highest priority pending interrupt is Group 1, a read of GICV_IAR or GICV_HPPIR returns the interrupt ID of the corresponding interrupt.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

EnableGrp1, bit [1]

Enables the signaling of Group 1 virtual interrupts by the virtual CPU interface to the virtual machine:

0 Signaling of Group 1 interrupts is disabled.

1 Signaling of Group 1 interrupts is enabled.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

EnableGrp0, bit [0]

Enables the signaling of Group 0 virtual interrupts by the virtual CPU interface to the virtual machine:

`0` Signaling of Group 0 interrupts is disabled.

`1` Signaling of Group 0 interrupts is enabled.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICV_CTLR:

GICV_CTLR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	`0x0000`

8.15.8 GICV_DIR, Virtual Machine Deactivate Interrupt Register

The GICV_DIR characteristics are:

Purpose

Deactivates a specified virtual interrupt in the GICH_LR<n> List registers.

This register corresponds to the physical CPU interface register GICC_DIR.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WO	WO	WO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_DIR provides equivalent functionality.
- For AArch64 implementations, ICC_DIR_EL1 provides equivalent functionality.

Writes to this register are valid only when GICV_CTLR.EOImode == 1. Writes to this register are otherwise UNPREDICTABLE.

When affinity routing is enabled, it is a programming error to use memory-mapped registers to access the GIC.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICV_DIR is a 32-bit register.

Field descriptions

The GICV_DIR bit assignments are:

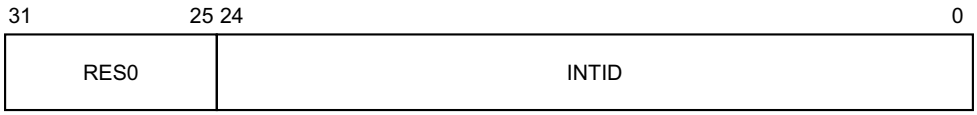

Bits [31:25]

Reserved, RES0.

INTID, bits [24:0]

The INTID of the signaled interrupt.

Note

INTIDs 1020-1023 are reserved and convey additional information such as spurious interrupts.

When affinity routing is not enabled:

- Bits [23:13] are RES0.
- For SGIs, bits [12:10] identify the CPU interface corresponding to the source PE. For all other interrupts these bits are RES0.

When the virtual machine writes to this register, the specified interrupt in the List registers is changed from active to inactive, or from active and pending to pending. If the specified interrupt is present in the List registers but is not in either the active or active and pending states, the effect is UNPREDICTABLE. If the specified interrupt is not present in the List registers, GICH_HCR.EOIcount is incremented, potentially generating a maintenance interrupt.

Note

If the specified interrupt is not present in the List registers, the virtual machine cannot recover the INTID. Therefore, the hypervisor must ensure that, when GICV_CTLR.EOImode == 1, no more than one active interrupt is transferred from the List registers into a software list. If more than one active interrupt that is not stored in the List registers exists, the hypervisor must handle accesses to GICV_DIR in software, typically by trapping these accesses.

If the corresponding GICH_LR<n>.HW == 1, indicating a hardware interrupt, then a deactivate request is sent to the physical Distributor, identifying the physical INTID from the corresponding field in the List register. This effect is identical to a Non-secure write to GICC_DIR from the PE having that physical INTID. This means that if the corresponding physical interrupt is marked as Group 0, the request is ignored.

Note

Interrupt deactivation using this register is based on the provided INTID, with no requirement to deactivate interrupts in any particular order. A single register is therefore used to deactivate both Group 0 and Group 1 interrupts.

Accessing the GICV_DIR:

GICV_DIR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	0x1000

8.15.9 GICV_EOIR, Virtual Machine End Of Interrupt Register

The GICV_EOIR characteristics are:

Purpose

A write to this register performs a priority drop for the specified Group 0 virtual interrupt and, if GICV_CTLR.EOImode == 0, also deactivates the interrupt.

This register corresponds to the physical CPU interface register GICC_EOIR.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WO	WO	WO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_EOIR0 provides equivalent functionality.
- For AArch64 implementations, ICC_EOIR0_EL1 provides equivalent functionality.

This register is used for Group 0 interrupts only. GICV_AEOIR provides equivalent functionality for Group 1 interrupts.

When affinity routing is enabled, it is a programming error to use memory-mapped registers to access the GIC.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICV_EOIR is a 32-bit register.

Field descriptions

The GICV_EOIR bit assignments are:

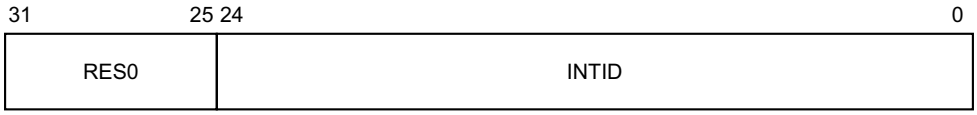

Bits [31:25]

Reserved, RES0.

INTID, bits [24:0]

The INTID of the signaled interrupt.

Note

INTIDs 1020-1023 are reserved and convey additional information such as spurious interrupts.

When affinity routing is not enabled:

- Bits [23:13] are RES0.
- For SGIs, bits [12:10] identify the CPU interface corresponding to the source PE. For all other interrupts these bits are RES0.

The behavior of this register depends on the setting of GICV_CTLR.EOImode:

GICV_CTLR.EOImode	Behavior
0	Both the priority drop and the deactivate interrupt effects occur.
1	Only the priority drop effect occurs.

A successful EOI request means that:

- The highest priority bit in GICH_APR<n> is cleared, causing the running priority to drop.
- If the appropriate GICV_CTLR.EOImode bit == 0, the interrupt is deactivated in the corresponding List register, GICH_LR<n>. If GICH_LR<n>.HW == 1, indicating the INTID corresponds to a hardware interrupt, a deactivate request is also sent to the physical Distributor, identifying the physical INTID from the corresponding field in the List register. This effect is identical to a Non-secure write to GICC_DIR from the PE having that physical INTID. This means that if the corresponding physical interrupt is marked as Group 0, and GICD_CTLR.DS == 0, the deactivation request is ignored. See GICC_EOIR for more information.

Note

Only Group 1 interrupts can target the hypervisor, and therefore only Group 1 interrupts are deactivated in the Distributor.

Accessing the GICV_EOIR:

GICV_EOIR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	0x0010

8.15.10 GICV_HPPIR, Virtual Machine Highest Priority Pending Interrupt Register

The GICV_HPPIR characteristics are:

Purpose

Provides the INTID of the highest priority pending Group 0 virtual interrupt in the List registers.

This register corresponds to the physical CPU interface register GICC_HPPIR.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_HPPIR0 provides equivalent functionality.
- For AArch64 implementations, ICC_HPPIR0_EL1 provides equivalent functionality.

This register is used for Group 0 interrupts only. GICV_AHPPIR provides equivalent functionality for Group 1 interrupts.

When affinity routing is enabled, it is a programming error to use memory-mapped registers to access the GIC.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICV_HPPIR is a 32-bit register.

Field descriptions

The GICV_HPPIR bit assignments are:

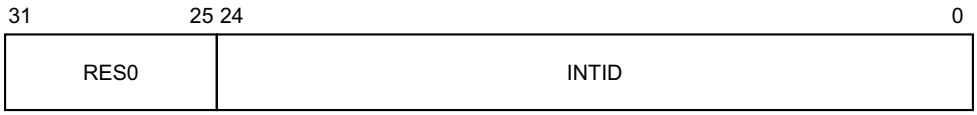

Bits [31:25]

Reserved, RES0.

INTID, bits [24:0]

The INTID of the signaled interrupt.

Note

INTIDs 1020-1023 are reserved and convey additional information such as spurious interrupts.

When affinity routing is not enabled:

- Bits [23:13] are RES0.
- For SGIs, bits [12:10] identify the CPU interface corresponding to the source PE. For all other interrupts these bits are RES0.

Reads of the GICC_HPPIR that do not return a valid INTID return a spurious INTID, 1022 or 1023, see *Special INTIDs* on page 3-40.

Highest priority pending interrupt Group	GICV_HPPIR read	GICV_CTLR.AckCtl	Returned INTID
1	Non-secure	x	ID of Group 1 interrupt
1	Secure	0	1022
1	Secure	1	ID of Group 1 interrupt
0	Non-secure	x	1023
0	Secure	x	ID of Group 0 interrupt
No pending interrupts	x	x	1023

If the CPU interface supports only a single Security state, the entries that apply to Secure reads describe the behavior.

Accessing the GICV_HPPIR:

GICV_HPPIR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	`0x0018`

8.15.11 GICV_IAR, Virtual Machine Interrupt Acknowledge Register

The GICV_IAR characteristics are:

Purpose

Provides the INTID of the signaled Group 0 virtual interrupt. A read of this register by the PE acts as an acknowledge for the interrupt.

This register corresponds to the physical CPU interface register GICC_IAR.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_IAR0 provides equivalent functionality.
- For AArch64 implementations, ICC_IAR0_EL1 provides equivalent functionality.

This register is used for Group 0 interrupts only. GICV_AIAR provides equivalent functionality for Group 1 interrupts.

When affinity routing is enabled, it is a programming error to use memory-mapped registers to access the GIC.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICV_IAR is a 32-bit register.

Field descriptions

The GICV_IAR bit assignments are:

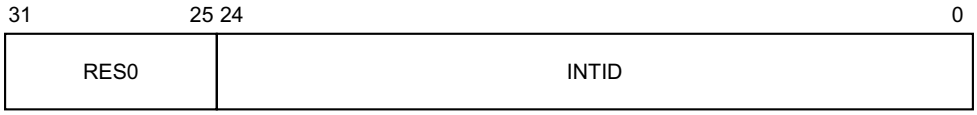

Bits [31:25]

Reserved, RES0.

INTID, bits [24:0]

The INTID of the signaled interrupt.

Note

INTIDs 1020-1023 are reserved and convey additional information such as spurious interrupts.

When affinity routing is not enabled:

- Bits [23:13] are RES0.
- For SGIs, bits [12:10] identify the CPU interface corresponding to the source PE. For all other interrupts these bits are RES0.

When the virtual machine writes to this register, the virtual CPU interface acknowledges the highest priority pending virtual interrupt and sets the state in the corresponding List register to active. The appropriate bit in the active priorities register GICH_APR<n> is set to 1.

If GICH_LR<n>.HW == 0, indicating that the interrupt is software-triggered, then bits [12:10] of GICH_LR<n> are returned in bits [12:10] of GICV_IAR. Otherwise bits [12:10] are RES0.

A read of this register returns the spurious INTID 1023 if either of the following is true:

- There are no pending interrupts of sufficiently high priority value to be signaled to the PE with the virtual CPU interface enabled and GICH_HCR.En == 1.
- Interrupt signaling by the virtual CPU interface is disabled.

A read of this register returns the spurious INTID 1022 if the highest priority pending interrupt is Group 1 and GICV_CTLR.AckCtl == 0.

Accessing the GICV_IAR:

GICV_IAR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	0x000C

8.15.12 GICV_IIDR, Virtual Machine CPU Interface Identification Register

The GICV_IIDR characteristics are:

Purpose

Provides information about the implementer and revision of the virtual CPU interface.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

Configurations

This register is available in all configurations of the GIC. If the GIC implementation supports two Security states this register is Common.

Attributes

GICV_IIDR is a 32-bit register.

Field descriptions

The GICV_IIDR bit assignments are:

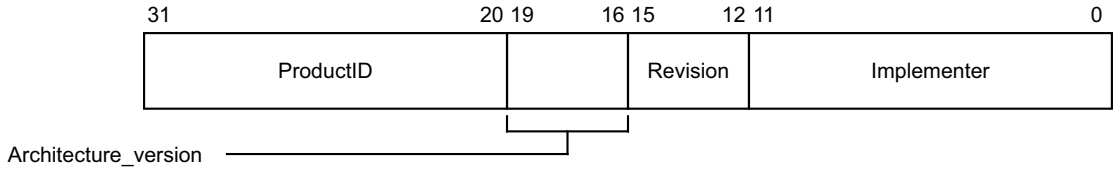

ProductID, bits [31:20]

An IMPLEMENTATION DEFINED product identifier.

Architecture_version, bits [19:16]

The version of the GIC architecture that is implemented.

`0001` GICv1.

`0010` GICv2.

`0011` GICv3.

`0100` GICv4 memory-mapped interface supported. Support for the System register interface is discoverable from PE registers ID_PFR1 and ID_AA64PFR0_EL1.

Other values are reserved.

Revision, bits [15:12]

An IMPLEMENTATION DEFINED revision number for the CPU interface.

Implementer, bits [11:0]

Contains the JEP106 code of the company that implemented the CPU interface.

- Bits [11:8] are the JEP106 continuation code of the implementer. For an ARM implementation, this field is `0x4`.
- Bit [7] is always 0.

- Bits [6:0] are the JEP106 identity code of the implementer. For an ARM implementation, bits [7:0] are therefore `0x3B`.

Accessing the GICV_IIDR:

GICV_IIDR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	`0x00FC`

8.15.13 GICV_PMR, Virtual Machine Priority Mask Register

The GICV_PMR characteristics are:

Purpose

This register provides a virtual interrupt priority filter. Only virtual interrupts with higher priority than the value in this register are signaled to the PE.

Note

Higher interrupt priority corresponds to a lower value of the Priority field.

This register corresponds to the physical CPU interface register GICC_PMR.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_PMR provides equivalent functionality.
- For AArch64 implementations, ICC_PMR_EL1 provides equivalent functionality.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

The Priority field of this register is aliased to GICH_VMCR.VPMR, to enable state to be switched easily between virtual machines during context-switching.

Attributes

GICV_PMR is a 32-bit register.

Field descriptions

The GICV_PMR bit assignments are:

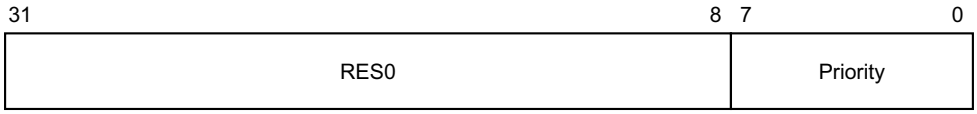

Bits [31:8]

Reserved, RES0.

Priority, bits [7:0]

The priority mask level for the virtual CPU interface. If the priority of the interrupt is higher than the value indicated by this field, the interface signals the interrupt to the PE.

If the GIC implementation supports fewer than 256 priority levels some bits might be RAZ/WI, as follows:

- For 128 supported levels, bit [0] = `0b0`.
- For 64 supported levels, bits [1:0] = `0b00`.
- For 32 supported levels, bits [2:0] = `0b000`.

- For 16 supported levels, bits [3:0] = `0b0000`.

See *Interrupt prioritization* on page 4-65 for more information.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICV_PMR:

GICV_PMR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	`0x0004`

8.15.14 GICV_RPR, Virtual Machine Running Priority Register

The GICV_RPR characteristics are:

Purpose

This register indicates the running priority of the virtual CPU interface.

This register corresponds to the physical CPU interface register GICC_RPR.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICC_RPR provides equivalent functionality.
- For AArch64 implementations, ICC_RPR_EL1 provides equivalent functionality.

Depending on the implementation, if no bits are set to 1 in GICH_APR<n>, indicating no active virtual interrupts in the virtual CPU interface, the priority reads as `0xFF` or `0xF8` to reflect the number of supported interrupt priority bits defined by GICH_VTR.PRIbits.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICV_RPR is a 32-bit register.

Field descriptions

The GICV_RPR bit assignments are:

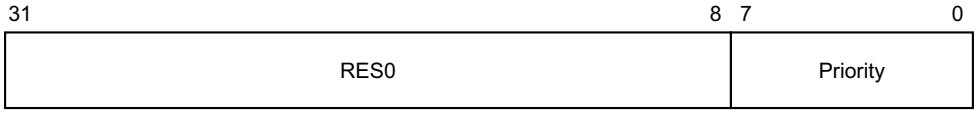

Bits [31:8]

Reserved, RES0.

Priority, bits [7:0]

The current running priority on the virtual CPU interface.

Accessing the GICV_RPR:

GICV_RPR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	`0x0014`

8.15.15 GICV_STATUSR, Virtual Machine Error Reporting Status Register

The GICV_STATUSR characteristics are:

Purpose

Provides software with a mechanism to detect:

- Accesses to reserved locations.
- Writes to read-only locations.
- Reads of write-only locations.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

This is an optional register. If the register is implemented, GICC_STATUSR must also be implemented. If the register is not implemented, the location is RAZ/WI.

This register is used only when System register access is not enabled. If System register access is enabled, this register is not updated. Equivalent function might be provided by appropriate traps and exceptions.

Configurations

In systems where this register is implemented, ARM expects that when a virtual machine is scheduled, the hypervisor ensures that this register is cleared to 0. The hypervisor might check for illegal accesses when the virtual machine is unscheduled.

Attributes

GICV_STATUSR is a 32-bit register.

Field descriptions

The GICV_STATUSR bit assignments are:

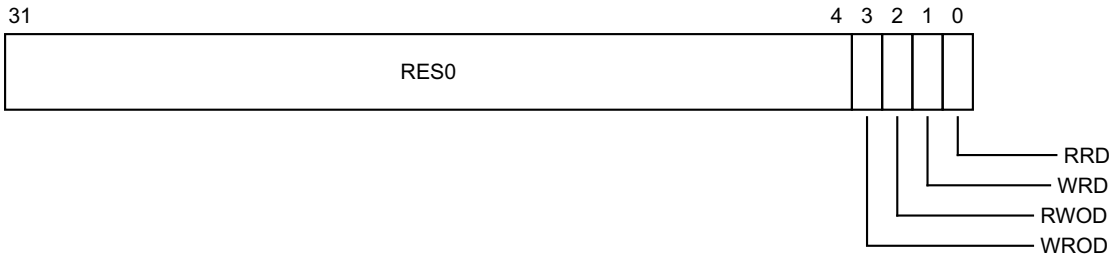

Bits [31:4]

Reserved, RES0.

WROD, bit [3]

Write to an RO location.

0 Normal operation.

1 A write to an RO location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

RWOD, bit [2]

Read of a WO location.

0 Normal operation.

1 A read of a WO location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

WRD, bit [1]

Write to a reserved location.

0 Normal operation.

1 A write to a reserved location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

RRD, bit [0]

Read of a reserved location.

0 Normal operation.

1 A read of a reserved location has been detected.

When a violation is detected, software must write 1 to this register to reset it.

Accessing the GICV_STATUSR:

GICV_STATUSR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual CPU interface	0x002C

8.16 The GIC virtual interface control register map

The GIC virtual interface control registers are management registers. Configuration software on the PE must ensure they are accessible only by a hypervisor, or similar software.

Note

Reserved register addresses are RAZ/WI.

Table 8-32 shows the register map for the GIC virtual interface control registers.

Table 8-32 GIC virtual interface control register map

Offset	Name	Type	Reset	Description
0x0000	GICH_HCR	RW	0x00000000	Hypervisor Control Register
0x0004	GICH_VTR	RO	IMPLEMENTATION DEFINED	VGIC Type Register
0x0008	GICH_VMCR	RW	-	Virtual Machine Control Register
0x000C	-	-	-	Reserved
0x0010	GICH_MISR	RO	0x00000000	Maintenance Interrupt Status Register
0x0014-0x001C	-	-	-	Reserved
0x0020	GICH_EISR	RO	0x00000000	End of Interrupt Status Register
0x0024-0x002C	-	-	-	Reserved
0x0030	GICH_ELRSR	RO	IMPLEMENTATION DEFINED[a]	Empty List Register Status Register
0x0034-0x00EC	-	-	-	Reserved
0x00F0-0x00FC	GICH_APR<n>	RW	0x00000000	Active Priorities Register
0x0100-0x013C	GICH_LR<n>	RW	0x00000000	List Registers 0-15 lower bits

a. Each bit that has a corresponding List register resets to 1, meaning that the reset value of the register depends on the number of List registers implemented.

Note

It is IMPLEMENTATION DEFINED whether an access to a GIC virtual interface control register using the memory-mapped interface accesses the same state as an access using the System register interface, or whether the two interfaces access different states.

8.17 The GIC virtual interface control register descriptions

This section describes each of the GIC virtual interface control registers in register name order.

8.17.1 GICH_APR<n>, Active Priorities Registers, n = 0 - 3

The GICH_APR<n> characteristics are:

Purpose

These registers track which preemption levels are active in the virtual CPU interface, and indicate the current active priority. Corresponding bits are set to 1 in this register when an interrupt is acknowledged, based on GICH_LR<n>.Priority, and the least significant bit set is cleared on EOI.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

These registers are used only when System register access is not enabled. When System register access is enabled the following registers provide equivalent functionality:

- In AArch64:
  - For Group 0, ICH_AP0R<n>_EL2.
  - For Group 1, ICH_AP1R<n>_EL2.
- In AArch32:
  - For Group 0, ICH_AP0R<n>.
  - For Group 1, ICH_AP1R<n>.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

The number of registers required depends on how many bits are implemented in GICH_LR<n>.Priority:

- When 5 priority bits are implemented, 1 register is required (GICH_APR0).
- When 6 priority bits are implemented, 2 registers are required (GICH_APR0, GICH_APR1).
- When 7 priority bits are implemented, 4 registers are required (GICH_APR0, GICH_APR1, GICH_APR2, GICH_APR3).

Unimplemented registers are RAZ/WI.

Attributes

GICH_APR<n> is a 32-bit register.

Field descriptions

The GICH_APR<n> bit assignments are:

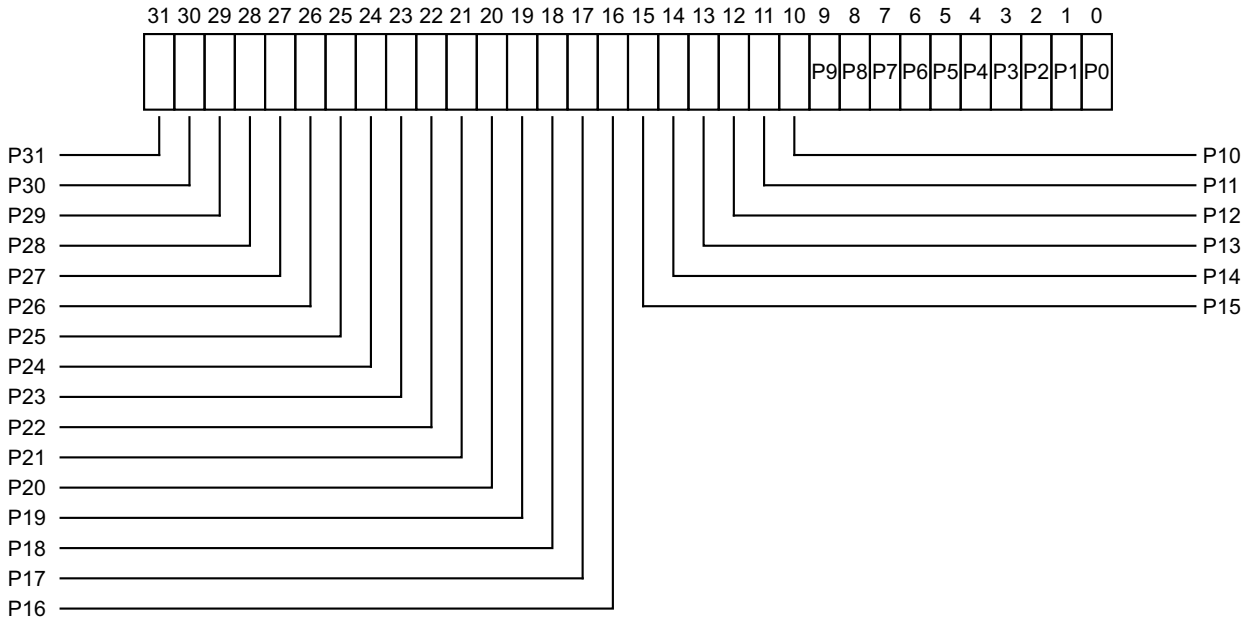

P<x>, bit [x], for x = 0 to 31

Active priorities. Possible values of each bit are:

`0` There is no interrupt active at the priority corresponding to that bit.

`1` There is an interrupt active at the priority corresponding to that bit.

The correspondence between priorities and bits depends on the number of bits of priority that are implemented.

If 5 bits of priority are implemented (bits [7:3] of priority), then there are 32 priority groups, and the active state of these priorities are held in GICH_APR0 in the bits corresponding to Priority[7:3].

If 6 bits of priority are implemented (bits [7:2] of priority), then there are 64 priority groups, and:

- The active state of priorities 0 - 124 are held in GICH_APR0 in the bits corresponding to 0:Priority[6:2].
- The active state of priorities 128 - 252 are held in GICH_APR1 in the bits corresponding to 1:Priority[6:2].

If 7 bits of priority are implemented (bits [7:1] of priority), then there are 128 priority groups, and:

- The active state of priorities 0 - 62 are held in GICH_APR0 in the bits corresponding to 00:Priority[5:1].
- The active state of priorities 64 - 126 are held in GICH_APR1 in the bits corresponding to 01:Priority[5:1].
- The active state of priorities 128 - 190 are held in GICH_APR2 in the bits corresponding to 10:Priority[5:1].
- The active state of priorities 192 - 254 are held in GICH_APR3 in the bits corresponding to 11:Priority[5:1].

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the GICH_APR<n>:

GICH_APR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual interface control	`0x00F0` + 4n

8.17.2 GICH_EISR, End Interrupt Status Register

The GICH_EISR characteristics are:

Purpose

Indicates which List registers have outstanding EOI maintenance interrupts.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICH_EISR provides equivalent functionality.
- For AArch64 implementations, ICH_EISR_EL2 provides equivalent functionality.

Bits corresponding to unimplemented List registers are RAZ.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICH_EISR is a 32-bit register.

Field descriptions

The GICH_EISR bit assignments are:

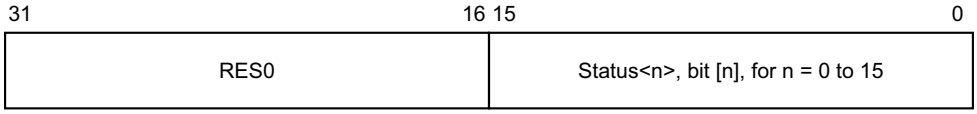

Bits [31:16]

Reserved, RES0.

Status<n>, bit [n], for n = 0 to 15

EOI maintenance interrupt status for List register <n>:

`0` GICH_LR<n> does not have an EOI maintenance interrupt.

`1` GICH_LR<n> has an EOI maintenance interrupt that has not been handled.

For any GICH_LR<n> register, the corresponding status bit is set to 1 if all of the following are true:

- GICH_LR<n>.State is `0b00`.
- GICH_LR<n>.HW == 0.
- GICH_LR<n>.EOI == 1.

See GICH_LR<n> for more information.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICH_EISR:

GICH_EISR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual interface control	0x0020

8.17.3 GICH_ELRSR, Empty List Register Status Register

The GICH_ELRSR characteristics are:

Purpose

Indicates which List registers contain valid interrupts.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICH_ELRSR provides equivalent functionality.
- For AArch64 implementations, ICH_ELRSR_EL2 provides equivalent functionality.

Bits corresponding to unimplemented List registers are RES0.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICH_ELRSR is a 32-bit register.

Field descriptions

The GICH_ELRSR bit assignments are:

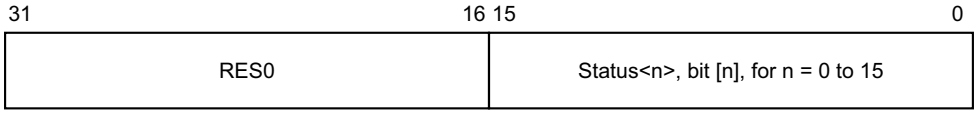

Bits [31:16]

Reserved, RES0.

Status<n>, bit [n], for n = 0 to 15

Status bit for List register <n>:

`0` GICH_LR<n>, if implemented, contains a valid interrupt. Using this List register can result in overwriting a valid interrupt.

`1` GICH_LR<n> does not contain a valid interrupt. The List register is empty and can be used without overwriting a valid interrupt or losing an EOI maintenance interrupt.

For any GICH_LR<n> register, the corresponding status bit is set to 1 if GICH_LR<n>.State is `0b00` and either:

- GICH_LR<n>.HW == 1.
- GICH_LR<n>.EOI == 0.

When this register has an architecturally-defined reset value, this field resets to `1`.

Accessing the GICH_ELRSR:

GICH_ELRSR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual interface control	0x0030

8.17.4 GICH_HCR, Hypervisor Control Register

The GICH_HCR characteristics are:

Purpose

Controls the virtual CPU interface.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICH_HCR provides equivalent functionality.
- For AArch64 implementations, ICH_HCR_EL2 provides equivalent functionality.

GICH_HCR.En must be set to 1 for any virtual or maintenance interrupt to be asserted.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICH_HCR is a 32-bit register.

Field descriptions

The GICH_HCR bit assignments are:

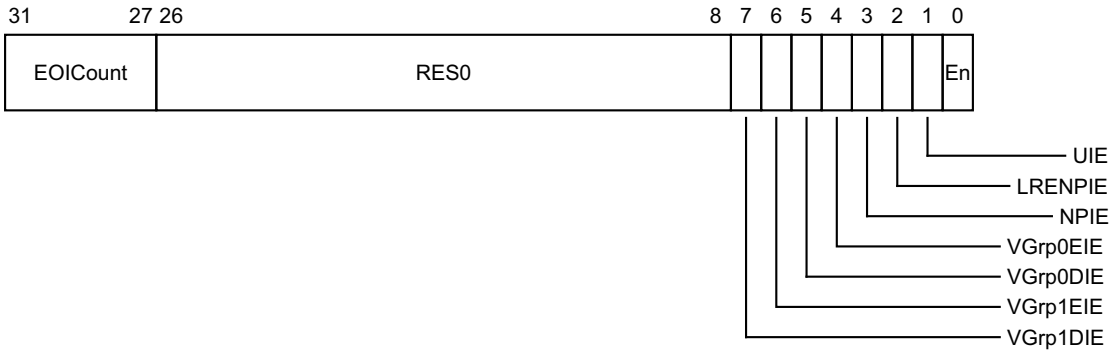

EOICount, bits [31:27]

Counts the number of EOIs received that do not have a corresponding entry in the List registers. The virtual CPU interface increments this field automatically when a matching EOI is received. EOIs that do not clear a bit in GICH_APR<n> do not cause an increment. If an EOI occurs when the value of this field is 31, then the field wraps to 0.

The maintenance interrupt is asserted whenever this field is nonzero and GICH_HCR.LRENPIE == 1.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Bits [26:8]

Reserved, RES0.

VGrp1DIE, bit [7]

VM Group 1 Disabled Interrupt Enable.

Enables the signaling of a maintenance interrupt while signaling of Group 1 interrupts from the virtual CPU interface to the connected virtual machine is disabled:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled when GICV_CTLR.EnableGrp1 == 0.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

VGrp1EIE, bit [6]

VM Group 1 Enabled Interrupt Enable.

Enables the signaling of a maintenance interrupt while signaling of Group 1 interrupts from the virtual CPU interface to the connected virtual machine is enabled:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled when GICV_CTLR.EnableGrp1 == 1.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

VGrp0DIE, bit [5]

VM Group 0 Disabled Interrupt Enable.

Enables the signaling of a maintenance interrupt while signaling of Group 0 interrupts from the virtual CPU interface to the connected virtual machine is disabled:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled when GICV_CTLR.EnableGrp0 == 0.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

VGrp0EIE, bit [4]

VM Group 0 Enabled Interrupt Enable.

Enables the signaling of a maintenance interrupt while signaling of Group 0 interrupts from the virtual CPU interface to the connected virtual machine is enabled:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled when GICV_CTLR.EnableGrp0 == 1.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

NPIE, bit [3]

No Pending Interrupt Enable.

Enables the signaling of a maintenance interrupt while no pending interrupts are present in the List registers:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled while the List registers contain no interrupts in the pending state.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

LRENPIE, bit [2]

List Register Entry Not Present Interrupt Enable.

Enables the signaling of a maintenance interrupt while the virtual CPU interface does not have a corresponding valid List register for an EOI request:

0 Maintenance interrupt disabled.

1 Maintenance interrupt signaled while GICH_HCR.EOICount is not 0.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

UIE, bit [1]

Underflow Interrupt Enable.

Enables the signaling of a maintenance interrupt when the List registers are either empty or hold only one valid entry.

0 Maintenance interrupt disabled.

1 A maintenance interrupt is signaled if zero or one of the List register entries are marked as a valid interrupt.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

En, bit [0]

Enable.

Global enable bit for the virtual CPU interface.

0 Virtual CPU interface operation is disabled.

1 Virtual CPU interface operation is enabled.

When this field is 0:

- The virtual CPU interface does not signal any maintenance interrupts.
- The virtual CPU interface does not signal any virtual interrupts.
- A read of GICV_IAR or GICV_AIAR returns a spurious interrupt ID.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

The VGrp1DIE, VGrp1EIE, VGrp0DIE, and VGrp0EIE fields permit the hypervisor to track the virtual CPU interfaces that are enabled. The hypervisor can then route interrupts that have multiple targets correctly and efficiently, without having to read the virtual CPU interface status.

See *Maintenance interrupts* on page 5-85 and GICH_MISR for more information.

Accessing the GICH_HCR:

GICH_HCR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual interface control	0x0000

8.17.5 GICH_LR<n>, List Registers, n = 0 - 15

The GICH_LR<n> characteristics are:

Purpose

These registers provide context information for the virtual CPU interface.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICH_LR<n> provides equivalent functionality.
- For AArch64 implementations, ICH_LR<n>_EL2 provides equivalent functionality.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

A maximum of 16 List registers can be provided. GICH_VTR.ListRegs defines the number implemented. Unimplemented List registers are RAZ/WI.

Attributes

GICH_LR<n> is a 32-bit register.

Field descriptions

The GICH_LR<n> bit assignments are:

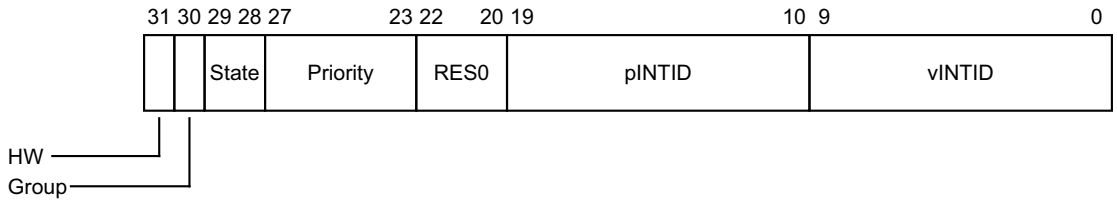

HW, bit [31]

Indicates whether this virtual interrupt is a hardware interrupt, meaning that it corresponds to a physical interrupt. Deactivation of the virtual interrupt also causes the deactivation of the physical interrupt corresponding to the INTID:

`0` This interrupt is triggered entirely in software. No notification is sent to the Distributor when the virtual interrupt is deactivated.

`1` A hardware interrupt. A deactivate interrupt request is sent to the Distributor when the virtual interrupt is deactivated, using GICH_LR<n>.pINTID to indicate the physical interrupt identifier.

If GICV_CTLR.EOImode == 0, this request corresponds to a write to GICV_EOIR or GICV_AEOIR, otherwise it corresponds to a write to GICV_DIR.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Group, bit [30]

Indicates whether the interrupt is Group 0 or Group 1:

`0` Group 0 virtual interrupt. GICV_CTLR.FIQEn determines whether it is signaled as a virtual IRQ or as a virtual FIQ, and GICV_CTLR.EnableGrp0 enables signaling of this interrupt to the virtual machine.

`1` Group 1 virtual interrupt, signaled as a virtual IRQ. GICV_CTLR.EnableGrp1 enables signaling of this interrupt to the virtual machine.

Note

GICV_CTLR.CBPR controls whether GICV_BPR or GICV_ABPR determines if a pending Group 1 interrupt has sufficient priority to preempt current execution.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

State, bits [29:28]

The state of the interrupt. This field has one of the following values:

`00` Inactive

`01` Pending

`10` Active

`11` Active and pending

The GIC updates these state bits as virtual interrupts proceed through the interrupt life cycle. Entries in the inactive state are ignored, except for the purpose of generating virtual maintenance interrupts.

Note

For hardware interrupts, the active and pending state is held in the Distributor rather than the virtual CPU interface. A hypervisor must only use the active and pending state for software originated interrupts, which are typically associated with virtual devices, or for SGIs.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Priority, bits [27:23]

The priority of this interrupt.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Bits [22:20]

Reserved, RES0.

pINTID, bits [19:10]

The function of this field depends on the value of GICH_LR<n>.HW.

When GICH_LR<n>.HW == 0:

- Bit [19] indicates whether the interrupt triggers an EOI maintenance interrupt. If this bit is 1, then when the interrupt identified by vINTID is deactivated, an EOI maintenance interrupt is asserted.
- Bits [18:13] are reserved, SBZ.
- If the vINTID field value corresponds to an SGI (that is, 0-15), bits [12:10] contain the number of the requesting PE. This appears in the corresponding field of GICV_IAR or GICV_AIAR. If the vINTID field value is not 0-15, this field must be cleared to 0.

When GICH_LR<n>.HW == 1:

- This field indicates the pINTID that the hypervisor forwards to the Distributor. This field is only required to implement enough bits to hold a valid value for the ID configuration. Any unused higher order bits are RAZ/WI.
- If the value of pINTID is 0-15 or 1020-1023, behavior is UNPREDICTABLE. If the value of pINTID is 16-31, this field applies to the PPI associated with this same PE as the virtual CPU interface requesting the deactivation.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

vINTID, bits [9:0]

This INTID is returned to the VM when the interrupt is acknowledged through GICV_IAR. Each valid interrupt stored in the List registers must have a unique vINTID for that virtual CPU interface. If the value of vINTID is 1020-1023, behavior is UNPREDICTABLE.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Accessing the GICH_LR<n>:

GICH_LR<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual interface control	`0x0100` + 4n

8.17.6 GICH_MISR, Maintenance Interrupt Status Register

The GICH_MISR characteristics are:

Purpose

Indicates which maintenance interrupts are asserted.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICH_MISR provides equivalent functionality.
- For AArch64 implementations, ICH_MISR_EL2 provides equivalent functionality.

A maintenance interrupt is asserted only if at least one bit is set to 1 in this register and if GICH_HCR.En == 1.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICH_MISR is a 32-bit register.

Field descriptions

The GICH_MISR bit assignments are:

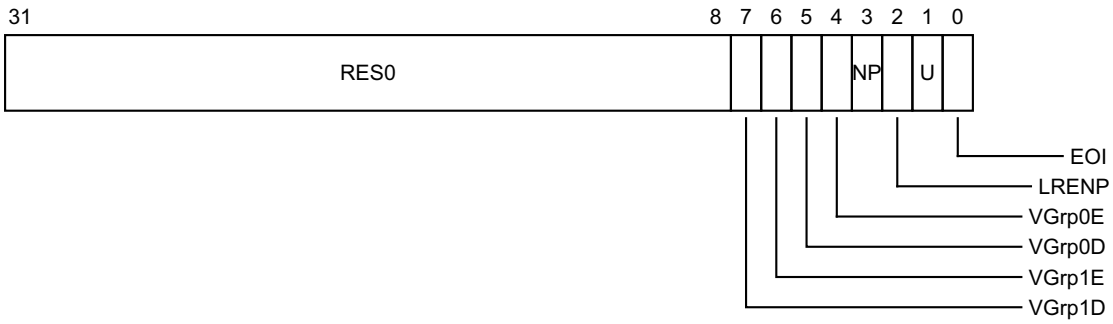

Bits [31:8]

Reserved, RES0.

VGrp1D, bit [7]

vPE Group 1 Disabled.

`0` vPE Group 1 Disabled maintenance interrupt not asserted.

`1` vPE Group 1 Disabled maintenance interrupt asserted.

This maintenance interrupt is asserted when GICH_HCR.VGrp1DIE == 1 and GICH_VMCR.VENG1 == 0.

When this register has an architecturally-defined reset value, this field resets to `0`.

VGrp1E, bit [6]

vPE Group 1 Enabled.

0 vPE Group 1 Enabled maintenance interrupt not asserted.

1 vPE Group 1 Enabled maintenance interrupt asserted.

This maintenance interrupt is asserted when GICH_HCR.VGrp1EIE == 1 and GICH_VMCR.VENG1 == 1.

When this register has an architecturally-defined reset value, this field resets to 0.

VGrp0D, bit [5]

vPE Group 0 Disabled.

0 vPE Group 0 Disabled maintenance interrupt not asserted.

1 vPE Group 0 Disabled maintenance interrupt asserted.

This maintenance interrupt is asserted when GICH_HCR.VGrp0DIE == 1 and GICH_VMCR.VENG0 == 0.

When this register has an architecturally-defined reset value, this field resets to 0.

VGrp0E, bit [4]

vPE Group 0 Enabled.

0 vPE Group 0 Enabled maintenance interrupt not asserted.

1 vPE Group 0 Enabled maintenance interrupt asserted.

This maintenance interrupt is asserted when GICH_HCR.VGrp0EIE == 1 and GICH_VMCR.VENG0 == 1.

When this register has an architecturally-defined reset value, this field resets to 0.

NP, bit [3]

No Pending.

0 No Pending maintenance interrupt not asserted.

1 No Pending maintenance interrupt asserted.

This maintenance interrupt is asserted when GICH_HCR.NPIE == 1 and no List register is in the pending state.

When this register has an architecturally-defined reset value, this field resets to 0.

LRENP, bit [2]

List Register Entry Not Present.

0 List Register Entry Not Present maintenance interrupt not asserted.

1 List Register Entry Not Present maintenance interrupt asserted.

This maintenance interrupt is asserted when GICH_HCR.LRENPIE == 1 and GICH_HCR.EOICount is nonzero.

When this register has an architecturally-defined reset value, this field resets to 0.

U, bit [1]

Underflow.

0 Underflow maintenance interrupt not asserted.

1 Underflow maintenance interrupt asserted.

This maintenance interrupt is asserted when GICH_HCR.UIE == 1 and zero or one of the List register entries are marked as a valid interrupt.

When this register has an architecturally-defined reset value, this field resets to 0.

EOI, bit [0]

End Of Interrupt.

`0` End Of Interrupt maintenance interrupt not asserted.

`1` End Of Interrupt maintenance interrupt asserted.

This maintenance interrupt is asserted when at least one bit in GICH_EISR == 1.

When this register has an architecturally-defined reset value, this field resets to `0`.

Note

A List register is in the pending state only if the corresponding GICH_LR<n> value is `01`, that is, pending. The active and pending state is not included.

Accessing the GICH_MISR:

GICH_MISR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual interface control	`0x0010`

8.17.7 GICH_VMCR, Virtual Machine Control Register

The GICH_VMCR characteristics are:

Purpose

Enables the hypervisor to save and restore the virtual machine view of the GIC state. This register is updated when a virtual machine updates the virtual CPU interface registers.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICH_VMCR provides equivalent functionality.
- For AArch64 implementations, ICH_VMCR_EL2 provides equivalent functionality.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICH_VMCR is a 32-bit register.

Field descriptions

The GICH_VMCR bit assignments are:

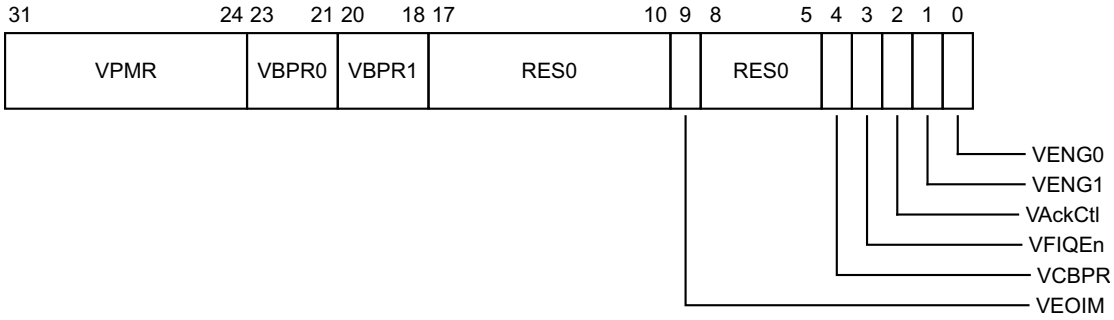

VPMR, bits [31:24]

Virtual priority mask. The priority mask level for the CPU interface. If the priority of an interrupt is higher than the value indicated by this field, the interface signals the interrupt to the PE.

This alias field is updated when a VM updates GICV_PMR.Priority.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

VBPR0, bits [23:21]

Virtual Binary Point Register, Group 0. Defines the point at which the priority value fields split into two parts, the Group priority field and the subpriority field. The Group priority field determines Group 0 interrupt preemption, and also determines Group 1 interrupt preemption if GICH_VMCR.VCBPR == 1.

This alias field is updated when a VM updates GICV_BPR.Binary_Point.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

VBPR1, bits [20:18]

Virtual Binary Point Register, Group 1. Defines the point at which the priority value fields split into two parts, the Group priority field and the subpriority field. The Group priority field determines Group 1 interrupt preemption if GICH_VMCR.VCBPR == 0.

This alias field is updated when a VM updates GICV_ABPR.Binary_Point.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Bits [17:10]

Reserved, RES0.

VEOIM, bit [9]

Virtual EOImode. Possible values of this bit are:

`0` A write of an INTID to GICV_EOIR or GICV_AEOIR drops the priority of the interrupt with that INTID, and also deactivates that interrupt.

`1` A write of an INTID to GICV_EOIR or GICV_AEOIR only drops the priority of the interrupt with that INTID. Software must write to GICV_DIR to deactivate the interrupt.

This alias field is updated when a VM updates GICV_CTLR.EOImode.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

Bits [8:5]

Reserved, RES0.

VCBPR, bit [4]

Virtual Common Binary Point Register. Possible values of this bit are:

`0` GICV_ABPR determines the preemption group for Group 1 interrupts.

`1` GICV_BPR determines the preemption group for Group 1 interrupts.

This alias field is updated when a VM updates GICV_CTLR.CBPR.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

VFIQEn, bit [3]

Virtual FIQ enable. Possible values of this bit are:

`0` Group 0 virtual interrupts are presented as virtual IRQs.

`1` Group 0 virtual interrupts are presented as virtual FIQs.

This alias field is updated when a VM updates GICV_CTLR.FIQEn.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

VAckCtl, bit [2]

Virtual AckCtl. Possible values of this bit are:

`0` If the highest priority pending interrupt is Group 1, a read of GICV_IAR or GICV_HPPIR returns an INTID of 1022.

`1` If the highest priority pending interrupt is Group 1, a read of GICV_IAR or GICV_HPPIR returns the INTID of the corresponding interrupt.

This alias field is updated when a VM updates GICV_CTLR.AckCtl.

This field is supported for backwards compatibility with GICv2. ARM deprecates the use of this field.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

VENG1, bit [1]

Virtual interrupt enable, Group 1. Possible values of this bit are:

`0` Group 1 virtual interrupts are disabled.

`1` Group 1 virtual interrupts are enabled.

This alias field is updated when a VM updates GICV_CTLR.EnableGrp1.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

VENG0, bit [0]

Virtual interrupt enable, Group 0. Possible values of this bit are:

`0` Group 0 virtual interrupts are disabled.

`1` Group 0 virtual interrupts are enabled.

This alias field is updated when a VM updates GICV_CTLR.EnableGrp0.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Note

A List register is in the pending state only if the corresponding GICH_LR<n> value is `01`, that is, pending. The active and pending state is not included.

Accessing the GICH_VMCR:

GICH_VMCR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual interface control	`0x0008`

8.17.8 GICH_VTR, Virtual Type Register

The GICH_VTR characteristics are:

Purpose

Indicates the number of implemented virtual priority bits and List registers.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

This register is used only when System register access is not enabled. When System register access is enabled:

- For AArch32 implementations, ICH_VTR provides equivalent functionality.
- For AArch64 implementations, ICH_VTR_EL2 provides equivalent functionality.

Configurations

This register is available when the GIC implementation supports interrupt virtualization.

Attributes

GICH_VTR is a 32-bit register.

Field descriptions

The GICH_VTR bit assignments are:

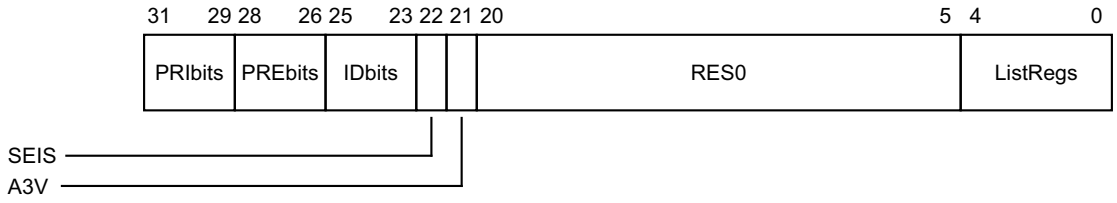

PRIbits, bits [31:29]

The number of virtual priority bits implemented, minus one.

An implementation must implement at least 32 levels of virtual priority (5 priority bits).

PREbits, bits [28:26]

The number of virtual preemption bits implemented, minus one.

An implementation must implement at least 32 levels of virtual preemption priority (5 preemption bits).

The value of this field must be less than or equal to the value of GICH_VTR.PRIbits.

IDbits, bits [25:23]

The number of virtual interrupt identifier bits supported:

`000` 16 bits.

`001` 24 bits.

All other values are reserved.

SEIS, bit [22]

SEI support. Indicates whether the virtual CPU interface supports generation of SEIs:

`0` The virtual CPU interface logic does not support generation of SEIs.

`1` The virtual CPU interface logic supports generation of SEIs.

A3V, bit [21]

Affinity 3 valid. Possible values are:

`0` The virtual CPU interface logic only supports zero values of the Aff3 field in ICC_SGI0R_EL1, ICC_SGI1R_EL1, and ICC_ASGI1R_EL1.

`1` The virtual CPU interface logic supports nonzero values of the Aff3 field in ICC_SGI0R_EL1, ICC_SGI1R_EL1, and ICC_ASGI1R_EL1.

Bits [20:5]

Reserved, RES0.

ListRegs, bits [4:0]

The number of implemented List registers, minus one.

Accessing the GICH_VTR:

GICH_VTR can be accessed through the memory-mapped interface:

Component	Offset
GIC Virtual interface control	`0x0004`

8.18 The ITS register map

The ITS address map consists of two separate 64KB frames starting from an IMPLEMENTATION DEFINED address specified in ITS_base. This base address must be aligned to a 64KB boundary. The two frames are:

- The control registers. which are located at ITS_base + 0x000000.
- The interrupt translation space, which is located at ITS_base + 0x010000.

Table 8-33 shows the GIC register map for the ITS control registers.

Table 8-33 ITS control register map

Offset	Name	Type	Reset	Description
0x0000	GITS_CTLR	RW	See the register description	ITS control register
0x0004	GITS_IIDR	RO	IMPLEMENTATION DEFINED	ITS Identification register
0x0008	GITS_TYPER	RO	IMPLEMENTATION DEFINED	ITS Type register
0x0010-0x001C	-	-	-	Reserved
0x0020-0x003C	-	-	IMPLEMENTATION DEFINED	IMPLEMENTATION DEFINED registers.
0x0040-0x007C	-	-	-	Reserved
0x0080	GITS_CBASER	RW	See the register description	ITS Command Queue Descriptor
0x0088	GITS_CWRITER	RW	See the register description	ITS Write register
0x0090	GITS_CREADR	RO	See the register description	ITS Read register
0x0098-0x00FC	-	-	-	Reserved
0x0100-0x0138	GITS_BASER<n>	RW	See the register description	ITS Translation Table Descriptors
0x0140-0xBFFC	-	-	-	Reserved
0xC000-0xFFCC	-	-	IMPLEMENTATION DEFINED	IMPLEMENTATION DEFINED registers.
0xFFD0-0xFFFC	-	RO	-	Reserved for ID registers, see *Identification registers* on page 8-173

Table 8-34 shows the GIC register map for the ITS translation registers.

Table 8-34 ITS translation register map

Offset	Name	Type	Reset	Description
0x0000-0x003C	-	-	-	Reserved
0x0040	GITS_TRANSLATER	WO	-	ITS Translation register
0x0044-0xFFFC	-	-	-	Reserved

8.19 The ITS register descriptions

This section describes each of the ITS registers in register name order.

8.19.1 GITS_BASER<n>, ITS Translation Table Descriptors, n = 0 - 7

The GITS_BASER<n> characteristics are:

Purpose

Specifies the base address and size of the ITS translation tables.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

Configurations

A copy of this register is provided for each ITS translation table.

Bits [63:32] and bits [31:0] are accessible independently.

A maximum of 8 GITS_BASER<n> registers can be provided. Unimplemented registers are RES0.

When GITS_CTLR.Enable == 1 or GITS_CTLR.Quiescent == 0, writing this register is UNPREDICTABLE.

Attributes

GITS_BASER<n> is a 64-bit register.

Field descriptions

The GITS_BASER<n> bit assignments are:

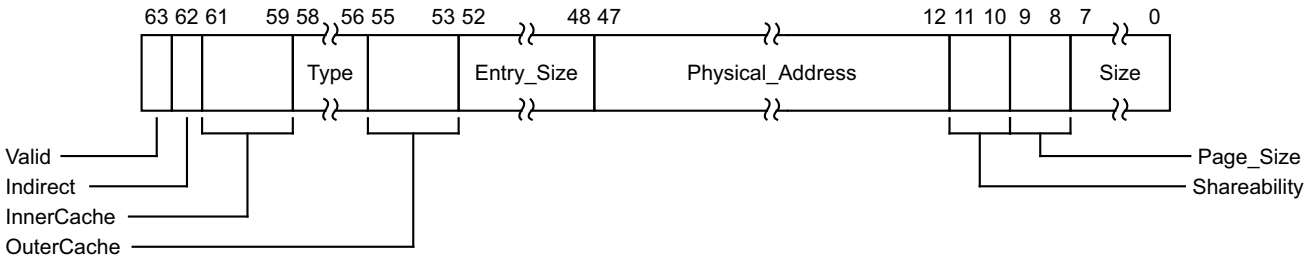

Valid, bit [63]

Indicates whether software has allocated memory for the translation table:

`0` No memory is allocated for the translation table. The ITS discards any writes to the interrupt translation page when either:

- GITS_BASER<n>.Type specifies any valid table entry type other than interrupt collections, that is, any value other than `100`.
- GITS_BASER<n>.Type specifies an interrupt collection and GITS_TYPER.HCC == 0.

`1` Memory is allocated to the translation table.

When this register has an architecturally-defined reset value, this field resets to `0`.

Indirect, bit [62]

This field indicates whether an implemented register specifies a single, flat table or a two-level table where the first level contains a list of descriptors.

Note

This field is RAZ/WI for implementations that only support flat tables.

`0` Single Level. The Size field indicates the number of pages used by the ITS to store data associated with each table entry.

`1` Two Level. The Size field indicates the number of pages which contain an array of 64-bit descriptors to pages that are used to store the data associated with each table entry. A little endian memory order model is used.

See *The ITS tables* on page 6-99 for more information.

This field is RAZ/WI for GIC implementations that only support flat tables.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

InnerCache, bits [61:59]

Indicates the Inner Cacheability attributes of accesses to the table. The possible values of this field are:

`000` Device-nGnRnE.

`001` Normal Inner Non-cacheable.

`010` Normal Inner Cacheable Read-allocate, Write-through.

`011` Normal Inner Cacheable Read-allocate, Write-back.

`100` Normal Inner Cacheable Write-allocate, Write-through.

`101` Normal Inner Cacheable Write-allocate, Write-back.

`110` Normal Inner Cacheable Read-allocate, Write-allocate, Write-through.

`111` Normal Inner Cacheable Read-allocate, Write-allocate, Write-back.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Type, bits [58:56]

Read only. Specifies the type of entity that requires entries in the corresponding translation table. The possible values of the field are:

`000` Unimplemented. This register does not correspond to a translation table.

`001` Devices. This register corresponds to a translation table that scales with the width of the DeviceID. Only a single GITS_BASER<n> register reports this type.

`010` vPEs. GICv4 only. This register corresponds to a translation table that scales with the number of vPEs in the system. The translation table requires (ENTRY_SIZE * N) bytes of memory, where N is the number of vPEs in the system. Only a single GITS_BASER<n> register reports this type.

`100` Interrupt collections. This register corresponds to a translation table that scales with the number of interrupt collections in the system. The translation table requires (ENTRY_SIZE * N) bytes of memory, where N is the number of interrupt collections. Not more than one GITS_BASER<n> register will report this type.

Other values are reserved.

Note

The minimum number of entries that an ITS must support is N+1, where N is the number of physical PEs in the system.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

OuterCache, bits [55:53]

Indicates the Outer Cacheability attributes of accesses to the table. The possible values of this field are:

000 Memory type defined in InnerCache field. For Normal memory, Outer Cacheability is the same as Inner Cacheability.

001 Normal Outer Non-cacheable.

010 Normal Outer Cacheable Read-allocate, Write-through.

011 Normal Outer Cacheable Read-allocate, Write-back.

100 Normal Outer Cacheable Write-allocate, Write-through.

101 Normal Outer Cacheable Write-allocate, Write-back.

110 Normal Outer Cacheable Read-allocate, Write-allocate, Write-through.

111 Normal Outer Cacheable Read-allocate, Write-allocate, Write-back.

It is IMPLEMENTATION DEFINED whether this field has a fixed value or can be programmed by software. Implementing this field with a fixed value is deprecated.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

Entry_Size, bits [52:48]

Read-only. Specifies the number of bytes per translation table entry, minus one.

Physical_Address, bits [47:12]

Physical Address. When Page_Size is 4KB or 16KB:

- Bits [51:48] of the base physical address are zero.
- This field provides bits[47:12] of the base physical address of the table.
- Bits[11:0] of the base physical address are zero.
- The address must be aligned to the size specified in the Page Size field. Otherwise the effect is CONSTRAINED UNPREDICTABLE, and can be one of the following:
  - Bits[X:12], where X is derived from the page size, are treated as zero.
  - The value of bits[X:12] are used when calculating the address of a table access.

When Page_Size is 64KB:

- Bits[47:16] of the register provide bits[47:16] of the base physical address of the table.
- Bits[15:12] of the register provide bits[51:48] of the base physical address of the table.
- Bits[15:0] of the base physical address are 0.

In implementations that support fewer than 52 bits of physical address, any unimplemented upper bits may be RAZ/WI.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Shareability, bits [11:10]

Indicates the Shareability attributes of accesses to the table. The possible values of this field are:

00 Non-shareable.

01 Inner Shareable.

10 Outer Shareable.

11 Reserved. Treated as 00.

It is IMPLEMENTATION DEFINED whether this field has a fixed value or can be programmed by software. Implementing this field with a fixed value is deprecated.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

Page_Size, bits [9:8]

The size of page that the translation table uses:

00 4KB.

01 16KB.

10 64KB.

11 Reserved. Treated as 10.

Note

If the GIC implementation supports only a single, fixed page size, this field might be RO.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

Size, bits [7:0]

The number of pages of physical memory allocated to the table, minus one. GITS_BASER<n>.Page_Size specifies the size of each page.

If GITS_BASER<n>.Type == 0, this field is RAZ/WI.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

Accessing the GITS_BASER<n>:

GITS_BASER<n> can be accessed through the memory-mapped interface:

Component	Offset
GIC ITS control	0x0100 + 8n

8.19.2 GITS_CBASER, ITS Command Queue Descriptor

The GITS_CBASER characteristics are:

Purpose

Specifies the base address and size of the ITS command queue.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

When GITS_CTLR.Enable == 1 or GITS_CTLR.Quiescent == 0, writing this register is UNPREDICTABLE.

Configurations

Bits [63:32] and bits [31:0] are accessible separately.

Attributes

GITS_CBASER is a 64-bit register.

Field descriptions

The GITS_CBASER bit assignments are:

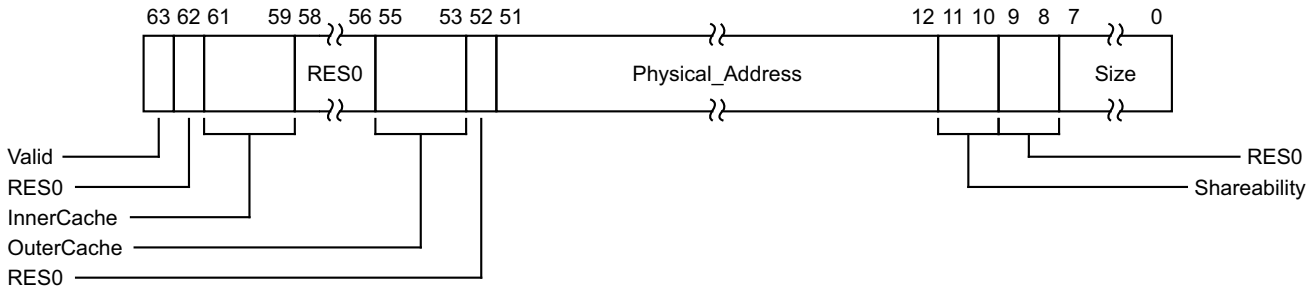

Valid, bit [63]

Indicates whether software has allocated memory for the command queue:

`0` No memory is allocated for the command queue.

`1` Memory is allocated to the command queue.

When this register has an architecturally-defined reset value, this field resets to `0`.

Bit [62]

Reserved, RES0.

InnerCache, bits [61:59]

Indicates the Inner Cacheability attributes of accesses to the command queue. The possible values of this field are:

`000` Device-nGnRnE.

`001` Normal Inner Non-cacheable.

`010` Normal Inner Cacheable Read-allocate, Write-through.

`011` Normal Inner Cacheable Read-allocate, Write-back.

100 Normal Inner Cacheable Write-allocate, Write-through.

101 Normal Inner Cacheable Write-allocate, Write-back.

110 Normal Inner Cacheable Read-allocate, Write-allocate, Write-through.

111 Normal Inner Cacheable Read-allocate, Write-allocate, Write-back.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Bits [58:56]

Reserved, RES0.

OuterCache, bits [55:53]

Indicates the Outer Cacheability attributes of accesses to the command queue. The possible values of this field are:

000 Memory type defined in InnerCache field. For Normal memory, Outer Cacheability is the same as Inner Cacheability.

001 Normal Outer Non-cacheable.

010 Normal Outer Cacheable Read-allocate, Write-through.

011 Normal Outer Cacheable Read-allocate, Write-back.

100 Normal Outer Cacheable Write-allocate, Write-through.

101 Normal Outer Cacheable Write-allocate, Write-back.

110 Normal Outer Cacheable Read-allocate, Write-allocate, Write-through.

111 Normal Outer Cacheable Read-allocate, Write-allocate, Write-back.

It is IMPLEMENTATION DEFINED whether this field has a fixed value or can be programmed by software. Implementing this field with a fixed value is deprecated.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

Bit [52]

Reserved, RES0.

Physical_Address, bits [51:12]

Bits [51:12] of the base physical address of the command queue. Bits [11:0] of the base address are 0.

In implementations supporting fewer than 52 bits of physical address, unimplemented upper bits are RAZ/WI.

If bits [15:12] are not all zeros, behavior is a CONSTRAINED UNPREDICTABLE choice:

- Bits [15:12] are treated as if all the bits are zero. The value read back from those bits is either the value written or zero.
- The result of the calculation of an address for a command queue read can be corrupted.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Shareability, bits [11:10]

Indicates the Shareability attributes of accesses to the command queue. The possible values of this field are:

00 Non-shareable.

01 Inner Shareable.

10 Outer Shareable.

11 Reserved. Treated as 00.

It is IMPLEMENTATION DEFINED whether this field has a fixed value or can be programmed by software. Implementing this field with a fixed value is deprecated.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.

Bits [9:8]

Reserved, RES0.

Size, bits [7:0]

The number of 4KB pages of physical memory allocated to the command queue, minus one.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

The command queue is a circular buffer and wraps at Physical Address [47:0] + (4096 * (Size + 1)).

Note

When this register is successfully written, the value of GITS_CREADR is set to zero.

Accessing the GITS_CBASER:

GITS_CBASER can be accessed through the memory-mapped interface:

Component	Offset
GIC ITS control	0x0080-0x0084

8.19.3 GITS_CREADR, ITS Read Register

The GITS_CREADR characteristics are:

Purpose

Specifies the offset from GITS_CBASER where the ITS reads the next ITS command.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

Configurations

This register is cleared to 0 when a value is written to GITS_CBASER.

Bits [63:32] and bits [31:0] are accessible separately.

Attributes

GITS_CREADR is a 64-bit register.

Field descriptions

The GITS_CREADR bit assignments are:

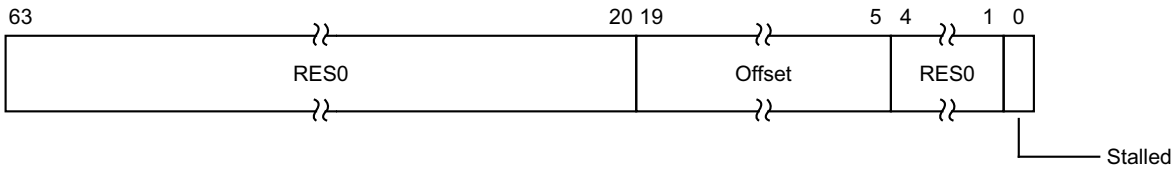

Bits [63:20]

Reserved, RES0.

Offset, bits [19:5]

Bits [19:5] of the offset from GITS_CBASER. Bits [4:0] of the offset are zero.

Bits [4:1]

Reserved, RES0.

Stalled, bit [0]

Reports whether the processing of commands is stalled because of a command error.

`0` ITS command queue is not stalled because of a command error.

`1` ITS command queue is stalled because of a command error.

See *The ITS command interface* on page 6-105.

Accessing the GITS_CREADR:

GITS_CREADR can be accessed through the memory-mapped interface:

Component	Offset
GIC ITS control	0x0090-0x0094

8.19.4 GITS_CTLR, ITS Control Register

The GITS_CTLR characteristics are:

Purpose

Controls the operation of an ITS.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

Configurations

The ITS_Number (bits [7:4]) and bit [1] fields apply only in GICv4 implementations, and are RES0 in GICv3 implementations.

Attributes

GITS_CTLR is a 32-bit register.

Field descriptions

The GITS_CTLR bit assignments are:

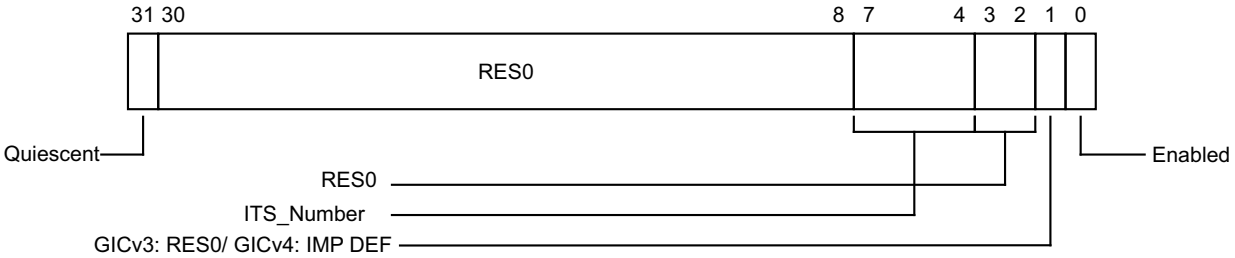

Quiescent, bit [31]

Read-only. Indicates completion of all ITS operations when GITS_CTLR.Enable == 0.

0 The ITS is not quiescent and cannot be powered down.

1 The ITS is quiescent and can be powered down.

For the ITS to be quiescent, there must be no transactions in progress. In addition, all operations required to ensure that mapping data is consistent with external memory must be complete.

Note

In distributed GIC implementations, this bit is set to 1 only after the ITS forwards any operations that have not yet been completed to the Redistributors and receives confirmation that all such operations have reached the appropriate Redistributor.

When this register has an architecturally-defined reset value, this field resets to 1.

Bits [30:8]

Reserved, RES0.

ITS_Number, bits [7:4]

In GICv3 implementations this field is RES0.

In GICv4:

- When GITS_TYPER.VMOVP == 0:
  - In implementations that have more than one ITS, ITS_Number indicates the number of the ITS that is used with VMOVP.
  - It is IMPLEMENTATION DEFINED whether this field is programmable or RO.
  - If this field is programmable, changing this field when GITS_CTLR.Quiescent==0 or GITS_CTLR.Enabled==1 is UNPREDICTABLE.
  - When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to a value that is architecturally UNKNOWN.
- When GITS_TYPER.VMOVP == 1:
  - This field is RES0.

Bits [3:2]

Reserved, RES0.

Bit [1]

In GICv3 implementations this bit is RES0.

In GICv4 implementations this bit is IMPLEMENTATION DEFINED.

When this register has an architecturally-defined reset value, if this field is implemented as an RW field, it resets to `0`.

Enabled, bit [0]

Controls whether an Interrupt Translation Space is enabled:

`0` The Interrupt Translation Space is not enabled. Writes to the Interrupt Translation Space are ignored and no further command queue entries are processed.

`1` The Interrupt Translation Space is enabled. Writes to the Interrupt Translation Space result in interrupt translations and the command queue is processed.

If a write to this register changes this field from 1 to 0, the Interrupt Translation Space must ensure that both:

- Any caches containing mapping data are made consistent with external memory.
- GITS_CTLR.Quiescent == 0 until all caches are consistent with external memory.

When this register has an architecturally-defined reset value, this field resets to `0`.

Accessing the GITS_CTLR:

GITS_CTLR can be accessed through the memory-mapped interface:

Component	Offset
GIC ITS control	`0x0000`

8.19.5 GITS_CWRITER, ITS Write Register

The GITS_CWRITER characteristics are:

Purpose

Specifies the offset from GITS_CBASER where software writes the next ITS command.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RW	RW	RW

Configurations

Bits [63:32] and bits [31:0] are accessible separately.

Attributes

GITS_CWRITER is a 64-bit register.

Field descriptions

The GITS_CWRITER bit assignments are:

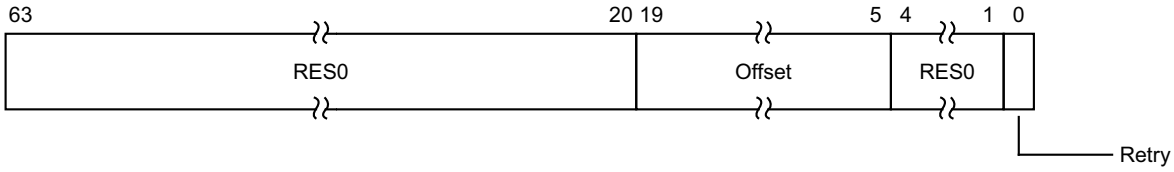

Bits [63:20]

Reserved, RES0.

Offset, bits [19:5]

Bits [19:5] of the offset from GITS_CBASER. Bits [4:0] of the offset are zero.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally UNKNOWN.

Bits [4:1]

Reserved, RES0.

Retry, bit [0]

Writing this bit has the following effects:

`0` No effect on the processing commands by the ITS.

`1` Restarts the procGessing of commands by the ITS if it stalled because of a command error.

Note

If the processing of commands is not stalled because of a command error, writing 1 to this bit has no effect.

When read, this bit is RES0.

See: *The ITS command interface* on page 6-105.

If GITS_CWRITER is written with a value outside of the valid range specified by GITS_CBASER.Physical_Address and GITS_CBASER.Size, behavior is a CONSTRAINED UNPREDICTABLE choice, as follows:

- The command queue is considered invalid, and no further commands are processed until GITS_CWRITER is written with a value that is in the valid range.
- The value is treated as a valid UNKNOWN value.

An implementation might choose to report a system error in an IMPLEMENTATION DEFINED manner.

Accessing the GITS_CWRITER:

GITS_CWRITER can be accessed through the memory-mapped interface:

Component	Offset
GIC ITS control	0x0088-0x008C

8.19.6 GITS_IIDR, ITS Identification Register

The GITS_IIDR characteristics are:

Purpose

Provides information about the implementer and revision of the ITS.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

Configurations

This register is available in all configurations of the GIC. If the GIC implementation supports two Security states, this register is Common.

Attributes

GITS_IIDR is a 32-bit register.

Field descriptions

The GITS_IIDR bit assignments are:

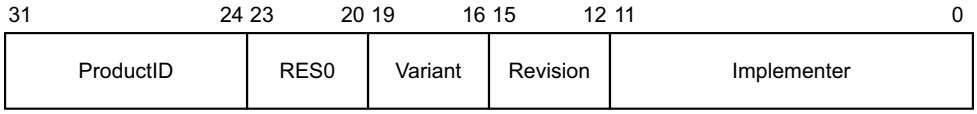

ProductID, bits [31:24]

An IMPLEMENTATION DEFINED product identifier.

Bits [23:20]

Reserved, RES0.

Variant, bits [19:16]

An IMPLEMENTATION DEFINED variant number. Typically, this field is used to distinguish product variants, or major revisions of a product.

Revision, bits [15:12]

An IMPLEMENTATION DEFINED revision number. Typically, this field is used to distinguish minor revisions of a product.

Implementer, bits [11:0]

Contains the JEP106 code of the company that implemented the ITS:

- Bits [11:8] are the JEP106 continuation code of the implementer. For an ARM implementation, this field is `0x4`.
- Bit [7] is always 0.
- Bits [6:0] are the JEP106 identity code of the implementer. For an ARM implementation, bits [7:0] are therefore `0x3B`.

Accessing the GITS_IIDR:

GITS_IIDR can be accessed through the memory-mapped interface:

Component	Offset
GIC ITS control	0x0004

8.19.7 GITS_TRANSLATER, ITS Translation Register

The GITS_TRANSLATER characteristics are:

Purpose

Written by a requesting a Device to signal an interrupt for translation by the ITS.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
WO	WO	WO

16-bit access to bits [15:0] of this register must be supported. When this register is written by a 16-bit transaction, bits [31:16] are written as zero.

Implementations must ensure that:

- A unique DeviceID is provided for each requesting device, and the DeviceID is presented to the ITS when a write to this register occurs in a manner that cannot be spoofed by any agent capable of performing writes.
- The DeviceID presented corresponds to the DeviceID field in the ITS commands.

Writes to this register are ignored if any of the following are true:

- GITS_CTLR.Enable == 0.
- The presented DeviceID is not mapped to an Interrupt Translation Table.
- The DeviceID is larger than the supported size.
- The DeviceID is mapped to an Interrupt Translation Table, but the EventID is outside the range specified by the MAPD command.
- The EventID is mapped to an Interrupt Translation Table and the EventID is within the range specified by the MAPD command, but the EventID is unmapped.

Translation requests that result from writes to this register are subject to certain ordering rules. See *Ordering of translations with the output to ITS commands* on page 6-107 for more information.

Configurations

This register is at the same offset as GICD_SETSPI_NSR in the Distributor, and is at the same offset as GICR_SETLPIR in the Redistributor.

Attributes

GITS_TRANSLATER is a 32-bit register.

Field descriptions

The GITS_TRANSLATER bit assignments are:

EventID, bits [31:0]

An identifier corresponding to the interrupt to be translated.

Note

The size of the EventID is DeviceID specific, and set when the DeviceID is mapped to an ITT (using a MAPD command).

The number of EventID bits implemented is reported by GITS_TYPER.IDbits. If a write specifies non-zero identifiers bits outside this range behavior is a CONSTRAINED UNPREDICTABLE choice between:

- Non-zero identifier bits outside the supported range are ignored.
- The write is ignored.

The DeviceID presented to an ITS is used to index a device table. The device table maps the DeviceID to an interrupt translation table for that device.

Accessing the GITS_TRANSLATER:

GITS_TRANSLATER can be accessed through the memory-mapped interface:

Component	Offset
GIC ITS translation	0x0040

8.19.8 GITS_TYPER, ITS Type Register

The GITS_TYPER characteristics are:

Purpose

Specifies the features that an ITS supports.

Usage constraints

This register is accessible as follows:

Security disabled	Secure	Non-secure
RO	RO	RO

Configurations

There are no configuration notes.

Attributes

GITS_TYPER is a 64-bit register.

Field descriptions

The GITS_TYPER bit assignments are:

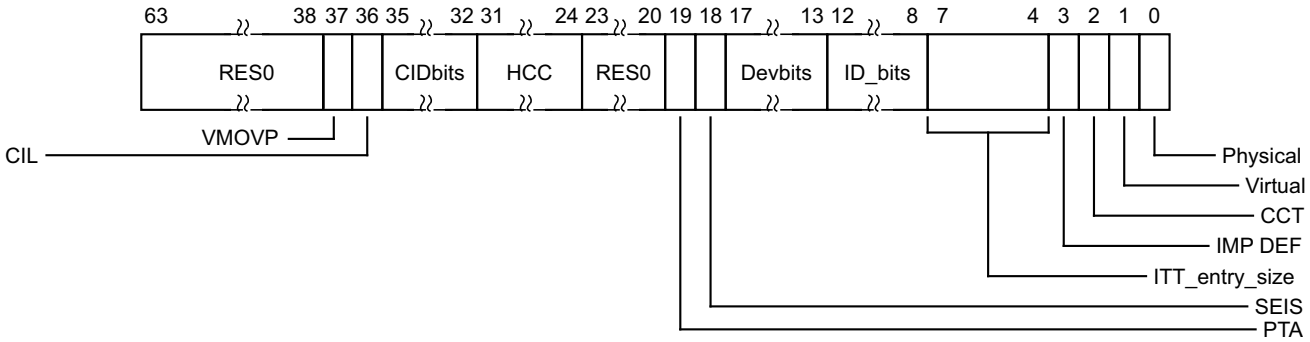

Bits [63:37]

Reserved, RES0.

VMOVP Indicates the form of the VMOVP command.

`0` When moving a vPE, software must issue a VMOVP on all ITSs that have mappings for that vPE. The `ITSList` and `Sequence Number` fields in the VMOVP command must ensure synchronization, otherwise behavior is UNPREDICTABLE.

`1` When moving a vPE, software must only issue a VMOVP on one of the ITSs that has a mapping for that vPE. The `ITSList` and `Sequence Number` fields in the VMOVP command are RES0.

CIL, bit [36]

Collection ID Limit.

`0` ITS supports 16-bit Collection ID, GITS_TYPER.CIDbits is RES0.

`1` GITS_TYPER.CIDbits indicates supported Collection ID size

In implementations that do not support Collections in external memory, this bit is RES0 and the number of Collections supported is reported by GITS_TYPER.HCC.

CIDbits, bits [35:32]

Number of Collection ID bits.

- The number of bits of Collection ID - 1.
- When GITS_TYPER.CIL==0, this field is RES0.

HCC, bits [31:24]

Hardware Collection Count. The number of interrupt collections supported by the ITS without provisioning of external memory.

Note

Collections held in hardware are unmapped at reset.

Bits [23:20]

Reserved, RES0.

PTA, bit [19]

Physical Target Addresses. Indicates the format of the target address:

0 The target address corresponds to the PE number specified by GICR_TYPER.Processor_Number.

1 The target address corresponds to the base physical address of the required Redistributor.

See *RDbase* on page 6-110 for more information.

SEIS, bit [18]

SEI support. Indicates whether the virtual CPU interface supports generation of SEIs:

0 The ITS does not support local generation of SEIs.

1 The ITS supports local generation of SEIs.

Devbits, bits [17:13]

The number of DeviceID bits implemented, minus one.

ID_bits, bits [12:8]

The number of EventID bits implemented, minus one.

ITT_entry_size, bits [7:4]

Read-only. Indicates the number of bytes per translation table entry, minus one.

See *MAPD* on page 6-118 for more information.

IMPLEMENTATION DEFINED, bit [3]

IMPLEMENTATION DEFINED.

CCT, bit [2]

Cumulative Collection Tables.

0 The total number of supported collections is determined by the number of collections held in memory only.

1 The total number of supported collections is determined by number of collections that are held in memory and the number indicated by GITS_TYPER.HCC.

If GITS_TYPER.HCC==0, or if memory backed collections are not supported (all GITS_BASER<n>.Type != 100), this bit is RES0.

Virtual, bit [1]

Indicates whether the ITS supports virtual LPIs and direct injection of virtual LPIs:

0 The ITS does not support virtual LPIs or direct injection of virtual LPIs.

1 The ITS supports virtual LPIs and direct injection of virtual LPIs.

This field is RES0 in GICv3 implementations.

Physical, bit [0]

Indicates whether the ITS supports physical LPIs:

0 The ITS does not support physical LPIs.

1 The ITS supports physical LPIs.

This field is RES1, indicating that the ITS supports physical LPIs.

Accessing the GITS_TYPER:

GITS_TYPER can be accessed through the memory-mapped interface:

Component	Offset
GIC ITS control	0x0008-0x000C

8.20 Pseudocode

AArch64 functions shows the pseudocode for the System registers when executing in AArch64 state. The same pseudocode can be used for the System registers when executing in AArch32 state by substituting the AArch64 register names with the equivalent AArch32 register names.

Note

An AArch64 register name includes the lowest Exception level that can access the register as a suffix to the register name. An AArch32 register name does not contain this suffix. For example the AArch64 Interrupt Controller Deactivate Interrupt Register is ICC_DIR_EL1, while the AArch32 equivalent is ICC_DIR.

Functions for memory-mapped registers on page 8-677 shows the pseudocode for the memory-mapped registers.

Note

Some variable names used the pseudocode differ from those used in the body text. For a list of the affected variables, see *Pseudocode terminology* on page B-734.

8.20.1 AArch64 functions

This subsection describes the AArch64 functions. The functions are indicated by the hierarchical path names, for example `aarch64/support`. The functions are:

- *aarch64/support/ICC_DIR_EL1*.
- *aarch64/support/ICC_EOIR0_EL1* on page 8-664.
- *aarch64/support/ICC_EOIR1_EL1* on page 8-665.
- *aarch64/support/ICC_HPPIR0_EL1* on page 8-666.
- *aarch64/support/ICC_HPPIR1_EL1* on page 8-666.
- *aarch64/support/ICC_IAR0_EL1* on page 8-666.
- *aarch64/support/ICC_IAR1_EL1* on page 8-667.
- *aarch64/support/ICC_PMR_EL1* on page 8-667.
- *aarch64/support/ICC_RPR_EL1* on page 8-668.
- *aarch64/support/ICH_EISR_EL2* on page 8-668.
- *aarch64/support/ICH_ELSR_EL2* on page 8-669.
- *aarch64/support/VirtualRead HPPIR0* on page 8-669.
- *aarch64/support/VirtualRead HPPIR1* on page 8-669.
- *aarch64/support/VirtualReadIAR0* on page 8-670.
- *aarch64/support/VirtualReadIAR1* on page 8-671.
- *aarch64/support/VirtualWriteDIR* on page 8-671.
- *aarch64/support/VirtualWriteEOIR0* on page 8-672.
- *aarch64/support/VirtualWriteEOIR1* on page 8-673.
- *aarch64/support/CheckGroup0ForSpecialIdentifiers* on page 8-675.
- *aarch64/support/CheckGroup1ForSpecialIdentifiers* on page 8-676.
- *aarch64/support/PRIMask* on page 8-677.
- *aarch64/support/VRIMask* on page 8-677.

aarch64/support/ICC_DIR_EL1

```
// ICC_DIR_EL1 - assignment form
// =============================

ICC_DIR_EL1[] = bits(64) data

    // First check if System Registers are enabled
    SystemRegisterAccessPermitted(2);
```

```
    // Check whether Non-secure EL1 writes of ICC_DIR_EL1 are trapped to EL2
    if !IsSecure() && PSTATE.EL == EL1 && ICH_HCR_EL2.TDIR =='1' then
        SystemRegisterTrap(EL2);

    // Check if the access is virtual
    if (HaveEL(EL2) && !IsSecure() && PSTATE.EL == EL1 &&
        (HCR_EL2.FMO == '1' || HCR_EL2.IMO == '1')) then
        VirtualWriteDIR(data);          // Access the Virtual DIR register
        return;

    // Check for spurious ID.  LPIs are not allowed and the access is physical
    if !InterruptIdentifierValid(data, FALSE) then
        return;

    ID = data<INTID_SIZE-1:0>;

    // Now start handling the interrupt
    if !EOImodeSet() then
        // EOI mode is not set, so don't deactivate
        IMPLEMENTATION_DEFINED "SError DIR_EOIMODE_NOT_SET";
    else
        route_fiq_to_el3 = HaveEL(EL3) && SCR_EL3.FIQ == '1';
        route_irq_to_el3 = HaveEL(EL3) && SCR_EL3.IRQ == '1';
        route_fiq_to_el2 = HaveEL(EL2) && !IsSecure() && HCR_EL2.FMO == '1';
        route_irq_to_el2 = HaveEL(EL2) && !IsSecure() && HCR_EL2.IMO == '1';
        if PSTATE.EL == EL3 then
            Deactivate(ID);
        elsif PSTATE.EL == EL2 then
            if SingleSecurityState() && IsGrp0Int(ID) && !route_fiq_to_el3 then
                Deactivate(ID);
            elsif !IsSecureInt(ID) && !IsGrp0Int(ID) && !route_irq_to_el3 then
                Deactivate(ID);
        elsif PSTATE.EL == EL1 && !IsSecure() then
            if SingleSecurityState() && IsGrp0Int(ID) && !route_fiq_to_el3 && !route_fiq_to_el2 then
                Deactivate(ID);
            elsif !IsSecureInt(ID) && !IsGrp0Int(ID) && !route_irq_to_el3 && !route_irq_to_el2 then
                Deactivate(ID);
        elsif PSTATE.EL == EL1 && IsSecure() then
            if IsGrp0Int(ID) && !route_fiq_to_el3 then
                Deactivate(ID);
            elsif (!IsGrp0Int(ID) && (!IsSecureInt(ID) || !SingleSecurityState()) &&
                   !route_irq_to_el3) then
                Deactivate(ID);

    return;
```

aarch64/support/ICC_EOIR0_EL1

```
// ICC_EOIR0_EL1 - assignment form
// ===============================

ICC_EOIR0_EL1[] = bits(64) data

    eoiID = data<INTID_SIZE-1:0>;

    // First check if System Registers are enabled
    SystemRegisterAccessPermitted(0);

    // Check if the access is virtual
    if HaveEL(EL2) && !IsSecure() && PSTATE.EL == EL1 && HCR_EL2.FMO == '1' then
        VirtualWriteEOIR0(data);        // Access the Virtual EOIR0 register

    // Check for spurious ID.  LPIs are allowed and the access is physical
    if !InterruptIdentifierValid(data, TRUE) then
        return;

    // Now start handling the interrupt
```

```
    // Is the highest priority G0S, G1S or G1NS
    pGroup = GetHighestActiveGroup(ICC_AP0R_EL1, ICC_AP1R_EL1NS, ICC_AP1R_EL1S);
    pPriority = GetHighestActivePriority(ICC_AP0R_EL1, ICC_AP1R_EL1NS, ICC_AP1R_EL1S);

    if pGroup == IntGroup_None then
        // There are no active priorities
        if GenerateLocalSError() then
            // Reporting of locally generated SEIs is supported
            IMPLEMENTATION_DEFINED "SError EOI0_NO_INTS_ACTIVE";

    elsif pGroup == IntGroup_G0 && (!HaveEL(EL3) || SingleSecurityState() || IsSecure()) then
        // Highest priority is Group 0
        // Drop the priority
        boolean dropped = PriorityDrop[ICC_AP0R_EL1];

        if !EOImodeSet() then           // If EOI mode is set, don't deactivate
            // Deactivate the interrupt unless it is an LPI
            if !IsLPI(eoiID) then Deactivate(eoiID);

    elsif GenerateLocalSError() then    // Locally generated SEIs are supported
        // Highest priority is Group 1
        IMPLEMENTATION_DEFINED "SError EOI0_HIGHEST_IS_G1";

        return;
```

aarch64/support/ICC_EOIR1_EL1

```
// ICC_EOIR1_EL1 - assignment form
// ===============================

ICC_EOIR1_EL1[] = bits(64) data

    eoiID = data<INTID_SIZE-1:0>;

    // First check if System Registers are enabled
    SystemRegisterAccessPermitted(1);

    // Check if the access is virtual
    if HaveEL(EL2) && !IsSecure() && PSTATE.EL == EL1 && HCR_EL2.IMO == '1' then
        VirtualWriteEOIR1(data);        // Access the Virtual EOIR1 register

    // Check for spurious ID.  LPIs are allowed and the access is physical
    if !InterruptIdentifierValid(data, TRUE) then
        return;

    // Now start handling the interrupt
    // Is the highest priority G0S, G1S or G1NS
    pGroup = GetHighestActiveGroup(ICC_AP0R_EL1, ICC_AP1R_EL1NS, ICC_AP1R_EL1S);
    pPriority = GetHighestActivePriority(ICC_AP0R_EL1, ICC_AP1R_EL1NS, ICC_AP1R_EL1S);

    if pGroup == IntGroup_None then
        // There are no active priorities
        if GenerateLocalSError() then
            // Reporting of locally generated SEIs is supported
            IMPLEMENTATION_DEFINED "SError EOI1_NO_INTS_ACTIVE";

    elsif pGroup == IntGroup_G1NS && (PSTATE.EL == EL3 || !IsSecure()) then
        // Highest priority is Non-Secure Group 1
        // Drop the priority
        boolean dropped = PriorityDrop[ICC_AP1R_EL1NS];

        if !EOImodeSet() then           // If EOI mode is set, don't deactivate
            // Deactivate the interrupt unless it is an LPI
            if !IsLPI(eoiID) then Deactivate(eoiID);

    elsif pGroup == IntGroup_G1S && IsSecure() then
        // Highest priority is Secure Group 1 and we are secure
```

```
        // Drop the priority
        boolean dropped = PriorityDrop[ICC_AP1R_EL1S];

        if !EOImodeSet() then           // If EOI mode is set, don't deactivate
            // Deactivate the interrupt unless it is an LPI
            if !IsLPI(eoiID) then Deactivate(eoiID);

    elsif GenerateLocalSError() then    // Locally generated SEIs are supported
        // Highest priority is Group 0 or Secure Group 1 and we are not secure
        IMPLEMENTATION_DEFINED "SError EOI1_HIGHEST_NOT_ACCESSIBLE";

        return;
```

aarch64/support/ICC_HPPIR0_EL1

```
// ICC_HPPIR0_EL1 - non-assignment form
// ====================================

bits(32) ICC_HPPIR0_EL1[]

    // First check if System Registers are enabled
    SystemRegisterAccessPermitted(0);

    // Check if the access is virtual
    if HaveEL(EL2) && !IsSecure() && PSTATE.EL == EL1 && HCR_EL2.FMO == '1' then
        return VirtualReadHPPIR0();       // Access the Virtual HPPIR0 register

    // Now start handling the interrupt
    pendID = HighestPriorityPendingInterrupt();
    pendID = CheckGroup0ForSpecialIdentifiers(pendID);

    return ZeroExtend(pendID);
```

aarch64/support/ICC_HPPIR1_EL1

```
// ICC_HPPIR1_EL1 - non-assignment form
// ====================================

bits(32) ICC_HPPIR1_EL1[]

    // First check if System Registers are enabled
    SystemRegisterAccessPermitted(1);

    // Check if the access is virtual
    if HaveEL(EL2) && !IsSecure() && PSTATE.EL == EL1 && HCR_EL2.IMO == '1' then
        return VirtualReadHPPIR1();       // Access the Virtual HPPIR1 register

    // Now start handling the interrupt
    pendID = HighestPriorityPendingInterrupt();
    pendID = CheckGroup1ForSpecialIdentifiers(pendID);

    return ZeroExtend(pendID);
```

aarch64/support/ICC_IAR0_EL1

```
// ICC_IAR0_EL1 - non-assignment form
// ==================================

bits(32) ICC_IAR0_EL1[]

    // First check if System Registers are enabled
    SystemRegisterAccessPermitted(0);

    // Check if the access is virtual
    if HaveEL(EL2) && !IsSecure() && PSTATE.EL == EL1 && HCR_EL2.FMO == '1' then
        return VirtualReadIAR0();                       // Access the Virtual IAR0 register
```

```
    // Now start handling the interrupt
    if !CanSignalInterrupt() then
        return ZeroExtend(INTID_SPURIOUS);

    // Gets the highest priority pending and enabled interrupt
    pendID = HighestPriorityPendingInterrupt();
    pendID = CheckGroup0ForSpecialIdentifiers(pendID);

    // Check that pendID is not a special interrupt ID
    if !IsSpecial(pendID) then
        AcknowledgeInterrupt(pendID);                    // Set active and attempt to clear pending

    return ZeroExtend(pendID);
```

aarch64/support/ICC_IAR1_EL1

```
// ICC_IAR1_EL1 - non-assignment form
// ==================================

bits(32) ICC_IAR1_EL1[]

    // First check if System Registers are enabled
    SystemRegisterAccessPermitted(1);

    // Check if the access is virtual
    if HaveEL(EL2) && !IsSecure() && PSTATE.EL == EL1 && HCR_EL2.IMO == '1' then
        return VirtualReadIAR1();                        // Access the Virtual IAR1 register

    // Now start handling the interrupt
    if !CanSignalInterrupt() then
        return ZeroExtend(INTID_SPURIOUS);

    // Gets the highest priority pending and enabled interrupt
    pendID = HighestPriorityPendingInterrupt();
    pendID = CheckGroup1ForSpecialIdentifiers(pendID);

    // Check that pendID is not a special interrupt ID
    if !IsSpecial(pendID) then
        AcknowledgeInterrupt(pendID);                    // Set active and attempt to clear pending

    return ZeroExtend(pendID);
```

aarch64/support/ICC_PMR_EL1

```
// ICC_PMR_EL1[] - non-assignment form
// ==================================

bits(32) ICC_PMR_EL1[]

    // First check if System Registers are enabled
    SystemRegisterAccessPermitted(2); // Set group to 2 so "TC" bit is checked

    if (HaveEL(EL2) && !IsSecure() && PSTATE.EL == EL1 &&
        (HCR_EL2.FMO == '1' || HCR_EL2.IMO == '1')) then
        // At least one interrupt is virtualized so return the virtual mask
        return ZeroExtend(ICH_VMCR_EL2.VPMR AND VPRIMask());

    pPriority = ICC_PMR_EL1.Priority;

    if HaveEL(EL3) && !IsSecure() && SCR_EL3.FIQ == '1' then
        // A non-secure GIC access and group 0 inaccessible to Non-secure.
        if pPriority<7> == '0' then
            // Priority is in Secure half and not visible to Non-secure
            pPriority<7:0> = Zeros();
        elsif pPriority != PRIMask() then
```

```
            // Non-secure access and not idle, so physical priority must be shifted
            pPriority<7:0> = (pPriority AND PRIMask())<6:0>:'0';

    return ZeroExtend(pPriority);

// ICC_PMR_EL1[] - assignment form
// ===============================

ICC_PMR_EL1 = bits(32) data

    // First check if System Registers are enabled
    SystemRegisterAccessPermitted(2); // Set group to 2 so "TC" bit is checked

    if (HaveEL(EL2) && !IsSecure() && PSTATE.EL == EL1 &&
        (HCR_EL2.FMO == '1' || HCR_EL2.IMO == '1')) then
        // At least one interrupt is virtualized so update the virtual mask
        ICH_VMCR_EL2.VPMR = data<7:0> AND VPRIMask();
        return;

    if HaveEL(EL3) && !IsSecure() && SCR_EL3.FIQ == '1' then
        // A Non-secure GIC access and Group 0 inaccessible to Non-secure.
        mod_write_val = ('1':data<7:1>) AND PRIMask();
        // Non-secure state can only update the Priority Mask Register if the current value is in
        // the range 0x80 to 0xFF
        if ICC_PMR_EL1.Priority<7> == '1' then
            ICC_PMR_EL1.Priority = mod_write_val;
        // Otherwise PMR is between 0x00 and 0x7F and the write is ignored
    else                                                    // A Secure GIC access
        ICC_PMR_EL1.Priority = data<7:0> AND PRIMask();

    return;
```

aarch64/support/ICC_RPR_EL1

```
// ICC_RPR_EL1 - non-assignment form
// =================================

bits(32) ICC_RPR_EL1[]

    // First check if System Registers are enabled
    SystemRegisterAccessPermitted(2); // Set group to 2 so "TC" bit is checked

    if (HaveEL(EL2) && !IsSecure() && PSTATE.EL == EL1 &&
        (HCR_EL2.FMO == '1' || HCR_EL2.IMO == '1')) then
        // At least one interrupt is virtualized so return the virtual priority
        return ZeroExtend(GetHighestActiveVPriority(ICH_AP0R_EL2, ICH_AP1R_EL2));

    // Get physical priority.
    pPriority = GetHighestActivePriority(ICC_AP0R_EL1, ICC_AP1R_EL1NS, ICC_AP1R_EL1S);

    if HaveEL(EL3) && !IsSecure() && SCR_EL3.FIQ == '1' then
        // A Non-secure GIC access and Group 0 inaccessible to Non-secure.
        if pPriority<7> == '0' then
            // Priority is in Secure half and not visible to Non-secure
            pPriority = Zeros();
        elsif !IsOnes(pPriority) then
            // Non-secure access and not idle, so physical priority must be shifted
            pPriority<7:0> = (pPriority AND PRIMask())<6:0>:'0';

    return ZeroExtend(pPriority);
```

aarch64/support/ICH_EISR_EL2

```
// ICH_EISR_EL2 - non-assignment form
// ==================================
```

```
bits(32) ICH_EISR_EL2[]
    bits(32) rval = Zeros();

    for i = 0 to NumListRegs() - 1
        if (ICH_LR_EL2[i].State == IntState_Invalid && ICH_LR_EL2[i].HW == '0' &&
            ICH_LR_EL2[i].EOI == '1') then
            rval<i> = '1';

    return rval;
```

aarch64/support/ICH_ELSR_EL2

```
// ICH_ELSR_EL2 - non-assignment form
// ==================================

bits(32) ICH_ELSR_EL2[]
    bits(32) rval = Zeros();

    for i = 0 to NumListRegs() - 1
        if (ICH_LR_EL2[i].State == IntState_Invalid &&
            (ICH_LR_EL2[i].HW == '1' || ICH_LR_EL2[i].EOI == '0')) then
            rval<i> = '1';

    return rval;
```

aarch64/support/VirtualRead HPPIR0

```
// VirtualReadHPPIR0()
// ===================

bits(32) VirtualReadHPPIR0()

    lrIndex = HighestPriorityVirtualInterrupt();

    if (GICH_VLPIR.State == IntState_Pending &&
        (lrIndex < 0 || PriorityIsHigher(GICH_VLPIR.Priority, ICH_LR_EL2[lrIndex].Priority))) then
        // A virtual LPI is the highest priority
        vID = GICH_VLPIR.VirtualID<INTID_SIZE-1:0>;
        if GICH_VLPIR.Group != '0' then
            vID = INTID_SPURIOUS;

    elsif lrIndex >= 0 then               // lrIndex is valid, that is, positive
        vID = ICH_LR_EL2[lrIndex].VirtualID<INTID_SIZE-1:0>;
        if (vID != INTID_SPURIOUS && (ICH_LR_EL2[lrIndex].Group != '0' ||
                                      ICH_LR_EL2[lrIndex].State == IntState_Invalid)) then
            vID = INTID_SPURIOUS;         // If the highest priority isn't group 0, then no interrupt

    else
        vID = INTID_SPURIOUS;

    return ZeroExtend(vID);
```

aarch64/support/VirtualRead HPPIR1

```
// VirtualReadHPPIR1()
// ===================

bits(32) VirtualReadHPPIR1()

    lrIndex = HighestPriorityVirtualInterrupt();

    if (GICH_VLPIR.State == IntState_Pending &&
        (lrIndex < 0 || PriorityIsHigher(GICH_VLPIR.Priority, ICH_LR_EL2[lrIndex].Priority))) then
        // A virtual LPI is the highest priority
        vID = GICH_VLPIR.VirtualID<INTID_SIZE-1:0>;
        if GICH_VLPIR.Group != '1' then
```

```
        vID = INTID_SPURIOUS;

    elsif lrIndex >= 0 then               // lrIndex is valid, that is, positive
        vID = ICH_LR_EL2[lrIndex].VirtualID<INTID_SIZE-1:0>;
        if (vID != INTID_SPURIOUS && (ICH_LR_EL2[lrIndex].Group != '1' ||
                                      ICH_LR_EL2[lrIndex].State == IntState_Invalid)) then
            vID = INTID_SPURIOUS;         // If the highest priority isn't group 1, then no interrupt

    else
        vID = INTID_SPURIOUS;

    return ZeroExtend(vID);
```

aarch64/support/VirtualReadIAR0

```
// VirtualReadIAR0()
// =================

bits(32) VirtualReadIAR0()

    integer lrIndex = HighestPriorityVirtualInterrupt();

    if !CanSignalVirtualInterrupt() then
        return ZeroExtend(INTID_SPURIOUS);

    if (GICH_VLPIR.State == IntState_Pending &&
        (lrIndex < 0 || PriorityIsHigher(GICH_VLPIR.Priority, ICH_LR_EL2[lrIndex].Priority))) then
        // A virtual LPI is the highest priority
        vID = GICH_VLPIR.VirtualID<INTID_SIZE-1:0>;
        if GICH_VLPIR.Group == '0' then
            vPriorityGroup = VPriorityGroup(GICH_VLPIR.Priority);
            ICH_AP0R_EL2<UInt(vPriorityGroup)> = '1';
            AcknowledgeVInterrupt(GICH_VLPIR.VirtualID<INTID_SIZE-1:0>);
            GICH_VLPIR.State = IntState_Invalid;         // Set the virtual LPI to Idle
        else
            vID = INTID_SPURIOUS;
        return ZeroExtend(vID);

    // lrIndex must be valid (i.e. non-negative)
    vID = ICH_LR_EL2[lrIndex].VirtualID<INTID_SIZE-1:0>;
    pID = ICH_LR_EL2[lrIndex].PhysicalID<INTID_SIZE-1:0>;

    if (vID != INTID_SPURIOUS &&
        (ICH_LR_EL2[lrIndex].Group != '0' || ICH_LR_EL2[lrIndex].State == IntState_Invalid)) then
        // If the highest priority isn't Group 0, then no interrupt
        return ZeroExtend(INTID_SPURIOUS);

    if !IsSpecial(vID) then                              // Check that it is not a spurious interrupt
        ICH_LR_EL2[lrIndex].State = IntState_Active;     // Set the list register state to Active
        vPriorityGroup = VPriorityGroup(ICH_LR_EL2[lrIndex].Priority);
        ICH_AP0R_EL2<UInt(vPriorityGroup)> = '1';        // Set the corresponding bit in APR
    else
        ICH_LR_EL2[lrIndex].State = IntState_Invalid;

        setEI = ICH_LR_EL2[lrIndex].EOI == '1';

        // Generate a maintenance interrupt if required
        if ICH_LR_EL2[lrIndex].HW == '0' && setEI then
            // Set the appropriate EISR bit to generate a maintenance interrupt
            ICH_EISR_EL2<lrIndex> = '1';
        else
            // Set the appropriate ELRSR bit
            ICH_ELRSR_EL2<lrIndex> = '1';

    return ZeroExtend(vID);
```

aarch64/support/VirtualReadIAR1

```
// VirtualReadIAR1()
// =================

bits(32) VirtualReadIAR1()

    integer lrIndex = HighestPriorityVirtualInterrupt();

    if !CanSignalVirtualInterrupt() then
        return ZeroExtend(INTID_SPURIOUS);

    if (GICH_VLPIR.State == IntState_Pending &&
        (lrIndex < 0 || PriorityIsHigher(GICH_VLPIR.Priority, ICH_LR_EL2[lrIndex].Priority))) then
        // A virtual LPI is the highest priority
        vID = GICH_VLPIR.VirtualID<INTID_SIZE-1:0>;
        if GICH_VLPIR.Group == '1' then
            vPriorityGroup = VPriorityGroup(GICH_VLPIR.Priority);
            ICH_AP1R_EL2<UInt(vPriorityGroup)> = '1';
            AcknowledgeVInterrupt(GICH_VLPIR.VirtualID<INTID_SIZE-1:0>);
            GICH_VLPIR.State = IntState_Invalid;         // Set the virtual LPI to Idle
        else
            vID = INTID_SPURIOUS;
        return ZeroExtend(vID);

    // lrIndex must be valid (i.e. non-negative)
    vID = ICH_LR_EL2[lrIndex].VirtualID<INTID_SIZE-1:0>;
    pID = ICH_LR_EL2[lrIndex].PhysicalID<INTID_SIZE-1:0>;

    if ICH_LR_EL2[lrIndex].Group != '1' || ICH_LR_EL2[lrIndex].State == IntState_Invalid then
        // If the highest priority isn't Group 1, then no interrupt
        return ZeroExtend(INTID_SPURIOUS);

    if !IsSpecial(vID) then                              // Check that it is not a spurious interrupt
        ICH_LR_EL2[lrIndex].State = IntState_Active;     // Set the list register state to Active
        vPriorityGroup = VPriorityGroup(ICH_LR_EL2[lrIndex].Priority);
        ICH_AP1R_EL2<UInt(vPriorityGroup)> = '1';        // Set the corresponding bit in APR
    else
        ICH_LR_EL2[lrIndex].State = IntState_Invalid;

        setEI = ICH_LR_EL2[lrIndex].EOI == '1';

        // Generate a maintenance interrupt if required
        if ICH_LR_EL2[lrIndex].HW == '0' && setEI then
            // Set the appropriate EISR bit to generate a maintenance interrupt
            ICH_EISR_EL2<lrIndex> = '1';
        else
            // Set the appropriate ELRSR bit
            ICH_ELRSR_EL2<lrIndex> = '1';

    return ZeroExtend(vID);
```

aarch64/support/VirtualWriteDIR

```
// VirtualWriteDIR()
// =================

VirtualWriteDIR(bits(64) data)

    // When an error is detected return to avoid unpredictable behaviour
    if ICH_VMCR_EL2.VEOIM == '0' then
        // EOI mode is not set so generate a maintenance interrupt
        if ICH_VTR_EL2.SEIS == '1' then
            // Reporting of locally generated virtual SEIs is supported
            IMPLEMENTATION_DEFINED "SError DIR_EOIMODE_NOT_SET";
        return;
```

```
    // Check for spurious ID.  LPIs are not allowed and the access is virtual
    if !VirtualIdentifierValid(data, FALSE) then
        return;

    ID = data<INTID_SIZE-1:0>;
    lrIndex = FindActiveVirtualInterrupt(ID);

    if lrIndex < 0 then
        // No valid list register corresponds to the EOI ID, increment EOI count
        ICH_HCR_EL2.EOIcount = ICH_HCR_EL2.EOIcount + 1;
        return;

    setEI = ICH_LR_EL2[lrIndex].EOI == '1';

    if ICH_LR_EL2[lrIndex].HW =='1' then
        // Deactivate the physical interrupt if EOI Mode is set
        pID = ICH_LR_EL2[lrIndex].PhysicalID;
        if UInt(pID) < 1020 then Deactivate(ZeroExtend(pID));
    else
        // Generate a maintenance interrupt if required
        if setEI then
            // Set the appropriate EISR bit to generate a maintenance interrupt
            ICH_EISR_EL2<lrIndex> = '1';

    // Clear the Active state
    if ICH_LR_EL2[lrIndex].State == IntState_ActivePending then
        ICH_LR_EL2[lrIndex].State = IntState_Pending;
    else
        ICH_LR_EL2[lrIndex].State = IntState_Invalid;
        if ICH_LR_EL2[lrIndex].HW == '1' || !setEI then
            ICH_ELRSR_EL2<lrIndex> = '1';

    return;
```

aarch64/support/VirtualWriteEOIR0

```
// VirtualWriteEOIR0()
// ===================

VirtualWriteEOIR0(bits(64) data)

    eoiID = data<24-1:0>;
    vPriority = GetHighestActiveVPriority(ICH_AP0R_EL2, ICH_AP1R_EL2);

    // Check the identifier is valid.  LPIs are allowed and the access is virtual
    if !VirtualIdentifierValid(data, TRUE) then
        return;

    // Now perform the priority drop
    drop = VPriorityBitsSet(ICH_AP0R_EL2, ICH_AP1R_EL2);

    if !drop then
        if ICH_VTR_EL2.SEIS == '1' then
            // Reporting of locally generated virtual SEIs is supported
            IMPLEMENTATION_DEFINED "SError EOI0_HIGHEST_IS_G1";
        return;

    // It is IMPLEMENTATION DEFINED whether the priority is dropped before the error checks
    if boolean IMPLEMENTATION_DEFINED "Drop before checks" then
        VPriorityDrop[ICH_AP0R_EL2, ICH_AP1R_EL2] = '0';
        dropped = TRUE;
    else
        dropped = FALSE;

    // Find the matching List Register
    lrIndex = FindActiveVirtualInterrupt(eoiID);
```

```
    if lrIndex < 0 then

        if IsLPI(eoiID) then
            // It is a virtual LPI not in the List Registers
            // so just priority drop and return without incrementing EOI count
            return;
        else
            // No valid list register corresponds to the EOI ID, increment EOI count
            if drop && ICH_VMCR_EL2.VEOIM == '0' then
                ICH_HCR_EL2.EOIcount = ICH_HCR_EL2.EOIcount + 1;
            return;

    // Start error checks
    // When an error is detected return to avoid unpredictable behaviour
    if ICH_LR_EL2[lrIndex].Group != '0' then
        // The EOI ID is not Group 0
        if ICH_VTR_EL2.SEIS == '1' then
            // Reporting of locally generated virtual SEIs is supported
            IMPLEMENTATION_DEFINED "SError EOI0_HIGHEST_IS_G1";
        return;

    if VPriorityGroup(ICH_LR_EL2[lrIndex].Priority, 0) != vPriority then
        // The EOI ID is not the highest priority
        if ICH_VTR_EL2.SEIS == '1' then
            // Reporting of locally generated virtual SEIs is supported
            IMPLEMENTATION_DEFINED "SError EOI0_NOT_HIGHEST_PRIORITY";
        return;
    // End of error checks

    if !dropped then VPriorityDrop[ICH_AP0R_EL2, ICH_AP1R_EL2] = '0';

    if ICH_VMCR_EL2.VEOIM == '0' || IsLPI(eoiID) then
        setEI = ICH_LR_EL2[lrIndex].EOI == '1';


        // EOI mode not set, or it is an LPI and no deactivate is expected
        // so clear the active state in the List Register
        if ICH_LR_EL2[lrIndex].HW =='1' then
            // Deactivate the physical interrupt
            pID = ICH_LR_EL2[lrIndex].PhysicalID;
            if UInt(pID) < 1020 then Deactivate(ZeroExtend(pID));
        else
            // Generate a maintenance interrupt if required
            if setEI then
                //Set the appropriate EISR bit to generate a maintenance interrupt
                ICH_EISR_EL2<lrIndex> = '1';

       // Clear the Active state
       if ICH_LR_EL2[lrIndex].State == IntState_ActivePending then
            ICH_LR_EL2[lrIndex].State = IntState_Pending;
       else
            ICH_LR_EL2[lrIndex].State = IntState_Invalid;
            if ICH_LR_EL2[lrIndex].HW == '1' || !setEI then
                ICH_ELRSR_EL2<lrIndex> = '1';


        return;
```

aarch64/support/VirtualWriteEOIR1

```
// VirtualWriteEOIR1()
// ===================

VirtualWriteEOIR1(bits(64) data)

    eoiID = data<24-1:0>;
    vPriority = GetHighestActiveVPriority(ICH_AP0R_EL2, ICH_AP1R_EL2);
```

```
// Check for spurious ID.  LPIs are allowed and the access is virtual
if !VirtualIdentifierValid(data, TRUE) then
    return;

// Now perform the priority drop
drop = VPriorityBitsSet(ICH_AP0R_EL2, ICH_AP1R_EL2);

if !drop then
    if ICH_VTR_EL2.SEIS == '1' then
        // Reporting of locally generated virtual SEIs is supported
        IMPLEMENTATION_DEFINED "SError EOI1_HIGHEST_IS_G0";
    return;

// It is IMPLEMENTATION DEFINED whether the priority is dropped before the error checks
if boolean IMPLEMENTATION_DEFINED "Drop before checks" then
    VPriorityDrop[ICH_AP0R_EL2, ICH_AP1R_EL2] = '0';
    dropped = TRUE;
else
    dropped = FALSE;

// Find the matching List Register
lrIndex = FindActiveVirtualInterrupt(eoiID);

if lrIndex < 0 then
    if IsLPI(eoiID) then
        // It is a virtual LPI not in the List Registers
        // so just priority drop and return without incrementing EOI count
        return;
    else
        // No valid list register corresponds to the EOI ID, increment EOI count
        if drop && ICH_VMCR_EL2.VEOIM == '0' then
            ICH_HCR_EL2.EOIcount = ICH_HCR_EL2.EOIcount + 1;
        return;

// Start error checks
// When an error is detected return to avoid unpredictable behaviour
if ICH_LR_EL2[lrIndex].Group != '1' then
    // The EOI ID is not Group 1
    if ICH_VTR_EL2.SEIS == '1' then
        // Reporting of locally generated virtual SEIs is supported
        IMPLEMENTATION_DEFINED "SError EOI1_HIGHEST_IS_G0";
    return;

if VPriorityGroup(ICH_LR_EL2[lrIndex].Priority, 1) != vPriority then
    // The EOI ID is not the highest priority
    if ICH_VTR_EL2.SEIS == '1' then
        // Reporting of locally generated virtual SEIs is supported
        IMPLEMENTATION_DEFINED "SError EOI1_NOT_HIGHEST_PRIORITY";
    return;
// End of error checks

if !dropped then VPriorityDrop[ICH_AP0R_EL2, ICH_AP1R_EL2] = '0';

if ICH_VMCR_EL2.VEOIM == '0' || IsLPI(eoiID) then
    setEI = ICH_LR_EL2[lrIndex].EOI == '1';

    // EOI mode not set, or it is an LPI and no deactivate is expected
    // so clear the active state in the List register
    if ICH_LR_EL2[lrIndex].HW =='1' then
        // Deactivate the physical interrupt
        pID = ICH_LR_EL2[lrIndex].PhysicalID;
        if UInt(pID) < 1020 then Deactivate(ZeroExtend(pID));
     else
        // Generate a maintenance interrupt if required
        if set EI then
            // Set the appropriate EISR bit to generate a maintenance interrupt
            ICH_EISR_EL2<lrIndex. = '1';
```

```
         //Clear the Active state
         if ICH_LR_EL2[lrIndex].State == IntState_ActivePending then
            ICH_LR_EL2[lrIndex].State = IntState_Pending;
         else
        ICH_LR_EL2[lrIndex].State = IntState_Invalid;
            if ICH_LR_EL2[lrIndex].HW == '1' || !setEI then
                ICH_ELRSR_EL2<lrIndex> = '1';

        return;
```

aarch64/support/CheckGroup0ForSpecialIdentifiers

```
// CheckGroup0ForSpecialIdentifiers()
// ==================================

bits(INTID_SIZE) CheckGroup0ForSpecialIdentifiers(bits(INTID_SIZE) pendID)

    if !IsGrp0Int(pendID) && PSTATE.EL != EL3 then
        // If the highest priority is Group 1, then no interrupt
        return INTID_SPURIOUS;

    if IsSecureInt(pendID) && !IsSecure() then
        // Secure interrupt not visible in Non-secure
        return INTID_SPURIOUS;

    if pendID != INTID_SPURIOUS && PSTATE.EL == EL3 then // An interrupt is pending
        if !IsGrp0Int(pendID) then
            if IsSecureInt(pendID) then              // Group 1 interrupt for the other state
                return INTID_SECURE;                 // Group 1 interrupt for Secure EL1 or IRQ/FIQ
            else
                return INTID_NONSECURE;              // Group 1 interrupt for Non-secure
        elsif ICC_CTLR_EL3.RM == '1' then
            return INTID_SECURE;                     // Group 0 Secure interrupt for Secure EL1

    return pendID;

// =================

bits(32) VirtualReadIAR0
()

    integer lrIndex = HighestPriorityVirtualInterrupt();

    if !CanSignalVirtualInterrupt() then
        return ZeroExtend(INTID_SPURIOUS);

    if (GICH_VLPIR.State == IntState_Pending &&
        (lrIndex < 0 || PriorityIsHigher(GICH_VLPIR.Priority, ICH_LR_EL2[lrIndex].Priority))) then
        // A virtual LPI is the highest priority
        vID = GICH_VLPIR.VirtualID<INTID_SIZE-1:0>;
        if GICH_VLPIR.Group == '0' then
            vPriorityGroup = VPriorityGroup(GICH_VLPIR.Priority);
            ICH_AP0R_EL2

<UInt(vPriorityGroup)> = '1';
            AcknowledgeVInterrupt(GICH_VLPIR.VirtualID<INTID_SIZE-1:0>);
            GICH_VLPIR.State = IntState_Invalid;         // Set the virtual LPI to Idle
        else
            vID = INTID_SPURIOUS;
        return ZeroExtend(vID);

    // lrIndex must be valid (i.e. non-negative)
    vID = ICH_LR_EL2[lrIndex].VirtualID<INTID_SIZE-1:0>;
    pID = ICH_LR_EL2[lrIndex].PhysicalID<INTID_SIZE-1:0>;

    if (vID != INTID_SPURIOUS &&
```

```
        (ICH_LR_EL2[lrIndex].Group != '0' || ICH_LR_EL2[lrIndex].State == IntState_Invalid)) then
        // If the highest priority isn't Group 0, then no interrupt
        return ZeroExtend(INTID_SPURIOUS);

    if !IsSpecial(vID) then                             // Check that it is not a spurious interrupt
        ICH_LR_EL2[lrIndex].State = IntState_Active;    // Set the list register state to Active
        vPriorityGroup = VPriorityGroup(ICH_LR_EL2[lrIndex].Priority);
        ICH_AP0R_EL2

<UInt(vPriorityGroup)> = '1';        // Set the corresponding bit in APR
    else
        ICH_LR_EL2[lrIndex].State = IntState_Invalid;

        setEI = ICH_LR_EL2[lrIndex].EOI == '1';

        // Generate a maintenance interrupt if required
        if ICH_LR_EL2[lrIndex].HW == '0' && setEI then
            // Set the appropriate EISR bit to generate a maintenance interrupt
            ICH_EISR_EL2<lrIndex> = '1';
        else
            // Set the appropriate ELRSR bit
            ICH_ELRSR_EL2<lrIndex> = '1';

    return ZeroExtend(vID);
```

aarch64/support/CheckGroup1ForSpecialIdentifiers

```
// CheckGroup1ForSpecialIdentifiers()
// ==================================

bits(INTID_SIZE) CheckGroup1ForSpecialIdentifiers(bits(INTID_SIZE) pendID)

    if IsGrp0Int(pendID) && PSTATE.EL != EL3 then
        // If the highest priority is Group 0 and not at EL3 then no interrupt
        return INTID_SPURIOUS;

    if UInt(pendID) != INTID_SPURIOUS then              // An enabled interrupt is pending
        if PSTATE.EL == EL3 && ICC_CTLR_EL3.RM == '1' then
            if !IsGrp0Int(pendID) then                  // Group 1 interrupt for Non-Secure EL1
                return INTID_NONSECURE;
            else                                        // Indicate a Group 0 interrupt is pending
                return INTID_SECURE;
        elsif !IsGrp0Int(pendID) then                   // IRQ is routed to EL1/2 or RM is zero
            if IsSecureInt(pendID) then                 // Group 1 Secure interrupt
                if !IsSecure() then
                    return INTID_SPURIOUS;              // Not visible in Non-secure
            elsif IsSecure() && PSTATE.EL != EL3 then       // Group 1 Non-secure interrupt
                return INTID_SPURIOUS;                  // Not visible at Secure EL1
        else
            return INTID_SPURIOUS;                      // Group 0 interrupt

    return pendID;

// ================

bits(32) VirtualReadIAR1
()

    integer lrIndex = HighestPriorityVirtualInterrupt();

    if !CanSignalVirtualInterrupt() then
        return ZeroExtend(INTID_SPURIOUS);

    if (GICH_VLPIR.State == IntState_Pending &&
        (lrIndex < 0 || PriorityIsHigher(GICH_VLPIR.Priority, ICH_LR_EL2[lrIndex].Priority))) then
        // A virtual LPI is the highest priority
        vID = GICH_VLPIR.VirtualID<INTID_SIZE-1:0>;
```

```
        if GICH_VLPIR.Group == '1' then
            vPriorityGroup = VPriorityGroup(GICH_VLPIR.Priority);
            ICH_AP1R_EL2

<UInt(vPriorityGroup)> = '1';
            AcknowledgeVInterrupt(GICH_VLPIR.VirtualID<INTID_SIZE-1:0>);
            GICH_VLPIR.State = IntState_Invalid;         // Set the virtual LPI to Idle
        else
            vID = INTID_SPURIOUS;
        return ZeroExtend(vID);

    // lrIndex must be valid (i.e. non-negative)
    vID = ICH_LR_EL2[lrIndex].VirtualID<INTID_SIZE-1:0>;
    pID = ICH_LR_EL2[lrIndex].PhysicalID<INTID_SIZE-1:0>;

    if ICH_LR_EL2[lrIndex].Group != '1' || ICH_LR_EL2[lrIndex].State == IntState_Invalid then
        // If the highest priority isn't Group 1, then no interrupt
        return ZeroExtend(INTID_SPURIOUS);

    if !IsSpecial(vID) then                              // Check that it is not a spurious interrupt
        ICH_LR_EL2[lrIndex].State = IntState_Active;     // Set the list register state to Active
        vPriorityGroup = VPriorityGroup(ICH_LR_EL2[lrIndex].Priority);
        ICH_AP1R_EL2

<UInt(vPriorityGroup)> = '1';        // Set the corresponding bit in APR
    else
        ICH_LR_EL2[lrIndex].State = IntState_Invalid;

        setEI = ICH_LR_EL2[lrIndex].EOI == '1';

        // Generate a maintenance interrupt if required
        if ICH_LR_EL2[lrIndex].HW == '0' && setEI then
            // Set the appropriate EISR bit to generate a maintenance interrupt
            ICH_EISR_EL2<lrIndex> = '1';
        else
            // Set the appropriate ELRSR bit
            ICH_ELRSR_EL2<lrIndex> = '1';

    return ZeroExtend(vID);
```

aarch64/support/PRIMask

```
// PRIMask()
// =========

bits(8) PRIMask()
    pri_bits = UInt(if HaveEL(EL3) then ICC_CTLR_EL3.PRIbits else ICC_CTLR_EL1.PRIbits);
    return Ones(pri_bits + 1):Zeros(7 - pri_bits);
```

aarch64/support/VRIMask

```
// VPRIMask()
// ==========

bits(8) VPRIMask()
    pri_bits = UInt(ICH_VTR_EL2.PRIbits);
    return Ones(pri_bits + 1):Zeros(7 - pri_bits);
```

8.20.2 Functions for memory-mapped registers

This subsection describes the functions that relate to the memory-mapped registers.The functions are indicated by the hierarchical path names, for example `shared/support`. The functions are:

- *shared/support/GICC_AIAR* on page 8-678.
- *shared/support/GICC_EOIR_NS* on page 8-678.

- *shared/support/GICC_EOIR_S*.
- *shared/support/GICC_IAR_NS* on page 8-679.
- *shared/support/GICC_IAR_S* on page 8-679.
- *shared/support/GICV_IAR* on page 8-679.

shared/support/GICC_AIAR

```
// GICC_AIAR[] - non-assignment form
// =================================

bits(32) GICC_AIAR[integer cpu_id]

    pendID = HighestPriorityPendingInterrupt(cpu_id);

    // If the highest priority isn't enabled then no interrupt
    if (!IsGrp0Int(pendID) && GICC_CTLR.EnableGrp1NS == '0') || IsGrp0Int(pendID) then
        pendID = INTID_SPURIOUS;

    if pendID != INTID_SPURIOUS then    // An enabled interrupt is pending
        if IsGrp0Int(pendID) then       // Highest priority is Secure
            pendID = INTID_SPURIOUS;

    // Check that it is not a spurious interrupt
    if !IsSpecial(pendID) then
        AcknowledgeInterrupt(pendID);   // Set active and attempt to clear pending

    return ZeroExtend(pendID);
```

shared/support/GICC_EOIR_NS

```
// GICC_EOIR_NS[] - assignment form
// ================================

GICC_EOIR_NS[integer cpu_id] = bits(32) data

    // Is the highest priority G0S, G1S or G1NS
    pGroup = GetHighestActiveGroup(GICC_APR0, GICC_APR1);
    pPriority = GetHighestActivePriority(GICC_APR0, GICC_APR1);

    if pGroup == IntGroup_None then            // There are no active interrupts
        IMPLEMENTATION_DEFINED "SError EOI1_NO_INTS_ACTIVE";

    elsif pGroup == IntGroup_G1NS && !IsSecure() then   // Non-secure Group 1
        // Drop the priority
        PriorityDrop(cpu_id, pPriority);
        // Deactivate the interrupt if EOI mode is not set
        if !EOImodeSet(cpu_id) then Deactivate(cpu_id, data<15:0>);

    else                                       // Group 0 or Secure Group 1 and access is Non-secure
        IMPLEMENTATION_DEFINED "SError EOI1_HIGHEST_NOT_ACCESSIBLE";

    return;
```

shared/support/GICC_EOIR_S

```
// GICC_EOIR_S[] - assignment form
// ===============================

GICC_EOIR_S[integer cpu_id] = bits(32) data

    // Is the highest priority G0S, G1S or G1NS
    pGroup = GetHighestActiveGroup(GICC_APR0, GICC_APR1);
    pPriority = GetHighestActivePriority(GICC_APR0, GICC_APR1);
```

```
    if pGroup == IntGroup_None then            // There are no active interrupts
        IMPLEMENTATION_DEFINED "SError EOI0_NO_INTS_ACTIVE";

    elsif pGroup == IntGroup_G0 && IsSecure() then    // Group 0 and the access is Secure
        // Drop the priority
        PriorityDrop(cpu_id, pPriority);
        // Deactivate the interrupt if EOI mode is not set
        if !EOImodeSet(cpu_id) then Deactivate(cpu_id, data<15:0>);

    elsif pGroup == 'IntGroup_G0' && !IsSecure() then   // Group 0 and the access is Non-secure
        if boolean IMPLEMENTATION_DEFINED "GICC_STATUSR implemented" then
            // Set the attempted security violation bit
            GICC_STATUSR.ASV = '1';

    else                                       // Group 1
        IMPLEMENTATION_DEFINED "SError EOI0_HIGHEST_IS_G1";

    return;
```

shared/support/GICC_IAR_NS

```
// GICC_IAR_NS[] - non-assignment form
// ===================================

bits(32) GICC_IAR_NS[integer cpu_id]

    pendID = HighestPriorityPendingInterrupt(cpu_id);

    // If the highest priority isn't enabled or is for the other security state then no interrupt
    if (!IsGrp0Int(pendID) && GICC_CTLR.EnableGrp1NS == '0') || IsGrp0Int(pendID) then
        pendID = INTID_SPURIOUS;

    // Check that it is not a spurious interrupt
    if !IsSpecial(pendID) then
        AcknowledgeInterrupt(pendID);   // Set active and attempt to clear pending

    return ZeroExtend(pendID);
```

shared/support/GICC_IAR_S

```
// GICC_IAR_S[] - non-assignment form
// ==================================

bits(32) GICC_IAR_S[integer cpu_id]

    pendID = HighestPriorityPendingInterrupt(cpu_id);

    // If the highest priority isn't enabled or is for the other security state then no interrupt
    if (IsGrp0Int(pendID) && GICC_CTLR.EnableGrp0 == '0') || !IsGrp0Int(pendID) then
        pendID = INTID_SPURIOUS;

    // Check that it is not a spurious interrupt
    if !IsSpecial(pendID) then
        AcknowledgeInterrupt(pendID);   // Set active and attempt to clear pending

    return ZeroExtend(pendID);
```

shared/support/GICV_IAR

```
// GICV_IAR[] - non-assignment form
// ================================

bits(32) GICV_IAR[integer cpu_id]

    lrIndex = HighestPriorityVirtualInterrupt(cpu_id);
```

```
vID = ICH_LR_EL2[lrIndex].VirtualID<INTID_SIZE-1:0>;

if ICH_LR_EL2[lrIndex].State == IntState_Invalid then
    vID = INTID_SPURIOUS;

if vID != INTID_SPURIOUS then
    if ICH_LR_EL2[lrIndex].Group == '1' then
        if GICV_CTLR.AckCtl == '1' then
            rval = ICV_IAR1_EL1[cpu_id];
        else
            rval = ZeroExtend(INTID_GROUP1);
    else
        rval = ICV_IAR0_EL1[cpu_id];
else
    rval = ZeroExtend(vID);

return rval;
```

Chapter 9
System Error Reporting

This chapter describes GIC support for System Error reporting. It contains the following section:

- *About System Error reporting on page 9-682*.

9.1 About System Error reporting

The GIC architecture provides optional support for locally-generated system error interrupts generated by the CPU interface. This support is IMPLEMENTATION DEFINED.

Whether a CPU interface supports locally-generated system error interrupts associated with physical interrupts is discoverable from either ICC_CTLR_EL1.SEIS or ICC_CTLR_EL3.SEIS. The GIC reports these using the ARMv8 SError exception. The ITS can also generate system errors, see the description of the GITS_TYPER.SEIS bit.

Whether the GIC supports locally-generated system error interrupts associated with virtual interrupts is discoverable from ICH_VTR_EL2.SEIS. The GIC reports these using either the SError exception or the virtual SError exception. Locally-generated System Error interrupts from Non-secure EL1 are reported:

- Using the SError exception when HCR_EL2.AMO == 0.
- Using the virtual SError exception, when HCR_EL2.AMO == 1. Where supported, a virtual SError exception is normally taken to Non-secure EL1. However, when HCR_EL2.AMO == 1, the hypervisor can intercept locally-generated system error interrupts using ICH_HCR_EL2.TSEI.

9.1.1 Pseudocode

The following pseudocode indicates whether a local system error is generated.

```
// GenerateLocalSError()
// =====================

boolean GenerateLocalSError()
   if HaveEL(EL3) then
       return ICC_CTLR_EL3.SEIS == '1';
   else
       return ICC_CTLR_EL1.SEIS == '1';
```

Chapter 10
Legacy Operation and Asymmetric Configurations

This chapter describes GIC support for legacy operation and asymmetric configurations. It contains the following sections:

10.1 Legacy support of interrupts and asymmetric configurations

Whether a GICv3 implementation includes a mechanism to support legacy operation of physical interrupts is IMPLEMENTATION DEFINED. Where supported, this mechanism is the same as in GICv2, with the following exceptions:

- GICC_CTLR.AckCtl is RAZ/WI, and separate registers must handle Group 0 and Group 1 physical interrupts.
- The GICv2 configuration lockdown feature and the associated **CFGSDISABLE** signal are not supported. GICD_TYPER.LSPI is RES0.
- For asymmetric operation, a routing modifier bit is used as part of the security context switch control mechanism that handles the highest priority pending interrupt. See *The asymmetric configuration* on page 10-688 for more information.

  In addition, software executing in Secure state in a system that is configured for asymmetric operation is not permitted to manage Non-secure interrupts:

  — When ICC_CTLR_EL3.RM == 1, it is a requirement that GICC_CTLR.FIQen == 1 or the behavior of Secure EL1 is UNPREDICTABLE.
- A hypervisor executing at EL2 can only control virtual interrupts for the PE that it is executing on, and cannot control virtual interrupts on other PEs.
- The individual enables for SGIs, GICD_ISENABLER<n> where n=0, always reset to zero.
- Interrupts that belong to a group that is disabled in GICD_CTLR cannot block interrupts that belong to a group that is enabled. This means that if the highest priority pending interrupt is in a group that is disabled, this does not prevent the GIC from forwarding interrupts that are in a group that is enabled to the CPU interfaces.

Note

Secure Group 1 interrupts are treated as Group 0 interrupts during legacy operation.

In GICv3, the following restrictions apply when the Non-secure state is using affinity routing and the Secure state is not using affinity routing:

- GICD_ITARGETSR<n> is RES0 for any SPI where affinity routing is enabled for the current Security state.

  Note

  — Legacy Secure software cannot re-route Non-secure interrupts because GICD_IROUTER<n> is inaccessible to Secure accesses, and might not be interpreted correctly.

  — Legacy Secure software can change the group of the interrupt.
- The mapping between the bit positions and the affinity is IMPLEMENTATION DEFINED, and is reported by GICR_TYPER.Processor_Number.
- If an SGI is generated in Non-secure state and GICD_CTLR.DS = 0 then a Group 0 SGI cannot be set as pending, irrespective of the value of GICD_NSACR<n>.
- If an SGI is generated in Secure state and routed using the Targeted list model, then the SGI must be delivered to those PEs whose number is indicated by the appropriate bit in GICR_TYPER.Processor_Number. When routed using the Targeted list model but excluding the originating PE, the SGI must be delivered to all PEs except the originating PE. This includes PEs with GICR_TYPER.Processor_Number >7.

  Note

  Software executing in Secure state that does not use affinity routing cannot use a Non-secure alias to GICD_SGIR to generate Non-secure SGIs, because this would result in a Non-secure write to GICD_SGIR, and GICD_SGIR is RAZ/WI when affinity routing is enabled for the Non-secure state.

When affinity routing is disabled for the Security state of an access, GICD_SGIR behaves as defined for GICv2, with the following exceptions:

- Writing to GICD_SGIR from a PE with GICR_TYPER.Processor_Number > 7 results in one of the following CONSTRAINED UNPREDICTABLE behaviors:
  - The write is ignored.
  - The originating PE ID is treated as having an UNKNOWN valid value.
- Writing to GICD_SGIR when the TargetListFilter field is 11 results in one of the following CONSTRAINED UNPREDICTABLE behaviors:
  - The write is ignored.
  - The TargetListFilter field is treated as having an UNKNOWN valid value.

In GICv2, pending SGIs were banked by the originating PE and by the target PE. In GICv3 this is simplified so that when affinity routing is enabled for a Security state, pending SGIs are only banked by the target PE:

- An originating PE ID is no longer provided when reading ICC_IAR0_EL1 or ICC_IAR1_EL1.
- A originating PE ID is no longer required when writing to an ICC_EOIR0_EL1 or ICC_EOIR1_EL1.
- Only 16 SGI pending bits are required for each Redistributor.

When the ARE bit in GICD_CTLR is set to 1 for a Security state, some Distributor registers that were banked for each PE are changed:

- GICD_SPENDSGIR<n>is RES0. In GICv3 SGIs are not pending by originating PE and the equivalent functionality is provided by GICR_ISPENDR0[0:15].
- GICD_CPENDSGIR<n> is RES0. In GICv3 SGIs are not pending by originating PE and the equivalent functionality is provided by GICR_ICPENDR0

GICD_SGIR is disabled when affinity routing is enabled for a Security state.

Writes to ICC_SGI0R_EL1, ICC_SGI1R_EL1, and ICC_ASGI1R_EL1 only generate SGIs for the other Security state when affinity routing is enabled for both Security states:

- When the Distributor supports two Security states, that is when GICD_CTLR.DS == 0, and affinity routing is disabled for the Secure state in the Distributor, then Non-secure writes to ICC_SGI0R_EL1 and ICC_ASGI1R_EL1 do not set any SGIs pending.
- When the Distributor supports only a single Security state, that is when GICD_CTLR.DS == 1, then Non-secure writes to both ICC_SGI0R_EL1 and ICC_ASGI1R_EL1 result in the generation of Group 0 SGIs.

For further information about the GICv2 architecture, see *ARM® Generic Interrupt Controller, Architecture version 2.0, Architecture Specification*.

10.1.1 Use of the special INTID 1022

INTID 1022 is only used for legacy operation, and is returned if all of the following conditions are true:

- The interrupt that is acknowledged is either:
  - A Secure read of GICC_IAR or GICC_HPPIR.
  - A Non-secure read of GICV_IAR or GICV_HPPIR.
- The highest priority pending interrupt is a Group 1 interrupt.
- For a read of GICV_IAR, GICV_CTLR.AckCtl == 0.
- The interrupt priority is sufficient for it to be signaled to the PE.

INTID 1022 indicates that there is a Group 1 interrupt of sufficient priority to be signalled to the PE, and that the interrupt must be acknowledged by a read of GICC_AIAR or GICV_AIAR, or observed by a read of GICC_AHPPIR or GICV_AHPPIR, as appropriate.

10.1.2 Legacy configurations

For physical interrupts, there are three possible configurations that can support legacy operation:

- GICD_CTLR.DS == 1, when the relevant ICC_SRE_EL3.SRE, ICC_SRE_EL2.SRE, and ICC_SRE_EL1.SRE are cleared to 0. In this case the GIC supports a single address space, and the behavior is the same as in GICv2 without the Security extensions.
- GICD_CTLR.DS == 0 and all of ICC_SRE_EL3.SRE, ICC_SRE_EL2.SRE, where implemented, and ICC_SRE_EL1.SRE are cleared to 0. In this case the GIC supports both Secure and Non-secure address spaces, and the behavior is the same as in GICv2 with the Security extensions.
- GICD_CTLR.DS == 0, and the system is using affinity routing for Non-secure physical interrupts. In this case, the Secure copy of ICC_SRE_EL1.SRE is cleared to 0. This configuration supports a legacy Secure operating system environment together with a Non-secure environment that supports affinity routing. This configuration is referred to as an *asymmetric configuration*.

Legacy operation is a deprecated feature. In an implementation that does not support legacy operation the following bits, where implemented, are RAO/WI:

- ICC_SRE_EL1.SRE.
- ICC_SRE_EL2.SRE.
- ICC_SRE_EL3.SRE.
- ICC_SRE.SRE.
- ICC_HSRE.SRE.
- ICC_MSRE.SRE.
- GICD_CTLR.ARE_NS.
- GICD_CTLR.ARE_S.

10.1.3 Legacy operation and bypass support

Interrupt bypass support during legacy operation is controlled using GICC_CTLR.

GICC_CTLR.{EnableGrp0, EnableGrp1} must have the value 0 when ICC_SRE_EL1.SRE == 1 and GICD_CTLR.DS == 1, otherwise GICv3 behavior is UNPREDICTABLE.

The following pseudocode defines the bypass behavior for an FIQ interrupt exception.

```
if GICC_CTLR.FIQEn == 0 then
  if (GICC_CTLR.FIQBypDisGrp0 && GICC_CTLR.FIQBypDisGrp1) == 0 then
    use BypassFIQsource
  else
    FIQ deasserted
else
  if GICC_CTRL.EnableGrp0 == 0 then
    if GICC_CTLR.FIQBypDisGrp0 == 0 then
      use BypassFIQsource
    else
      FIQ deasserted
   else
    use GICv3 FIQ output
```

The following pseudocode defines the bypass behavior for an IRQ interrupt exception.

```
if FIQEn == 0 then
  if (GICC_CTLR.EnableGrp1 || GICC_CTLR.EnableGrp0) == 0 then
    if (GICC_CTLR.IRQBypDisGrp0 && GICC_CTLR.IRQBypDisGrp1) == 0 then
      use BypassIRQsource
    else
      IRQ deasserted
  else
    use GICv3 IRQ Output
else
  if GICC_CTLR.EnableGrp1 == 0 then
    if GICC_CTLR.IRQBypDisGrp1 == 0 then
```

```
    use BypassIRQsource
  else
    IRQ Deasserted
else
  Use GICv3 IRQ Output
```

10.2 The asymmetric configuration

In a system that implements EL3,and where EL3 is using AArch64 state, the GIC architecture supports asymmetric configuration. A GICv3 system is configured for asymmetric operations when:

- GICD_CTLR.ARE_NS == 1.
- GICD_CTLR.ARE_S == 0.
- ICC_SRE_EL3.SRE == 1:
  - The Secure monitor is using System register access.
- If Secure EL1 is using AArch64 state, ICC_SRE_EL1(S).SRE == 0. If Secure EL1 is using AArch32 state, ICC_SRE(S).SRE == 0.
  - The Secure OS uses legacy GIC support.

For execution in Non-secure AArch64 state:

- If EL2 is implemented, ICC_SRE_EL2.SRE == 1.
- If EL2 is not implemented, ICC_SRE_EL1(NS).SRE == 1.

For execution in Non-secure AArch32 state:

- If EL2 is implemented, ICC_HSRE.SRE == 1 when EL2 is executing in AArch32 state. Otherwise, ICC_SRE_EL2 == 1.
- If EL2 is not implemented, ICC_SRE (NS).SRE== 1.

Note

If EL2 is implemented and using the System register interface, a vPE can access the memory-mapped interface.

When EL3 is using AArch64 state

The Secure Monitor software, executing at EL3 in AArch64 state, uses the System register interface.

The Secure OS, executing at Secure EL1 in either AArch32 state or AArch64 state, uses the legacy memory-mapped interface.

The Non-secure hypervisor or OS handling physical interrupts, executing at Non-secure EL2 or EL1 in either AArch32 state or AArch64 state, uses the System register interface.

When EL3 is using AArch32 state

Asymmetric operation is UNPREDICTABLE.

In this situation, ARM expects Group 0 interrupts to be handled by a Secure OS, and Non-secure Group 1 interrupts to be handled by the Non-secure hypervisor or OS.

Note

This situation is not compatible with the use of Secure Group 1 interrupts, as this concept is new in GICv3 and is therefore not understood by legacy Secure OS code.

10.2.1 Asymmetric operation and the use of ICC_CTLR_EL3.RM

ICC_CTLR_EL3.RM controls whether software executing at EL3 can acknowledge or observe Secure Group 0 and Non-secure Group 1 interrupts as the highest priority pending interrupt.

When ICC_CTLR_EL3.RM = = 1:

- Secure Group 0 interrupts return a special INTID value of 1020. This affects accesses to ICC_IAR0_EL1 and ICC_HPPIR0_EL1.
- Non-secure Group 1 interrupts return a special INTID value of 1021. This affects accesses to ICC_IAR1_EL1 and ICC_HPPIR1_EL1.

For more information about special INTIDs, see *Special INTIDs* on page 3-40.

10.3 Support for legacy operation of VMs

To support legacy operation for virtual interrupts, the GIC must support the GICV_* memory-mapped register interface. Whether this support is provided is IMPLEMENTATION DEFINED. All VM accesses to the GICD_* Distributor registers must trap to the hypervisor, which is responsible for running a virtual Distributor associated with the legacy VM.

The following constraints apply to virtual interrupts that are handled as part of legacy operation:

- The GICv2 configuration lockdown feature is not supported. This means that a hypervisor must virtualize GICD_TYPER.LSPI as a RAZ/WI bit to the scheduled legacy VM.
- A multiprocessing VM can support a maximum of eight vPEs, which is the maximum number of PEs that are supported in GICv2. These vPEs are independently associated with the same Redistributor or with different Redistributors.

Note

Legacy operation for virtual interrupts supports GICV_CTLR.AckCtl. Legacy operation for physical interrupts does not support GICV_CTLR.AckCtl.

During legacy operation, GICV_CTLR controls the signaling of interrupts by the CPU interface to the PE, as follows:

- GICV_CTLR.EnableGrp0 bit controls the signaling of Group 0 interrupts.
- GICV_CTLR.EnableGrp1 bit controls the signaling of Group 1 interrupts.

For detailed information about the control and configuration of Group 0 and Group 1 PPI, SGI, and SPI interrupts, and their virtualization during legacy operation, see *ARM® Generic Interrupt Controller, Architecture version 2.0, Architecture Specification*.

10.3.1 Accessing GIC virtual CPU interface registers using the memory-mapped register interface

The virtual CPU interface is in the Non-secure memory map. A hypervisor uses the Non-secure stage 2 address translations to ensure that the vPE cannot access other memory-mapped GIC registers.

Figure 10-1 on page 10-690 shows a GICv3 configuration executing in AArch64 state where:

- Affinity routing and System register access are enabled for Non-secure accesses, that is GICD_CTLR.ARE_NS == 1 and ICC_SRE_EL2.SRE == 1.
- Virtualization is supported, that is ICH_HCR_EL2.En == 1.
- EL1 is configured to support legacy operation, that is ICC_SRE_EL1(NS).SRE == 0.
- The PE is configured to handle virtual interrupts, using HCR_EL2.{IMO, FMO}.

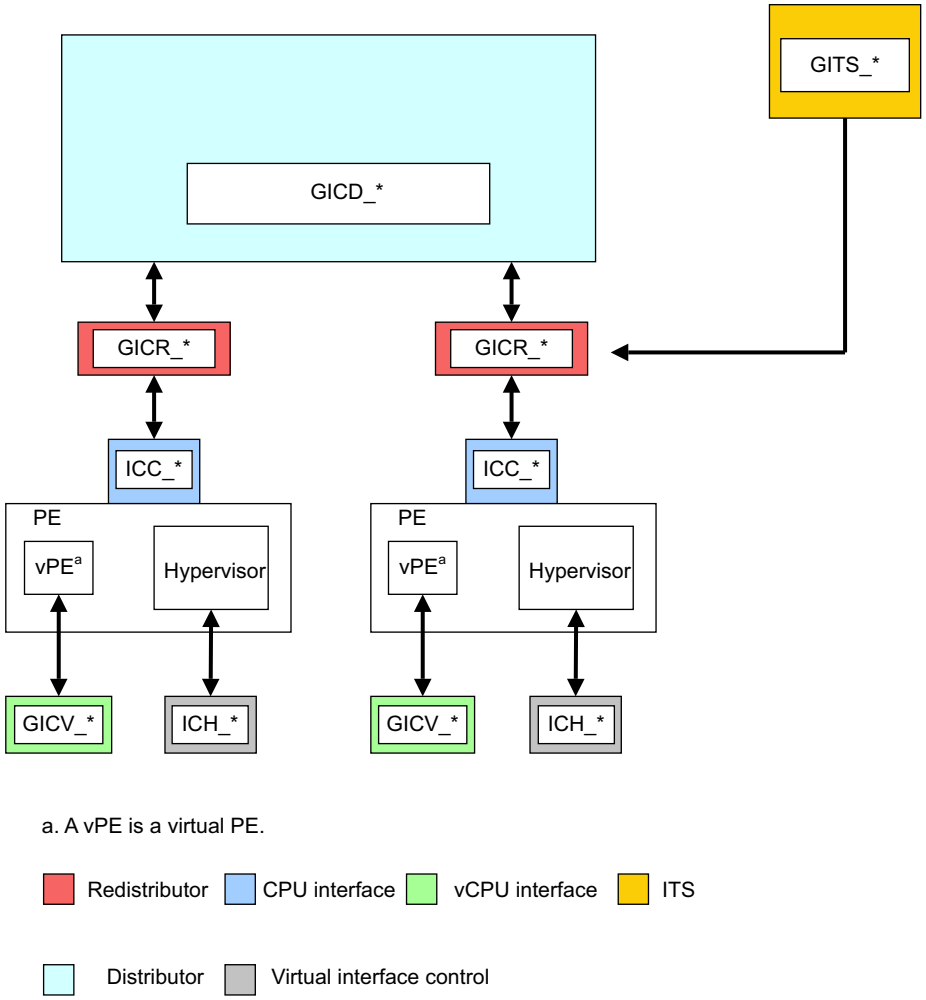

Figure 10-1 GICv3 register interfaces with legacy support

Appendix A
GIC Stream Protocol interface

This appendix describes the AXI4-Stream protocol standard message-based interface that the optional GIC Stream Protocol interface uses. It contains the following sections:

A.1 Overview

The GIC Stream Protocol interface describes the optional interface between the IRI and the PE, more specifically that between the Redistributor and the associated CPU interface. The interface supports independent development of an IRI and a PE, that includes System register support for the CPU interface. ARM recommends that a GIC implementation uses this stream protocol interface.

A communication channel that provides a packet interface, based on the AMBA 4 AXI-4 Stream Protocol, is required for each direction:

- From the Redistributor to the CPU interface.
- From the CPU interface to the Redistributor.

See *Signals and the GIC Stream Protocol on page A-693* for more information.

A.1.1 Terminology

The direction of communication for commands is referred to as downstream or upstream, where:

- Downstream is the direction associated with a command that is initiated by a Redistributor and sent to its associated CPU interface.
- Upstream is the direction associated with a command initiated by a CPU interface and sent to its associated Redistributor.

Note

This terminology can also be applied to communication within an IRI, that is between the Distributor and Redistributor. In this case:

- An upstream transfer is a transfer from a Redistributor to the Distributor.
- A downstream transfer is a transfer from the Distributor to a Redistributor.

A.2 Signals and the GIC Stream Protocol

The GIC Stream Protocol interface is based on the unidirectional AXI4-Stream Interface. Therefore, to support bidirectional communication, the GIC Stream Protocol interface consists of an AXI4-Stream Protocol Interface in each direction, that is:

- A downstream AXI4-Stream Interface containing connections from one or more Redistributors to an equivalent number of CPU interfaces. On this interface, the Redistributor is the master and the CPU interface is the slave.
- An upstream AXI4-Stream Interface containing connections from one or more CPU interface to an equivalent number of Redistributors. On this interface, the CPU interface is the master and the Redistributor is the slave.

Multiple packets on an AXI4-Stream Interface cannot be interleaved, that is, only one packet can be transferred in each direction at a given time.

Figure A-1 shows an example implementation of the GIC Stream Protocol interface.

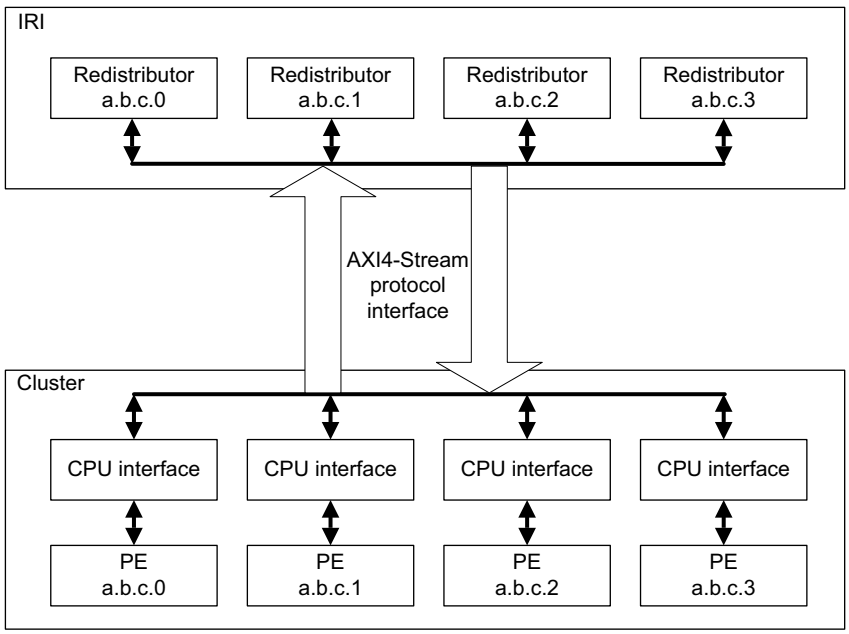

Figure A-1 Example of a GIC Stream Protocol interface

The GIC architecture requires a GIC implementation to include a Redistributor corresponding to each connected CPU interface, and defines an enumeration notation for identifying PEs. On any AXI4-Stream Interface, each Redistributor must only communicate with its corresponding CPU interface.

The *AMBA® 4 AXI4-Stream Protocol Specification* defines a packet as a group of bytes that are transported together across an AXI4-Stream interface.

An interconnect between an IRI and a CPU interface must ensure that the stream packet sequence is transferred over the stream protocol interface in the same order in which it was created.

A.2.1 Signals

The interface requires a global clock, **ACLK**, and a reset signal, **ARESETn**.

For the GIC Stream Protocol, each stream interface is identified by a prefix to the AXI-4 signal names:

- Downstream signals from a Redistributor to the CPU interface are prefixed with the letters **IRI**.
- Upstream signals from the CPU interface to a Redistributor are prefixed with the letters **ICC**.

Table A-1 shows the GIC Stream Protocol interface from the Redistributor to the downstream CPU interface.

Table A-1 Redistributor to downstream CPU interface

Signal [a]	Description
IRITVALID	When set to 1, this signal indicates that the master is driving a valid transfer.
IRITREADY	When set to 1, this signal indicates that the slave can accept a transfer in the current cycle.
IRITDATA[BN:0]	The interface data path.
IRITLAST	When set to 1, this signal indicates the final transfer of a packet.
IRITDEST[N:0]	When more than one PE is supported by the stream interface, this signal identifies the target CPU interface to provide routing information for the stream.Otherwise this signal is not required.

a. These signals were previously prefixed with **ICD** in the preliminary architecture information.

Table A-2 shows the GIC Stream Protocol interface from the CPU interface to the upstream Redistributor.

Table A-2 CPU interface to upstream Redistributor interface

Signal	Description
ICCTVALID	When set to 1, this signal indicates that the master is driving a valid transfer.
ICCTREADY	When set to 1, this signal indicates that the slave can accept a transfer in the current cycle.
ICCTDATA[BN:0]	The interface data path.
ICCTLAST	When set to 1, this signal indicates the final transfer of a packet.
ICCTID[N:0]	When more than one PE is supported by the stream interface, this signal identifies the originating CPU interface, to provide routing information for the stream. Otherwise this signal is not required.

In Table A-1 and Table A-2:

- BN is the number associated with the most significant bit on a datapath that is required to be an integral number of bytes wide.
- N is the value log(base2) M rounded up to the nearest integer, where M is the number of PEs supported by the interface.
- Values of **TDEST** and **TID** must be allocated sequentially without gaps, in order of ascending affinity.

For further information about the signals used by the GIC Stream Protocol interface, and for details about handshaking, see *AMBA® 4 AXI4-Stream Protocol Specification*.

A.2.2 Packet format

The GIC architecture issues packets across the GIC Stream Protocol interface where the initial half-byte of a packet indicates the packet type.

The declared size of a packet is always a multiple of the implemented datapath width used for the stream transfer. Where the number of bytes required by a packet is less than the overall packet size, the unused bytes are marked as reserved and filled with the value zero.

Supported INTID sizes

The GIC architecture supports 16-bit and 24-bit INTID fields. Where the INTID is an argument within a packet, the ID length field in the packet header defines which ID format is used, as follows:

- ID length == 0 for 16-bit INTIDs.
- ID length == 1 for 24-bit INTIDs.

Note

Reserved fields must be transmitted.

A downstream control command is used during interface initialization to inform a CPU interface whether the IRI supports 24-bit INTIDs.

For a 24-bit INTID where bits[23:16] have the value zero, the stream interface is allowed to identify and transfer the packet with a 16-bit INTID field.

A protocol error occurs when a PE generates a packet using a 24-bit INTID with nonzero bits[23:16] and the IRI only supports 16-bit INTIDs.

The Downstream Control Acknowledge command from the CPU interface returns the maximum INTID lengths supported by both the Redistributor and the CPU interface. The Redistributor and the CPU interface must not send a command that contains an INTID exceeding this length.

When both the Redistributor and the CPU interface support an INTID length larger than 16 bits, but the value of an INTID in a particular packet can only be encoded using 16 bits, it is permissible to send a 16-bit value, where ID length == `0b00`.

Software generation of protocol errors and packet errors

Software programming must never be permitted to cause a hardware *protocol error* because this might result in the PE or GIC becoming non-operational.

Where errors exist in the values of fields within packets, this is called a *packet error*. Protocol and packet errors can cause UNPREDICTABLE behavior.The manner in which these errors are reported is IMPLEMENTATION DEFINED.

Hardware generation of packet errors

If packets are sent that do not correspond to the architected format described in this specification, and the incorrect format is not the result of software misprogramming, then the architecture makes no guarantees about the behavior of the GIC. In such cases, the GIC can become non-operational in many ways, and this might result in the system hanging, data corruption, or any other effect. Physical damage to the system cannot be precluded in some implementations.

In high reliability systems, implementations might choose to report such cases using IMPLEMENTATION DEFINED system errors, but this is outside the scope of the architecture.

A.3 The GIC Stream Protocol

The GIC Stream Protocol supports two types of packet:

- Command packets for control actions.
- Response packets that acknowledge command packets.

Note

The Activate and Release commands are designated as a command, but also provide a response semantic to the Set and VSet commands. See the Activate and Release commands for more details.

Table A-3 shows a summary of the downstream Redistributor commands.

Table A-3 Redistributor commands

Command	ID	Parameters in the first 16-bit transfer	Data in subsequent transfers	Description
Clear	0x3	Bits[7:6]: ID length	INTID	Resets a specified pending physical interrupt.
Downstream Control	0x8	Bit[15:12]: Length Bits[11:4]: Identifier	Length bytes of data	Writes data to the CPU interface. Length must be greater than 0 and less than 9.
Quiesce	0x4	-	-	Requests that the CPU interface enters the quiescent state.
Set	0x1	Bits[15:8]: Priority Bits[7:6]: ID length Bit[5] GrpMod Bit[4]: Group	INTID	Sets the highest priority pending physical interrupt for a PE.
VClear	0x7	Bits[7:6]: ID length	Virtual INTID	Resets a specified pending virtual interrupt. This command is provided in GICv4 only.
VSet	0x6	Bits[15:8]: Priority Bits[7:6]: ID length Bit[4]: Group	Virtual INTID	Sets the highest priority pending virtual interrupt for a VM. This command is provided in GICv4 only.

Table A-4 shows a summary of the upstream Redistributor responses.

Table A-4 Redistributor responses

Response	ID	Parameters in the first 16-bit transfer	Data in subsequent transfers	Description
Activate Acknowledge	0xC	Bit[4]: V	-	Acknowledges that the Redistributor received an Activate command, and confirms that the effects of the activate are visible.

Table A-4 Redistributor responses (continued)

Response	ID	Parameters in the first 16-bit transfer	Data in subsequent transfers	Description
Deactivate Acknowledge	0xA	-	-	Acknowledges that the Redistributor received a Deactivate command, and confirms that the effects of the deactivate are visible.
Generate SGI Acknowledge	0x9	-	-	Acknowledges that the Redistributor received a Generate SGI command, and that the effects of the command are guaranteed to become visible to other PEs.
Upstream Control Acknowledge	0xB	-	-	Acknowledges receipt of an Upstream Control command, and confirms that the effects of the write operation are visible.

All other command and response IDs are reserved. If the Redistributor receives a reserved ID, this constitutes a protocol error, see *Software generation of protocol errors and packet errors on page A-695*.

Table A-5 shows a summary of all the CPU interface commands.

Table A-5 CPU interface commands

Command	ID	Parameters in the first 16-bit transfer	Data in subsequent transfers	Description
Activate	0x1	Bits[7:6]: ID length Bit[4]: V	INTID	A pending to active notification request as a result of an interrupt acknowledge on the CPU interface.
Deactivate	0x6	Bits[10:8]: Groups Bits[7:6]: ID length	INTID	Deactivate request for a specified interrupt.
Generate SGI	0x7	Bits[15:12] SGInum Bit[8]: A3V Bits[7]: IRM Bit[6]: NS Bit[5:4]: SGT	Affinity Routing Values (A0 to A3)	Requests that the Redistributor issues an SGI.
Upstream Control	0x8	Bits[15:12]: Length Bits[11:4]: Identifier	Length bytes of data	A system control command that might, for example, pass the configuration status to the Redistributor. Length must be greater than 0 and less than 9.

Table A-6 shows a summary of all the CPU interface responses.

Table A-6 CPU interface responses

Response	ID	Parameters in the first 16-bit transfer	Data in subsequent transfers	Description
Clear Acknowledge	0x4	Bit [4]: V	-	Acknowledges that the CPU interface received a Clear command for a specified interrupt.
Downstream Control Acknowledge	0xB	-	-	Acknowledges that the CPU interface received a Downstream Control command from the Redistributor.
Quiesce Acknowledge	0x9	-	-	Acknowledges a Quiesce command, and confirms that the Redistributor to CPU interface is in the quiescent state.
Release	0x3	Bits[7:6] ID length Bit[4]: V	INTID	Releases control of an interrupt when the CPU interface cannot handle the interrupt, and provides a reason for the release.

All other command and response IDs are reserved. If the CPU interface receives a reserved ID, this constitutes a protocol error, see *Software generation of protocol errors and packet errors on page A-695*.

A command packet has an equivalent handshake response packet that acknowledges the command. There are two exceptions to this rule:

- The Clear and VClear commands are both acknowledged by a Clear Acknowledge response with a bitfield in the header that indicates which command is acknowledged.
- The Set and VSet command packets are acknowledged by a Release response or an Activate command. This means that the Activate command also has the semantics of a response with respect to the Set and VSet commands. When an Activate command is used, the Redistributor acknowledges that command with an Activate Acknowledge response. Release, Activate, and Activate Acknowledge packets all have a bitfield in the header that indicates whether the response is to a Set or VSet command.

  Responses to a Set or VSet command are dependent on system contexts and events on the CPU interface. A Release response occurs when a pending interrupt that has been forwarded to the CPU interface cannot be maintained as pending or activated by the CPU interface. This can occur, for example, when:
  - The interrupt group of the INTID is disabled.
  - The highest pending physical interrupt is updated by a Set command before it is activated.
  - The highest pending virtual interrupt is updated by a VSet command before it is activated.

A.3.1 Rules associated with the downstream Redistributor commands

The following rules affect the generation of Redistributor commands:

- When GICR_WAKER.ProcessorSleep == 0, the first packet that is issued to the CPU interface must be a Downstream Control packet. This packet communicates the number of supported Security states, together with the physical and virtual INTID lengths that the GIC Stream Protocol interface supports.
- There can never be more than one outstanding Downstream Control command, and a Redistributor must only generate response packets until the Downstream Control command is acknowledged.
- On receipt of a Set command, a CPU interface is required to release the previous pending physical interrupt back to the Redistributor.
- Unless restricted by another rule in this section, two Set commands can be generated and outstanding at the same time, and the Redistributor must be able to accept an Activate command for a physical interrupt when a Set command is transferred.

- On receipt of a VSet command, a CPU interface is required to release the previous pending virtual interrupt back to the Redistributor
- A VSet command must only be generated when the Redistributor is able to accept an Activate command for a virtual interrupt while the VSet command is being transferred.
- There can never be more than one outstanding Clear command, and a Redistributor can generate further packets, other than Set or Clear packets, until the Clear command is acknowledged.
- There can never be more than one outstanding VClear command, and a Redistributor can generate further packets, other than VSet or VClear packets, until the VClear command is acknowledged.
- There can never be more than one outstanding Quiesce command, and a Redistributor must only generate response packets until the Quiesce command is acknowledged with a Quiesce Acknowledge.

A.3.2 Rules associated with the upstream CPU interface commands

The following rules affect the generation of CPU interface commands:

- There can never be more than one outstanding Upstream Control command, and a CPU interface must wait for an Upstream Control Acknowledge before issuing another Upstream Control command.
- There can never be more than one outstanding Deactivate command. This means that a CPU interface must wait for a Deactivate command to be acknowledged before issuing another Deactivate command. The CPU interface can continue to send other commands before receiving the Deactivate Acknowledge response.
- There can never be more than one outstanding Generate SGI command. This means that a CPU interface must wait for a Generate SGI command to be acknowledged before issuing another Generate SGI command. The CPU interface can continue to send other commands before receiving the Generate SGI Acknowledge response.
- Before issuing a Clear Acknowledge response with the V bit set to 0, the CPU interface must issue any Release commands that are required to move the physical interrupt specified in the Clear command to the inactive state on the CPU interface.
- Before issuing a Clear Acknowledge response with the V bit set to 1, the CPU interface must issue any Release commands that are required to move the virtual interrupt specified in the VClear command to the inactive state on the CPU interface.
- Before issuing a Quiesce Acknowledge response, all other outstanding commands from the Redistributor must be acknowledged, and a Release command must remove any pending interrupts on the CPU interface.

A.4 Alphabetic list of command and response packet formats

This subsection lists all the command and response packet formats in alphabetical order. The heading for each command or response subsection includes a label, ICC or IRI, that indicates the agent that generated the packet.

A.4.1 Activate (ICC)

The CPU interface sends an Activate command when acknowledging an interrupt. When the Redistributor receives the Activate command, it sets the interrupt to the active state. The CPU interface must send an Activate command only when Redistributor action is required, as follows:

- For SPIs, SGIs, and PPIs, where the Redistributor must clear the pending bit for edge-triggered interrupts, and set the active bit.
- For LPIs where the Redistributor must clear the pending bit.

The Activate command generated by the CPU interface, unlike other commands, also acts as a response to a Set or VSet command:

- A Set or VSet command results in a Release response or Activate command in finite time. The amount of time is determined by when the pending interrupt changes its state within the CPU interface. An Activate command acknowledges the original Set command. The Activate command is itself acknowledged using an Activate Acknowledge response from the Redistributor.

Figure A-2 shows the Activate command format.

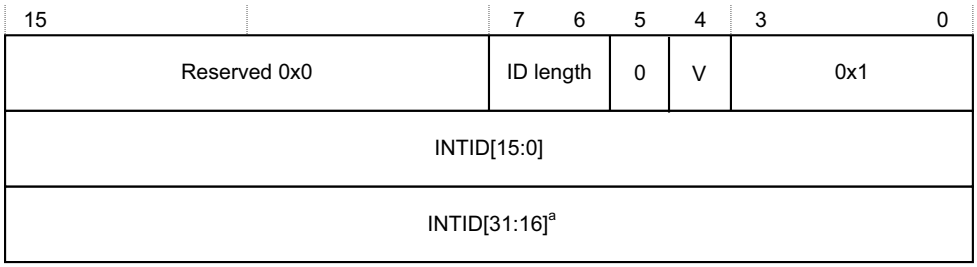

a. If the command includes this field, bits[31:24] are 0.

Figure A-2 Activate

In Figure A-2:

- ID length indicates the number of INTID bits the Activate command includes. See *Supported INTID sizes on page A-695* for more information.
- V indicates the original command the to which the Activate command corresponds:
  - **0** The Activate corresponds to a Set command.
  - **1** The Activate corresponds to a VSet command.
- INTID is the value that the CPU interface returns after a valid read of ICC_IAR0_EL1, ICC_IAR1_EL1, or GICC_IAR.

Note

During legacy operation, the INTID that is returned for SGIs includes the source PE in the GICC_IAR.Source_CPU_ID field.

A.4.2 Activate Acknowledge (IRI)

The Redistributor sends an Activate Acknowledge response to confirm receipt of an Activate command, and confirms that the effects of the activate operation are visible to the Redistributor and other PEs. Figure A-3 on page A-702 shows the Activate Acknowledge response format.

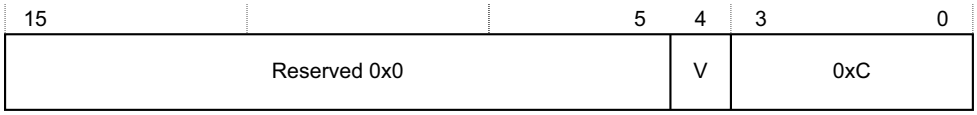

Figure A-3 Activate Acknowledge

In Figure A-3, V indicates the original command to which the Activate Acknowledge corresponds:

`0b0` The Activate Acknowledge corresponds to a Set command.

`0b1` The Activate Acknowledge corresponds to a VSet command.

Note

There is no requirement for ActivateAcknowledge commands to be issued in the same order as the Activate command to which they are responding.

A.4.3 Clear (IRI)

The Clear command clears the specified pending interrupt. Figure A-4 shows the Clear command format.

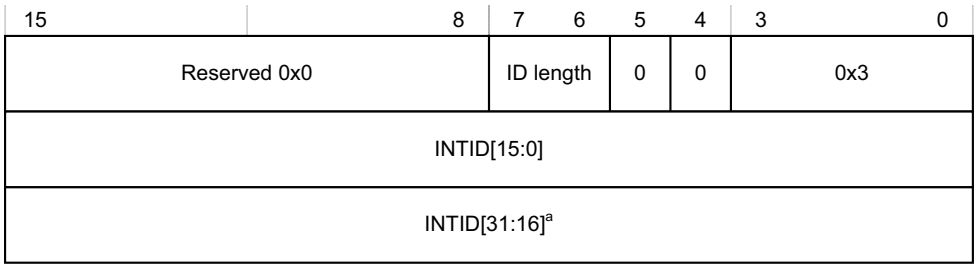

a. If the command includes this field, bits[31:24] are 0.

Figure A-4 Clear

In Figure A-4:

- ID length indicates the number of INTID bits that the Clear command includes. See *Supported INTID sizes on page A-695* for more information.
- INTID identifies the interrupt to be cleared.

The CPU interface must always respond to a Clear command with a Clear Acknowledge response where V == 0.

If the interrupt is pending in the CPU interface, the CPU interface must issue a Release response, or an Activate response that remains outstanding for the interrupt before it issues a Clear Acknowledge command.

If the interrupt is not pending or present on the CPU interface, the Clear command has no effect. However, the CPU interface must still issue a Clear Acknowledge response.

A.4.4 Clear Acknowledge (ICC)

The CPU interface sends a Clear Acknowledge response to acknowledge the receipt of a Clear or VClear command. Figure A-5 shows the Clear Acknowledge response format.

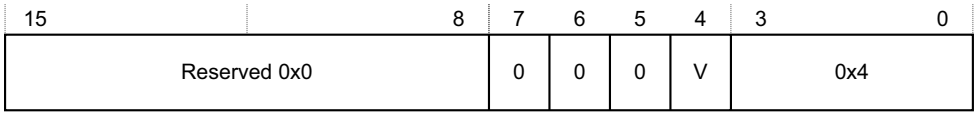

Figure A-5 Clear Acknowledge

In Figure A-5, V indicates the original command to which the Clear Acknowledge corresponds:

0 The Clear Acknowledge corresponds to a Clear command.

1 The Clear Acknowledge corresponds to a VClear command.

Note

No INTID field is required for this command because only a single Clear can be outstanding for a CPU interface at any time.

A.4.5 Deactivate (ICC)

The Deactivate command deactivates an interrupt, provided the initiating Exception level and Security state can access the interrupt group to which the INTID belongs. The Redistributor sends a Deactivate Acknowledge in response to a Deactivate command. Figure A-6 on page A-704 shows the Deactivate command format.

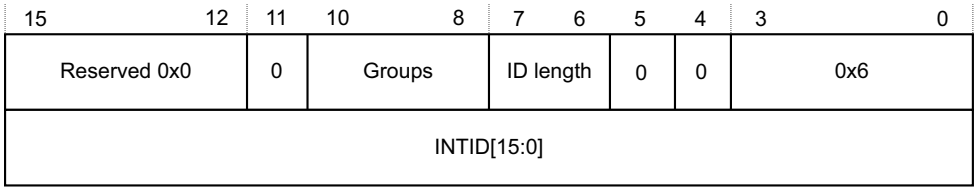

Figure A-6 Deactivate

In Figure A-6:

- `Groups` indicates the interrupt groups that the initiating Exception level and Security state are permitted to modify:

  Bit[10] When this bit is set to 1, Secure Group 1 interrupts can be modified.

  Bit[9] When this bit is set to 1, Non-secure Group 1 interrupts can be modified.

  Bit[8] When this bit is set to 1, Group 0 interrupts can be modified.

  Note

  When sending a Deactivate command, at least one of the Groups bits must be set to 1. A protocol error occurs if none of these bits are set to 1.

- `ID length` indicates the number of INTID bits the Deactivate command includes. See *Supported INTID sizes on page A-695* for more information. The Deactivate command applies only to SPIs, PPIs, and SGIs, each of which has INTIDs no higher than 8192. This field must therefore be set to `0b00`, indicating a 16-bit INTID.
- `INTID` is the 32-bit value read from the corresponding interrupt acknowledge cycle that is presented in the write to ICC_EOIR0_EL1 or ICC_EOIR1_EL1.

When System register access is enabled for the initiating Exception level and Security state, one of the Groups bits is set according to the rules in *Groups field when System register access is enabled on page A-705*.

Note

In an implementation that supports two Security states, for Secure EL1 to be permitted to handle Group 1 interrupts, that is, IRQs not taken to EL3, both bit[9] and bit[10] must be set to 1.

When System register access is not enabled for the initiating Exception level and Security state, the Groups field is set according to the Security state of the initiating Exception level. That is, bit [9] is set to 1 for Non-secure write access, and bits [10:8] are all set to 1 for Secure write access. In an implementation that supports only a single Security state, write accesses that result in the generation of a Deactivate command are treated as Secure writes.

In an implementation that supports two Security states, Group 0 and Secure Group 1 interrupts can be modified only from a Secure initiating Exception level. This includes EL3, regardless of the setting of SCR_EL3.NS. In an implementation that supports only a single Security state, the Redistributor can ignore bit[10].

Note

The Redistributor must send a Deactivate Acknowledge in response to a Deactivate command.

- If affinity routing is enabled for an interrupt group, the Redistributor must acknowledge, but otherwise ignore, any Deactivate command with an ID in the range 1019 < INTID < 8192.
- If affinity routing is not enabled for an interrupt group, the Redistributor must acknowledge, but otherwise ignore any Deactivate command with an ID in the range 1019 < INTID < 8192 where bits [9:4] are not 0. That is, it might issue Deactivate packets for SGIs with a non-zero CPU number in bits[12:10] of GICC_IAR.
- If affinity routing is not enabled for an interrupt group and the ID specifies an SGI, and the PE specified by the CPU number in bits[12:10] does not support operation when affinity routing is not enabled, the Redistributor must acknowledge but otherwise ignore the Deactivate command.

Groups field when System register access is enabled

When System register access is enabled for the initiating Exception level and Security state, the following pseudocode describes the rules for specifying the value of the Groups field:

```
// DeactivateGroups_SRE()
// ======================

(boolean,boolean,boolean) DeactivateGroups_SRE(bits(2) effective_EL)
    boolean groups_G0S  = FALSE;
    boolean groups_G1NS = FALSE;
    boolean groups_G1S  = FALSE;

    if effective_EL == EL3 then
        groups_G0S  = TRUE;
        groups_G1NS = TRUE;
        groups_G1S  = GICD_CTLR.DS == '0';

    elsif effective_EL == EL2 then
        // This also covers the case when the HW bit is one in a List Register
        // corresponding to a write at Non-secure EL1
        groups_G0S  = (!HaveEL(EL3) || SCR_EL3.FIQ == '0') && GICD_CTLR.DS == '1';
        groups_G1NS =  !HaveEL(EL3) || SCR_EL3.IRQ == '0';

    elsif effective_EL == EL1 && IsSecure() then
        // Secure EL1
        groups_G0S  =  !HaveEL(EL3) || SCR_EL3.FIQ == '0';
        groups_G1NS =  !HaveEL(EL3) || SCR_EL3.IRQ == '0';
        groups_G1S  = (!HaveEL(EL3) || SCR_EL3.IRQ == '0') && GICD_CTLR.DS == '0';

    elsif effective_EL == EL1 && !IsSecure() then
        // Non-secure EL1
        groups_G0S  = (!HaveEL(EL3) || SCR_EL3.FIQ == '0') && (!HaveEL(EL2) || HCR_EL2.FMO == '0') &&
GICD_CTLR.DS == '1';
        groups_G1NS = (!HaveEL(EL3) || SCR_EL3.IRQ == '0') && (!HaveEL(EL2) || HCR_EL2.IMO == '0');

    return (groups_G1S, groups_G1NS, groups_G0S);
```

If the Deactivate command relates to a virtual interrupt that has a corresponding physical interrupt in the List registers, that is ICH_LR<n>_EL2.HW is set to 1, a virtual write caused the deactivation of the physical interrupt.

The bits are set as if an equivalent write had been performed at EL2. That is, `effective_EL` == 2.

A.4.6 Deactivate Acknowledge (IRI)

The Redistributor sends a Deactivate Acknowledge response to confirm receipt of a Deactivate command, and to confirm that the effects of the deactivate operation are visible to the Redistributor and other PEs. Figure A-7 shows the Deactivate Acknowledge response format.

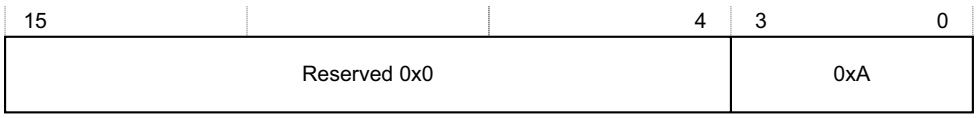

Figure A-7 Deactivate Acknowledge

A.4.7 Downstream Control (IRI)

The Downstream Control command transfers a specified number of bytes of data to the CPU interface. Figure A-8 on page A-706 shows the Downstream Control command format.

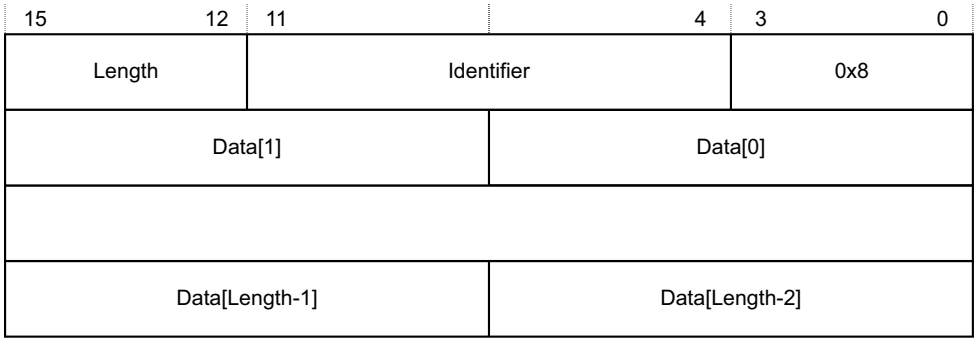

Figure A-8 Downstream Control

In Figure A-8:

- Length indicates the number of bytes of valid data appended to the 2 byte header.

  If this field specifies a number of bytes that is not exactly divisible by the interface width, as *Signals and the GIC Stream Protocol on page A-693* describes, any surplus bytes beyond this specified length in the last transfer must be zero. The CPU interface must ignore such bytes.

- Identifier is a value that specifies the format of the data provided, and can have the values shown in Table A-7.

Table A-7 Downstream Control Identifier values

Data value name	Identifier value	Length	Contents
Settings (configure interface)	`0x00`	`0x1`	Data[0] holds the Redistributor global settings, and these bits have the following meanings: **[7:6]** VL. Indicates the supported vINTID length. **[5:4]** PL. Indicates the supported pINTID length. **[3:1]** Reserved. RES0. **[0]** DS. Disable Security. Indicates the value of GICD_CTLR.DS. **Note** Bit[0] is set to 1 if the GIC supports only a single Security state.
Reserved	`0x01 - 0x7F`	-	-
IMPLEMENTATION DEFINED	`0x80 - 0xFF`	-	Reserved for IMPLEMENTATION DEFINED variables.

Note

Each identifier value can have a different length, but a particular identifier value must always have the same length.

The CPU interface must always respond to a Downstream Control command with a Downstream Control Acknowledge response.

After the CPU interface receives a Downstream Control command where DS == 1, a packet protocol violation occurs if it receives a subsequent Downstream Control command where DS == 0, before an intervening hardware reset.

If a CPU interface receives an IMPLEMENTATION DEFINED value that it cannot interpret, this constitutes a protocol error. See *Software generation of protocol errors and packet errors on page A-695*.

Note

The IMPLEMENTATION DEFINED values of the Downstream Write Command must only be used where the Distributor and the CPU interface interpret the IMPLEMENTATION DEFINED values to mean the same thing. This is typically the case where both components have been produced as part of the same system design.

A.4.8 Downstream Control Acknowledge (ICC)

The CPU interface sends a Downstream Control Acknowledge response to confirm receipt of a Downstream Control command. Figure A-9 shows the Downstream Control Acknowledge response format.

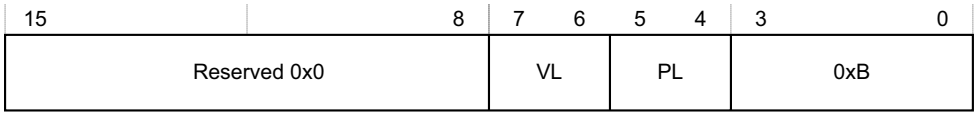

Figure A-9 Downstream Control Acknowledge

In Figure A-9:

- VL indicates the virtual INTID length, that is, the supported number of INTID bits.
- PL indicates the physical INTID length, that is, the supported number of INTID bits.

See *Supported INTID sizes on page A-695* for more information.

VL must be set to the minimum of:

- The value of VL contained in the first Downstream Control command received after reset.
- The value that ICH_VTR_EL2.IDbits specifies.

PL must be set to the minimum of:

- The value of PL contained in the first Downstream Control command received after reset.
- The value that ICC_CTLR_EL3.IDbits or ICC_CTLR_EL1.IDbits, as appropriate, specifies.

A.4.9 Generate SGI (ICC)

The CPU interface sends a Generate SGI command to the Redistributor to generate an SGI. The Redistributor sends a Generate SGI Acknowledge in response to a Generate SGI command. Figure A-10 shows the Generate SGI command format.

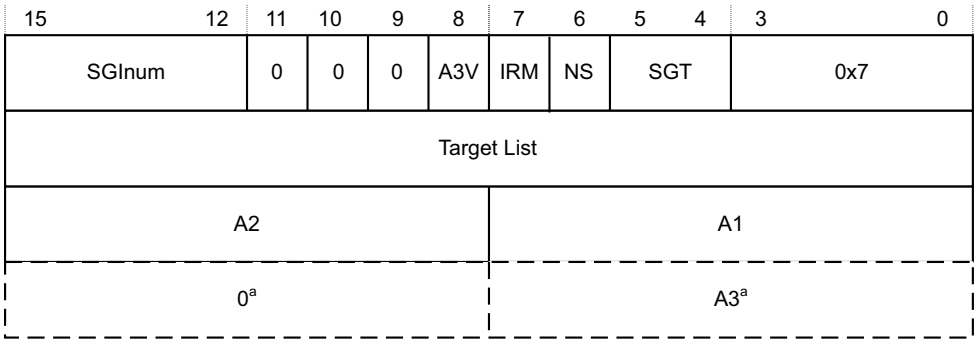

a. Whether this part of the packet is transmitted depends on the value of A3V.

Figure A-10 Generate SGI

In Figure A-10:

- SGInum indicates the INTID of the SGI to be generated.

- A3V indicates whether the command includes an A3 field. When A3V is 0, the packet does not include an A3 field, and the Redistributor must use 0 as the value of A3.When a CPU interface supports the Aff3 field and a write to ICC_SGI0R, ICC_SGI1R or ICC_ASGI1R specifies Aff3 == 0, the resulting packet must clear A3V to zero.
- IRM indicates the Interrupt Routing Mode to be used.When IRM is set to 1, Target List, A1, A2, and A3 are ignored. The A3V field is RES0.
- NS indicates whether the Generate SGI command originates from Non-secure state:

  0 The command originates from a Secure Execution state.

  1 The command originates from a Non-secure Execution state.
- SGT specifies the register access that caused the Generate SGI command:

  `0b00` ICC_SGI0R_EL1.

  `0b01` ICC_SGI1R_EL1.

  `0b10` ICC_ASGI1R_EL1.

  `0b11` Reserved.

  When the Redistributor supports two Security states and affinity routing is not enabled for the Secure state in the Redistributor, Generate SGI commands that correspond to Non-secure writes to ICC_SGI0R_EL1 and ICC_ASGI1R_EL1 must be acknowledged and discarded, and must not set an SGI pending.

  When the Redistributor supports a single Security state, that is, GICD_CTLR.DS == 1, Generate SGI commands that correspond to Non-secure writes to ICC_SGI0R_EL1 or ICC_ASGI1R_EL1 generate a Group 0 SGI.
- Target List is the group of target PEs defined by the routing mode. For SGIs, the GIC routing mode defines a group of target PEs, `targetlist`. This field is treated as defined in ICC_SGI0R_EL1, ICC_SGI1R_EL1, and ICC_ASGI1R_EL1.
- A1, A2, and A3 are the affinity level values used for generating the set of target PEs. These fields are treated in the same way as the Affinity value fields in ICC_SGI0R_EL1, ICC_SGI1R_EL1, and ICC_ASGI1R_EL1 Whether the A3 field is supported is IMPLEMENTATION DEFINED.

  Note

  In systems where the Redistributor only supports the zero value for A3, the Redistributor must acknowledge any Generate SGI commands where A3V == 1 with a Generate SGI Acknowledge response, but must otherwise ignore the command.

A.4.10 Generate SGI Acknowledge (IRI)

The Redistributor sends a Generate SGI Acknowledge response to confirm that is has received a Generate SGI command from the CPU interface, and that the effects of that command are guaranteed to become visible to other PEs. Figure A-11 shows the Generate SGI Acknowledge response format.

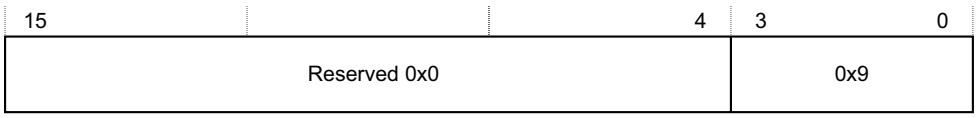

Figure A-11 Generate SGI Acknowledge

Note

Receipt of a Generate SGI Acknowledge response by a CPU interface does not guarantee that the corresponding SGI pending state is set, but it does guarantee that the pending state will become set.

A.4.11 Quiesce (IRI)

The Redistributor sends a Quiesce command to request that the CPU interface enters the quiescent state. Figure A-12 shows the Quiesce command format.

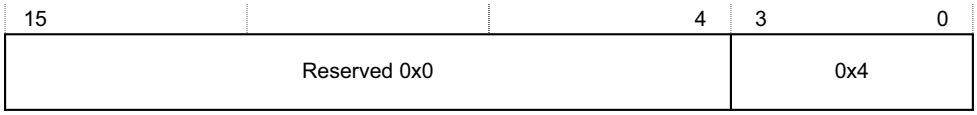

Figure A-12 Quiesce

A CPU interface is quiescent when there are no pending interrupts and all outstanding operations are complete. To ensure quiescence, a CPU interface must:

- Respond to any outstanding Clear and VClear commands by sending a Clear Acknowledge command.
- Release any pending virtual or physical interrupts.
- Ensure it receives an acknowledge response from the Redistributor to indicate completion of all outstanding:
  - Generate SGI requests.
  - Activate requests.
  - Deactivate requests.
  - Upstream Control.
- Respond to the Quiesce commands by sending a Quiesce Acknowledge response as the final transfer.

In addition, software must ensure that the Redistributor receives no traffic after the CPU interface sends the Quiesce Acknowledge response. Failure to adhere to this results in UNPREDICTABLE behavior. In practice, because such timing is not predictable, software must ensure that no traffic is generated after the GICR_WAKER.ProcessorSleep bit is set to 1, see Chapter 7 *Power Management*.

A CPU interface cannot receive a Quiesce command if a Downstream Control Acknowledge response is outstanding. See *Rules associated with the downstream Redistributor commands on page A-698* for more information.

A.4.12 Quiesce Acknowledge (ICC)

The CPU interface sends a Quiesce Acknowledge response to confirm receipt of a Quiesce command, and to confirm that it is quiescent. Figure A-13 shows the Quiesce Acknowledge response format.

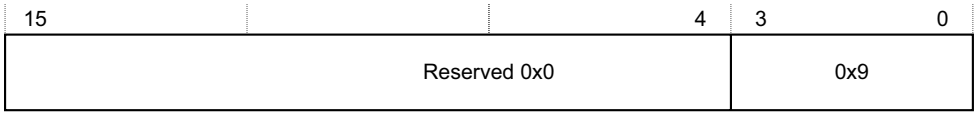

Figure A-13 Quiesce Acknowledge

The Quiesce command acts as a form of barrier. Before sending a Quiesce Acknowledge response, the CPU interface must be quiescent, that is, it must fulfil the requirements for quiescence specified in *Quiesce (IRI)*.

A.4.13 Release (ICC)

The CPU interface logic sends a Release response when it cannot handle a particular interrupt. Figure A-14 on page A-710 shows the Release response format.

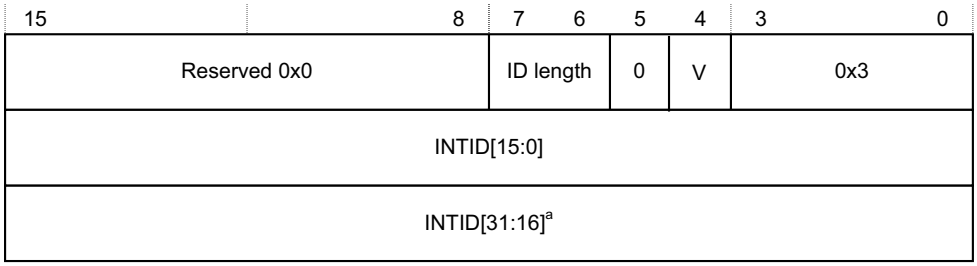

a. If the command includes this field, bits[31:24] are 0.

Figure A-14 Release

In Figure A-14:

- ID length indicates the number of INTID bits the Release response includes. See *Supported INTID sizes on page A-695* for more information.
- V indicates the original command to which the Release response corresponds:
  - **0** The Release corresponds to a Set command.
  - **1** The Release corresponds to a VSet command.
- INTID is the value that the CPU interface returns after a valid read of ICC_IAR0_EL1 or ICC_IAR1_EL1.

Note

— During legacy operation, the INTID that is returned for SGIs includes the source PE in the GICC_IAR.Source_CPU_ID field.

— If the INTID corresponds to an interrupt that uses the 1 of N model, the Redistributor might forward the interrupt to a different PE or it might send the interrupt to the same PE again. See ICC_CTLR_EL3 for information about how the PMHE field might affect the 1 of N selection.

If the CPU interface issues a Release response as a result of disabling an interrupt group, ARM recommends that it sends the Upstream Control command that contains the revised interrupt group enable information before issuing the Release response.

A.4.14 Set (IRI)

The Set command sets the highest priority pending interrupt for a PE. The PE has control of the interrupt and might respond to a read of ICC_IAR0_EL1 or ICC_IAR1_EL1 with the INTID. Figure A-15 shows the Set command format.

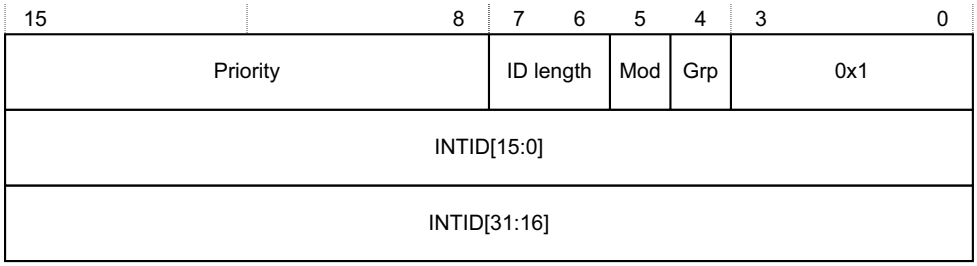

Figure A-15 Set

In Figure A-15:

- Priority indicates the actual priority of the interrupt, that is, the Secure, unshifted view. Bits corresponding to unimplemented priority bits in the CPU interface are RES0.
- ID length indicates the number of INTID bits the Set command includes. See *Supported INTID sizes on page A-695* for more information.

- Mod represents the value of the GICD_IGRPMODR<n>.Group status bit for the interrupt.
- Grp represents the interrupt group, as indicated by the corresponding GICD_IGROUPR<n>.Group status bit.
- INTID is the value that the CPU interface returns after a valid read of an ICC_IAR0_EL1 or ICC_IAR1_EL1.

Note

During legacy operation, the INTID that is returned for SGIs includes the source PE in the GICC_IAR.Source_CPU_ID field.

If the Redistributor sends a Set command, the interrupt specified in the command replaces any outstanding highest pending interrupt, that is, the command sets a new highest priority pending interrupt. Where a pending interrupt is replaced, the CPU interface must Release it back to the Redistributor.

The Redistributor must:

- Ensure that no more than two Set commands that are waiting for a response are outstanding per PE at any time.
- Send a Set command only if it can accept an Activate command where V == 0.

Note

An implementation can guarantee this by treating the Set command as outstanding until either a Release command is received for the Set command, or an Activate Acknowledge response is sent for the corresponding Activate.

- Never send a Set command when any of the following conditions apply:
  - The INTID is a special interrupt number, that is, 1020-1023.
  - Affinity routing is enabled for an interrupt group and 1023 < INTID < 8192.
  - Affinity routing is not enabled for an interrupt group, 1023 < INTID < 8192, and bits[9:4] are non-zero. That is, the Redistributor is permitted to send Set commands for SGIs where bits[12:10] of ICC_IAR0_EL1 or ICC_IAR1_EL1 specify the CPUID of the source PE.
  - Affinity routing is not enabled for an interrupt group, and INTID > 8191
  - The Set command has the same INTID as a previous Set command, unless the Redistributor has received an Activate command or Release response.

Note

The Redistributor must not send a Set command for the same interrupt again until the interrupt is no longer active or is released. Where INTID < 8192, typically this requires the CPU interface to send a Deactivate command, although clearing the active state explicitly in the Redistributor is an alternative method.

If the interrupt group is disabled, the CPU interface cannot handle the interrupt, and must Release the interrupt.

A.4.15 Upstream Control (ICC)

This command communicates data to the Redistributor. Figure A-16 shows the Upstream Control command format.

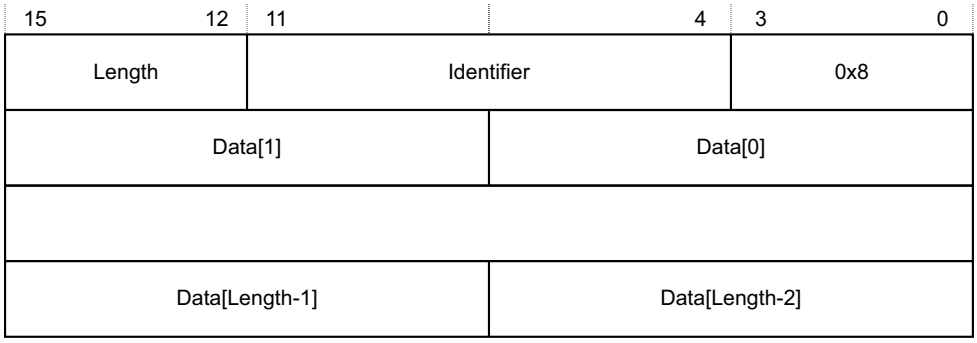

Figure A-16 Upstream Control

In Figure A-16 on page A-711:

- Length indicates the number of bytes of valid data.

  If this field specifies a number of bytes that is not exactly divisible by the interface width, as *Signals and the GIC Stream Protocol on page A-693* describes, any surplus bytes beyond this specified length in the last transfer must be zero. The Redistributor must ignore such bytes.

- Identifier is a value that specifies the format of the data provided.

Table A-8 shows the possible Identifier values.

Table A-8 Upstream Control Identifier values

Data value	Identifier	Length	Contents of Data[0] field
Physical interface enables	0x00	0x1	This value contains the physical CPU interface enable bit values that must be communicated to the Redistributor. Bits [2:0] of Data[0] have the following meanings: **[2]** EnableGrp1, secure. The value of the Secure copy of ICC_IGRPEN1_EL1.Enable. **[1]** EnableGrp1, Non-secure. The value of the Non-secure copy of ICC_IGRPEN1_EL1.Enable. **[0]** EnableGrp0, Secure. The value of ICC_IGRPEN0_EL1.Enable. For PEs that do not include EL3, or when the GIC supports only a single Security state, see the individual register descriptions for more information about the value of these bits. To ensure the state of the enable bits can be communicated easily to the Redistributor after powerup, this command must be generated by any write to a physical enable bit. If multiple writes to a physical enable bit occur before the CPU interface issues the command, the GIC can combine these writes into a single command.
Virtual interface enables	0x01	0x1	This value contains the virtual CPU interface enable bit values that must be communicated to the Redistributor. Bits [1:0] of Data[0] have the following meanings: **[1]** EnableGrp1. The value of ICH_VMCR_EL2.VENG1. **[0]** EnableGrp0. The value of ICH_VMCR_EL2.VENG0. To ensure the state of the enable bits can be communicated easily to the Redistributor after powerup, this command must be generated by any write to a virtual enable bit. If multiple writes to a virtual enable bit occur before the CPU interface issues the command, the GIC can combine these writes into a single command. **Note** If EL2 accesses memory-mapped registers, and uses GICH_VMCR, the VM must access GICV_* registers. If the GIC shares state between the GICH_* registers and the ICH_* System registers, it might communicate any change to the virtual enable bits.

Table A-8 Upstream Control Identifier values (continued)

Data value	Identifier	Length	Contents of Data[0] field
Physical priority	0x02	0x1	This value contains the current value of the Priority Mask Register (PMR): **[7:0]** The value written to ICC_PMR_EL1. The CPU interface must issue this command when the PE successfully writes to ICC_PMR_EL1 and ICC_CTLR_EL3.PMHE bit is set to 1. The command must be generated by any successful write that changes the value of ICC_PMR_EL1. If multiple writes to ICC_PMR_EL1 occur before the CPU interface issues the command, the GIC can combine these writes into a single command. **Note** • In GIC implementations that use this value, the Redistributor copy of the value must reset to the idle priority, that is, 0xF8 in cases where only 5 bits of priority are implemented. • If the CPU interface receives a Set command with a priority lower than the current value in ICC_PMR_EL1 before the Upstream Control Acknowledge is received, the GIC might Release that Set command.
-	0x03 - 0x07	-	Reserved.
-	0x80 - 0xFF	-	IMPLEMENTATION DEFINED

If a Redistributor receives an IMPLEMENTATION DEFINED value that it cannot interpret, this constitutes a protocol error. See *Software generation of protocol errors and packet errors on page A-695*.

A.4.16 Upstream Control Acknowledge (IRI)

The Redistributor sends an Upstream Control Acknowledge response to confirm receipt of an Upstream Control command, and to confirm that the effects of the write operation are visible to the Distributor. Figure A-17 shows the Upstream Control Acknowledge response format.

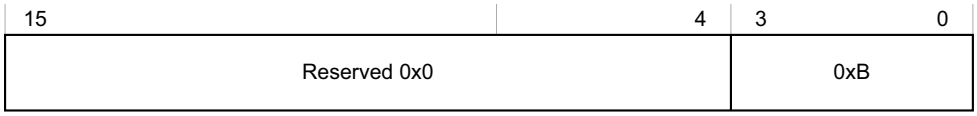

Figure A-17 Upstream Control Acknowledge

A.4.17 VClear (IRI)

The VClear command resets the highest priority pending virtual interrupt.

This command is provided in GICv4 only.

Figure A-18 shows the VClear command format.

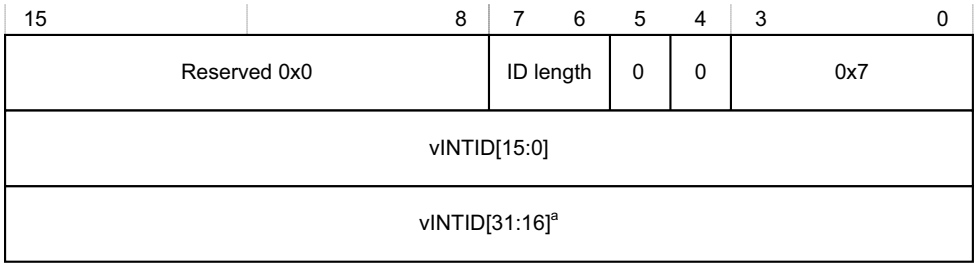

a. If the command includes this field, bits[31:24] are 0.

Figure A-18 VClear

In Figure A-18:

- ID length indicates the number of vINTID bits the VClear command includes. See *Supported INTID sizes on page A-695* for more information.
- vINTID identifies the virtual interrupt to be cleared.

The CPU interface must always respond to a VClear command by sending a Clear Acknowledge response where V==1.

If the interrupt is pending in the CPU interface, the CPU interface must issue a Release response, or an Activate response that remains outstanding for the interrupt before it issues a Clear Acknowledge command.

If the interrupt is not pending or present on the CPU interface, the VClear command has no effect. However, the CPU interface must still issue a Clear Acknowledge response.

Note

This command does not affect LPIs in the List registers.

A.4.18 VSet (IRI)

The Redistributor sends a VSet command to set a virtual interrupt pending for a VM. The PE has control of the interrupt and can respond to a read of ICC_IAR0_EL1 or ICC_IAR1_EL1 with the vINTID.

This command is provided in GICv4 only.

Figure A-19 on page A-715 shows the VSet command format.

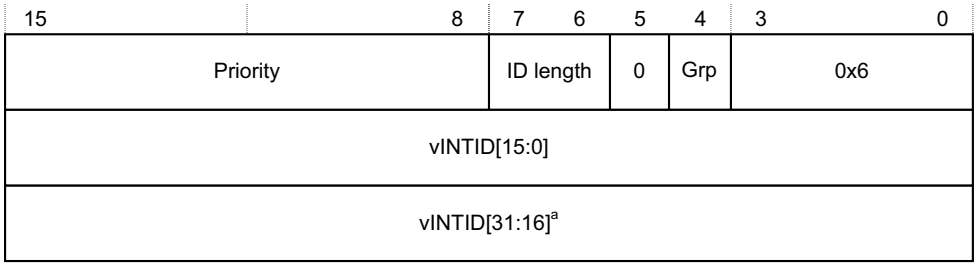

a. If the command includes this field, bits[31:24] are 0.

Figure A-19 VSet

In Figure A-19:

- Priority indicates the actual priority of the interrupt, that is, the Secure, unshifted view. Bits corresponding to unimplemented priority bits in the CPU interface are RES0.
- ID length indicates the number of v INTID bits the VSet command includes. See *Supported INTID sizes on page A-695* for more information.
- Grp represents the interrupt group.
- vINTID is the value that the CPU interface returns after a valid read of ICC_IAR0_EL1 or ICC_IAR1_EL1. When affinity routing is not enabled for a Security state, the CPUID field in ICC_IAR0_EL1 and ICC_IAR1_EL1 identifies the source PE for SGIs.

The Redistributor sends a VSet command when the virtual interrupt specified by vINTID is set as pending in the resident virtual LPI Pending table. The CPU interface must either activate the virtual interrupt by sending an Activate command where V == 1, or Release the virtual interrupt to the Redistributor.

If the Redistributor sends a VSet command, the interrupt specified in the command always replaces any previous interrupt, that is, the command sets a new highest priority pending interrupt. If the replaced interrupt is still valid and pending, the CPU interface must Release it back to the Redistributor.

The Distributor must:

- Ensure no more than two VSet commands that are waiting for a response are outstanding per PE at any time.
- Send a VSet command only if it can accept an Activate command where V == 1.

  Note

  An implementation can guarantee this by treating the VSet command as outstanding until either a Release response is received for the VSet command, or an Activate Acknowledge response is sent for the corresponding Activate command.

- Never send a VSet command when Virtual INTID < 8192.
- Send a VSet command for an interrupt after receipt of either an Activate or a Release command for that particular interrupt.

The CPU interface must Release an interrupt, ensuring that V == 1, if it cannot handle the interrupt for either of the following reasons:

- The interrupt group is disabled. This includes when the VM interface is disabled, that is, when GICH_HCR.En or ICH_HCR.En, as appropriate, is cleared to 0.
- The hypervisor is not using the System register interface, that is, when either of the following applies:
  - During legacy operation, when ICC_SRE_EL2.SRE == 0.
  - EL2 is not present.

When the Non-secure copy of ICC_SRE_EL1.SRE == 0, it is UNPREDICTABLE whether the specified virtual interrupt is factored into virtual priority calculations and reads of the GICV_* registers.

Appendix B
Pseudocode Definition

This appendix provides a definition of the pseudocode used in this specification, and lists the helper procedures and support functions used by pseudocode to perform useful architecture-specific jobs. For functions that are referenced in this specification but that are not defined in this appendix, see *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*.

This appendix contains the following sections:

B.1 About ARM pseudocode

ARM pseudocode provides precise descriptions of some areas of the architecture. The following sections describe the pseudocode in detail:

- *Data types* on page B-719.
- *Expressions* on page B-723.
- *Operators and built-in functions* on page B-725.
- *Statements and program structure* on page B-730.

Miscellaneous helper procedures and support functions on page B-735 describes some pseudocode helper functions, that are used by the pseudocode functions that are described elsewhere in this document.

B.1.1 General limitations of ARM pseudocode

The pseudocode statements `IMPLEMENTATION_DEFINED`, `SEE`, `UNDEFINED`, and `UNPREDICTABLE` indicate behavior that differs from that indicated by the pseudocode being executed. If one of them is encountered:

- Earlier behavior indicated by the pseudocode is only specified as occurring to the extent required to determine that the statement is executed.
- No subsequent behavior indicated by the pseudocode occurs. This means that these statements terminate pseudocode execution.

For more information, see *Simple statements* on page B-730.

B.2 Data types

This section describes:

- *General data type rules*.
- *Bitstrings*.
- *Integers* on page B-720.
- *Reals* on page B-720.
- *Booleans* on page B-720.
- *Enumerations* on page B-720.
- *Lists* on page B-721.
- *Arrays* on page B-722.

B.2.1 General data type rules

ARM architecture pseudocode is a strongly-typed language. Every constant and variable is of one of the following types:

- Bitstring.
- Integer.
- Boolean.
- Real.
- Enumeration.
- List.
- Array.

The type of a constant is determined by its syntax. The type of a variable is normally determined by assignment to the variable, with the variable being implicitly declared to be of the same type as whatever is assigned to it. For example, the assignments `x = 1`, `y = '1'`, and `z = TRUE` implicitly declare the variables `x`, `y`, and `z` to have types integer, bitstring of length 1, and Boolean, respectively.

Variables can also have their types declared explicitly by preceding the variable name with the name of the type. This is most often done in function definitions for the arguments and the result of the function.

The remaining subsections describe each data type in more detail.

B.2.2 Bitstrings

A bitstring is a finite-length string of 0s and 1s. Each length of bitstring is a different type. The minimum permitted length of a bitstring is 1.

The type name for bitstrings of length `N` is `bits(N)`. A synonym of `bits(1)` is `bit`.

Bitstring constants are written as a single quotation mark, followed by the string of 0s and 1s, followed by another single quotation mark. For example, the two constants of type `bit` are `'0'` and `'1'`. Spaces can be included in bitstrings for clarity.

A special form of bitstring constant with `'x'` bits is permitted in bitstring comparisons. See *Equality and non-equality testing* on page B-725.

Every bitstring value has a left-to-right order, with the bits being numbered in standard *little-endian* order. That is, the leftmost bit of a bitstring of length *N* is bit (*N*–1) and its right-most bit is bit 0. This order is used as the most-significant-to-least-significant bit order in conversions to and from integers. For bitstring constants and bitstrings derived from encoding diagrams, this order matches the way they are printed.

Bitstrings are the only concrete data type in pseudocode, in the sense that they correspond directly to the contents of registers, memory locations, and instructions. All of the remaining data types are abstract.

B.2.3 Integers

Pseudocode integers are unbounded in size and can be either positive or negative. That is, they are mathematical integers rather than what computer languages and architectures commonly call integers. Computer integers are represented in pseudocode as bitstrings of the appropriate length, associated with suitable functions to interpret those bitstrings as integers.

The type name for integers is `integer`.

Integer constants are normally written in decimal, such as `0`, `15`, `-1234`. They can also be written in C-style hexadecimal, such as `0x55` or `0x80000000`. Hexadecimal integer constants are treated as positive unless they have a preceding minus sign. For example, `0x80000000` is the integer $+2^{31}$. If -2^{31} must be written in hexadecimal, it must be written as `-0x80000000`.

B.2.4 Reals

Pseudocode reals are unbounded in size and precision. That is, they are mathematical real numbers, not computer floating-point numbers. Computer floating-point numbers are represented in pseudocode as bitstrings of the appropriate length, associated with suitable functions to interpret those bitstrings as reals.

The type name for reals is `real`.

Real constants are written in decimal with a decimal point. This means `0` is an integer constant but `0.0` is a real constant.

B.2.5 Booleans

A Boolean is a logical TRUE or FALSE value.

The type name for Booleans is `boolean`. This is not the same type as `bit`, which is a length–1 bitstring. Boolean constants are `TRUE` and `FALSE`.

B.2.6 Enumerations

An enumeration is a defined set of symbolic constants, such as:

```
enumeration InstrSet {InstrSet_A32, InstrSet_T32, InstrSet_A64};
```

An enumeration always contains at least one symbolic constant, and a symbolic constant must not be shared between enumerations.

Enumerations must be declared explicitly, although a variable of an enumeration type can be declared implicitly by assigning one of the symbolic constants to it. By convention, each of the symbolic constants starts with the name of the enumeration followed by an underscore. The name of the enumeration is its *type name*, or *type*, and the symbolic constants are its possible *constants*.

Note

A Boolean is a pre-declared enumeration that does not follow the normal naming convention and it has a special role in some pseudocode constructs, such as `if` statements. This means the enumeration of a `boolean` is:

```
enumeration boolean {FALSE, TRUE};
```

B.2.7 Lists

A list is an ordered set of other data items, separated by commas and enclosed in parentheses, for example:

```
(bits(32) shifter_result, bit shifter_carry_out)
```

A list always contains at least one data item.

Lists are often used as the return type for a function that returns multiple results. For example, this list at the start of this section is the return type of the function `Shift_C()` that performs a standard ARM shift or rotation, when its first operand is of type `bits(32)`.

Some specific pseudocode operators use lists surrounded by other forms of bracketing than the `(…)` parentheses. These are:

- Bitstring extraction operators, that use lists of bit numbers or ranges of bit numbers surrounded by angle brackets `<…>`.
- Array indexing, that uses lists of array indexes surrounded by square brackets `[…]`.
- Array-like function argument passing, that uses lists of function arguments surrounded by square brackets `[…]`.

Each combination of data types in a list is a separate type, with type name given by listing the data types. This means that the example list at the start of this section is of type `(bits(32), bit)`. The general principle that types can be declared by assignment extends to the types of the individual list items in a list. For example:

```
(shift_t, shift_n) = ('00', 0);
```

implicitly declares `shift_t`, `shift_n`, and `(shift_t, shift_n)` to be of types `bits(2)`, `integer`, and `(bits(2), integer)`, respectively.

A list type can also be explicitly named, with explicitly named elements in the list. For example:

```
type ShiftSpec is (bits(2) shift, integer amount);
```

After this definition and the declaration:

```
ShiftSpec abc;
```

the elements of the resulting list can then be referred to as `abc.shift`, and `abc.amount`. This qualified naming of list elements is only permitted for variables that have been explicitly declared, not for those that have been declared by assignment only.

Explicitly naming a type does not alter what type it is. For example, after the above definition of `ShiftSpec`, `ShiftSpec`, and `(bits(2), integer)` are two different names for the same type, not the names of two different types. To avoid ambiguity in references to list elements, it is an error to declare a list variable multiple times using different names of its type or to qualify it with list element names not associated with the name by which it was declared.

An item in a list that is being assigned to can be written as `"-"` to indicate that the corresponding item of the assigned list value is discarded. For example:

```
(shifted, -) = LSL_C(operand, amount);
```

List constants are written as a list of constants of the appropriate types, for example the `('00', 0)` in the earlier example.

B.2.8 Arrays

Pseudocode arrays are indexed by either enumerations or integer ranges. An integer range is represented by the lower inclusive end of the range, then `..`, then the upper inclusive end of the range.

For example:

```
// The names of the Banked core registers.

enumeration RName {RName_0usr, RName_1usr, RName_2usr, RName_3usr, RName_4usr, RName_5usr,
                   RName_6usr, RName_7usr, RName_8usr, RName_8fiq, RName_9usr, RName_9fiq,
                   RName_10usr, RName_10fiq, RName_11usr, RName_11fiq, RName_12usr, RName_12fiq,
                   RName_SPusr, RName_SPfiq, RName_SPirq, RName_SPsvc,
                   RName_SPabt, RName_SPund, RName_SPmon, RName_SPhyp,
                   RName_LRusr, RName_LRfiq, RName_LRirq, RName_LRsvc,
                   RName_LRabt, RName_LRund, RName_LRmon,
                   RName_PC};

array bits(8) _Memory[0..0xFFFFFFFF];
```

Arrays are always explicitly declared, and there is no notation for a constant array. Arrays always contain at least one element, because:

- Enumerations always contain at least one symbolic constant.
- Integer ranges always contain at least one integer.

Arrays do not usually appear directly in pseudocode. The items that syntactically look like arrays in pseudocode are usually array-like functions such as `R[i]`, `MemU[address, size]` or `Elem[vector, i, size]`. These functions package up and abstract additional operations normally performed on accesses to the underlying arrays, such as register banking, memory protection, endian-dependent byte ordering, exclusive-access housekeeping and Advanced SIMD element processing.

B.3 Expressions

This section describes:

- *General expression syntax*.
- *Operators and functions - polymorphism and prototypes* on page B-724.
- *Precedence rules* on page B-724.

B.3.1 General expression syntax

An expression is one of the following:

- A constant.
- A variable, optionally preceded by a data type name to declare its type.
- The word `UNKNOWN` preceded by a data type name to declare its type.
- The result of applying a language-defined operator to other expressions.
- The result of applying a function to other expressions.

Variable names normally consist of alphanumeric and underscore characters, starting with an alphabetic or underscore character.

Each register described in the text is to be regarded as declaring a correspondingly named bitstring variable, and that variable has the stated behavior of the register. For example, if a bit of a register is defined as RAZ/WI, then the corresponding bit of its variable reads as 0 and ignore writes.

An expression like `bits(32) UNKNOWN` indicates that the result of the expression is a value of the given type, but the architecture does not specify what value it is and software must not rely on such values. The value produced must not:

- Return information that cannot be accessed at the current or a lower level of privilege using instructions that are not UNPREDICTABLE and do not return UNKNOWN values.
- Be promoted as providing any useful information to software.

Note

Some earlier documentation describes this as an UNPREDICTABLE value. UNKNOWN values are similar to the definition of UNPREDICTABLE, but do not indicate that the entire architectural state becomes unspecified.

Only the following expressions are assignable. This means that these are the only expressions that can be placed on the left-hand side of an assignment.

- Variables.
- The results of applying some operators to other expressions.

  The description of each language-defined operator that can generate an assignable expression specifies the circumstances under which it does so. For example, those circumstances might require that one or more of the expressions the operator operates is an assignable expression.
- The results of applying array-like functions to other expressions. The description of an array-like function specifies the circumstances under which it can generate an assignable expression.

Every expression has a data type:

- For a constant, this data type is determined by the syntax of the constant.
- For a variable, there are the following possible sources for the data type:
  - An optional preceding data type name.
  - A data type the variable was given earlier in the pseudocode by recursive application of this rule.
  - A data type the variable is being given by assignment, either by direct assignment to the variable, or by assignment to a list of which the variable is a member.

It is a pseudocode error if none of these data type sources exists for a variable, or if more than one of them exists and they do not agree about the type.

- For a language-defined operator, the definition of the operator determines the data type.
- For a function, the definition of the function determines the data type.

B.3.2 Operators and functions - polymorphism and prototypes

Operators and functions in pseudocode can be polymorphic, producing different functionality when applied to different data types. Each resulting form of an operator or function has a different prototype definition. For example, the operator + has forms that act on various combinations of integers, reals, and bitstrings.

One particularly common form of polymorphism is between bitstrings of different lengths. This is represented by using `bits(N)`, `bits(M)`, or similar, in the prototype definition.

B.3.3 Precedence rules

The precedence rules for expressions are:

1. Constants, variables, and function invocations are evaluated with higher priority than any operators using their results.
2. Expressions on integers follow the normal operator precedence rules of *exponentiation before multiply/divide before add/subtract*, with sequences of multiply/divides or add/subtracts evaluated left-to-right.
3. Other expressions must be parenthesized to indicate operator precedence if ambiguity is possible, but do not have to be if all permitted precedence orders under the type rules necessarily lead to the same result. For example, if `i`, `j`, and `k` are integer variables, `i > 0 && j > 0 && k > 0` is acceptable, but `i > 0 && j > 0 || k > 0` is not.

B.4 Operators and built-in functions

This section describes:

- *Operations on generic types*.
- *Operations on Booleans*.
- *Bitstring manipulation*.
- *Arithmetic* on page B-728.

B.4.1 Operations on generic types

The following operations are defined for all types.

Equality and non-equality testing

Any two values `x` and `y` of the same type can be tested for equality by the expression `x == y` and for non-equality by the expression `x != y`. In both cases, the result is of type `boolean`.

A special form of comparison is defined with a bitstring constant that includes `'x'` bits in addition to `'0'` and `'1'` bits. The bits corresponding to the `'x'` bits are ignored in determining the result of the comparison. For example, if opcode is a 4-bit bitstring, `opcode == '1x0x'` is equivalent to `opcode<3> == '1' && opcode<1> == '0'`.

Note

This special form is permitted in the implied equality comparisons in `when` parts of `case … of …` structures.

Conditional selection

If `x` and `y` are two values of the same type and `t` is a value of type `boolean`, then `if t then x else y` is an expression of the same type as `x` and `y` that produces `x` if `t` is `TRUE` and `y` if `t` is `FALSE`.

B.4.2 Operations on Booleans

If `x` is a Boolean, then `!x` is its logical inverse.

If `x` and `y` are Booleans, then `x && y` is the result of ANDing them together. As in the C language, if `x` is `FALSE`, the result is determined to be `FALSE` without evaluating `y`.

If `x` and `y` are Booleans, then `x || y` is the result of ORing them together. As in the C language, if `x` is `TRUE`, the result is determined to be `TRUE` without evaluating `y`.

If `x` and `y` are Booleans, then `x ^ y` is the result of exclusive-ORing them together.

B.4.3 Bitstring manipulation

The following bitstring manipulation functions are defined:

Bitstring length and most significant bit

If `x` is a bitstring:

- The bitstring length function `Len(x)` returns the length of `x` as an integer.
- `TopBit(x)` is the leftmost bit of x. Using bitstring extraction, this means:

  `TopBit(x)= x<Len(x)-1>`.

Bitstring concatenation and replication

If `x` and `y` are bitstrings of lengths `N` and `M` respectively, then `x:y` is the bitstring of length `N+M` constructed by concatenating `x` and `y` in left-to-right order.

If x is a bitstring and n is an integer with `n > 0`:

- `Replicate(x, n)` is the bitstring of length `n*Len(x)` consisting of n copies of x concatenated together
- `Zeros(n)` = `Replicate('0', n)`, `Ones(n)` = `Replicate('1', n)`.

Bitstring extraction

The bitstring extraction operator extracts a bitstring from either another bitstring or an integer. Its syntax is `x<integer_list>`, where x is the integer or bitstring being extracted from, and `<integer_list>` is a list of integers enclosed in angle brackets rather than the usual parentheses. The length of the resulting bitstring is equal to the number of integers in `<integer_list>`. In `x<integer_list>`, each of the integers in `<integer_list>` must be:

- `>= 0`
- `< Len(x)` if x is a bitstring.

The definition of `x<integer_list>` depends on whether `integer_list` contains more than one integer:

- If `integer_list` contains more than one integer, `x<i, j, k,…, n>` is defined to be the concatenation:

  `x<i> : x<j> : x<k> : … : x<n>`.

- If `integer_list` consists of one integer `i`, `x<i>` is defined to be:
  - If x is a bitstring, `'0'` if bit `i` of x is a zero and `'1'` if bit `i` of x is a one.
  - If x is an integer, let y be the unique integer in the range `0` to `2^(i+1)-1` that is congruent to x modulo `2^(i+1)`. Then `x<i>` is `'0'` if `y < 2^i` and `'1'` if `y >= 2^i`.

    Loosely, this definition treats an integer as equivalent to a sufficiently long two's complement representation of it as a bitstring.

In `<integer_list>`, the notation `i:j` with `i >= j` is shorthand for the integers in order from `i` down to `j`, with both end values included. For example, `instr<31:28>` is shorthand for `instr<31, 30, 29, 28>`.

The expression `x<integer_list>` is assignable provided x is an assignable bitstring and no integer appears more than once in `<integer_list>`. In particular, `x<i>` is assignable if x is an assignable bitstring and `0 <= i < Len(x)`.

Encoding diagrams for registers frequently show named bits or multi-bit fields. For example, the encoding diagram for the ICC_SGI1R shows its bit<28> as IS. In such cases, the syntax `ICC_SGI1R.IS` is used as a more readable synonym for `ICC_SGI1R<28>`.

Logical operations on bitstrings

If x is a bitstring, `NOT(x)` is the bitstring of the same length obtained by logically inverting every bit of x.

If x and y are bitstrings of the same length, `x AND y`, `x OR y`, and `x EOR y` are the bitstrings of that same length obtained by logically ANDing, ORing, and exclusive-ORing corresponding bits of x and y together.

Bitstring count

If x is a bitstring, `BitCount(x)` produces an integer result equal to the number of bits of x that are ones.

Testing a bitstring for being all zero or all ones

If x is a bitstring:

- `IsZero(x)` produces `TRUE` if all of the bits of x are zeros and `FALSE` if any of them are ones
- `IsZeroBit(x)` produces `'1'` if all of the bits of x are zeros and `'0'` if any of them are ones.

`IsOnes(x)` and `IsOnesBit(x)` work in the corresponding ways. This means:

```
IsZero(x)    = (BitCount(x) == 0)
IsOnes(x)    = (BitCount(x) == Len(x))
IsZeroBit(x) = if IsZero(x) then '1' else '0'
IsOnesBit(x) = if IsOnes(x) then '1' else '0'
```

Lowest and highest set bits of a bitstring

If x is a bitstring, and `N = Len(x)`:

- `LowestSetBit(x)` is the minimum bit number of any of its bits that are ones. If all of its bits are zeros, `LowestSetBit(x) = N`.
- `HighestSetBit(x)` is the maximum bit number of any of its bits that are ones. If all of its bits are zeros, `HighestSetBit(x) = -1`.
- `CountLeadingZeroBits(x)` is the number of zero bits at the left end of x, in the range 0 to N. This means:

  `CountLeadingZeroBits(x) = N - 1 - HighestSetBit(x)`.
- `CountLeadingSignBits(x)` is the number of copies of the sign bit of x at the left end of x, excluding the sign bit itself, and is in the range 0 to N–1. This means:

  `CountLeadingSignBits(x) = CountLeadingZeroBits(x<N-1:1> EOR x<N-2:0>)`.

Zero-extension and sign-extension of bitstrings

If x is a bitstring and i is an integer, then `ZeroExtend(x, i)` is x extended to a length of i bits, by adding sufficient zero bits to its left. That is, if `i == Len(x)`, then `ZeroExtend(x, i) = x`, and if `i > Len(x)`, then:

```
ZeroExtend(x, i) = Replicate('0', i-Len(x)) : x
```

If x is a bitstring and i is an integer, then `SignExtend(x, i)` is x extended to a length of i bits, by adding sufficient copies of its leftmost bit to its left. That is, if `i == Len(x)`, then `SignExtend(x, i) = x`, and if `i > Len(x)`, then:

```
SignExtend(x, i) = Replicate(TopBit(x), i-Len(x)) : x
```

It is a pseudocode error to use either `ZeroExtend(x, i)` or `SignExtend(x, i)` in a context where it is possible that `i < Len(x)`.

Converting bitstrings to integers

If x is a bitstring, `SInt(x)` is the integer whose two's complement representation is x:

```
// SInt()
// ======

integer SInt(bits(N) x)
    result = 0;
    for i = 0 to N-1
        if x<i> == '1' then result = result + 2^i;
    if x<N-1> == '1' then result = result - 2^N;
    return result;
```

`UInt(x)` is the integer whose unsigned representation is x:

```
// UInt()
// ======
```

```
integer UInt(bits(N) x)
    result = 0;
    for i = 0 to N-1
        if x<i> == '1' then result = result + 2^i;
    return result;
```

`Int(x, unsigned)` returns either `SInt(x)` or `UInt(x)` depending on the value of its second argument:

```
// Int()
// =====

integer Int(bits(N) x, boolean unsigned)
    result = if unsigned then UInt(x) else SInt(x);
    return result;
```

B.4.4 Arithmetic

Most pseudocode arithmetic is performed on integer or real values, with operands being obtained by conversions from bitstrings and results converted back to bitstrings afterwards. As these data types are the unbounded mathematical types, no issues arise about overflow or similar errors.

Unary plus, minus, and absolute value

If `x` is an integer or real, then `+x` is `x` unchanged, `-x` is `x` with its sign reversed, and `Abs(x)` is the absolute value of `x`. All three are of the same type as `x`.

Addition and subtraction

If `x` and `y` are integers or reals, `x+y` and `x-y` are their sum and difference. Both are of type `integer` if `x` and `y` are both of type `integer`, and `real` otherwise.

Addition and subtraction are particularly common arithmetic operations in pseudocode, and so it is also convenient to have definitions of addition and subtraction acting directly on bitstring operands.

If `x` and `y` are bitstrings of the same length `N`, so that `N = Len(x) = Len(y)`, then `x+y` and `x-y` are the least significant `N` bits of the results of converting them to integers and adding or subtracting them. Signed and unsigned conversions produce the same result:

```
x+y = (SInt(x) + SInt(y))<N-1:0>
    = (UInt(x) + UInt(y))<N-1:0>
x-y = (SInt(x) - SInt(y))<N-1:0>
    = (UInt(x) - UInt(y))<N-1:0>
```

If `x` is a bitstring of length `N` and `y` is an integer, `x+y` and `x-y` are the bitstrings of length `N` defined by `x+y = x + y<N-1:0>` and `x-y = x - y<N-1:0>`. Similarly, if `x` is an integer and `y` is a bitstring of length `M`, `x+y` and `x-y` are the bitstrings of length `M` defined by `x+y = x<M-1:0> + y` and `x-y = x<M-1:0> - y`.

Comparisons

If `x` and `y` are integers or reals, then `x == y`, `x != y`, `x < y`, `x <= y`, `x > y`, and `x >= y` are equal, not equal, less than, less than or equal, greater than, and greater than or equal comparisons between them, producing Boolean results. In the case of `==` and `!=`, this extends the generic definition applying to any two values of the same type to also act between integers and reals.

Multiplication

If `x` and `y` are integers or reals, then `x * y` is the product of `x` and `y`. It is of type `integer` if `x` and `y` are both of type `integer`, and `real` otherwise.

Division and modulo

If `x` and `y` are integers or reals, then `x/y` is the result of dividing `x` by `y`, and is always of type `real`.

If x and y are integers, then `x DIV y` and `x MOD y` are defined by:

```
x DIV y = RoundDown(x/y)
x MOD y = x - y * (x DIV y)
```

It is a pseudocode error to use any of x/y, `x MOD y`, or `x DIV y` in any context where y can be zero.

Square root

If x is an integer or a real, `Sqrt(x)` is its square root, and is always of type `real`.

Rounding and aligning

If x is a real:

- `RoundDown(x)` produces the largest integer n so that `n <= x`
- `RoundUp(x)` produces the smallest integer n so that `n >= x`
- `RoundTowardsZero(x)` produces `RoundDown(x)` if `x > 0.0`, `0 if x == 0.0`, and `RoundUp(x) if x < 0.0`.

If x and y are both of type `integer`, `Align(x, y) = y * (x DIV y)` is of type `integer`.

If x is of type `bitstring` and y is of type `integer`, `Align(x, y) = (Align(UInt(x), y))<Len(x)-1:0>` is a `bitstring` of the same length as x.

It is a pseudocode error to use either form of `Align(x, y)` in any context where y can be 0. In practice, `Align(x, y)` is only used with y a constant power of two, and the bitstring form used with `y = 2^n` has the effect of producing its argument with its n low-order bits forced to zero.

Scaling

If n is an integer, `2^n` is the result of raising `2` to the power n and is of type `real`.

If x and n are of type `integer`, then:

- `x << n = RoundDown(x * 2^n)`
- `x >> n = RoundDown(x * 2^(-n))`.

Maximum and minimum

If x and y are integers or reals, then `Max(x, y)` and `Min(x, y)` are their maximum and minimum respectively. Both are of type `integer` if x and y are both of type `integer`, and `real` otherwise.

B.5 Statements and program structure

The following sections describe the control statements used in the pseudocode:

- *Simple statements*.
- *Compound statements* on page B-731.
- *Comments* on page B-733.

B.5.1 Simple statements

Each of the following simple statements must be terminated with a semicolon, as shown.

Assignments

An assignment statement takes the form:

```
<assignable_expression> = <expression>;
```

Procedure calls

A procedure call takes the form:

```
<procedure_name>(<arguments>);
```

Return statements

A procedure return takes the form:

```
return;
```

and a function return takes the form:

```
return <expression>;
```

where `<expression>` is of the type declared in the function prototype line.

UNDEFINED

This subsection describes the statement:

```
UNDEFINED;
```

This statement indicates a special case that replaces the behavior defined by the current pseudocode, apart from behavior required to determine that the special case applies. The replacement behavior is that the Undefined Instruction exception is taken.

UNPREDICTABLE

This subsection describes the statement:

```
UNPREDICTABLE;
```

This statement indicates a special case that replaces the behavior defined by the current pseudocode, apart from behavior required to determine that the special case applies. The replacement behavior is UNPREDICTABLE.

SEE…

This subsection describes the statement:

```
SEE <reference>;
```

This statement indicates a special case that replaces the behavior defined by the current pseudocode, apart from behavior required to determine that the special case applies. The replacement behavior is that nothing occurs as a result of the current pseudocode because some other piece of pseudocode defines the required behavior. The `<reference>` indicates where that other pseudocode can be found.

It usually refers to another instruction but can also refer to another encoding or note of the same instruction.

IMPLEMENTATION_DEFINED

This subsection describes the statement:

```
IMPLEMENTATION_DEFINED {<text>};
```

This statement indicates a special case that replaces the behavior defined by the current pseudocode, apart from behavior required to determine that the special case applies. The replacement behavior is IMPLEMENTATION DEFINED. An optional `<text>` field can give more information.

B.5.2 Compound statements

Indentation normally indicates the structure in compound statements. The statements contained in structures such as `if … then … else …` or procedure and function definitions are indented more deeply than the statement itself, and their end is indicated by returning to the original indentation level or less.

Indentation is normally done by four spaces for each level.

if … then … else …

A multi-line `if … then … else …` structure takes the form:

```
if <boolean_expression> then
    <statement 1>
    <statement 2>
    …
    <statement n>
elsif <boolean_expression> then
    <statement a>
    <statement b>
    …
    <statement z>
else
    <statement A>
    <statement B>
    …
    <statement Z>
```

The block of lines consisting of `elsif` and its indented statements is optional, and multiple such blocks can be used.

The block of lines consisting of `else` and its indented statements is optional.

Abbreviated one-line forms can be used when there are only simple statements in the `then` part and in the `else` part, if it is present, such as:

```
if <boolean_expression> then <statement 1>
if <boolean_expression> then <statement 1> else <statement A>
if <boolean_expression> then <statement 1> <statement 2> else <statement A>
```

Note

In these forms, `<statement 1>`, `<statement 2>` and `<statement A>` must be terminated by semicolons. This and the fact that the `else` part is optional are differences from the `if … then … else …` expression.

repeat … until …

A `repeat … until …` structure takes the form:

```
repeat
    <statement 1>
    <statement 2>
    …
    <statement n>
until <boolean_expression>;
```

while … do

A `while … do` structure takes the form:

```
while <boolean_expression> do
    <statement 1>
    <statement 2>
    …
    <statement n>
```

for …

A `for …` structure takes the form:

```
for <assignable_expression> = <integer_expr1> to <integer_expr2>
    <statement 1>
    <statement 2>
    …
    <statement n>
```

case … of …

A `case … of …` structure takes the form:

```
case <expression> of
    when <constant values>
        <statement 1>
        <statement 2>
        …
        <statement n>
        … more "when" groups …
    otherwise
        <statement A>
        <statement B>
        …
        <statement Z>
```

In this structure, `<constant values>` consists of one or more constant values of the same type as `<expression>`, separated by commas. Abbreviated one line forms of `when` and `otherwise` parts can be used when they contain only simple statements.

If `<expression>` has a bitstring type, `<constant values>` can also include bitstring constants containing `'x'` bits. For details see *Equality and non-equality testing* on page B-725.

Procedure and function definitions

A procedure definition takes the form:

```
<procedure name>(<argument prototypes>)
    <statement 1>
    <statement 2>
    …
    <statement n>
```

where `<argument prototypes>` consists of zero or more argument definitions, separated by commas. Each argument definition consists of a type name followed by the name of the argument.

Note

This first prototype line is not terminated by a semicolon. This helps to distinguish it from a procedure call.

A function definition is similar but also declares the return type of the function:

```
<return type> <function name>(<argument prototypes>)
    <statement 1>
    <statement 2>
    …
    <statement n>
```

An array-like function is similar but with square brackets:

```
<return type> <function name>[<argument prototypes>]
    <statement 1>
    <statement 2>
    …
    <statement n>
```

An array-like function also usually has an assignment prototype:

```
<function name>[<argument prototypes>] = <value prototypes>
    <statement 1>
    <statement 2>
    …
    <statement n>
```

B.5.3 Comments

Two styles of pseudocode comment exist:

- // starts a comment that is terminated by the end of the line
- /* starts a comment that is terminated by */.

B.6 Pseudocode terminology

Table B-1 lists the terms used in the body text and the terms used in pseudocode to denote the same concept, throughout this specification.

Table B-1 Pseudocode terms

Term used in the body text	Term used in pseudocode
vPE	`VCPU`
vPEID	`VCPUID`
pINTID	`pID`
vINTID	`vID`

B.7 Miscellaneous helper procedures and support functions

The functions described in this section are not part of the pseudocode specification. They are miscellaneous *helper* procedures and functions used by pseudocode that are not described elsewhere in this document. Each has a brief description and a pseudocode prototype, except that the prototype is omitted where it is identical to the section title.

Note

Some variable names used the pseudocode differ from those used in the body text. For a list of the affected variables, see *Pseudocode terminology* on page B-734.

B.7.1 Helper functions

The functions listed in the following sections are indicated by the hierarchical path names, for example `shared/gic/helper:`

shared/gic/helper/AcknowledgeInterrupt

```
// AcknowledgeInterrupt()
// ==============================
// Acknowledges the INTID and sets the appropriate ICC_AP{0,1}R_EL1 active priority bit


AcknowledgeInterrupt(bits(INTID_SIZE) ID);
```

shared/gic/helper/AcknowledgeVInterrupt

```
// AcknowledgeVInterrupt()
// =======================
// Acknowledges vINTID

AcknowledgeVInterrupt(bits(INTID_SIZE) ID);
```

shared/gic/helper/AlwaysUsingSysRegs

```
// AlwaysUsingSysRegs()
// ====================

boolean AlwaysUsingSysRegs();
```

shared/gic/helper/Deactivate

```
// Deactivate()
```

```
// =================

// Deactivates the INTID


Deactivate(bits(INTID_SIZE) INTID);
```

shared/gic/helper/INTID_SIZE

```
// INTID_SIZE
// ==========
// The number of interrupt ID bits implemented at the Distributor and Redistributor.
// This value is IMPLEMENTATION DEFINED and discoverable from GICD_TYPER.IDbits.

constant integer INTID_SIZE = integer IMPLEMENTATION_DEFINED “Distributor INTID size”;
```

shared/gic/helper/IntGroup

```
// IntGroup()
// ==================

enumeration IntGroup { IntGroup_None, IntGroup_G0, IntGroup_G1NS, IntGroup_G1S };
LRType GICH_VLPIR;                              // Holds virtual LPIs received from the Distributor
```

shared/gic/helper/Interrupt

```
// Interrupt()
// ================

constant bits(2) IntState_Invalid = ‘00’;
constant bits(2) IntState_Pending = ‘01’;
constant bits(2) IntState_Active = ‘10’;
constant bits(2) IntState_ActivePending = ‘11’;
```

shared/gic/helper/IsGrp0Int

```
// IsGrp0Int()
// ======================
// Returns TRUE if the  INTID is in Group 0

boolean IsGrp0Int(bits(INTID_SIZE) ID);
```

shared/gic/helper/IsSecureInt

```
// IsSecureInt()
// =============

boolean IsSecureInt(bits(INTID_SIZE) ID);
```

shared/gic/helper/PermissionFailureException

```
 //PermissionFailureException()
// ========================

PermissionFailureException(integer targetEL);
```

shared/gic/helper/PriorityIsHigher

```
// PriorityIsHigher()
// ==========================

boolean PriorityIsHigher(bits(8) first, bits(8) second);
```

shared/gic/helper/SingleSecurityState

```
// SingleSecurityState()
// =====================

// Returns TRUE if the Distributor supports  a single Security state, for example when GICD_CTLR_DS == 1.


boolean SingleSecurityState();
```

shared/gic/helper/Special

```
// IsSpecial()
// ===========

boolean IsSpecial(bits(INTID_SIZE) intID)

    return UInt(intID) >= 1020 && UInt(intID) <= 1023;

boolean IsLPI(bits(INTID_SIZE) intID)

    return UInt(intID) >= 8192;
```

The definition for this function is as follows:

```
constant bits(INTID_SIZE) INTID_SECURE    = 1020<INTID_SIZE-1:0>;
constant bits(INTID_SIZE) INTID_NONSECURE = 1021<INTID_SIZE-1:0>;
constant bits(INTID_SIZE) INTID_GROUP1    = 1022<INTID_SIZE-1:0>;
constant bits(INTID_SIZE) INTID_SPURIOUS  = 1023<INTID_SIZE-1:0>;
type INTID = bits(INTID_SIZE);
```

shared/gic/helper/SystemRegisterTrap

```
// SystemRegisterTrap()
// ====================

SystemRegisterTrap(bits(2) target_el);
```

B.7.2 Support functions

This section lists the support functions that are not listed elsewhere in this specification.The functions listed in the following sections are indicated by the hierarchical path names, for example `shared/support`:

- *shared/support/ActivePRIBits*.
- *shared/support/CanSignalInterrupt* on page B-738.
- *shared/support/CanSignalVirtualInt* on page B-738.
- *shared/support/CanSignalVirtualInterrupt* on page B-738.
- *shared/support/ClearPendingState* on page B-739.
- *shared/support/HighestPriorityPendingInterrupt* on page B-739.
- *shared/support/HighestPriorityVirtualInterrupt* on page B-739.
- *shared/support/PRIBits* on page B-740.
- *shared/support/PriorityDrop* on page B-740.
- *shared/support/PriorityGroup* on page B-740.
- *shared/support/SetPendingState* on page B-740.
- *shared/support/SystemRegisterAccessPermitted* on page B-741.

shared/support/ActivePRIBits

```
// ActivePRIBits()
// ===============
```

```
integer ActivePRIBits()
    pri_bits = PRIBits();
    return 2^(pri_bits - 1);
```

shared/support/CanSignalInterrupt

```
// CanSignalInterrupt()
// ====================

boolean CanSignalInterrupt()

    // Get the priority group of the current “Set” using the BPR appropriate to the group
    setPriorityGroup = GroupBits(GICC_SETR.Priority, GICC_SETR.Group);
    runningPriority = GetHighestActivePriority(ICC_AP0R_EL1, ICC_AP1R_EL1NS, ICC_AP1R_EL1S);

    // Get the priority group of highest APR using the BPR appropriate to the SET packet
    preemptionLevel = GroupBits(runningPriority<7:1>:'0', GICC_SETR.Group);

    if (GICC_SETR.State == IntState_Pending &&
        UInt(GICC_SETR.Priority) < UInt(ICC_PMR_EL1.Priority)) then
        // The “Set” is higher priority than PMR
        if (runningPriority == 0x7F) || (setPriorityGroup<preemptionLevel) then
            return TRUE;                                // The Set can preempt

    return FALSE;                                       // Can’t preempt so no interrupt
```

shared/support/CanSignalVirtualInt

```
// CanSignalVirtualInt()
// =====================

boolean CanSignalVirtualInt(bits(64) listReg)

    LRType vInt = listReg;

    // First check whether the virtual interface is enabled
    if ICH_HCR_EL2.En == ‘0’ then
        return FALSE;

    // Get the priority group of “vInt” using the BPR appropriate to the group
    vintPriorityGroup = VGroupBits(vInt.Priority, vInt.Group);
    runningPriority = GetHighestActivePriority(ICH_AP0R_EL2, ICH_AP1R_EL2, Zeros());

    // Get the priority group of highest APR using the BPR appropriate to the APR group
    preemptionLevel = VGroupBits(runningPriority<7:1>:'0', vInt.Group);

    if vInt.State == IntState_Pending && UInt(vInt.Priority) < UInt(ICH_VMCR_EL2.VPMR) then
        // “vInt” is higher priority than PMR
        if (runningPriority == 0x7F) || (UInt(vintPriorityGroup) < UInt(preemptionLevel) then    // The
“vInt” can preempt
            return TRUE;

    return FALSE;                                                   // Can’t preempt so no interrupt
```

shared/support/CanSignalVirtualInterrupt

```
// CanSignalVirtualInterrupt()
// ===========================

boolean CanSignalVirtualInterrupt()
    integer lrIndex = HighestPriorityVirtualInterrupt();

    if (GICH_VLPIR.State == IntState_Pending &&
        (lrIndex < 0 || PriorityIsHigher(GICH_VLPIR.Priority, ICH_LR_EL2[lrIndex].Priority))) then
        // A virtual LPI is the highest priority
        return CanSignalVirtualInt(GICH_VLPIR);
```

```
    elsif lrIndex >= 0 then
        // A list register is the highest priority
        return CanSignalVirtualInt(ICH_LR_EL2[lrIndex]);

    // There are no valid and enabled interrupts
    return FALSE;
```

shared/support/ClearPendingState

```
// ClearPendingState()
// ===================

boolean ClearPendingState(InterruptTableEntry ite)

    if ite.Type == physical_interrupt then
        CollectionTableEntry cte = ReadCollectionTable(UInt(ite.ICID));

        if (!cte.Valid) then
            return FALSE;

        bits(32) rd_base = cte.RDbase;

        ClearPendingStateLocal(GICR_PENDBASER[rd_base], ite.OutputID);

    else
        VCPUTableEntry vte = ReadVCPUTable(UInt(ite.VCPUID));

        if (!vte.Valid) then
            return FALSE;

        bits(32) rd_base = vte.RDbase;
        Address vpt = vte.VPT_base;

        ClearPendingStateLocal(vpt, ite.OutputID);

    return TRUE;
```

shared/support/HighestPriorityPendingInterrupt

```
// HighestPriorityPendingInterrupt()
// =================================

bits(INTID_SIZE) HighestPriorityPendingInterrupt()
    if GICC_SETR.State != IntState_Pending then        // No interrupt pending
        return INTID_SPURIOUS;

    case GICC_SETR.Group of
        when IntGroup_G1NS
            if ICC_IGRPEN1_EL1NS.Enable == '0' then return INTID_SPURIOUS;
        when IntGroup_G1S
            if ICC_IGRPEN1_EL1S.Enable == '0' then return INTID_SPURIOUS;
        when IntGroup_G0
            if ICC_IGRPEN0_EL1.Enable == '0' then return INTID_SPURIOUS;
        otherwise                                      // Reserved
            return INTID_SPURIOUS;

    return GICC_SETR.ID;
```

shared/support/HighestPriorityVirtualInterrupt

```
// HighestPriorityVirtualInterrupt()
// =================================
// Returns -1 if there are no pending virtual interrupts

integer HighestPriorityVirtualInterrupt()
```

```
    integer lrIndex = -1;
    bits(8) priority = Ones();

    // Find the List Register with the highest priority enabled pending interrupt
    for i = 0 to NumListRegs() - 1
        if (ICH_LR_EL2[i].State == IntState_Pending &&
            ((ICH_LR_EL2[i].Group == '0' && ICH_VMCR_EL2.VENG0 == '1') ||
             (ICH_LR_EL2[i].Group == '1' && ICH_VMCR_EL2.VENG1 == '1')) &&
            PriorityIsHigher(ICH_LR_EL2[i].Priority,  priority)) then
            // Found an enabled pending list register with a higher priority
            priority = ICH_LR_EL2[i].Priority;
            lrIndex = i;

    return lrIndex;
```

shared/support/PRIBits

```
// PRIBits()
// =========

integer PRIBits()
    pri_bits = UInt(if HaveEL(EL3) then ICC_CTLR_EL3.PRIbits else ICC_CTLR_EL1.PRIbits);
    return pri_bits + 1;
```

shared/support/PriorityDrop

```
// PriorityDrop
// ============
// Clears the highest active priority in the supplied register; returns FALSE if no priorities were
active.

boolean PriorityDrop[bits(128) &ap]
```

shared/support/PriorityGroup

```
// PriorityGroup()
// ===============
// Returns the priority group field for the minimum BPR value for the group

bits(8) PriorityGroup(bits(8) priority, IntGroup group)
    p_bits = PRIBits();

    if p_bits == 8 then
        mask = Ones(7):'0';
    else
        mask = Ones(p_bits):Zeros(8 - p_bits);

    return (priority AND mask);
```

shared/support/SetPendingState

```
// SetPendingState()
// =================

boolean SetPendingState(InterruptTableEntry ite)

    if ite.Type == physical_interrupt then
        CollectionTableEntry cte = ReadCollectionTable(UInt(ite.ICID));

        if !cte.Valid then
            return FALSE;

        bits(32) rd_base = cte.RDbase;
```

```
        SetPendingStateLocal(GICR_PENDBASER[rd_base], ite.OutputID);

    else
        VCPUTableEntry vte = ReadVCPUTable(UInt(ite.VCPUID));

        if !vte.Valid then
            return FALSE;

        bits(32) rd_base = vte.RDbase;
        Address vpt = vte.VPT_base;

        SetVirutalPendingStateLocal(vpt, ite.OutputID);

        if (GICR_VPENDBASER[rd_base].Valid == '1' &&
            GICR_VPENDBASER[rd_base].PhysicalAddress != vpt<47:16>) then
            if ite.DoorbellID != ZeroExtend(INTID_SPURIOUS, 32) then
                // Not resident so set the doorbell interrupt pending as well
                SetPendingStateLocal(GICR_PENDBASER[rd_base], ite.DoorbellID);

    return TRUE;
```

shared/support/SystemRegisterAccessPermitted

```
// SystemRegisterAccessPermitted()
// ===============================

SystemRegisterAccessPermitted(integer group)

    // The "group" parameter indicates which set of registers is being accessed
    // 0   FIQ (Group 0) registers
    // 1   IRQ (Group 1) registers
    // 2   Common registers
    // First check if any System Registers are enabled
    if PSTATE.EL == EL0 || (HaveEL(EL3) && ICC_SRE_EL3.SRE == '0') then
        // System registers aren't enabled.
        UndefinedFault();

    // Check that whether the access is to virtual or physical state
    if HaveEL(EL2) && !IsSecure() && PSTATE.EL == EL1 then
        accessIsVirtual = ((group IN {0,2} && HCR_EL2.FMO == '1') ||
                           (group IN {1,2} && HCR_EL2.IMO == '1'));
    else
        accessIsVirtual = FALSE;

    sreEL1S = (HaveEL(EL3) && ICC_SRE_EL1S.SRE == '1') || AlwaysUsingSysRegs();
    sreEL2 = HaveEL(EL2) && (ICC_SRE_EL2.SRE == '1' || sreEL1S);

    // Check whether Non-secure EL1 is using system registers or not
    if HaveEL(EL2) && (HCR_EL2.FMO == '1' || HCR_EL2.IMO == '1' || HCR_EL2.AMO == '1') then
        // EL2 is implemented and at least one interrupt exception is virtualized
        sreEL1NS = ((sreEL2  && ICC_SRE_EL1NS.SRE == '1') ||
                    (sreEL1S && (HCR_EL2.FMO == '0' || HCR_EL2.IMO == '0' || HCR_EL2.AMO == '0')));
    elsif HaveEL(EL2) then
        sreEL1NS = (ICC_SRE_EL2.SRE == '1' && ICC_SRE_EL1NS.SRE == '1') || sreEL1S;
    else
        sreEL1NS = ICC_SRE_EL1NS.SRE == '1' || sreEL1S;

    // Check if System Registers are enabled for the EL and security state
    if ((PSTATE.EL == EL2 && !IsSecure() && !sreEL2) ||
        (PSTATE.EL == EL1 && IsSecure() && ICC_SRE_EL1S.SRE == '0') ||
        (PSTATE.EL == EL1 && !IsSecure() && !sreEL1NS)) then
        UndefinedFault();              // System registers aren't enabled.

    // Check if the access should trap to the hypervisor
    if HaveEL(EL2) && !IsSecure() && PSTATE.EL == EL1 then
        if ((group == 0 && ICH_HCR_EL2.TALL0 == '1') ||
```

```
        (group == 1 && ICH_HCR_EL2.TALL1 == '1') ||
        (group == 2 && ICH_HCR_EL2.TC == '1')) then
        SystemRegisterTrap(EL2);

// Check that access is allowed given the routing
if (!HaveEL(EL3) ||
    (group IN {0,2} && SCR_EL3.FIQ == '0') ||
    (group IN {1,2} && SCR_EL3.IRQ == '0')) then
    lowestPhysicalEL = EL1;
else
    lowestPhysicalEL = EL3;

if !accessIsVirtual && UInt(PSTATE.EL) < UInt(lowestPhysicalEL) then
    if ELUsingAArch32(EL3) then
        UndefinedFault();
    else
        SystemRegisterTrap(EL3);

return;
```

Appendix C
Revisions

This appendix describes the technical changes between released issue A and issue B of this specification.

Table C-1 Issue B

Change	Location
Changed ICC_APxR_EL1[31:16] to ICC_APxR_EL1[31:16] in the tables showing the implementation of ICC_APxR_EL1	• Table 4-10 on page 4-70 • Table 4-11 on page 4-70
Changed CPU interface INTID size from 32 to 24 in `CPUInterfaceIDSize()`	*INTIDs* on page 3-39
Added a note the visibility of the effect of clearing GICD_CTLR.EnableGrp* when changing GICD_CTLR.ARE_S or GICD_CTLR.ARE_NS from 0 to 1.	*Affinity routing* on page 3-43
Clarified the description of the Priority drop in the interrupt lifecycle	*Interrupt lifecycle* on page 4-46
Clarified the description of the PE acknowledging SGIs, PPIs, and SPIs at the CPU interface	*Physical CPU interface* on page 4-46
Inserted statement that the effects of reading ICC_IAR0_EL1 and ICC_IAR1_EL1 on the state of the returned INTID are only guaranteed to be visible after executing a `DSB`	*Activation* on page 4-47
Added a table summarizing the signaling of interrupts in systems that support only a single Security state	Table 4-5 on page 4-60
Added additional information about the effects of disabling interrupts	*Effect of disabling interrupts* on page 4-64

Table C-1 Issue B (continued)

Change	Location
Added a new section to describe the ITS tables	*The ITS tables* on page 6-99
Added summary encoding tables for the System registers	• *AArch64 System register descriptions* on page 8-179 • *AArch64 virtualization control System registers* on page 8-272 • *AArch32 System register descriptions* on page 8-298 • *AArch32 virtualization control System registers* on page 8-399
Added statement that a number of functions referred to in this specification are defined in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile*	Appendix B *Pseudocode Definition*
Clarified when the Secure and Non-secure copy of ICC_SRE_EL1.SRE might be RAO/WI	*Implementations with fixed System register enables* on page 8-167
Added a note clarifying that a number of GIC System registers, unlike other ARMv8 AArch64 registers, can be banked by Security state	*Register banking* on page 8-173
Added a new section about using SGIs in asymmetric implementations	Chapter 10 *Legacy Operation and Asymmetric Configurations*
Added choice of CONSTRAINED UNPREDICTABLE behaviors that result from ITS command errors	*Command errors* on page 6-111
Clarified comment in `EOImodeSet()`	*Deactivation* on page 4-49
Clarified comment in pseudocode for ICC_EOIR0_EL1 and ICC_EOIR1_EL1	*aarch64/support/ICC_EOIR0_EL1* on page 8-664 *aarch64/support/ICC_EOIR1_EL1* on page 8-665
Added pseudocode for `VPRIMask()`	*aarch64/support/VRIMask* on page 8-677
Added additional information about legacy operations	*Legacy support of interrupts and asymmetric configurations* on page 10-684
Clarified what happens when single ARE bit is changed from 0 to 1	*Changing affinity routing enables* on page 3-44
Clarified that size of level 2 tables is determined by GIRS_BASER<n>.Page_Size	*The ITS tables* on page 6-99
Clarified that use of register offsets `0xFD0 -0xFFC` in GICD_*, GICR_*, and GITS_* is IMPLEMENTATION DEFINED	*Identification registers* on page 8-173
Clarified that AArch64 Op0=3 CRn=c12,Op1=4 CRm=c9 Op2=4 and AARch32 CRn=c12 Op1=4 CRm=c9 Op2=4 are IMPLEMENTATION DEFINED	*AArch64 System register descriptions* on page 8-179 *AArch32 System register descriptions* on page 8-298
Clarified unpredictable cases in pseudocode for `VirtualInterruptIdentifierValid()`	*Valid interrupt ID check pseudocode* on page 3-41
Clarified that ActivateAcknowledge commands do not have to be issued in the same order as the Activate command to which they are responding	*Activate Acknowledge (IRI)* on page A-701
Clarified that writing to GICD_SGIR from a PE where GICD_TYPER.ProcessorNumber is greater than 7 results in CONSTRAINED UNPREDICTABLE behavior	*Legacy support of interrupts and asymmetric configurations* on page 10-684

Table C-1 Issue B (continued)

Change	Location
Clarified that writing 0b11 to the Target List Filter is CONSTRAINED UNPREDICTABLE	*Legacy support of interrupts and asymmetric configurations* on page 10-684
Clarified comment in LPIOutOfRange()	*shared/gic/its/its_helper/LPIOutOfRange* on page 6-143
Added CONSTRAINED UNPREDICTABLE behavior for ITS command errors	*Command errors* on page 6-111
Clarified behavior when enabling ITS without memory having been allocated for private tables	*The ITS tables* on page 6-99
Added section on implementations with mixed INTID sizes	*Implementations with mixed INTD sizes* on page 3-41
Added relaxation that allows 52 bits of physical address space	*MAPC* on page 6-117 *MAPD* on page 6-118 *MOVALL* on page 6-122 *SYNC* on page 6-125 *VMAPP* on page 6-128 *VMOVP* on page 6-133 *IMPLEMENTATION DEFINED sizes in ITS command parameters* on page 6-110 *GITS_CBASER, ITS Command Queue Descriptor* on page 8-647 *GITS_BASER<n>, ITS Translation Table Descriptors, n = 0 - 7* on page 8-643 *GICR_PROPBASER, Redistributor Properties Base Address Register* on page 8-526 *GICR_PENDBASER, Redistributor LPI Pending Table Base Address Register* on page 8-523 *GICR_VPROPBASER, Virtual Redistributor Properties Base Address Register* on page 8-542 *GICR_VPENDBASER, Virtual Redistributor LPI Pending Table Base Address Register* on page 8-538
Clarified that memory-mapped registers use little-endian memory order model	*GIC memory-mapped register access* on page 8-157 *The ITS command interface* on page 6-105 *GITS_BASER<n>, ITS Translation Table Descriptors, n = 0 - 7* on page 8-643
Clarified use of Valid bit in MAPD	*MAPD* on page 6-118
Clarified that Deactivate command must be acknowledged	*Deactivate (ICC)* on page A-703
Clarified statement about virtual priority calculations and reads of the GICV_* registers for VSet command	*VSet (IRI)* on page A-714
Clarified that writes that clear GICD_CTLR Enable bits must be visible	*Changing affinity routing enables* on page 3-44
Clarified numerous statements in the register descriptions	Chapter 8 *Programmers' Model*
Clarified that an interrupt that is in a disabled group cannot block an interrupt that is in an enabled group	*Legacy support of interrupts and asymmetric configurations* on page 10-684
Clarified that software writes to GICD_IPRIORITYR<n> and GICR_IPRIORITY<n> must be visible in finite time	*Changing the priority of enabled PPIs, SGIs, and SPIs* on page 4-76

Table C-1 Issue B (continued)

Change	Location
Clarified that message-based SPIs can be generated by a write to GICD_SETSPI_NSR or GICD_SETSPI_SR and cleared by a write to GICD_CLRSPI_NSR or GICD_CLRSPI_SR.	*Shared Peripheral Interrupts* on page 4-56
Clarified that some registers cannot be used in an implementation that includes one or more ITSs	*LPIs* on page 6-92
Clarified that GICD_IPRIORITYR<n> resets to an IMPLEMENTATION DEFINED value	*The GIC Distributor register map* on page 8-429
Clarified that on receipt of a VSet command a CPU interface is required to release the previous pending virtual interrupt	*Rules associated with the downstream Redistributor commands* on page A-698
Clarified that a Quiesce Acknowledge command must be preceded by a Release command that removes any pending interrupts on the CPU interface	*Rules associated with the upstream CPU interface commands* on page A-699
Clarified description of Clear command	*Clear (IRI)* on page A-703
Clarified description of VClear command	*VClear (IRI)* on page A-714
Clarified that a Release response can occur when the highest pending virtual interrupt is updated by a VSet command	*The GIC Stream Protocol* on page A-696
Added alternative form of the VMOVP command	*VMOVP* on page 6-133
Clarified that values for TDEST and TID must be allocated sequentially	*Signals* on page A-693
Clarified `CanSignalInterrupt()` and `CanSignalVirtualInt()` pseudocode	*shared/support/CanSignalInterrupt* on page B-738 *shared/support/CanSignalVirtualInt* on page B-738
Clarified allocation of memory to LPI Pending tables	*LPI Pending tables* on page 6-97
Added information about when GICR_VPENDBASER.IDAI can be cleared to 0	*Virtual LPI Configuration tables and virtual LPI Pending tables* on page 6-97
Clarified behavior on wake-up	Chapter 7 *Power Management*

Glossary

Activate

An interrupt is activated when its state changes either:

- From pending to active.
- From pending to active and pending.

For more information see *Interrupt handling state machine on page 4-50*.

Affinity level

Provides an indication of relative locality in a multiprocessor system, by defining a particular level within the system hierarchy. The affinity levels that the GIC uses correspond to those defined in the *Multiprocessor Affinity Register* (MPIDR), an ARM processor system control register.

Banked register

A register that has multiple instances, with the instance that is in use depending on the PE mode, Security state, or other PE state.

For more information about register banking in the GIC see *Register banking on page 8-173*.

Big-endian memory

Means that, for example:

- A byte or halfword at a word-aligned address is the most significant byte or halfword in the word at that address.
- A byte at a halfword-aligned address is the most significant byte in the halfword at that address.

CONSTRAINED UNPREDICTABLE

Where an instruction can result in UNPREDICTABLE behavior, the ARMv8 architecture specifies a narrow range of permitted behaviors. This range is the range of CONSTRAINED UNPREDICTABLE behavior. All implementations that are compliant with the architecture must follow the CONSTRAINED UNPREDICTABLE behavior.

Execution at Non-secure EL1 or EL0 of an instruction that is CONSTRAINED UNPREDICTABLE can be implemented as generating a trap exception that is taken to EL2, provided that at least one instruction that is not UNPREDICTABLE and is not CONSTRAINED UNPREDICTABLE causes a trap exception that is taken to EL2.

In body text, the term CONSTRAINED UNPREDICTABLE is shown in SMALL CAPITALS.

See also UNPREDICTABLE.

Context switch The saving and restoring of computational state when switching between different threads or processes. In this manual, the term context switch describes any situation where the context is switched by an operating system and might or might not include changes to the address space.

Deactivate An interrupt is deactivated when its state changes either:

- From active to inactive.
- From active and pending to pending.

For more information see *Interrupt handling state machine on page 4-50*.

Deprecated Something that is present in the ARM architecture for backwards compatibility. Whenever possible software must avoid using deprecated features. Features that are deprecated but are not optional are present in current implementations of the ARM architecture, but might not be present, or might be deprecated and OPTIONAL, in future versions of the ARM architecture.

See also OPTIONAL.

Distributor A logical component in the GIC that receives interrupts, and determines the priority and distribution of SPIs and SGIs. The Distributor forwards the interrupt with the highest priority to the corresponding Redistributor and CPU interface, for priority masking and preemption handling.

See also Distributor on page 2-32.

Doubleword A 64-bit data item. Doublewords are normally at least word-aligned in ARM systems.

Endianness An aspect of the system memory mapping.

See also Big-endian memory and Little-endian memory.

Exception Handles an event. For example, an exception could handle an external interrupt or an undefined instruction.

GIC Stream Protocol interface

A optional interface between the IRI and the PE, specifically between the Redistributor and the associated CPU interface, that conforms to the AMBA AXI4-Stream protocol. The protocol defines a set of packets that can be sent between the CPU and the Distributor, together with ordering and flow control rules.

See also Appendix A *GIC Stream Protocol interface*.

Halfword A 16-bit data item. Halfwords are normally halfword-aligned in ARM systems.

IMPLEMENTATION DEFINED

Means that the behavior is not architecturally defined, but must be defined and documented by individual implementations.

In body text, the term IMPLEMENTATION DEFINED is shown in SMALL CAPITALS.

Implementation specific

Behavior that is not architecturally defined, and might not be documented by an individual implementations. Used when there are a number of implementation options available, and the option chosen does not affect software compatibility.

Interrupt grouping

This is a mechanism to align interrupt handling with the ARMv8 Exception model and Security model. Interrupts are configured as belonging to either Group 0 or Group 1. In a system with two Security states interrupts are configured as being in Group 0, Non-secure Group 1, or Secure Group 1.

See also Interrupt grouping on page 4-58.

Interrupt Translation Service (ITS)

An optional hardware mechanism that routes LPIs to the appropriate Redistributor. Software uses a command queue to configure an ITS.

See also The ITS on page 6-99.

Little-endian memory

Means that:

- A byte or halfword at a word-aligned address is the least significant byte or halfword in the word at that address.
- A byte at a halfword-aligned address is the least significant byte in the halfword at that address.

List registers

The List registers are a subset of the GIC virtual interface control registers that define the active and pending virtual interrupts for the virtual CPU interface. List registers indicate whether an interrupt is in Group 0 or Group 1, and therefore whether it is assigned to a virtual IRQ signal or virtual FIQ signal. The scheduled virtual machine accesses these interrupt indirectly, using the virtual CPU interface.

See also List register usage resulting in UNPREDICTABLE behavior on page 5-81.

Locality-specific Peripheral interrupt (LPI)

LPIs are optional message-based interrupts that target a specific PE. They can be routed using an optional ITS. LPIs are always Non-secure Group 1 interrupts, and have edge-triggered behavior.

See also LPIs on page 6-92.

OPTIONAL

When applied to a feature of the architecture, OPTIONAL indicates a feature that is not required in an implementation of the ARM architecture:

- If a feature is OPTIONAL and deprecated, this indicates that the feature is being phased out of the architecture. ARM expects such a features to be included in a new implementation only if there is a known backwards-compatibility reason for the inclusion of the feature.

  A feature that is OPTIONAL and deprecated might not be present in future versions of the architecture.
- A feature that is OPTIONAL but not deprecated is, typically, a feature added to a version of the ARM architecture after the initial release of that version of the architecture. ARM recommends that such features are included in all new implementations of the architecture.

In body text, these meanings of the term OPTIONAL are shown in SMALL CAPITALS.

See also Deprecated.

PE

See Processing element (PE).

Peripheral interrupt

An interrupt generated by the assertion of an interrupt request signal input to the GIC. The GIC architecture defines the following types of peripheral interrupt:

Private Peripheral Interrupt (PPI)

A peripheral interrupt that targets a single, specific PE. PPIs can be either Group 0 or Group 1 interrupts, and they have edge-triggered or level-sensitive behavior.

Shared Peripheral Interrupt (SPI)

A peripheral interrupt that the Distributor can route to a specified PE or to combination of PEs. SPIs can be either Group 0 or Group 1 interrupts, and they have edge-triggered or level-sensitive behavior.

See also *Shared Peripheral Interrupts on page 4-56.*

PPI

See Peripheral interrupt.

Preemption level

A preemption level is a supported group priority.

See also Preemption on page 4-71.

Priority drop A priority drop occurs when the PE signals to the GIC that the highest priority active interrupt has been handled to the point where the priority can be dropped to the priority that the interrupt had prior to being handled.

See also Interrupt lifecycle on page 4-46.

Processing element (PE)
The abstract machine defined in the ARM architecture, as documented in an ARM Architecture Reference Manual. A PE implementation compliant with the ARM architecture must conform with the behaviors described in the corresponding ARM Architecture Reference Manual.

See also ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile.

Quadword A 128-bit data item. Quadwords are normally at least word-aligned in ARM systems.

RAO *See* Read-As-One (RAO).

RAO/WI Read-As-One, Writes Ignored.

Hardware must implement the field as Read-As-One, and must ignore writes to the field.

Software can rely on the field reading as all 1s, and on writes being ignored.

This description can apply to a single bit that reads as 1, or to a field that reads as all 1s.

See also Read-As-One (RAO).

RAZ *See* Read-As-Zero (RAZ).

RAZ/WI Read-As-Zero, Writes Ignored.

Hardware must implement the field as Read-As-Zero, and must ignore writes to the field.

Software can rely on the field reading as all 0s, and on writes being ignored.

This description can apply to a single bit that reads as 0, or to a field that reads as all 0s.

See also Read-As-Zero (RAZ).

Read-As-One (RAO)
Hardware must implement the field as reading as all 1s.

Software:
- Can rely on the field reading as all 1s.
- Must use a SBOP policy to write to the field.

This description can apply to a single bit that reads as 1, or to a field that reads as all 1s.

Read-As-Zero (RAZ)
Hardware must implement the field as reading as all 0s.

Software:
- Can rely on the field reading as all 0s
- Must use a SBZP policy to write to the field.

This description can apply to a single bit that reads as 0, or to a field that reads as all 0s.

Redistributor A logical component in the GIC at affinity level 0 that is part of the Interrupt Routing Infrastructure (IRI). It connects the IRI to the CPU interface. Each PE in the system has a connected Redistributor that routes interrupts to the appropriate PEs. Every PE in the system has a corresponding Redistributor.

See also Redistributor on page 2-33.

RES0 A reserved bit or field with *Should-Be-Zero-or-Preserved (SBZP)* behavior.

Note

The following definition is consistent with that provided in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile,* and therefore has a broad scope.

This term is used for fields in register descriptions, and for fields in architecturally-defined data structures that are held in memory, for example in translation table descriptors.

Note

RES0 is not used in descriptions of instruction encodings.

Within the architecture, there are some cases where a register bit or bitfield:

- Is RES0 in some defined architectural context.
- Has different defined behavior in a different architectural context.

This means the definition of RES0 for register fields is:

If a bit is RES0 in all contexts

It is IMPLEMENTATION DEFINED whether:

1. The bit is hardwired to 0. In this case:
   - Reads of the bit always return 0.
   - Writes to the bit are ignored.
2. The bit can be written. In this case:
   - An indirect write to the register sets the bit to 0.
   - A read of the bit returns the last value successfully written, by either a direct or an indirect write, to the bit.

     If the bit has not been successfully written since reset, then the read of the bit returns the reset value if there is one, or otherwise returns an UNKNOWN value.
   - A direct write to the bit must update a storage location associated with the bit.
   - The value of the bit must have no effect on the operation of the PE, other than determining the value read back from the bit, unless the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* explicitly defines additional properties for the bit.

Whether RES0 bits or fields follow behavior 1 or behavior 2 is IMPLEMENTATION DEFINED on a field-by-field basis.

If a bit is RES0 only in some contexts

When the bit is described as RES0:

- An indirect write to the register sets the bit to 0.
- A read of the bit must return the value last successfully written to the bit, by either a direct or an indirect write, regardless of the use of the register when the bit was written.

  If the bit has not been successfully written since reset, then the read of the bit returns the reset value if there is one, or otherwise returns an UNKNOWN value.
- A direct write to the bit must update a storage location associated with the bit.
- While the use of the register is such that the bit is described as RES0, the value of the bit must have no effect on the operation of the PE, other than determining the value read back from that bit, unless the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* explicitly defines additional properties for the bit.

A bit that is RES0 in a context is reserved for possible future use in that context. To preserve forward compatibility, software:

- Must not rely on the bit reading as 0.
- Must use an SBZP policy to write to the bit.

The RES0 description can apply to bits or bitfields that are read-only, or are write-only:

- For a read-only bit, RES0 indicates that the bit reads as 0, but software must treat the bit as UNKNOWN.
- For a write-only bit, RES0 indicates that software must treat the bit as SBZ.

This RES0 description can apply to a single bit that should be written as its preserved value or as 0, or to a field that should be written as its preserved value or as all 0s.

In body text, the term RES0 is shown in SMALL CAPITALS.

See also Read-As-Zero (RAZ), Should-Be-Zero-or-Preserved (SBZP), UNKNOWN.

RES1 A reserved bit or field with *Should-Be-One-or-Preserved (SBOP)* behavior.

Note

The following definition is consistent with that provided in the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile,* and therefore has a broad scope.

This term is used for fields in register descriptions, and for fields in architecturally-defined data structures that are held in memory, for example in translation table descriptors.

Note

RES1 is not used in descriptions of instruction encodings.

Within the architecture, there are some cases where a register bit or bitfield:

- Is RES1 in some defined architectural context.
- Has different defined behavior in a different architectural context.

This means the definition of RES1 for register fields is:

If a bit is RES1 in all contexts

It is IMPLEMENTATION DEFINED whether:

1. The bit is hardwired to 1. In this case:
   - Reads of the bit always return 1.
   - Writes to the bit are ignored.
2. The bit can be written. In this case:
   - An indirect write to the register sets the bit to 1.
   - A read of the bit returns the last value successfully written, by either a direct or an indirect write, to the bit.

     If the bit has not been successfully written since reset, then the read of the bit returns the reset value if there is one, or otherwise returns an UNKNOWN value.
   - A direct write to the bit must update a storage location associated with the bit.
   - The value of the bit must have no effect on the operation of the PE, other than determining the value read back from the bit, unless the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* explicitly defines additional properties for the bit.

Whether RES1 bits or fields follow behavior 1 or behavior 2 is IMPLEMENTATION DEFINED on a field-by-field basis.

If a bit is RES1 only in some contexts

When the bit is described as RES1:

- An indirect write to the register sets the bit to 1.
- A read of the bit must return the value last successfully written to the bit, regardless of the use of the register when the bit was written.

  Note

  As indicated in this list, this value might be written by an indirect write to the register.

  If the bit has not been successfully written since reset, then the read of the bit returns the reset value if there is one, or otherwise returns an UNKNOWN value.

- A direct write to the bit must update a storage location associated with the bit.
- While the use of the register is such that the bit is described as RES1, the value of the bit must have no effect on the operation of the PE, other than determining the value read back from that bit, unless the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* explicitly defines additional properties for the bit.

A bit that is RES1 in a context is reserved for possible future use in that context. To preserve forward compatibility, software:

- Must not rely on the bit reading as 1.
- Must use an SBOP policy to write to the bit.

The RES1 description can apply to bits or bitfields that are read-only, or are write-only:

- For a read-only bit, RES1 indicates that the bit reads as 1, but software must treat the bit as UNKNOWN.
- For a write-only bit, RES1 indicates that software must treat the bit as SBO.

This RES1 description can apply to a single bit that should be written as its preserved value or as 1, or to a field that should be written as its preserved value or as all 1s.

In body text, the term RES1 is shown in SMALL CAPITALS.

See also Read-As-One (RAO), Should-Be-One-or-Preserved (SBOP), UNKNOWN.

Reserved

Unless otherwise stated:

- Instructions that are reserved or that access reserved registers have UNPREDICTABLE behavior.
- Bit positions described as reserved are:
  - In an RW register, RES0.
  - In an RO register, UNK.
  - In a WO register, RES0.

SBO

See Should-Be-One (SBO).

SBOP

See Should-Be-One-or-Preserved (SBOP).

SBZ

See Should-Be-Zero (SBZ).

SBZP

See Should-Be-Zero-or-Preserved (SBZP).

Should-Be-One (SBO)

Hardware must ignore writes to the field.

ARM strongly recommends that software writes the field as all 1s. If software writes a value that is not all 1s, it must expect an UNPREDICTABLE result.

This description can apply to a single bit that should be written as 1, or to a field that should be written as all 1s.

Should-Be-One-or-Preserved (SBOP)

From the introduction of the ARMv8 architecture, the description of *Should-Be-One-or-Preserved (SBOP)* is superseded by *RES1*.

Hardware must ignore writes to the field.

If software has read the field since the PE implementing the field was last reset and initialized, it must preserve the value of the field by writing the value that it previously read from the field. Otherwise, it must write the field as all 1s.

If software writes a value to the field that is not a value previously read for the field and is not all 1s, it must expect an UNPREDICTABLE result.

This description can apply to a single bit that should be written as its preserved value or as 1, or to a field that should be written as its preserved value or as all 1s.

Should-Be-Zero (SBZ)

Hardware must ignore writes to the field.

ARM strongly recommends that software write the field as all 0s. If software writes a value that is not all 0s, it must expect an UNPREDICTABLE result.

This description can apply to a single bit that should be written as 0, or to a field that should be written as all 0s.

Should-Be-Zero-or-Preserved (SBZP)

From the introduction of the ARMv8 architecture, the description *Should-Be-Zero-or-Preserved (SBZP)* is superseded by *RES0*.

Hardware must ignore writes to the field.

If software has read the field since the PE implementing the field was last reset and initialized, it must preserve the value of the field by writing the value that it previously read from the field. Otherwise, it must write the field as all 0s.

If software writes a value to the field that is not a value previously read for the field and is not all 0s, it must expect an UNPREDICTABLE result.

This description can apply to a single bit that should be written as its preserved value or as 0, or to a field that should be written as its preserved value or as all 0s.

SGI

See Software-generated interrupt (SGI).

Software-generated interrupt (SGI)

An interrupt generated by the GIC in response to software writing to an SGI register in the GIC. SGIs are typically used for inter-processor communication. SGIs can be either Group 0 or Group 1 interrupts, and have edge-triggered behavior.

See also Software Generated Interrupts on page 4-55.

SPI

See Peripheral interrupt

Spurious interrupt

An interrupt that does not require servicing. Usually, refers to an INTID returned by a GIC to a request from a connected PE. Returning a spurious INTID indicates that there is no pending interrupt on the CPU interface that the requesting PE can service.

See also Special INTIDs on page 3-40.

UNDEFINED

Indicates an instruction that is not architecturally defined and generates an Undefined Instruction exception. See the *ARM® Architecture Reference Manual, ARMv7-A and ARMv7-R edition* or the *ARM® Architecture Reference Manual, ARMv8, for ARMv8-A architecture profile* for more information.

UNK

An abbreviation indicating that software must treat a field as containing an UNKNOWN value.

Hardware must implement the bit as read as 0, or all 0s for a multi-bit field. Software must not rely on the field reading as zero.

See also UNKNOWN.

UNK/SBOP

Hardware must implement the field as Read-As-One, and must ignore writes to the field.

Software must not rely on the field reading as all 1s, and except for writing back to the register it must treat the value as if it is UNKNOWN. Software must use an SBOP policy to write to the field.

This description can apply to a single bit that should be written as its preserved value or as 1, or to a field that should be written as its preserved value or as all 1s.

See also Read-As-One (RAO), Should-Be-One-or-Preserved (SBOP), UNKNOWN.

UNK/SBZP

Hardware must implement the bit as Read-As-Zero, and must ignore writes to the field.

Software must not rely on the field reading as all 0s, and except for writing back to the register must treat the value as if it is UNKNOWN. Software must use an SBZP policy to write to the field.

This description can apply to a single bit that should be written as its preserved value or as 0, or to a field that should be written as its preserved value or as all 0s.

See also Read-As-Zero (RAZ), Should-Be-Zero-or-Preserved (SBZP), UNKNOWN.

UNKNOWN An UNKNOWN value does not contain valid data, and can vary from moment to moment, instruction to instruction, and implementation to implementation. An UNKNOWN value must not return information that cannot be accessed at the current or a lower level of privilege using instructions that are not UNPREDICTABLE, are not CONSTRAINED UNPREDICTABLE, and do not return UNKNOWN values.

An UNKNOWN value must not be documented or promoted as having a defined value or effect.

In body text, the term UNKNOWN is shown in SMALL CAPITALS.

See also CONSTRAINED UNPREDICTABLE, UNDEFINED, UNK, UNPREDICTABLE.

UNPREDICTABLE

Means the behavior cannot be relied on. UNPREDICTABLE behavior must not perform any function that cannot be performed at the current or a lower level of privilege using instructions that are not UNPREDICTABLE.

UNPREDICTABLE behavior must not be documented or promoted as having a defined effect.

An instruction that is UNPREDICTABLE can be implemented as UNDEFINED.

In body text, the term UNPREDICTABLE is shown in SMALL CAPITALS.

Valid interrupt ID

An interrupt ID, as returned by a read of ICC_IAR0_EL1 or ICC_IAR1_EL1, that is not a spurious interrupt ID.

See also Interrupt lifecycle on page 4-46.

WI Writes Ignored. In a register that software can write to, a WI attribute applied to a bit or field indicates that the bit or field ignores the value written by software and retains the value it had before that write.

See also RAO/WI, RAZ/WI, RES0, RES1.

Word A 32-bit data item. Words are normally word-aligned in ARM systems.