



Arm[®] Generic Interrupt Controller Architecture Specification

GIC architecture version 5

Document number	ARM-AES-0070
Document quality	BET
Document version	00bet0
Document confidentiality	Non-confidential
Document build information	0b7a9f48 doctool 0.55.1

Copyright © 2022-2025 Arm Limited or its affiliates. All rights reserved.

Arm® Generic Interrupt Controller Architecture Specification, GIC architecture version 5

Release information

Date	Version	Changes
2025/Mar/28	00bet0	<ul style="list-style-type: none">• First BET release.

Non-Confidential Proprietary Notice

This document is protected by copyright and other related rights and the use or implementation of the information contained in this document may be protected by one or more patents or pending patent applications. No part of this document may be reproduced in any form by any means without the express prior written permission of Arm Limited (“Arm”). **No license, express or implied, by estoppel or otherwise to any intellectual property rights is granted by this document unless specifically stated.**

Your access to the information in this document is conditional upon your acceptance that you will not use or permit others to use the information for the purposes of determining whether the subject matter of this document infringes any third party patents.

The content of this document is informational only. Any solutions presented herein are subject to changing conditions, information, scope, and data. This document was produced using reasonable efforts based on information available as of the date of issue of this document. The scope of information in this document may exceed that which Arm is required to provide, and such additional information is merely intended to further assist the recipient and does not represent Arm’s view of the scope of its obligations. You acknowledge and agree that you possess the necessary expertise in system security and functional safety and that you shall be solely responsible for compliance with all legal, regulatory, safety and security related requirements concerning your products, notwithstanding any information or support that may be provided by Arm herein. In addition, you are responsible for any applications which are used in conjunction with any Arm technology described in this document, and to minimize risks, adequate design and operating safeguards should be provided for by you.

This document may include technical inaccuracies or typographical errors. THIS DOCUMENT IS PROVIDED “AS IS”. ARM PROVIDES NO REPRESENTATIONS AND NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT OR FITNESS FOR A PARTICULAR PURPOSE WITH RESPECT TO THE DOCUMENT. For the avoidance of doubt, Arm makes no representation with respect to, and has undertaken no analysis to identify or understand the scope and content of, any patents, copyrights, trade secrets, trademarks, or other rights.

TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL ARM BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF ARM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Reference by Arm to any third party’s products or services within this document is not an express or implied approval or endorsement of the use thereof.

This document consists solely of commercial items. You shall be responsible for ensuring that any permitted use, duplication, or disclosure of this document complies fully with any relevant export laws and regulations to assure that this document or any portion thereof is not exported, directly or indirectly, in violation of such export laws. Use of the word “partner” in reference to Arm’s customers is not intended to create or refer to any partnership relationship with any other company. Arm may make changes to this document at any time and without notice.

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

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

The Arm corporate logo and words marked with ® or ™ are registered trademarks or trademarks of Arm Limited (or its affiliates) in the US and/or elsewhere. Please follow Arm’s trademark usage guidelines at <https://www.arm.com/company/policies/trademarks>. All rights reserved. Other brands and names mentioned in this document may be the trademarks of their respective owners.

Copyright © 2022-2025 Arm Limited or its affiliates. All rights reserved.

Arm Limited. Company 02557590 registered in England.

110 Fulbourn Road, Cambridge, England CB1 9NJ.

PRE-20349

8 March 2024

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 at Beta quality. Beta quality means that all major features of the specification are described, but some details might be missing.

Web Address

<https://www.arm.com>

Contents

Arm® Generic Interrupt Controller Architecture Specification, GIC architecture version 5

Arm® Generic Interrupt Controller Architecture Specification, GIC architecture version 5	ii
Release information	ii
Non-Confidential Proprietary Notice	iii

Preface

Conventions	xvi
Typographical conventions	xvi
Numbers	xvi
Pseudocode descriptions	xvi
Assembler syntax descriptions	xvi
Rules-based writing	xvii
Content item identifiers	xvii
Content item rendering	xvii
Content item classes	xvii
Additional reading	xviii
Feedback	xx
Feedback on this book	xx
Inclusive terminology commitment	xx

Chapter 1

Introduction

Chapter 2

PE architecture

2.1	Architecture features and extensions	23
2.2	The GICv5 CPU interface	25
2.3	Interrupt Domains	27
2.4	Interrupt types and identifiers	29
2.4.1	PE-Private Peripheral Interrupts (PPIs)	30
2.4.2	Logical Peripheral Interrupts (LPIs)	31
2.4.3	Shared Peripheral Interrupts (SPIs)	31
2.5	Inter-Processor Interrupts	33
2.6	GIC System instructions	35
2.6.1	LPI and SPI configuration	37
2.7	Interrupt Prioritization	40
2.8	Interrupt handling	42
2.8.1	Interrupt life cycle	43
2.9	The physical CPU interface	46
2.9.1	Physical PPIs	46
2.9.2	Physical priority masking	48
2.9.3	Preemptive interrupts	49
2.9.4	Physical interrupt signaling	50
2.9.5	Physical non-maskable interrupts	52
2.9.6	Doorbell PPIs	53
2.10	The virtual CPU interface	56
2.10.1	Virtual PPIs	56
2.10.2	Virtual priority masking	59
2.10.3	Virtual interrupt signaling	60
2.10.4	Virtual non-maskable interrupts	62

2.10.5	Selecting the resident VPE	62
2.10.6	Requesting VPE doorbells	64
2.10.7	Legacy virtual CPU interface	64
2.11	GIC synchronous exception priorities	67
2.12	Interrupt ordering model and synchronization requirements	68
2.12.1	GIC and GICR ordering semantics	69
2.12.2	GSB instruction semantics	71
2.12.3	GIC Ordering Model	71
2.13	Effects on the Transactional Memory Extension	76

Chapter 3

GICv5 system architecture

3.1	Interrupt Domains	81
3.2	Communication between GIC system components	83
3.3	Coherency considerations for GIC data structures	84

Chapter 4

Interrupt routing service (IRS)

4.1	Communication between the IRS and the CPU interface	87
4.2	Signaling interrupts	88
4.3	IRS Domains	90
4.4	IRS Configuration	92
4.4.1	Enabling and disabling the IRS	94
4.5	IRS synchronization requests	95
4.6	Interrupt configuration and state	97
4.7	The interrupt state table (IST)	99
4.7.1	Level 2 IST management	101
4.7.2	Initialization of level 2 IST entries	103
4.7.3	INTID state and configuration	103
4.7.4	IST metadata	104
4.7.5	Example IST structures	105
4.8	Physical interrupts	107
4.8.1	Physical LPIs	107
4.8.2	Physical SPIs	108
4.8.3	Physical interrupt routing	112
4.8.4	Physical interrupt signaling	114
4.9	Virtualization data structures	116
4.9.1	The VM table	117
4.9.2	The VPE table	122
4.10	Virtual interrupts	126
4.10.1	Virtual LPIs	126
4.10.2	Virtual SPIs	129
4.10.3	Virtual interrupt routing	132
4.10.4	Virtual interrupt signaling	133
4.10.5	VPE selection and configuration	135
4.10.6	VPE residency	136
4.10.7	VPE doorbells	136
4.10.8	1ofN doorbells	138
4.10.9	Save and restore of virtual interrupts	138
4.11	IRS power management	141
4.12	IRS memory access rules	143
4.13	IRS support for MPAM	146
4.14	IRS support for Memory Encryption Contexts	148
4.15	IRS support for software error reporting	149

Chapter 5

Interrupt translation service (ITS)

5.1	ITS Domains	155
-----	-----------------------	-----

5.1.1	Supporting Realm interrupts from Non-secure writes	156
5.2	Operation	158
5.2.1	Enabling and disabling the ITS	158
5.2.2	Interrupt event types	159
5.2.3	Software generated ITS events	159
5.2.4	ITS synchronization requests	160
5.3	Translation structures	161
5.3.1	The Device Table (DT)	162
5.3.2	The Interrupt Translation Table (ITT)	164
5.4	ITS cache management	167
5.4.1	ITS cache management for EventIDs	168
5.4.2	ITS cache management for DeviceIDs	169
5.5	ITS memory access rules	171
5.6	ITS support for MPAM	172
5.7	ITS support for Memory Encryption Contexts	173
5.8	ITS support for software error reporting	174
Chapter 6	Interrupt Wire Bridge (IWB)	
6.1	IWB wire control registers	180
6.2	IWB support for multiple Interrupt Domains	183
Chapter 7	GIC Performance Monitoring Unit (PMU)	
7.1	CoreSight PMU extensions	187
7.2	GIC PMU Overflow interrupt	189
7.3	GIC PMU event types	190
7.4	Event filtering	191
7.5	IRS PMU events	192
7.5.1	IRS PMU events filtering	194
7.6	ITS PMU events	196
7.6.1	ITS PMU events filtering	196
Chapter 8	System instructions	
8.1	System instructions for the Current Interrupt Domain	199
8.1.1	GIC CDAFF, Interrupt Set Target in the Current Interrupt Domain	200
8.1.2	GIC CDDI, Interrupt Deactivate in the Current Interrupt Domain	202
8.1.3	GIC CDDIS, Interrupt Disable in the Current Interrupt Domain	204
8.1.4	GIC CDEN, Interrupt Enable in the Current Interrupt Domain	206
8.1.5	GIC CDEOI, Priority Drop in the Current Interrupt Domain	208
8.1.6	GIC CDHM, Interrupt Handling mode state in the Current Interrupt Domain	210
8.1.7	GIC CDPEND, Interrupt Set/Clear Pending state in the Current Interrupt Domain	212
8.1.8	GIC CDPRI, Interrupt Set priority in the Current Interrupt Domain	214
8.1.9	GIC CDRCFG, Request Interrupt Configuration in the Current Interrupt Domain	216
8.1.10	GICR CDIA, Interrupt Acknowledge in the Current Interrupt Domain	218
8.1.11	GICR CDNMA, Non-maskable Interrupt Acknowledge in the Current Interrupt Domain	220
8.2	System instructions for the Virtual Interrupt Domain	222
8.2.1	GIC VDAFF, Interrupt Set Target in the Virtual Interrupt Domain	223
8.2.2	GIC VDDI, Interrupt Deactivate in the Virtual Interrupt Domain	225
8.2.3	GIC VDDIS, Interrupt Disable in the Virtual Interrupt Domain	227
8.2.4	GIC VDEN, Interrupt Enable in the Virtual Interrupt Domain	229
8.2.5	GIC VDHM, Interrupt Handling mode in the Virtual Interrupt Domain	231
8.2.6	GIC VDPEND, Interrupt Set/Clear Pending state in the Virtual Interrupt Domain	233

8.2.7	GIC VDPRI, Interrupt Set priority in the Virtual Interrupt Domain	235
8.2.8	GIC VDRCFG, Request Interrupt Configuration in the Virtual Interrupt Domain	237
8.3	System instructions for the Logical Interrupt Domain	239
8.3.1	GIC LDAFF, Interrupt Set Target in the Logical Interrupt Domain	240
8.3.2	GIC LDDI, Interrupt Deactivate in the Logical Interrupt Domain	242
8.3.3	GIC LDDIS, Interrupt Disable in the Logical Interrupt Domain	244
8.3.4	GIC LDEN, Interrupt Enable in the Logical Interrupt Domain	246
8.3.5	GIC LDHM, Interrupt Handling mode in the Logical Interrupt Domain . .	248
8.3.6	GIC LDPEND, Interrupt Set/Clear Pending state in the Logical Interrupt Domain	250
8.3.7	GIC LDPRI, Interrupt Set priority in the Logical Interrupt Domain	252
8.3.8	GIC LDRCFG, Request Interrupt Configuration in the Logical Interrupt Domain	254
8.4	GIC synchronization barrier instructions	256
8.4.1	GSB SYS, GIC Synchronization Barrier System	257
8.4.2	GSB ACK, GIC Synchronization Barrier Interrupt Acknowledge	258

Chapter 9

System registers

9.1	Synchronization requirements for GICv5 System registers	260
9.2	CPU interface registers	261
9.2.1	ICC_APR_EL1, Interrupt Controller Physical Active Priorities Register .	262
9.2.2	ICC_APR_EL3, Interrupt Controller Physical Active Priorities Register for EL3	265
9.2.3	ICC_CR0_EL1, Interrupt Controller EL1 Physical Control Register . . .	267
9.2.4	ICC_CR0_EL3, Interrupt Controller EL3 Physical Control Register . . .	271
9.2.5	ICC_DOMHPPIR_EL3, Interrupt Controller Domain Highest Priority Pending Interrupt Register	273
9.2.6	ICC_HAPR_EL1, Interrupt Controller Physical Highest Active Priority Register	275
9.2.7	ICC_HPPIR_EL1, Interrupt Controller Physical Highest Priority Pending Interrupt Register	277
9.2.8	ICC_HPPIR_EL3, Interrupt Controller Physical Highest Priority Pending Interrupt Register	279
9.2.9	ICC_IAFFIDR_EL1, Interrupt Controller PE Interrupt Affinity ID Register	281
9.2.10	ICC_ICSR_EL1, Interrupt Controller Interrupt Configuration and State Register	282
9.2.11	ICC_IDR0_EL1, Interrupt Controller ID Register 0	286
9.2.12	ICC_PCR_EL1, Interrupt Controller Physical Interrupt Priority Control Register	288
9.2.13	ICC_PCR_EL3, Interrupt Controller Interrupt Priority Control Register for EL3	291
9.2.14	ID_AA64PFR2_EL1, AArch64 Processor Feature Register 2	293
9.3	Virtual CPU interface registers	295
9.3.1	ICV_APR_EL1, Interrupt Controller Virtual Active Priorities Register . .	296
9.3.2	ICV_CR0_EL1, Interrupt Controller EL1 Virtual Control Register	299
9.3.3	ICV_HAPR_EL1, Interrupt Controller Virtual Highest Active Priority Register	302
9.3.4	ICV_HPPIR_EL1, Interrupt Controller Virtual Highest Priority Pending Interrupt Register	304
9.3.5	ICV_PCR_EL1, Interrupt Controller Virtual Interrupt Priority Control Register	306
9.4	PPI registers	309
9.4.1	ICC_PPI_CACTIVER<n>_EL1, Interrupt Controller Physical PPI Clear Active Registers, n = 0 - 1	310

9.4.2	ICC_PPI_CPENDR<n>_EL1, Interrupt Controller Physical PPI Clear Pending State Registers, n = 0 - 1	313
9.4.3	ICC_PPI_DOMAINR<n>_EL3, Interrupt Controller PPI Domain Registers, n = 0 - 3	316
9.4.4	ICC_PPI_ENABLER<n>_EL1, Interrupt Controller Physical PPI Enable Registers, n = 0 - 1	318
9.4.5	ICC_PPI_HMR<n>_EL1, Interrupt Controller Physical PPI Handling mode Registers, n = 0 - 1	321
9.4.6	ICC_PPI_PRIORITYR<n>_EL1, Interrupt Controller Physical PPI Priority Registers, n = 0 - 15	323
9.4.7	ICC_PPI_SACTIVER<n>_EL1, Interrupt Controller Physical PPI Set Active Registers, n = 0 - 1	325
9.4.8	ICC_PPI_SPENDR<n>_EL1, Interrupt Controller Physical PPI Set Pending State Registers, n = 0 - 1	328
9.5	Virtual PPI registers	331
9.5.1	ICV_PPI_CACTIVER<n>_EL1, Interrupt Controller Virtual PPI Clear Active Registers, n = 0 - 1	332
9.5.2	ICV_PPI_CPENDR<n>_EL1, Interrupt Controller Virtual PPI Clear Pending State Registers, n = 0 - 1	335
9.5.3	ICV_PPI_ENABLER<n>_EL1, Interrupt Controller Virtual PPI Clear Enable Registers, n = 0 - 1	338
9.5.4	ICV_PPI_HMR<n>_EL1, Interrupt Controller Virtual PPI Handling mode Registers, n = 0 - 1	341
9.5.5	ICV_PPI_PRIORITYR<n>_EL1, Interrupt Controller Virtual PPI Priority Registers, n = 0 - 15	343
9.5.6	ICV_PPI_SACTIVER<n>_EL1, Interrupt Controller Virtual PPI Set Active Registers, n = 0 - 1	345
9.5.7	ICV_PPI_SPENDR<n>_EL1, Interrupt Controller Virtual PPI Set Pending State Registers, n = 0 - 1	348
9.6	Hypervisor control registers	351
9.6.1	ICH_APR_EL2, Interrupt Controller Active Virtual Priorities Register	352
9.6.2	ICH_CONTEXTR_EL2, Interrupt Controller Virtual Context Register	354
9.6.3	ICH_HFGITR_EL2, Hypervisor GIC Fine-Grained Instruction Trap Register	358
9.6.4	ICH_HFGRTR_EL2, Hypervisor GIC Fine-Grained Read Trap Register	363
9.6.5	ICH_HFGWTR_EL2, Hypervisor GIC Fine-Grained Write Trap Register	369
9.6.6	ICH_HPPIR_EL2, Interrupt Controller Hypervisor Highest Priority Pending Interrupt Register	374
9.6.7	ICH_PPI_ACTIVER<n>_EL2, Interrupt Controller Virtual Interrupt Active Registers, n = 0 - 1	376
9.6.8	ICH_PPI_DVIR<n>_EL2, Interrupt Controller PPI Direct-inject Virtual Interrupt Registers, n = 0 - 1	379
9.6.9	ICH_PPI_ENABLER<n>_EL2, Interrupt Controller Virtual Interrupt Enable Registers, n = 0 - 1	382
9.6.10	ICH_PPI_PENDR<n>_EL2, Interrupt Controller Virtual Interrupt Pending State Registers, n = 0 - 1	384
9.6.11	ICH_PPI_PRIORITYR<n>_EL2, Interrupt Controller Virtual Interrupt Priority Registers, n = 0 - 15	386
9.6.12	ICH_VCTLR_EL2, Interrupt Controller Virtual CPU interface Control Register	388
9.6.13	ICH_VMCR_EL2, Interrupt Controller Virtual Machine Control Register	391
9.6.14	Nested virtualization	393
9.7	Legacy hypervisor control registers	395
9.7.1	ICH_AP0R<n>_EL2, Interrupt Controller Active Virtual Priorities Registers 0, n = 0 - 3	396

9.7.2	ICH_AP1R<n>_EL2, Interrupt Controller Active Virtual Priorities Registers 1, n = 0 - 3	398
9.7.3	ICH_EISR_EL2, Interrupt Controller End of Interrupt Status Register . .	401
9.7.4	ICH_ELRSR_EL2, Interrupt Controller Empty List Register Status Register	403
9.7.5	ICH_HCR_EL2, Interrupt Controller Hyp Control Register	405
9.7.6	ICH_LR<n>_EL2, Interrupt Controller List Registers, n = 0 - 15	409
9.7.7	ICH_MISR_EL2, Interrupt Controller Maintenance Interrupt State Register	412
9.7.8	ICH_VTR_EL2, Interrupt Controller VGIC Type Register	415
9.7.9	Nested virtualization	417
9.8	Legacy virtual CPU interface registers	418
9.8.1	AArch64 Legacy virtual CPU interface registers	418

Chapter 10

Registers and memory maps

10.1	Memory-mapped programmer's model	421
10.2	IRS register frames	424
10.2.1	IRS_CONFIG_FRAME, IRS configuration register frame	425
10.2.2	IRS_SETLPI_FRAME, IRS SETLPI register frame	521
10.3	ITS register frames	523
10.3.1	ITS_CONFIG_FRAME, ITS configuration register frame	524
10.3.2	ITS_TRANSLATE_FRAME, ITS translate register frame	575
10.4	IWB register frames	578
10.4.1	IWB_CONFIG_FRAME, IWB configuration registers frame	579
10.5	GIC PMU register frame	595
10.5.1	GIC_PMU_FRAME, GIC PMU register frame	596
10.6	Identification registers	614

Chapter 11

Data structures

11.1	ITS Data Structures	617
11.1.1	L1_DTE, Level 1 device table entry	618
11.1.2	L2_DTE, Level 2 device table entry	619
11.1.3	L1_ITTE, Level 1 interrupt translation table entry	622
11.1.4	L2_ITTE, Level 2 interrupt translation table entry	623
11.2	IRS Data Structures	625
11.2.1	L1_VMTE, Level 1 VM table entry	626
11.2.2	L2_VMTE, Level 2 VM table entry	627
11.2.3	L1_ISTE, Level 1 interrupt state table entry	633
11.2.4	L2_ISTE, Level 2 interrupt state table entry	634
11.2.5	VPETE, VPE table entry	636
11.2.6	VM_DESC, VM descriptor	637
11.2.7	VPE_DESC, VPE descriptor	638

Part A GICv5 Stream Protocol interface

Chapter A1 GICv5 Stream Protocol overview

Chapter A2 AMBA AXI5-Stream Transport Layer

A2.1	Signals	644
A2.2	Channel identification	646
A2.3	Link status	648

Chapter A3 Common behaviors

Chapter A4 Interrupt Handling channel

A4.1	Command summary	651
A4.2	Outstanding commands	653

A4.3	Connection management	655
A4.4	Managing the resident VPE	657
A4.4.1	Interrupt Handling Channel behaviors when there is a resident VPE	657
A4.4.2	Interrupt Handling Channel behaviors when there is no resident VPE	658
A4.5	Forwarding, recalling, and releasing interrupts	659
A4.6	INTID configuration	661
A4.7	IRS and CPU interface capabilities	662

Chapter A5

Interrupt Signaling channel

A5.1	Command summary	664
A5.2	Outstanding commands	665
A5.3	Signaling interrupts to the IRS	666
A5.4	Connection management	667

Chapter A6

Alphabetical list of commands

A6.1	Interrupt Handling channel	669
A6.1.1	Activate, Activate command (CPUIF -> IRS)	670
A6.1.2	ActivateAck, Activate Acknowledge command (IRS -> CPUIF)	671
A6.1.3	Deactivate, Deactivate interrupt command (CPUIF -> IRS)	672
A6.1.4	DeactivateAck, Deactivate interrupt Acknowledge command (IRS -> CPUIF)	674
A6.1.5	DownstreamControl, Downstream Control (IRS -> CPUIF)	675
A6.1.6	DownstreamControlAck, Downstream Control Acknowledge command (CPUIF -> IRS)	676
A6.1.7	Forward, Forward command (IRS -> CPUIF)	677
A6.1.8	Recall, Recall command (IRS -> CPUIF)	679
A6.1.9	Release, Release command (CPUIF -> IRS)	680
A6.1.10	RequestConfig, Request Interrupt Configuration command (CPUIF -> IRS)	682
A6.1.11	RequestConfigAck, Request Interrupt Configuration Acknowledge command (IRS -> CPUIF)	684
A6.1.12	SetAck, Set interrupt configuration acknowledge command (IRS -> CPUIF)	686
A6.1.13	SetEnabled, Set interrupt Enabled command (CPUIF -> IRS)	687
A6.1.14	SetHandling, Set Interrupt Handling Mode command (CPUIF -> IRS)	689
A6.1.15	SetPending, Set interrupt Pending command (CPUIF -> IRS)	691
A6.1.16	SetPriority, Set Interrupt Priority command (CPUIF -> IRS)	693
A6.1.17	SetResident, Set Resident command (CPUIF -> IRS)	695
A6.1.18	SetResidentAck, Set Resident acknowledge command (IRS -> CPUIF)	697
A6.1.19	SetTarget, Set Interrupt Target command (CPUIF -> IRS)	698
A6.1.20	Sync, synchronizes previously sent configuration changes (CPUIF -> IRS)	700
A6.1.21	SyncAck, synchronizes previously sent configuration changes (IRS -> CPUIF)	701
A6.1.22	UpstreamControl, Upstream Control (CPUIF -> IRS)	702
A6.1.23	UpstreamControlAck, Upstream Control Acknowledge command (IRS -> CPUIF)	703
A6.1.24	WakeRequest, Wake request (IRS -> CPUIF)	704
A6.2	Interrupt Signaling channel	705
A6.2.1	INT, Interrupt command (Interrupt Source -> IRS)	706
A6.2.2	Flush, Flush command (IRS -> Interrupt Source)	707
A6.2.3	FlushAck, Flush acknowledge command (Interrupt Source -> IRS)	708
A6.2.4	Quiesce, Quiesce command (Interrupt Source -> IRS)	709
A6.2.5	QuiesceAck, Quiesce acknowledge command (IRS -> Interrupt Source)	710
A6.2.6	Resample, Resample request command (IRS -> Interrupt Source)	711
A6.2.7	ResampleAck, Resample request acknowledge command (Interrupt Source -> IRS)	712
A6.2.8	Reset, Reset command (Interrupt Source -> IRS)	713

A6.2.9	ResetAck, Reset acknowledge command (IRS -> Interrupt Source) . . .	714
--------	---	-----

Chapter A7

Example sequences

A7.1	Bringing the Interrupt Handling channel online and taking it offline	716
A7.2	Simple interrupt life-cycle	721
A7.3	Replacing the candidate HPPI for an Interrupt Domain, or resident VPE	722
A7.4	Making a VPE resident	725
A7.5	Interrupt configuration	728
A7.6	Sending IPIs	729
A7.7	1 of N interrupts	730

Part B Litmus tests

Chapter B1

Interrupt ordering litmus tests

B1.1	Interrupt litmus test assumptions	733
B1.2	Atomicity of interrupt updates by GIC system instructions	734
B1.2.1	Notes	734
B1.2.2	Litmus test	734
B1.3	Multiple updates of the same Interrupt Location	735
B1.3.1	Notes	735
B1.3.2	Litmus test with two configuration updates	735
B1.3.3	Litmus test with an interrupt disable and an interrupt deactivate	735
B1.3.4	Litmus test with an interrupt deactivate and an interrupt disable	736
B1.4	Reading back interrupt writes on a single PE	737
B1.4.1	Notes	737
B1.4.2	Litmus test with a configuration update	737
B1.4.3	Litmus test with deactivate	737
B1.4.4	Litmus test with acknowledgement	738
B1.4.5	Litmus test with two configuration updates	738
B1.5	Reading interrupt configurations and subsequent updates	740
B1.5.1	Notes	740
B1.5.2	Litmus test with update to priority	740
B1.5.3	Litmus test with deactivate	740
B1.5.4	Litmus test with acknowledge	741
B1.6	Configuration and acknowledgement	742
B1.6.1	Notes	742
B1.6.2	Litmus test using disable without explicit synchronization	742
B1.6.3	Litmus test using disable with explicit synchronization	742
B1.6.4	Litmus test with deactivate	743
B1.6.5	Litmus test using priority without explicit synchronization	743
B1.6.6	Litmus test using priority with explicit synchronization.	744
B1.7	Acknowledge followed by interrupt changes	745
B1.7.1	Notes	745
B1.7.2	Litmus test with deactivate	745
B1.7.3	Litmus test with make pending	745
B1.8	Multiple updates with interleaved read	747
B1.8.1	Notes	747
B1.8.2	Litmus test	747
B1.9	Configuration write and IRQ unmask in PSTATE	748
B1.9.1	Notes	748
B1.9.2	Litmus test	748
B1.9.3	Litmus test with initially masked IRQs	748
B1.9.4	Litmus test with disable of a PPI	749
B1.10	Configuration write and exception status on a single PE	750

B1.10.1	Notes	750
B1.10.2	Litmus test without wait for IRQ exception to be signaled	750
B1.10.3	Litmus test with wait for IRQ exception to be signaled	750
B1.11	IPI and acknowledgement	752
B1.11.1	Notes	752
B1.11.2	Litmus test without explicit synchronization	752
B1.11.3	Litmus test with explicit synchronization	752
B1.12	Observing multiple writes on a different PE	754
B1.12.1	Notes	754
B1.12.2	Litmus test	754
B1.13	Read of the configuration of an interrupt	755
B1.13.1	Notes	755
B1.13.2	Litmus test without ISB	755
B1.13.3	Litmus test with ISB	755
B1.14	Multiple reads of the same config	757
B1.14.1	Notes	757
B1.14.2	Litmus test with ISBs	757
B1.15	GIC coherence order	758
B1.15.1	Notes	758
B1.15.2	Litmus test	758
B1.16	Independent reads of independent writes	760
B1.16.1	Notes	760
B1.16.2	Litmus test	760
B1.17	Message passing via flag in memory	761
B1.17.1	Notes	761
B1.17.2	Litmus test	761
B1.18	Message passing via interrupt priority configuration	762
B1.18.1	Notes	762
B1.18.2	Litmus test	762
B1.19	Message passing with an LPI and a device read	763
B1.19.1	Notes	763
B1.19.2	Litmus test	763
B1.19.3	Litmus test with address dependency	763
B1.20	Message passing with an LPI and a GSB	765
B1.20.1	Notes	765
B1.20.2	Litmus test	765
B1.21	Message passing with an IPI and a GSB	767
B1.21.1	Notes	767
B1.21.2	Litmus test	767
B1.22	Message passing using deactivate	768
B1.22.1	Notes	768
B1.22.2	Litmus test with explicit synchronization	768
B1.22.3	Litmus test with address dependency	769
B1.22.4	Litmus test with control dependency	769
B1.22.5	Litmus test without explicit synchronization	770
B1.22.6	Litmus test with a DSB but without a GSB	771
B1.22.7	Litmus test without a DSB but with a GSB	772
B1.23	IPI edge merging and message passing	773
B1.23.1	Notes	773
B1.23.2	Litmus test	773
B1.24	Device edge merging with GSB ACK	775
B1.24.1	Notes	775
B1.24.2	Litmus test	775
B1.25	Configuration read and interrupt acknowledge	777
B1.25.1	Notes	777

B1.25.2	Single-threaded litmus test	777
B1.25.3	Multi-threaded litmus test	777
B1.26	Atomicity of interrupt acknowledge	779
B1.26.1	Notes	779
B1.26.2	Litmus test	779
B1.27	Atomicity of interrupt disable and acknowledge	780
B1.27.1	Notes	780
B1.27.2	Litmus test	780
B1.28	Interrupt handler completion	781
B1.28.1	Notes	781
B1.28.2	Litmus test	781
B1.29	Configuration update while disabled	782
B1.29.1	Notes	782
B1.29.2	Litmus test	782
B1.30	Atomicity of interrupt acknowledge and retarget	784
B1.30.1	Notes	784
B1.30.2	Litmus test	784
B1.31	1ofN interrupt acknowledge	785
B1.31.1	Notes	785
B1.31.2	Litmus test	785
B1.32	Retargeting interrupts without synchronization	786
B1.32.1	Notes	786
B1.32.2	Litmus test	786
B1.33	Reading interrupt configuration and exception status	787
B1.33.1	Notes	787
B1.33.2	Litmus test with ISB before reading IRQ pending status	787
B1.33.3	Litmus test with ISB after reading IRQ pending status	787
B1.34	Reading interrupt configuration and IRQ unmask in PSTATE	789
B1.34.1	Notes	789
B1.34.2	Explanation	789
B1.34.3	Litmus test with ISB before unmasking IRQ in PSTATE	789
B1.34.4	Litmus test with ISB after unmasking IRQ in PSTATE	789
B1.35	PPI activate and system register read	790
B1.35.1	Notes	790
B1.35.2	Litmus test without ISB	790
B1.35.3	Litmus test with GSB	790
B1.35.4	Litmus test with ISB	791
B1.36	PPI disable and acknowledge	792
B1.36.1	Notes	792
B1.36.2	Litmus test with ISB	792
B1.37	PPI acknowledgement	793
B1.37.1	Notes	793
B1.37.2	Litmus test	793
B1.37.3	Litmus test	793
B1.38	Write after changing resident VM	795
B1.38.1	Notes	795
B1.38.2	Litmus test	795
B1.39	Write before changing resident VM	796
B1.39.1	Notes	796
B1.39.2	Litmus test	796
B1.40	Completion of GIC and GICR instructions in finite time	797
B1.40.1	Notes	797
B1.40.2	Litmus test with a priority update observed by a configuration read	797
B1.40.3	Litmus test with an interrupt acknowledge observed by a configuration read	797
B1.40.4	Litmus test with an interrupt becoming pending and then acknowledged	798

Chapter B2	Effects of disabling a PPI source on the PPI Pending state	
	B2.0.1 Notes	799
	B2.0.2 Litmus test with disable of the timer state	799

Part C Model

Chapter C1	Operational model
-------------------	--------------------------

Glossary

Preface

Conventions

Typographical 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 some common terms such as IMPLEMENTATION DEFINED.

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

Red text

Indicates an open issue.

Blue text

Indicates a link. This can be

- A cross-reference to another location within the document
- A URL, for example <http://developer.arm.com>

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. To improve readability, long numbers can be written with an underscore separator between every four characters, for example 0xFFFF_0000_0000_0000. Ignore any underscores when interpreting the value of a number.

Pseudocode descriptions

This book uses a form of pseudocode to provide precise descriptions of the specified functionality. This pseudocode is written in a monospace font. The pseudocode language is described in the Arm Architecture Reference Manual.

Assembler syntax descriptions

This book contains numerous syntax descriptions for assembler instructions and for components of assembler instructions. These are shown in a monospace font.

Rules-based writing

This specification consists of a set of individual *content items*. A content item is classified as one of the following:

- Declaration
- Rule
- Goal
- Information
- Rationale
- Implementation note
- Software usage

Declarations and Rules are normative statements. An implementation that is compliant with this specification must conform to all Declarations and Rules in this specification that apply to that implementation.

Declarations and Rules must not be read in isolation. Where a particular feature is specified by multiple Declarations and Rules, these are generally grouped into sections and subsections that provide context. Where appropriate, these sections begin with a short introduction.

Arm strongly recommends that implementers read *all* chapters and sections of this document to ensure that an implementation is compliant.

Content items other than Declarations and Rules are informative statements. These are provided as an aid to understanding this specification.

Content item identifiers

A content item may have an associated identifier which is unique among content items in this specification.

After this specification reaches beta status, a given content item has the same identifier across subsequent versions of the specification.

Content item rendering

In this document, a content item is rendered with a token of the following format in the left margin: L_{iiii}

- L is a label that indicates the content class of the content item.
- $iiii$ is the identifier of the content item.

Content item classes

Declaration

A Declaration is a statement that does one or more of the following:

- Introduces a concept
- Introduces a term
- Describes the structure of data
- Describes the encoding of data

A Declaration does not describe behavior.

A Declaration is rendered with the label D .

Rule

A Rule is a statement that describes the behavior of a compliant implementation.

A Rule explains what happens in a particular situation.

A Rule does not define concepts or terminology.

A Rule is rendered with the label *R*.

Goal

A Goal is a statement about the purpose of a set of rules.

A Goal explains why a particular feature has been included in the specification.

A Goal is comparable to a “business requirement” or an “emergent property.”

A Goal is intended to be upheld by the logical conjunction of a set of rules.

A Goal is rendered with the label *G*.

Information

An Information statement provides information and guidance as an aid to understanding the specification.

An Information statement is rendered with the label *I*.

Rationale

A Rationale statement explains why the specification was specified in the way it was.

A Rationale statement is rendered with the label *X*.

Implementation note

An Implementation note provides guidance on implementation of the specification.

An Implementation note is rendered with the label *U*.

Software usage

A Software usage statement provides guidance on how software can make use of the features defined by the specification.

A Software usage statement is rendered with the label *S*.

Additional reading

This section lists publications by Arm and by third parties.

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

- [1] *Arm® Architecture Reference Manual, for A-profile architecture*. (ARM DDI 0487) Arm Ltd.
- [2] *Arm® Base System Architecture 1.0C*. (ARM DEN 0094C) Arm Ltd.
- [3] *Arm® Generic Interrupt Controller Architecture Specification, GIC architecture version 3 and version 4*. (ARM IHI 0069) Arm Ltd.
- [4] *Arm® Architecture Reference Manual Supplement, Transactional Memory Extension (TME), for A-profile architecture*. (ARM DDI 0617) Arm Ltd.
- [5] *Arm® System Memory Management Unit Architecture Specification, SMMU architecture version 3*. (ARM IHI 0070) Arm Ltd.
- [6] *Arm® Realm Management Extension (RME) System Architecture*. (ARM DEN 0029) Arm Ltd.
- [7] *AMBA® AXI Protocol Specification*. (ARM IHI 0022) Arm Ltd.
- [8] *Arm® Power State Coordination Interface*. (ARM DEN 0022) Arm Ltd.
- [9] *PCI Express® Base Specification Revision 6.0*. PCI-SIG.

- [10] *VIRTIO (Virtual I/O) Specification*. See <https://github.com/oasis-tcs/virtio-spec>
- [11] *Arm® CoreSight™ Architecture Performance Monitoring Unit Architecture*. (ARM IHI 0091) Arm Ltd.
- [12] *Arm® CoreSight™ Architecture Specification*. (ARM IHI 0029) Arm Ltd.
- [13] *AMBA® AXI-Stream Protocol Specification*. (ARM IHI 0051) Arm Ltd.
- [14] *Arm® Specification Language Reference Manual*. (ARM DDI 0612) Arm Ltd.
- [15] *Arm® Reliability, Availability, and Serviceability (RAS) System Architecture for A-profile architecture*. (ARM IHI 0100) Arm Ltd.

Feedback

Arm welcomes feedback on its documentation.

Feedback on this book

If you have any comments or suggestions for additions and improvements create a ticket at <https://support.developer.arm.com>. As part of the ticket include:

- The title (Arm® Generic Interrupt Controller Architecture Specification, GIC architecture version 5).
- The number (ARM-AES-0070 00bet0).
- The page numbers to which your comments apply.
- The rule identifiers to which your comments apply, if applicable.
- A concise explanation of your comments.

Arm also welcomes general suggestions for additions and improvements.

Note

Arm tests PDFs only in Adobe Acrobat and Acrobat Reader, and cannot guarantee the appearance or behavior of any document when viewed with any other PDF reader.

Inclusive terminology commitment

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

Previous issues of this document included language that can be offensive. We have replaced this language. To report offensive language in this document, email terms@arm.com.

Chapter 1

Introduction

I _{SBNXT}	The Generic Interrupt Controller version 5 (GICv5) architecture defines requirements for interrupt sources and a common programming interface for translating, configuring, and handling interrupts.
I _{JCSXY}	Sources of interrupts can be peripherals (including PE functions, system components, and the GIC itself), and software running on PEs.
I _{JCVZT}	<p>The GICv5 CPU interface provides a common programming interface to manage and handle physical and virtual interrupts.</p> <p>Virtual interrupts are handled using the same programming interface as physical interrupts, providing binary compatible virtualization support.</p> <p>The GICv5 CPU interface also supports backwards compatibility for GICv3 virtual machines by providing EL2 controls as well as the GICv3 virtual CPU interface and GICv3 hypervisor control interface. Backward compatibility support in the GICv5 CPU interface is an optional feature.</p>
I _{NHXNK}	<p>To support virtualization, the GICv5 architecture provides mechanisms for managing interrupts for Virtual Machines (VMs) and virtual PEs (VPEs). The architecture supports multiple VMs, and each VM can contain one or more VPEs.</p> <p>Software can configure when a VPE becomes resident on a PE through a system register interface. It can also request doorbell notification interrupts when making a VPE non-resident at a later time.</p> <p>The GIC architecture does not define or mandate how a VM boundary is managed by software.</p> <p>The following are examples of VMs that the GICv5 architecture supports:</p> <ul style="list-style-type: none">• A VM running in Non-secure state at EL1 and EL0, managed by a hypervisor running in Non-secure state at EL2.

- A Realm running in Realm state at EL1 and EL0, and is managed by the Realm Management Monitor (RMM) running in Realm state at EL2.
- A Secure partition running in Secure state at EL1 and EL0, and is managed by the Secure Partition Manager (SPM) running in Secure state at EL2.

A VPE represents an execution context within a VM.

D_{QJTJG}

A VM is identified by a VM ID.

A VPE is identified by a VPE ID in the context of a VM.

I_{NBJJF}

VMs are currently specified using a 16-bit identifier. The architecture reserves space to expand to additional bits of VM ID across all registers and data structures in a future version of the architecture. Due to the need to support a VPE doorbell in the LPI INTID space for each VPE in each VM, expanding the number of VMs beyond 24 bits may also require expanding the number of INTID.IDs supported.

D_{BJBPZ}

A GICv5 CPU interface is connected to the *Interrupt Routing Infrastructure (IRI)*.

The IRI uses either the GICv5 system architecture as described in [Chapter 3 GICv5 system architecture](#), or is IMPLEMENTATION DEFINED.

If the IRI is IMPLEMENTATION DEFINED, it must meet the requirements for SPIs and LPIs specified in [Chapter 2 PE architecture](#), including the requirements specified in [2.12 Interrupt ordering model and synchronization requirements](#).

D_{GMPXJ}

In the context of this specification, a GICv5 *system* comprises one or more PEs with support for the GICv5 architecture, each connected to the IRI.

This specification expects that the same OS or hypervisor manages all the PEs in a GICv5 system.

The word system can be used in different contexts to describe a different set of components, for example to include PEs that do not implement support for the GICv5 architecture.

Chapter 2

PE architecture

2.1 Architecture features and extensions

R _P NBFY	GICv5 defines a new Armv9-A architecture feature, GICv5 CPU interface Extension, represented by FEAT_GCIE. This feature is supported in AArch64 only.
I _G PWSJ	Support for FEAT_GCIE is reported in ID_AA64PFR2_EL1.GCIE.
R _J MHXV	FEAT_GCIE is optional from Armv9.3-A[1].
I _S HDMS	FEAT_NMI is mandatory from Armv8.8-A and Armv9.3-A. This means that when a PE implements FEAT_GCIE, it also implements FEAT_NMI.
R _K TBJL	A PE that implements FEAT_GCIE does not implement FEAT_GICv3.
R _B QYTD	GICv5 defines an optional Armv9-A architecture feature, GICv5 legacy VPE support, represented by FEAT_GCIE_LEGACY.
I _B KBQV	FEAT_GCIE_LEGACY is implemented by a PE if ICC_IDR0_EL1.GCIE_LEGACY is >= 1. Otherwise, FEAT_GCIE_LEGACY is not implemented.
I _C XTVF	FEAT_GCIE_LEGACY provides a virtual CPU interface compatible with GICv3.3, including: <ul style="list-style-type: none">• GICv3.3 ICV registers.• GICv3.3 ICH registers, with some GICv5 updates.• GICv3.3 support for virtual Non-maskable interrupts.• GICv3.1 extended INTID range.• GICv3 maintenance interrupt.• GICv3 support for separate trapping of ICV_DIR_EL1.

See [2.10.7 Legacy virtual CPU interface](#) for more information.

$R_{QCWB\bar{Y}}$ A PE that implements FEAT_GCIE_LEGACY also implements FEAT_GCIE and EL2.

R_{VVMRJ} When a PE implements FEAT_GCIE, SCR_EL3.IRQ is RES0.

R_{JMXDP} When FEAT_GCIE and FEAT_EL3 are implemented on a PE, SCR_EL3.FIQ is RES1.

See also:

- [2.7 Interrupt Prioritization](#)
- [2.9 The physical CPU interface](#)
- [2.9.3 Preemptive interrupts](#)
- [3.1 Interrupt Domains](#)

2.2 The GICv5 CPU interface

I_{GNEVQ}

When FEAT_GCIE is implemented, the PE implements the GICv5 CPU interface. The GICv5 CPU interface is responsible for:

- Managing configuration and state of *Private Peripheral Interrupts* (PPIs).
- Masking and signaling of interrupts.
- Providing an interface for software to handle interrupts.
- Providing an interface for software to configure *Shared Peripheral Interrupts* (SPIs) and *Logical Peripheral Interrupts* (LPIs).
- If EL2 is implemented, the CPU interface supports both virtual interrupt and physical interrupts.

A PE with an integrated GICv5 CPU interface is illustrated in [Figure 2.1](#).

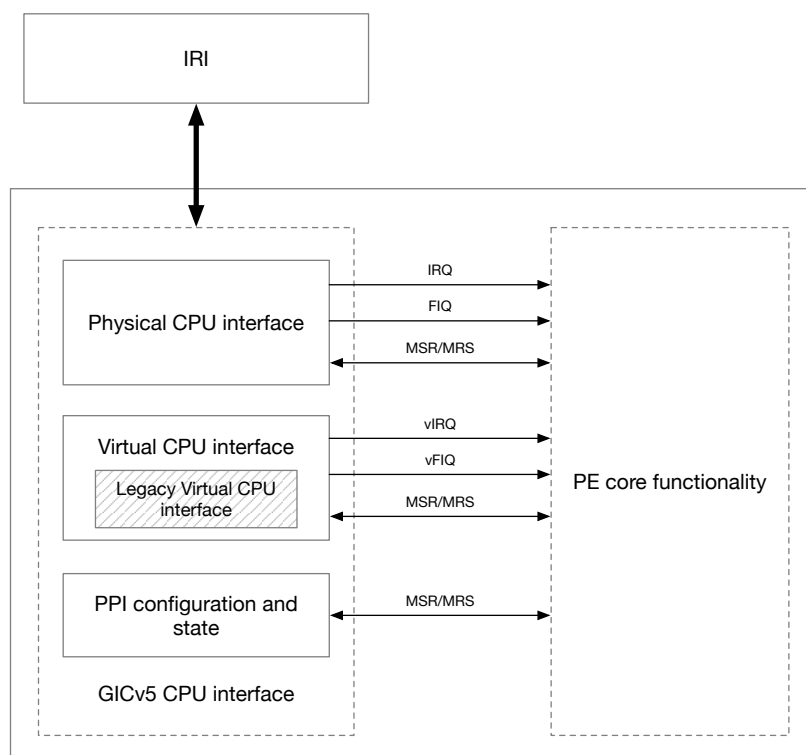


Figure 2.1: PE with an integrated GICv5 CPU interface.

Note

vFIQ is only signaled by the Legacy virtual CPU interface.

I_{CMPTD}

The GICv5 CPU interface defines when an interrupt is signaled to the PE.

The PE architecture supports the IRQ, FIQ, vIRQ, and vFIQ interrupt signals. These interrupts signals can have Superpriority as an additional attribute.

The PE architecture defines additional rules detailing whether an interrupt signal results in an exception being taken. See *Arm® Architecture Reference Manual, for A-profile architecture*[1] for more information.

I_{PYKTM}

The physical CPU interface is always active and determines whether the IRQ and FIQ interrupts are signaled. See [2.9 The physical CPU interface](#) for more information.

When executing at Exception levels below EL2, HCR_EL2.IMO is 1, and Legacy operation is not enabled, the virtual CPU interface determines whether the vIRQ interrupt is signaled. See [2.10 The virtual CPU interface](#) for more information.

When executing at Exception levels below EL2, HCR_EL2.IMO is 1, and Legacy operation is enabled, the Legacy virtual CPU interface determines whether the vIRQ and vFIQ interrupts are signaled. See [2.10.7 Legacy virtual CPU interface](#) for more information.

The GICv5 CPU interface never signals the vIRQ or vFIQ interrupts when executing at EL2 or EL3.

D_{MJHKK} The GICv5 CPU interface connects to the IRI that manages SPIs and LPIs. The IRI *presents* a candidate HPPI for each Interrupt Domain to the CPU interface and the CPU interface determines if the candidate HPPI is the HPPI.

R_{MCVBY} The mechanism used by the IRI to *present* a candidate HPPI is IMPLEMENTATION DEFINED.

I_{BTNJQ} Arm recommends that an implementation uses the GICv5 system architecture. See [Chapter 3 GICv5 system architecture](#) for more information.

Arm recommends that an implementation uses the GICv5 Stream Protocol for communication between PEs and the IRI. See [Chapter A1 GICv5 Stream Protocol overview](#) for more information.

R_{WKGBT} Every PE is assigned a unique PE interrupt Affinity ID in the system.

I_{XDCZN} The interrupt Affinity ID value for each PE is reported in ICC_IAFFIDR_EL1.IAFFID.

R_{GXVVJ} The PE interrupt Affinity ID is the same across all Physical Interrupt Domains.

R_{PXKWK} Support for bypassing the GICv5 CPU interface to signal interrupts to the PE is IMPLEMENTATION DEFINED. The behavior of System instruction and System registers, if the GICv5 CPU interface is bypassed, is IMPLEMENTATION DEFINED.

R_{TYFBB} All GIC System instructions and GIC System registers defined in this specification are available in Debug state.
See also:

- [2.8 Interrupt handling](#)

2.3 Interrupt Domains

D _{HJYDR}	An <i>Interrupt Domain</i> provides isolation of interrupt configuration and handling by its association to a Security state.
D _{XQLTP}	The architecture defines the following <i>Physical Interrupt Domains</i> , each corresponding to an implemented Security state and Exception level: <ul style="list-style-type: none"> • The Non-secure Interrupt Domain • The Realm Interrupt Domain • The Secure Interrupt Domain • The EL3 Interrupt Domain
D _{TTZCX}	When an interrupt belonging to a <i>Physical Interrupt Domain</i> is signaled, it is signaled as a physical interrupt. See 2.9 The physical CPU interface and ‘Asynchronous exception types’ in [1] for more information.
D _{HXMSH}	The architecture defines the term <i>Non-EL3 Interrupt Domain</i> as any Physical Interrupt Domain apart from the EL3 Interrupt Domain.
I _{NHCWZ}	The Non-secure, Secure, and Realm Interrupt Domains are Non-EL3 Interrupt Domains.
I _{WBZXZ}	Physical PPIs have a per-PE namespace that is shared across the Physical Interrupt Domains.
I _{VWGBP}	Physical SPIs have a single shared namespace across PEs and the Physical Interrupt Domains.
I _{NFQYW}	Physical LPIs have a separate namespace for each Physical Interrupt Domain that is shared across PEs.
R _{ZMRKX}	The implemented Physical Interrupt Domains depend on the implemented Security states as follows:

Implemented Security states	Implemented Physical Interrupt Domains
Non-secure	Non-secure
Secure	Secure
Non-secure, Secure	Non-secure, Secure, EL3
Non-secure, Realm, Root	Non-secure, Realm, EL3
Non-secure, Realm, Secure, Root	Non-secure, Realm, Secure, EL3

R _{ZFCXM}	The Physical Interrupt Domain associated with a Security state and Exception level is as follows:
--------------------	---

Exception level	Security state	Physical Interrupt Domain
EL3	Secure or Root ^a	EL3
EL2/1/0	Secure	Secure
EL2/1/0	Non-secure	Non-secure
EL2/1/0	Realm	Realm

- a. The Security state of EL3 depends on whether FEAT_RME is implemented. However, EL3 always uses the EL3 Interrupt Domain.

D _{XKSVB}	The architecture defines the <i>Current Physical Interrupt Domain</i> of a PE as the Physical Interrupt Domain corresponding to the current Exception level and Security state of that PE.
--------------------	--

I _{FDP} CN	For example, if a PE is running in EL3 on a system with the Secure and Non-secure Security states, the <i>Current Physical Interrupt Domain</i> is the EL3 Interrupt Domain.
D _{QWVZJ}	When EL2 is implemented, the <i>Virtual Interrupt Domain</i> defines the namespace for the virtual interrupts on a PE. When EL2 is not implemented, there is no Virtual Interrupt Domain.
D _{QHMVR}	The interrupts managed by the IRI for the Virtual Interrupt Domain are selected based on the <i>resident VPE</i> . The VM that the VPE belongs to is referred to as the <i>resident VM</i> .
I _{VHGT} D	The resident VPE and resident VM are programmed using ICH_CONTEXTR_EL2. See 2.10.5 Selecting the resident VPE for more information.
D _{ZWDFY}	When an interrupt belonging to the <i>Virtual Interrupt Domain</i> is signaled, it is signaled as a virtual interrupt. See 2.10 The virtual CPU interface and ‘Asynchronous exception types’ in [1] for more information.
D _{NSTKG}	The <i>Current Interrupt Domain</i> is one of the following: <ul style="list-style-type: none"> • The Virtual Interrupt Domain, when all of the following are true: <ul style="list-style-type: none"> – The current Exception level is EL1. – HCR_EL2.IMO is 1. – ICH_VCTLR_EL2.V3 is 0. • Otherwise, the Current Physical Interrupt Domain. <p>See 2.6 GIC System instructions for more information about the use of the Current Interrupt Domain for GIC System instructions.</p>

Note

The Current Interrupt Domain is not defined when ICH_VCTLR_EL2.V3 is 1 and executing at Exception levels below EL2. See [2.10.7 Legacy virtual CPU interface](#) for more information.

2.4 Interrupt types and identifiers

D _{SHNLW}	FEAT_GCIE supports the following types of interrupts: <ul style="list-style-type: none"> • Private Peripheral Interrupts (PPIs). • Logical Peripheral Interrupts (LPIs). • Shared Peripheral Interrupts (SPIs).
D _{QDWHN}	<p>An <i>interrupt</i> is identified by an INTID which comprises the following:</p> <ul style="list-style-type: none"> • INTID.TYPE The type of an interrupt is PPI, LPI, or SPIs. • INTID.ID The numeric ID of an interrupt. <p>Each interrupt comprises the following states:</p> <ul style="list-style-type: none"> • Pending An interrupt is either Pending or Idle. An interrupt is only signaled and acknowledged when it is Pending. • Active An interrupt is either Active or Inactive. An interrupt is only signaled and acknowledged when it is Inactive. <p>Each interrupt comprises the following configuration values:</p> <ul style="list-style-type: none"> • Enabled An interrupt is either Enabled or Disabled. An interrupt is only signaled and acknowledged when it is Enabled. • Priority The priority value of an interrupt. The priority values determine the order in which interrupts are acknowledged and allow masking interrupts of lower priority. • Handling mode The Handling mode of an interrupt is either Edge or Level. When the Handling mode is Edge, the Pending state becomes Idle when an interrupt is acknowledged. When the Handling mode is Level, the Pending state is unchanged when an interrupt is acknowledged. • Routing mode Targeted or 1ofN. • Affinity Specifies the destination of an interrupt.
D _{LDFFT}	Each interrupt type has a separate INTID.ID namespace.
I _{GTDKF}	For example, PPI 15 is a separate interrupt from LPI 15.
R _{PSTJC}	The architecture supports up to 24 bits of INTID.ID namespace for each interrupt type.
I _{PBBTX}	An implementation may support fewer than 24 bits of INTID.ID namespace.
I _{CTQKK}	The implemented INTID.ID namespace width may vary across interrupt types.
R _{GYVWB}	<p>Where relevant, the interrupt type is encoded using a 3-bit field as follows:</p> <ul style="list-style-type: none"> • 0b001: Private Peripheral Interrupts (PPIs) • 0b010: Logical Peripheral Interrupts (LPIs) • 0b011: Shared Peripheral Interrupts (SPIs) <p>All other values are reserved.</p>
R _{TJPHS}	<p>The INTID is a numerical value constructed from both the interrupt ID and the type as follows:</p> <ul style="list-style-type: none"> • Bits[31:29]: Interrupt type • Bits[28:24]: RES0 • Bits[23:0]: Interrupt ID
I _{VCLSH}	For example, an LPI with interrupt ID 1024 can be identified using INTID 0x40000400.

I _{D_{LGBQ}}	<p>The Handling mode of an interrupt specifies whether the Pending state becomes Idle when that interrupt is acknowledged.</p> <p>When using the GICv5 system architecture, the IRS specifies a separate Trigger mode for SPIs and the IWB specifies a Trigger mode for input wires, and interrupt events generated by an IWB or an IRS may update the Handling mode of the corresponding interrupt to match the Trigger mode.</p> <p>See Chapter 4 Interrupt routing service (IRS) and Chapter 6 Interrupt Wire Bridge (IWB) for more information.</p>
D _{CHFFT}	LPIs and SPIs are routed to appropriate PEs based on the Routing mode and Affinity. PPIs are private to a PE and therefore always signaled locally on a PE.
I _{ZPQDY}	The Routing mode of each LPI and SPI is either Targeted or 1ofN.
I _{LLQVD}	Support for 1ofN is optional for the IRI. When the GICv5 system architecture is used, the IRS reports whether 1ofN is supported in IRS_IDR0.ONE_N.
D _{QHNYP}	<p>When the Routing mode of an interrupt is 1ofN, the IRI dynamically selects a PE to present the interrupt to, based on information available to the IRI at any point in time.</p> <p>When the Routing mode of an interrupt is Targeted, the IRI is only permitted to present the interrupt to the PE corresponding to the interrupt Affinity.</p> <p>See 2.6.1 LPI and SPI configuration for more information about configuring the routing of LPIs and SPIs.</p>

2.4.1 PE-Private Peripheral Interrupts (PPIs)

G _{BGLBP}	GICv5 enables PE local events to be signaled as interrupts without recourse to the IRI.
D _{PDXVM}	The GIC CPU Interface Extension defines PPI IDs 0 through 127.
I _{QTYLP}	PPIs are generated by the CPU interface, typically by a function within the PE.
I _{YXRMV}	For example, the Arm Generic Timers generate a PPI when a timer condition is met.
R _{LWYBR}	<p>The mechanism by which PPI sources signal the CPU interface is IMPLEMENTATION DEFINED.</p> <p>Updates to the PPI Pending state from a PPI source are <i>autonomous asynchronous events</i> that complete in finite time. See <i>Arm® Architecture Reference Manual, for A-profile architecture</i>[1] for more information.</p>
I _{DQVMP}	When a PPI source is programmed using System register writes, even if a direct System register write causes a PPI signal to become asserted or de-asserted, this does not cause an indirect System register write to the ICC_PPI_xPENDR registers. Also in this case, the update from the PPI source to the PPI Pending state is an autonomous asynchronous event that completes in finite time.
D _{HQXDV}	Each PPI is private to the PE.
D _{BWMGN}	Each PE has a private namespace for physical PPIs. The namespace is common across all Physical Interrupt Domains on the PE.
I _{KPGFB}	<p>Because each PPI is private to the PE, and each PE has a private namespace for PPIs, when separate PEs use the same PPI INTID, they specify separate interrupts.</p> <p>For example, physical PPI 15 on PE 0 is a separate interrupt from physical PPI 15 on PE 1.</p>
D _{LHMTB}	If EL2 is implemented, the CPU interface supports virtual PPIs.
D _{RGQHS}	Each PE has a separate namespace for virtual PPIs.
I _{ZTTQT}	<p>Each virtual PPI is a separate interrupt from the physical PPI with the same INTID on the same PE.</p> <p>For example, physical PPI 15 on PE 0 is a separate interrupt from virtual PPI 15 on PE 0.</p> <p>However, a physical PPI can be assigned to the virtual PPI with the same INTID.</p> <p>See 2.10.1 Virtual PPIs for more information.</p>

<code>I_ZVRNF</code>	The Handling mode of a PPI is not software programmable but is discoverable through a read-only system register.
<code>I_WQBLP</code>	In GICv5 the configuration and state of PPIs is held in the CPU Interface. This is different to earlier versions of the GIC architecture, where the GIC IRI is responsible for PPI state and configuration.

See also:

- [2.9.1 Physical PPIs](#)
- [2.10.1 Virtual PPIs](#)

2.4.2 Logical Peripheral Interrupts (LPIs)

<code>G_PTPSK</code>	LPIs are expected to be provisioned by Hypervisor and OS software according to the needs of the software agent. Therefore, LPIs have a private namespace for each Interrupt Domain and typically require memory to be allocated for storing their configuration and state in the IRI.
<code>I_ZBQLG</code>	LPIs are so called because of the similarity between how LPIs in GICv3 and LPIs in GICv5 are translated by an ITS and often used for interrupt sources that signal their interrupts as MSIs. The word logical is chosen to reflect the INTID private namespace for each Interrupt Domain. GICv3 used the words locality-specific for historical reasons.
<code>D_BBVWN</code>	Each Physical Interrupt Domain has a private namespace for physical LPIs. The namespace is common across all PEs for the same Physical Interrupt Domain.
<code>I_PPFTN</code>	For example, LPI 15 as seen by PE 0 is the same interrupt as LPI 15 seen by PE 1, in the same Physical Interrupt Domain.
<code>D_BSNNZ</code>	The Virtual Interrupt Domain has a separate namespace for virtual LPIs. The namespace for virtual LPIs is determined by the resident VPE and resident VM . Each VM has a separate namespace for virtual LPIs. The virtual LPI namespace for a VM is common for all VPEs in a VM.

Note

This means that when separate VPEs for the same VM are [resident](#) across multiple PEs, the Virtual Interrupt Domains across those PEs share the same virtual LPI namespace.

See [2.10.5 Selecting the resident VPE](#) for more information.

<code>I_FSHHB</code>	For example, LPI 15 as seen by VPE 0 is the same interrupt as LPI 15 seen by VPE 1, in the same VM.
----------------------	---

2.4.3 Shared Peripheral Interrupts (SPIs)

<code>G_KCZPX</code>	SPIs are designed to function without being provisioned by software and without requiring memory to be allocated for storing their configuration and state in the IRI. SPIs are expected to be configured and provisioned when the system is designed and therefore have a common namespace across the Physical Interrupt Domains which allows firmware to configure the Interrupt Domain that each SPI is assigned to. SPIs are expected to be used in the following scenarios: <ul style="list-style-type: none"> • During early-boot where memory is not available to store interrupt state and configuration. • When supporting RAS or other error handling interrupts which must work correctly even in the presence of memory system failures.
----------------------	--

- In real-time sensitive systems, where the worst-case latency of interrupt delivery and handling must meet pre-determined timing requirements.

D_{NNDX}

The Physical Interrupt Domains have a common namespace for physical SPIs. The namespace is common across all PEs for all Physical Interrupt Domains.

This means that there is a single common namespace for physical SPIs in the system.

D_{VZHQW}

The Virtual Interrupt Domain has a separate namespace for virtual SPIs.

The namespace for virtual SPIs is determined by the [resident VPE](#) and [resident VM](#).

Each VM has a separate namespace for virtual SPIs. The virtual SPI namespace for a VM is common for all VPEs in a VM.

Note

This means that when separate VPEs for the same VM are [resident](#) across multiple PEs, the Virtual Interrupt Domains across those PEs share the same virtual SPI namespace.

See [2.10.5 Selecting the resident VPE](#) for more information.

2.5 Inter-Processor Interrupts

D _{XHRZC}	<p>The architecture supports inter-processor communications by using <i>inter-processor interrupts</i> (IPIs). IPIs are either SPIs or LPIs configured as Targeted with the Affinity specifying the destination of the IPI. Arm recommends that interrupts used for IPIs are not signaled by other interrupt sources.</p>
S _{BGNX}	<p>In a Physical Interrupt Domain, an IPI is configured and sent from a source PE to a destination PE as follows:</p> <ol style="list-style-type: none"> 1. A physical interrupt is selected and configured to target the destination PE. 2. The interrupt Handling mode is configured to Edge. 3. Software running on the source PE issues a GIC CDPEND instruction specifying the INTID of the interrupt. <p>For virtual machines, an IPI is configured and sent from a source VPE to a destination VPE as follows:</p> <ol style="list-style-type: none"> 1. A virtual interrupt is selected and configured to target the destination VPE. 2. The interrupt Handling mode is configured to Edge. 3. Software running on the source VPE issues a GIC CDPEND instruction specifying the INTID of the interrupt.
I _{SZFGV}	<p>The effect of making an interrupt Pending completes in finite time.</p> <p>This means that when an IPI is made pending, if the interrupt meets the conditions to be signaled by the IRI, it is signaled to the target PE or VPE in finite time.</p>

Note

When an IPI is signaled to its destination, there is no guarantee that it will generate an interrupt exception and be handled by the destination (V)PE. For example, the destination PE could be masking interrupt exceptions or continuously be handling higher priority interrupts.

I _{BXHQK}	<p>Arm strongly recommends that the IRI implements a sufficient number of interrupts that are not signaled by other interrupt sources, to send IPIs for each Interrupt Domain.</p> <p>The GICv5 system architecture allows software in each Interrupt Domain to allocate a sufficient number of LPIs and allows implementations to provide a sufficient number of SPIs for IPIs when appropriate.</p>
I _{WHWNY}	<p>Arm expects software to use LPIs allocated by system software to send IPIs.</p> <p>The architecture also allows the use of SPIs to send IPIs if a sufficient number of SPIs that are not connected to any interrupt sources are available.</p>
S _{PJWMR}	<p>Software may allocate the number of required LPIs for each PE and establish a mapping from an LPI INTID.ID to a PE and IPI number.</p> <p>If software requires n IPIs for each PE, and the system supports a maximum of m PEs, software would allocate $m * n$ LPIs for IPIs.</p> <p>The LPI IDs used for IPIs is managed by software.</p> <p>However, if software requires a fixed arithmetic mapping for the IPI number on a given PE to an LPI ID, then software can allocate the LPIs 0 through $(m * n) - 1$ for IPIs at system boot, and IPI number x on PE y is then given by $(y * n) + x$.</p> <p>If SPIs are used to send IPIs, whether it is possible to create a fixed arithmetic mapping depends on the INTIDs for implemented SPIs that can be used for this purpose.</p>
S _{FYVVP}	<p>Software can send IPIs by using the GIC <domain>AFF, <Xt>, GIC <domain>HM, <Xt> and GIC <domain>→PEND, <Xt> instructions to configure interrupts and to make them Pending.</p> <p>Software can configure interrupts for each potential destination PE during boot as follows:</p> <ul style="list-style-type: none"> • The interrupt Handling mode is configured to be Edge using the GIC <domain>HM, <Xt> instruction.

- The interrupt Routing mode is configured to be Targeted and the Affinity can be programmed to the destination PE using the GIC <domain>AFF, <Xt> instruction.

At runtime, software can send an IPI to a target PE by issuing the GIC <domain>PEND, <Xt> instruction for one of the interrupts targeted to the PE and allocated for IPIs.

See [2.6 GIC System instructions](#) for more information about the GIC System instructions.

2.6 GIC System instructions

I_{RSVHM}

FEAT_GCIE defines the GIC System instructions for the following purposes:

- Configure LPIs and SPIs managed by the IRI.
- Handle pending interrupts in the CPU interface.

I_{FVJSG}

The A64 assembly language syntax for GIC System instructions is one of the following:

- GIC <operation>, <Xt>.
- GICR <Xt>, <operation>.

The GIC variant is used when the instruction has the semantics of a write and is an alias of the SYS instruction.

The GICR variant is used when the instruction has the semantics of a read and is an alias of the SYSL instruction.

I_{ZCVGT}

The <operation> in GIC <operation>, <Xt> and GICR <Xt>, <operation> has the structure of <domain><command>.

Each instruction operates in the Interrupt Domain identified by the domain parameter as described below:

Domain parameter	Interrupt Domain
CD	Current Interrupt Domain
LD	Logical Interrupt Domain
VD	Virtual Interrupt Domain

Each instruction executes a command specified by a command parameter listed below in the identified Interrupt Domain.

Variant	Command	Name
GIC	AFF	Interrupt set Affinity
GIC	RCFG	Request interrupt configuration
GIC	DI	Deactivate interrupt
GIC	DIS	Interrupt clear Enable
GIC	EN	Interrupt set Enable
GIC	EOI	Interrupt Priority drop
GIC	PEND	Interrupt set/clear Pending state
GIC	PRI	Interrupt set Priority
GICR	IA	Interrupt acknowledge
GICR	NMIA	NMI acknowledge

- The commands for interrupt configuration are described in [2.6.1 LPI and SPI configuration](#).
- The commands for interrupt handling are described in [2.8 Interrupt handling](#).

<Xt> encodes additional parameters specific to the requested operation.

R_{XCLJC}

The effects of a GIC System instruction complete in finite time.

I_{BMZKC}

GIC System instructions and GSB instructions complete in finite time.

I _{MFDQT}	<p>When the GIC System instructions are executed with the <domain> parameter as CD, the operation applies to the Current Interrupt Domain.</p> <p>The GIC System instructions are not available when ICH_VCTLR_EL2.V3 is 1 and executing at Exception levels below EL2.</p>
I _{VNKRX}	<p>When the GIC System instructions are executed with the <domain> parameter as LD, the operation applies to the Physical Interrupt Domain selected by the current value of the SCR_EL3.{NSE,NS} fields.</p> <p>The LD parameter is only available at EL3.</p>
D _{GBYBF}	An interrupt specified by an INTID is either <i>reachable</i> or <i>unreachable</i> .
R _{PCDRJ}	<p>GIC System instructions are only permitted to access the configuration and state of reachable interrupts.</p> <p>When a GIC System instruction specifies an unreachable interrupt, it is treated as a NOP.</p>
R _{JRDTL}	<p>A physical PPI is reachable when all of the following are true:</p> <ul style="list-style-type: none"> • The PPI is implemented. • The PPI is assigned to the Physical Interrupt Domain that the instruction operates in. <p>A virtual PPI is reachable when the PPI is implemented.</p> <p>Otherwise, the PPI is unreachable.</p>
R _{GRSYS}	<p>A physical SPI is reachable when all of the following are true:</p> <ul style="list-style-type: none"> • The SPI is implemented by the IRI. • The SPI is assigned to the Physical Interrupt Domain that the instruction operates in. <p>A virtual SPI for a VM is reachable when the SPI for the resident VM is reachable by the IRI. See 4.10.2 Virtual SPIs for more information.</p> <p>Otherwise, the SPI is unreachable.</p>
R _{TTJYH}	<p>A physical LPI is reachable when all of the following are true:</p> <ul style="list-style-type: none"> • The LPI is within the configured range for the corresponding Physical Interrupt Domain in the IRI. See 4.8.1 Physical LPIs for more information. <p>A virtual LPI for a VM is reachable when the LPI for the resident VM is reachable by the IRI. See 4.10.1 Virtual LPIs for more information.</p> <p>Otherwise, the LPI is unreachable.</p>
I _{HDTHV}	<p>When a GIC System instruction specifies a virtual SPI or LPI, the VM used to define the namespace for the INTID is determined in one of the following ways:</p> <ul style="list-style-type: none"> • If the instruction operation is VDPEND, the VM is specified by the register argument passed to the instruction. • If the instruction domain parameter specifies the Current Interrupt Domain or the Virtual Interrupt Domain, the VM is the resident VM. <p>When a GIC System instruction operates in the Virtual Interrupt Domain and there is no resident VPE, all virtual LPIs and SPIs are unreachable.</p> <p>See 2.10.5 Selecting the resident VPE for more information about managing the resident VPE.</p>
I _{HZLFG}	<p>The following GIC System instructions are related to interrupt handling and access PPI, LPI, or SPI configuration and state:</p> <ul style="list-style-type: none"> • GICR CDIA. • GICR CDNMI. • GIC <domain>DI. <p>The GIC <domain>DI is treated as a NOP if it specifies an unreachable INTID.</p>

R_{WSSNV} When a GICR CDIA or GICR CDNMA instruction is executed, its architectural side effects are synchronized by the execution of a GSB ACK instruction.

The side effects of this synchronization include updates to the pending state of the IRQ, FIQ, vIRQ, and vFIQ asynchronous exceptions. An instruction executed after the GSB ACK is guaranteed to observe the updated pending status of these exceptions.

Note

The architectural side effects of a GICR CDIA or GICR CDNMA instruction are synchronized by the execution of a GSB ACK instruction, regardless of whether a PPI, LPI, or SPI is acknowledged.

I_{JNRPR} If HCR_EL2.NV is 1, execution at EL1 of a GIC System instruction that specifies VD as the <domain> parameter is trapped to EL2 with ESR_EL2.EC reporting 0x18.

I_{ZFDJZ} A GIC System instruction never successfully executes at EL1 if it specifies the VD as the <domain> parameter. The instruction is either UNDEFINED or trapped to EL2 depending on the value of HCR_EL2.NV.

I_{GMHGR} The architecture supports two barrier instructions used to define ordering between effects of GIC System instructions and other effects in the system:

- **GSB ACK:** Provides ordering between the effects of a GICR CDIA instruction in program order before the GSB ACK and other effects.
- **GSB SYS:** Provides ordering between the effects of any GIC System instruction in program order before the GSB SYS and other effects.

See [2.12.2 GSB instruction semantics](#) for more information.

See also:

- [2.12 Interrupt ordering model and synchronization requirements](#)

2.6.1 LPI and SPI configuration

I_{BFFCW} Configuration and state for LPIs and SPIs are managed by the IRI. GIC System instructions are used to update the configuration of an LPIs and SPIs.

I_{MBNRN} When an INTID whose Type is LPI or SPI specifies an ID which is beyond the range reported in ICC_IDR0_EL1.ID_BITS, the unimplemented upper identifier bits are RES0.

I_{KFJNB} The GIC System instruction commands that are used to update the configuration and state of interrupts managed by the IRI are listed in the following table:

Command	Name	Description
AFF	Interrupt set Affinity	Set the Affinity for the INTID.
DI	Interrupt clear Active	Clear Active for the INTID.
DIS	Interrupt clear Enable	Clear Enable for the INTID.
EN	Interrupt set Enable	Set Enable for the INTID.
PEND	Interrupt set/clear Pending state	Generate SET or CLEAR events for the INTID.
PRI	Interrupt set Priority	Set the Priority for the INTID.
HM	Interrupt set Handling mode	Set the Handling mode for the INTID.

The resulting GIC System instructions are:

- GIC <domain>AFF, <Xt>.
- GIC <domain>EN, <Xt>.
- GIC <domain>DI, <Xt>.
- GIC <domain>DIS, <Xt>.
- GIC <domain>PEND, <Xt>.
- GIC <domain>PRI, <Xt>.
- GIC <domain>HM, <Xt>.

If <Xt> specifies an unreachable LPI or SPI INTID, the instruction is treated as a NOP.

I_{WNZLQ} An INTID might be within the implemented INTID width but be unreachable. For example, the system might support 16 bits of INTID but software might have configured a smaller LPI space in the IRI. This is different to specifying an INTID beyond the implemented range, as unimplemented bits of INTID are treated as RES0.

I_{JGZYC} The Affinity of an interrupt is configured by issuing the GIC <domain>AFF, <Xt> instruction.

The IRM field in <Xt> controls the interrupt routing mode. The interrupt routing mode is either Targeted or 1ofN. When the interrupt is Targeted, the target PE is specified in the IAFFID field in <Xt>.

If the IRI supports fewer than 16 bits of IAFFID for the PE, unimplemented upper bits are RES0.

R_{JKYGN} When the Affinity of an interrupt is configured by the GIC <domain>AFF, <Xt> instruction and the IRM field configures the interrupt to be Targeted, all of the following are true:

- For a Physical Interrupt Domain, the IAFFID field specifies the PE with the corresponding interrupt Affinity ID.
- For the Virtual Interrupt Domain, the IAFFID field specifies the VPE in the VM using the corresponding VPE ID.

S_{WVYLC} Arm expects that Hypervisor software emulates a virtual GICv5 implementation where the emulated ICC_IAFFIDR_EL1.IAFFID value is the VPE ID of the corresponding VPE. This means that to software running in a VM, configuring the Affinity of a virtual interrupt using the virtualized ICC_IAFFIDR_EL1.IAFFID value will have the desired effect.

I_{WLBYY} The GIC <domain>PEND, <Xt> instruction has the following effects on the interrupt state:

- If the Pending field is cleared to 0, the IRI is requested to clear the Pending state of the interrupt to 0.
- If the Pending field is set to 1, the IRI is requested to set the Pending state of the interrupt to 1.

I_{PKRLN} The Handling mode of an interrupt can be configured by issuing the GIC <domain>HM, <Xt> instruction.

R_{YCDBY} It is IMPLEMENTATION DEFINED whether the Handling mode of an SPI is RO.

If the Handling mode of an SPI is RO, all of the following are true:

- GIC <domain>HM, <Xt> instructions targeting the SPI have no effect.
- The corresponding Trigger mode in the IRI is RO.
- The Handling mode corresponds to the Trigger mode.

See [4.8.2 Physical SPIs](#) for more information.

I_{TNGVR} Interrupt configuration managed by an IRI is queried using the System instruction GIC <domain>RCFG and read from the system register ICC_ICSR_EL1.

R_{NPDVH} On executing a GIC <domain>RCFG System instruction, the PE requests the configuration and state of the specified INTID from the IRI for the Interrupt Domain identified by the <domain> parameter.

If the INTID is reachable, ICC_ICSR_EL1.F is set to 0 and the other fields are populated with the configuration and state of the INTID.

If the INTID is unreachable, ICC_ICSR_EL1.F is set to 1 and the other fields become UNKNOWN.

I_{GBYBZ}

Where an instruction results in an update to a System register, as is the case with the GIC <domain>RCFG System instruction, explicit synchronization must be performed before the result is guaranteed to be visible. This means that the result of a GIC <domain>RCFG System instruction is not guaranteed to be visible in ICC_ICSR_EL1 until after an ISB or other context synchronization event.

See also:

- [2.12 Interrupt ordering model and synchronization requirements](#)
- [4.6 Interrupt configuration and state](#)

2.7 Interrupt Prioritization

<code>I_DQGTJ</code>	<p>The priority of an interrupt is an unsigned value that is used for the following purposes:</p> <ul style="list-style-type: none"> • Selecting which interrupt to signal among multiple interrupts that can be signaled for an Interrupt Domain. • Priority-based masking of interrupts. • When NMIs are enabled for an Interrupt Domain, configuring an interrupt with a priority value of 0 indicates that the interrupt is an NMI and should be signaled with Superpriority. See 2.9.5 Physical non-maskable interrupts for more information about NMIs. • When there are interrupts that can be signaled for multiple Interrupt Domains, the priority may be used to decide whether an interrupt should be signaled using the IRQ or FIQ interrupt exception. See 2.9.3 Preemptive interrupts for more information.
<code>I_ZCBZT</code>	<p>The number of implemented priority bits is reported in <code>ICC_IDR0_EL1.PRI_BITS</code>.</p> <p>The maximum number of implemented priority bits is 5, and the priority of an interrupt is always interpreted as a 5 bit unsigned value. Only bits [4:N] are implemented where $N = (4 - \text{ICC_IDR0_EL1.PRI_BITS})$. Unimplemented bits are RES0.</p> <p>This means that when fewer than 5 bits of priority is implemented, bits [3 - <code>ICC_IDR0_EL1.PRI_BITS:0</code>] are RES0.</p> <p>For example, if 4 bits of priority are implemented, only priority levels 0, 2, 4, 6, ..., 30 are implemented.</p>
<code>I_VWYWB</code>	<p>The number of implemented priority bits may vary between the CPU interface and the IRI. In this case, unimplemented lower order bits are treated as 0 by the component implementing the lowest number of bits. Arm recommends that software only programs priority values within the range supported by both the CPU interface and the IRI.</p>
<code>I_CNQYT</code>	<p>Some operating systems have requirements for a minimum number of interrupt priority levels. These requirements will be captured as part of the <i>Arm® Base System Architecture 1.0C</i>[2].</p>
<code>I_YQCPV</code>	<p>The priority value of an interrupt is only interpreted within an Interrupt Domain. Priorities of interrupts belonging to separate Interrupt Domains are never compared against one another.</p> <p>Therefore, the full implemented priority space is available to prioritize interrupts within each Interrupt Domain.</p>
<code>I_GDMTN</code>	<p>The priority value of a virtual interrupt is only interpreted to select higher priority interrupts for that VM. The priority value of a virtual interrupt is never compared against the priority value of a physical interrupt. Priorities of interrupts belonging to separate VMs are never compared against one another.</p> <p>Therefore, the full implemented priority space is available to prioritize interrupts within each VM.</p>
<code>R_XBQTT</code>	<p>In the GIC prioritization scheme, lower numbers have higher priority.</p> <p>For example, an interrupt with priority value 0 is of higher priority than an interrupt with priority value 16.</p>
<code>D_LDCKK</code>	<p>The <i>highest priority pending interrupt</i> (HPPI) for an Interrupt Domain is the highest priority interrupt selected among the candidate HPPIs.</p> <p>An HPPI is defined for each Physical Interrupt Domain and the Virtual Interrupt Domain:</p> <ul style="list-style-type: none"> • The EL3 physical HPPI. • The Secure physical HPPI. • The Realm physical HPPI. • The Non-secure physical HPPI. • The virtual HPPI for the Virtual Interrupt Domain.
<code>D_RKWNs</code>	<p>A <i>candidate HPPI</i> is a Pending, Inactive, and Enabled interrupt selected among a subset of interrupts for an Interrupt Domain.</p> <p>The PE considers the following candidate HPPIs for each Interrupt Domain:</p>

- The highest priority Pending, Inactive, and Enabled PPI for the Interrupt Domain.
- The candidate HPPI [presented](#) by the IRI for the Interrupt Domain.

R_{JXKKZ}

When the Pending, Active, or Enabled values of a PPI change, the effects on the HPPI determination complete in finite time.

See [2.9.4 Physical interrupt signaling](#) and [2.10.3 Virtual interrupt signaling](#) for more information.

R_{KXFKH}

The IRI selects the highest priority Pending, Inactive, and Enabled interrupt for each Interrupt Domain in finite time.

See [4.8.4 Physical interrupt signaling](#) and [4.10.4 Virtual interrupt signaling](#) for more information about the guarantees provided by the GICv5 system architecture.

See also:

- [2.8 Interrupt handling](#)
- [2.8.1 Interrupt life cycle](#)
- [2.9 The physical CPU interface](#)

2.8 Interrupt handling

D _{QWBZR}	When the CPU interface acknowledges an interrupt, the interrupt's priority becomes an <i>active priority</i> .
D _{DWSGL}	Each Interrupt Domain has a <i>running priority</i> , which is the highest active priority for that Interrupt Domain.
I _{FGFPH}	It is possible to have multiple active priorities at the same time. For example, consider the following scenario where an interrupt handler is interrupted by an NMI: <ol style="list-style-type: none"> 1. Interrupt X with priority 16 is acknowledged. Priority 16 becomes an active priority and the <i>running priority</i> of the Interrupt Domain becomes 16. 2. Another interrupt Y with priority 0 is signaled to the PE. 3. Interrupt Y is acknowledged. Priority 0 becomes an active priority and the <i>running priority</i> of the Interrupt domain becomes 0. At the end of the sequence above, priorities 16 and 0 are active priorities and the <i>running priority</i> for the Interrupt Domain is 0.
D _{MQKDW}	When there are no active priorities, the <i>running priority</i> is the <i>Idle priority</i> . The <i>Idle priority</i> is 0xFF.
I _{CFYGY}	Only interrupts with a priority higher than the <i>running priority</i> are signaled by the CPU interface. This prevents a high priority interrupt from being preempted by a low priority interrupt.
D _{VHCMF}	The CPU interface performs a <i>Priority drop</i> when software issues the GIC CDEOI instruction, and the <i>running priority</i> is not the Idle priority.
R _{KWMNP}	On a Priority drop, the highest active priority for the Interrupt Domain stops being active.
I _{FGKMJ}	After a Priority drop, because the <i>running priority</i> is defined as the highest active priority, the <i>running priority</i> becomes one of the following: <ol style="list-style-type: none"> 1. The highest active priority for which there has not been a Priority drop. 2. The Idle priority.
I _{XZFHX}	This allows pending interrupts with the previous <i>running priority</i> to be signaled by the CPU interface.
I _{QFZWR}	Active priorities stop being active in priority order, from the highest priority to the lowest priority.
I _{JFQSN}	GICv5 removes support for the binary point registers used in GICv3 to split the priority field into a preemption level and subpriority.
I _{FYZVD}	Software handles an interrupt signaled by the CPU interface as follows: <ol style="list-style-type: none"> 1. Software acknowledges the HPPI for the Current Interrupt Domain and obtains its INTID. 2. Software uses the INTID to locate the handler function for the acknowledged interrupt. 3. Software performs a Priority drop when it is ready to receive other interrupts of the same priority as the acknowledged interrupt. 4. Software deactivates the acknowledged interrupt once it is ready to receive that interrupt again.
I _{DFWSM}	FEAT_GCIE defines GIC System instructions to handle interrupts in the CPU interface. Each instruction executes a command that applies to all interrupt types in the Current Interrupt Domain. The commands are listed below:

Command	Name
IA	Interrupt acknowledge
NMIA	NMI acknowledge
EOI	Interrupt priority drop
DI	Interrupt deactivate

The System instructions for interrupt handling are listed below and described in [8.1 System instructions for the Current Interrupt Domain](#).

- GICR CDIA.
- GICR CDNMI.
- GIC CDEOI.
- GIC CDDI.

<code>I_ZYDVS</code>	The architecture does not define the order in which the Priority Drop and Interrupt deactivate are performed by software.
<code>I_WDLPR</code>	<p>The GICR CDIA System instruction acknowledges the HPPI in the Current Interrupt Domain if there is an HPPI with Sufficient priority and it is not an NMI.</p> <p>The result of the instruction is returned in the <Xt> register.</p> <p>The VALID field in <Xt> register indicates whether an interrupt was acknowledged by the GICR CDIA instruction.</p> <p>When the VALID field is 1, the HPPI was acknowledged.</p> <p>When the VALID field is 0, the HPPI was not acknowledged.</p>
<code>I_XBBNJ</code>	<p>The GICR CDNMI system instruction acknowledges the HPPI in the Current Interrupt Domain if there is an HPPI with Sufficient priority and it is an NMI.</p> <p>The result of the instruction is returned in the <Xt> register.</p> <p>The VALID field in <Xt> register indicates whether the HPPI was acknowledged.</p> <p>When the VALID field is 1, the HPPI was acknowledged.</p> <p>When the VALID field is 0, the HPPI was not acknowledged.</p>
<code>R_WNKKW</code>	<p>When an interrupt is acknowledged, all of the following are true:</p> <ul style="list-style-type: none"> • The interrupt becomes Active. • If the interrupt Handling mode is Edge, its Pending state is cleared. • The interrupt's priority becomes active and the running priority for the Current Interrupt Domain is updated. • The INTID of the interrupt is returned in the TYPE and ID fields of the <Xt> register.
<code>I_HBGJV</code>	The GIC CDEOI System instruction performs a Priority drop of the running priority in the Current Interrupt Domain.
<code>I_RYBDB</code>	Unlike earlier versions of the GIC architecture, the GIC CDEOI System instruction does not take an INTID as an argument.
<code>I_WPSSR</code>	The GIC CDDI System instruction performs an Interrupt deactivate for the specified INTID in the Current Interrupt Domain.

2.8.1 Interrupt life cycle

`R_TCZMV` [Figure 2.2](#) shows the interrupt state machine for interrupts whose Handling mode is Edge.

[Figure 2.3](#) shows the interrupt state machine for interrupts whose Handling mode is Level.

The interrupt Handling mode determines the effects on the interrupt state when the CPU Interface acknowledges the interrupt as follows:

- For Edge interrupts, all of the following are true:
 - The Pending state is atomically cleared to Idle.
 - The Active state is set to Active.
- For Level interrupts, The Active state is set to Active.

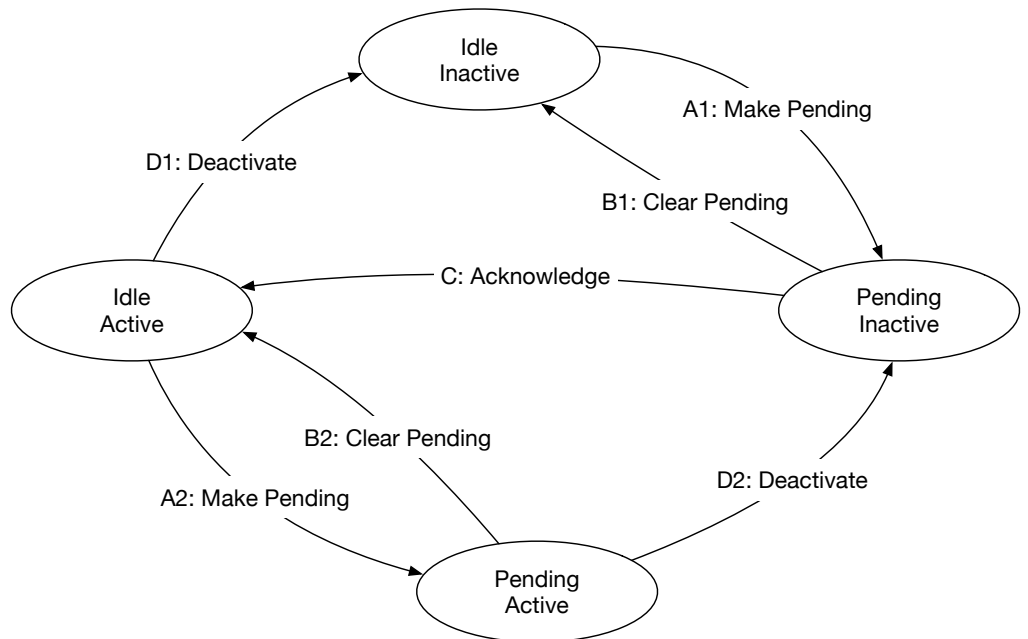


Figure 2.2: Interrupt state machine for Edge interrupts.

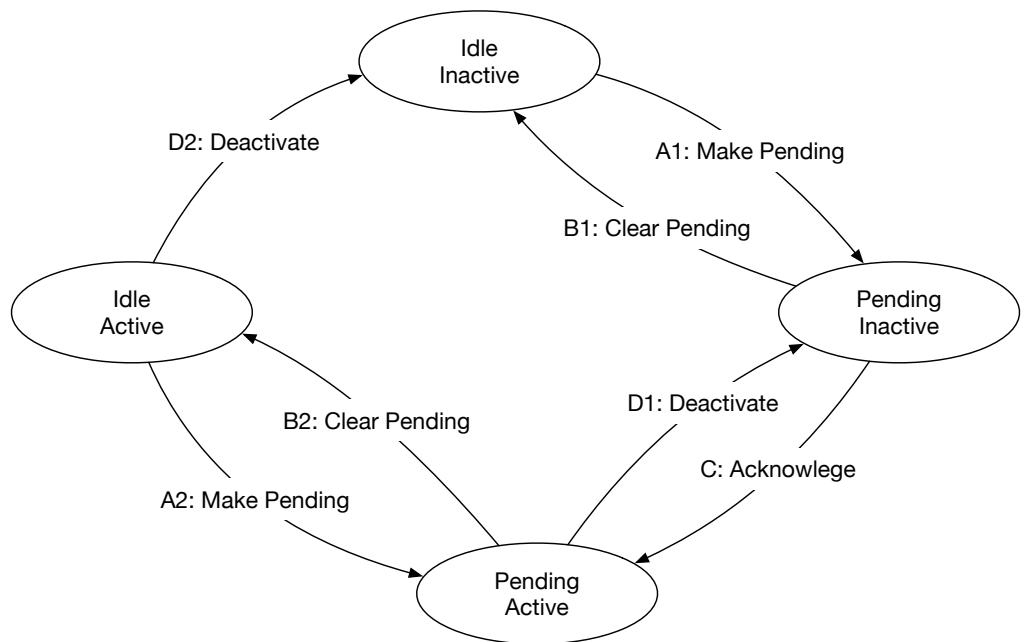


Figure 2.3: Interrupt state machine for Level interrupts.

The transitions shown in the figures can be caused by the following reasons:

- **Transition A1 or A2: Make Pending**

This transition occurs in one of the following situations:

- An LPI or SPI is signaled in the IRI or requested to be made Pending by software.

- **Transition B1 or B2: Clear Pending**

This transition occurs in one of the following situations:

- A level-sensitive LPI or SPI signal is de-asserted in the IRI or the Pending state is requested to be cleared by software.

- **Transition C: Acknowledge**

This transition occurs when the CPU Interface acknowledges the Interrupt.

For Edge interrupts, the Pending state is cleared, also known as *consumed*, as part of acknowledging the interrupt.

For Level interrupts, the Pending state is not affected as part of acknowledging the interrupt.

- **Transition D1 or D2: Deactivate**

This transition occurs when the CPU Interface deactivates the Interrupt.

I_{DMXLN}

An interrupt cannot be set to Active using any other GIC instruction than GICR CDIA and GICR CDNMA. An interrupt can be deactivated from any PE. Deactivating an interrupt which is already Inactive has no effect.

2.9 The physical CPU interface

2.9.1 Physical PPIs

R_{RVCZW}	A physical PPI is only considered for being selected as the candidate HPPI in the Physical Interrupt Domain it is assigned to.
I_{SHHCL}	The assignment of an architected PPI to a PPI source is defined by the architecture. Unassigned architected PPIs are Reserved.
D_{HVJHM}	PPIs in the range of IDs 0 through 63 are referred to as <i>architected PPIs</i> .
R_{FSFSL}	The assignment of PPIs in the range of IDs 64 through 127 to their PPI sources is IMPLEMENTATION DEFINED.
I_{YGJWQ}	The assignment of a physical PPI to a Physical Interrupt Domain is done as follows: <ul style="list-style-type: none"> • Architected PPIs are assigned by programming ICC_PPI_DOMAINR0_EL3 and ICC_PPI_DOMAINR1_EL3. • IMPLEMENTATION DEFINED PPIs are assigned by programming ICC_PPI_DOMAINR2_EL3 and ICC_PPI_DOMAINR3_EL3.
R_{RHCDS}	When EL3 is not implemented, each physical PPI is assigned to the Interrupt Domain identified by the Effective value of SCR_EL3.NS.
R_{XDVCM}	The assignments of architected PPIs to their PPI sources are listed in Table 2.7 .

Table 2.7: Architected PPI ID assignment

ID	Name	Source
31	TRBIRQ	Trace Buffer Unit
30	CNTP	EL1 Physical Timer
29	CNTPS	EL3 Physical Timer
28	CNTHV	Non-secure EL2 Virtual Timer
27	CNTV	EL1 Virtual Timer
26	CNTHP	Non-secure EL2 Physical Timer
25	GICMNT	GIC maintenance interrupt
24	CTIIRQ	Generic CTI interrupt trigger event
23	PMUIRQ	PMU overflow interrupt request
22	COMMIRQ	Debug communication channel
21	PMBIRQ	Profiling Buffer management interrupt request
20	CNTHPS	Secure EL2 Physical Timer
19	CNTHVS	Secure EL2 Virtual Timer
15	HACDBSIRQ	Hardware accelerator for cleaning Dirty state interrupt
3	SW_PPI	Reserved for software usage.
2	NS_DB_PPI	Doorbell PPI for a Non-secure Non-preemptive interrupt
1	RL_DB_PPI	Doorbell PPI for a Realm Non-preemptive interrupt
0	S_DB_PPI	Doorbell PPI for a Secure Non-preemptive interrupt

R_{CF}SKX The Handling modes of architected PPIs are listed in [Table 2.8](#).

Table 2.8: Architected PPI ID Handling modes

ID	Name	Handling Mode
31	TRBIRQ	Level
30	CNTP	Level
29	CNTPS	Level
28	CNTHV	Level
27	CNTV	Level
26	CNTHP	Level
25	GICMNT	Level
24	CTIIRQ	IMPLEMENTATION DEFINED
23	PMUIRQ	Level
22	COMMIRQ	Level
21	PMBIRQ	Level
20	CNTHPS	Level
19	CNTHVS	Level
15	HACDBSIRQ	Level
3	SW_PPI	Edge
2	NS_DB_PPI	Level
1	RL_DB_PPI	Level
0	S_DB_PPI	Level

I_{RR}GLL The configuration and state of PPIs is held in the PE System registers.
For physical PPIs, the following registers are used:

Table 2.9: Physical PPI system registers

Register	Name
ICC_PPI_SACTIVER<n>_EL1	Set interrupt Active
ICC_PPI_CACTIVER<n>_EL1	Clear interrupt Active
ICC_PPI_SPENDR<n>_EL1	Set interrupt Pending
ICC_PPI_CPENDR<n>_EL1	Clear interrupt Pending
ICC_PPI_ENABLER<n>_EL1	Interrupt Enable and Disable
ICC_PPI_PRIORITYR<n>_EL1	Interrupt Priority
ICC_PPI_DOMAINR<n>_EL3	Physical Interrupt Domain allocation
ICC_PPI_HMR<n>_EL1	Interrupt Handling mode (Level or Edge)

Register	Name
I _{RBPKR}	<p>The Active state of a PPI can also be updated as a side-effect of executing one of the following GIC System instructions:</p> <ul style="list-style-type: none"> • GICR CDIA. • GICR CDNMA. • GIC CDDI. • GIC VDDI. • GIC LDDI. <p>These System instructions are used during interrupt handling and are common across interrupt types.</p>
I _{ZLTKB}	<p>A Context Synchronization event is required before software can rely on the effects of direct writes to a PPI System register to affect instructions appearing in program order after the direct write. See <i>Arm® Architecture Reference Manual, for A-profile architecture</i>[1] for more information. This applies to the PPI system registers listed in the following tables:</p> <ul style="list-style-type: none"> • Table 2.9 • Table 2.15 • Table 2.16
R _{YRGMH}	<p>If the individual enable of a PPI is cleared, a Context synchronization event is sufficient to guarantee that the PPI is no longer the HPPI.</p>
S _{XNGGQ}	<p>A PPI might be shared between multiple Security states, and in this case, EL3 is expected to save and restore the configuration of that PPI during a change of Security state.</p>
I _{CPDGZ}	<p>A physical PPI is permitted to be a candidate HPPI for a Physical Interrupt Domain if and only if all of the following are true:</p> <ul style="list-style-type: none"> • ICC_PPI_ENABLER<n>_EL1.EN is 1. • ICC_PPI_SPENDR<n>_EL1.PEND is 1. • ICC_PPI_SACTIVER<n>_EL1.ACTIVE is 0.
I _{HTMYK}	<p>The CPU interface selects the candidate HPPI among the PPIs for each Physical Interrupt Domain by selecting the highest priority PPI that meets the physical candidate HPPI selection criteria.</p>
R _{JSQBW}	<p>If more than one PPI meets the physical candidate HPPI selection criteria for a Physical Interrupt Domain, it is IMPLEMENTATION DEFINED which of those PPIs becomes the candidate HPPI.</p>
I _{QMBYG}	<p>Any PPI source may generate an IMPLEMENTATION DEFINED wake event for the PE that they are connected to. The architecture does not specify when a PPI source may generate a wake event.</p>
I _{HYVBT}	<p>Examples of when a PPI source may generate an IMPLEMENTATION DEFINED wake event include the following:</p> <ul style="list-style-type: none"> • A generic timer is kept powered on while the PE is in low-power state and generates a wake event when the timer condition is asserted. • The Cross-Trigger interface is kept powered on while the PE is in a low-power state and asserts an output trigger event.

2.9.2 Physical priority masking

D _{MLTJM}	<p>The <i>Physical Priority Mask</i> is defined as follows:</p> <ul style="list-style-type: none"> • The Physical Priority Mask for a Non-EL3 Interrupt Domain is the value stored in the banked copy of ICC_PCR_EL1.PRIORITY for that Interrupt Domain. • The Physical Priority Mask for the EL3 Interrupt Domain is the value stored in ICC_PCR_EL3.PRIORITY.
--------------------	---

I _{RGRBD}	See 9.1 Synchronization requirements for GICv5 System registers for more information about synchronization requirements for updates to the Physical Priority Mask.
D _{XMBQZ}	The <i>physical running priority</i> is defined as follows: <ul style="list-style-type: none"> • The physical running priority for a Non-EL3 Interrupt Domain is the highest active priority stored in the banked copy of ICC_APR_EL1 for that Interrupt Domain. • The physical running priority for the EL3 Interrupt Domain is the highest active priority stored in ICC_APR_EL3.
I _{FWVXZ}	The physical running priority for the Current Physical Interrupt Domain is reported in ICC_HAPR_EL1.PRIORITY.
D _{MSQKF}	A physical interrupt has <i>Sufficient priority</i> to be signaled when all of the following are true: <ul style="list-style-type: none"> • The priority of the interrupt is higher than the physical running priority for the Physical Interrupt Domain. • The priority of the interrupt is equal to or higher than the Physical Priority Mask for the Physical Interrupt Domain.
I _{QYFVP}	The physical running priority is defined for each Physical Interrupt Domain even though ICC_HAPR_EL1 only shows the running priority for the Current Physical Interrupt Domain. The physical running priority is a factor in determining whether an interrupt has Sufficient priority in an Interrupt Domain for each Physical Interrupt Domain, regardless of whether an Interrupt Domain is the Current Physical Interrupt Domain. For example, when the Current Physical Interrupt Domain is the Non-secure Interrupt Domain, an interrupt in the Secure Interrupt Domain will only have Sufficient priority if the interrupt's Priority is higher than the highest active priority in the Secure banked copy of ICC_APR_EL1.
I _{WLVKM}	ICC_DOMHPPIR_EL3 reports if there is an HPPI of Sufficient priority for each Non-EL3 Interrupt Domain.

2.9.3 Preemptive interrupts

G _{LTfJK}	When the Current Physical Interrupt Domain is not the EL3 Interrupt Domain, and there is an HPPI of Sufficient priority for another Physical Interrupt Domain, the PE may be configured to take one of the following actions: <ul style="list-style-type: none"> • The Current Physical Interrupt Domain is preempted in order to switch to the Physical Interrupt Domain of the HPPI so that it can be handled immediately. See also 2.9.4 Physical interrupt signaling. • The Current Physical Interrupt Domain is notified through a separate interrupt that there is an HPPI for a different Physical Interrupt Domain, and the Current Physical Interrupt Domain can switch to the other Interrupt Domain when appropriate as defined by a software policy. See also 2.9.6 Doorbell PPIs.
D _{MQRSE}	The Interrupt Domain selected by ICC_CR0_EL3.PID is called the <i>Preemptive Interrupt Domain</i> .
I _{PJKYP}	When the value of ICC_CR0_EL3.PID corresponds to an Interrupt Domain that is not implemented, the PE behaves as if there is no Preemptive Interrupt Domain.
I _{VFWFW}	ICC_CR0_EL1.PID indicates whether the Physical Interrupt Domain associated with the Security state selected by SCR_EL3.{NSE,NS} is the Preemptive Interrupt Domain.
D _{NNBML}	The interrupt priority range of the Preemptive Interrupt Domain is split into two. The priority value used for the split is called the <i>Interrupt Preemptive Priority Threshold (IPPT)</i> . ICC_CR0_EL1.IPPT defines the Interrupt Preemptive Priority Threshold for each Physical Interrupt Domain.
I _{RLTDB}	The value of ICC_CR0_EL1.IPPT has no effect on the signaling of interrupts for an Interrupt Domain which is not the Preemptive Interrupt Domain.
D _{KXXGB}	A physical interrupt is a <i>Preemptive interrupt</i> when one of the following is true: <ul style="list-style-type: none"> • The interrupt belongs to the EL3 Interrupt Domain and the Current Physical Interrupt Domain is not the EL3 Interrupt Domain. • All of the following are true: <ul style="list-style-type: none"> – The Current Physical Interrupt Domain is not the Preemptive Interrupt Domain.

- The interrupt belongs to the Preemptive Interrupt Domain.
- The priority of the interrupt is higher than the IPPT of the Preemptive Interrupt Domain.

Otherwise, the interrupt is a *Non-preemptive interrupt*.

SYFVLN

For example, on a system with Secure and Non-secure Security states, the Secure Interrupt Domain can configure a subset of its interrupts as Preemptive interrupts and the remaining subset as Non-preemptive interrupts:

- The Preemptive interrupts preempt execution in the Non-secure Interrupt Domain before being handled in the Secure Interrupt Domain.
- The Non-preemptive interrupts allow the Non-secure Interrupt Domain to yield execution control to the Secure Interrupt Domain where they are handled.

2.9.4 Physical interrupt signaling

IWKDZ

The CPU interface determines the physical HPPI for each Physical Interrupt Domain and performs the following actions:

1. Applies priority masking based on software configured priority masks and the [running priority](#).
2. Determines whether the physical HPPIs signals the IRQ or FIQ interrupt to the PE.
3. Determines whether the interrupt is signaled with Superpriority.

This process is illustrated in [Figure 2.4](#):

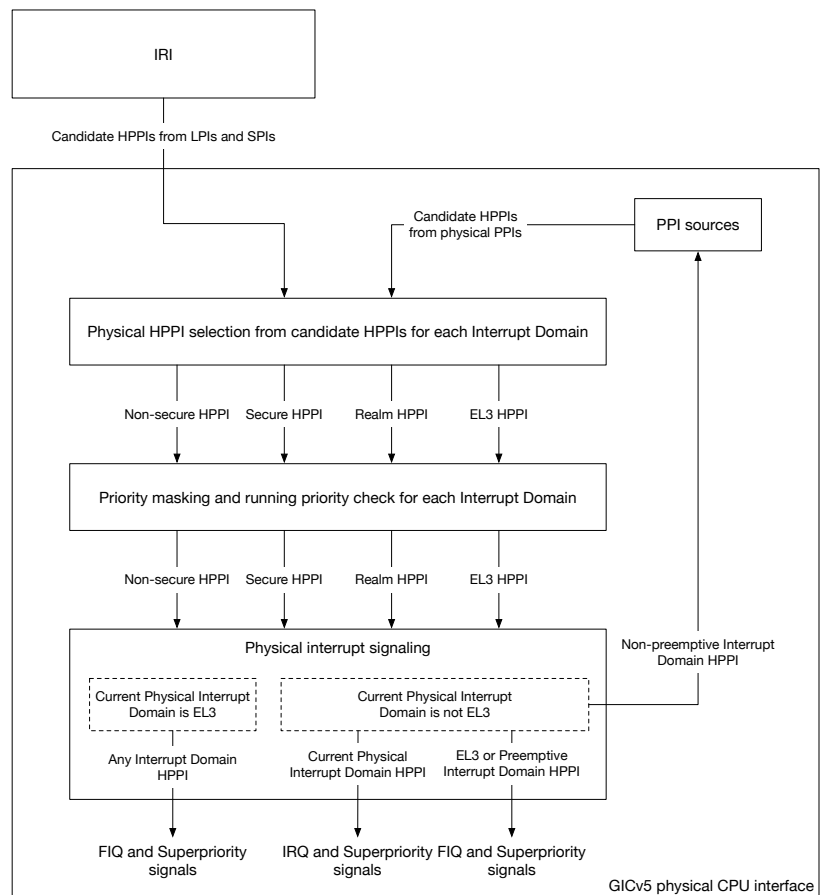


Figure 2.4: Physical interrupt determination

RZSWWN

The CPU interface determines whether there is an HPPI for a Physical Interrupt Domain when the value of ICC_CR0_ELx.EN is 1 for the Interrupt Domain.

The HPPI for an Interrupt Domain is determined by selecting the interrupt with the highest priority from the following:

- The candidate HPPI **presented** by the IRI for the Interrupt Domain.
- The candidate HPPI selected among the physical PPIs for the Interrupt Domain.

If ICC_CR0_ELx.EN is 0, the CPU interface behaves as if there are no candidate HPPIs for the Interrupt Domain.

There is no HPPI for an Interrupt Domain if any of the following are true:

- ICC_CR0_ELx.EN is 0 for the Interrupt Domain.
- There are no candidate HPPIs for the Interrupt Domain.

R_{VVBPS}

If the candidate HPPI **presented** by the IRI has the same priority as the candidate HPPI selected among the PPIs for a Physical Interrupt Domain, it is IMPLEMENTATION DEFINED which of the two candidate HPPIs is selected as the HPPI for the Physical Interrupt Domain.

R_{XNDSJ}

When there is at least one candidate HPPI for a Physical Interrupt Domain, the CPU interface determines the HPPI in finite time.

When there is a change to the candidate HPPIs due to updates to the configuration and state of the physical PPIs, or due to the IRI **presenting** a new candidate HPPI for the Interrupt Domain, the CPU interface redetermines the HPPI for the Physical Interrupt Domain in finite time.

R_{QLGBG}

The CPU interface signals the FIQ interrupt when one of the following is true:

- There is a Preemptive interrupt of Sufficient priority.
- All of the following are true:
 - The Current Physical Interrupt Domain is the EL3 Interrupt Domain.
 - There is an HPPI of Sufficient priority for any Interrupt Domain.

R_{ZGHMN}

The CPU interface signals the IRQ interrupt when all of the following are true:

- The Current Physical Interrupt Domain is a Non-EL3 Interrupt Domain.
- There is no Preemptive interrupt of Sufficient priority.
- There is an HPPI of Sufficient priority for the Current Physical Interrupt Domain.

I_{SXXCJ}

The CPU interface never signals the IRQ and FIQ interrupts at the same time.

I_{RYDXC}

The following table illustrates the physical interrupt signal asserted by the CPU interface when the PE is executing at any Exception level below EL3:

EL3 HPPI of Sufficient priority exists	Non-Current Preemptive Domain HPPI of Sufficient priority exists	Current Domain HPPI of Sufficient priority exists	Interrupt signal
Yes	x	x	FIQ
No	Yes	x	FIQ
No	No	Yes	IRQ
No	No	No	n/a

x: The value does not have any impact on which interrupt is signaled.

I_{LTYGD}

The following table illustrates the type of physical interrupt signaled by the CPU interface when the PE is executing at EL3:

EL3 HPPI of Sufficient priority exists	Non-EL3 Preemptive HPPI of Sufficient priority exists	Any non-Preemptive HPPI of Sufficient priority exists	Interrupt signal
Yes	x	x	FIQ
No	Yes	x	FIQ
No	No	Yes	FIQ
No	No	No	n/a

x: The value does not have any impact on which interrupt is signaled.

S_{MKKWV}

Software executing at EL3 can read ICC_DOMHPPIR_EL3 to determine if there are HPPIs of Sufficient priority for other Interrupt Domains. For example, this can be useful after taking the FIQ exception to determine what may have caused the exception when execution of the GICR CDIA System instruction returns VALID as 0 in the <Xt> register.

This may also be useful in certain power management scenarios to determine the idle state of the system.

I_{HDXLH}

Whether there is an HPPI of Sufficient priority is reported by the following registers based on the current Exception level:

- When the current Exception level is EL2 or EL1, ICC_HPPIR_EL1 reports whether there is an HPPI of Sufficient priority for the Current Physical Interrupt Domain.
- When the current Exception level is EL3, ICC_HPPIR_EL1 reports whether there is an HPPI of Sufficient priority for the Logical Interrupt Domain as selected by SCR_EL3.{NSE,NS}.

If there is an HPPI of Sufficient priority, it also reports the INTID of that HPPI.

S_{DWXXH}

EL3 firmware can use ICC_DOMHPPIR_EL3 to determine if a Security state has an HPPI of Sufficient priority. By switching SCR_EL3.{NSE,NS} to that Security state and subsequently reading ICC_HPPIR_EL1, the firmware can identify the INTID of that HPPI. EL3 firmware can use this information to switch to the software component responsible for handling the HPPI in the Security state.

I_{GVRTZ}

In GICv3, the ICx_HPPIRx_EL1 registers indicate the highest priority pending interrupt for the corresponding interrupt group irrespective of whether the interrupt has sufficient priority to be signaled to the PE.

GICv5 changes this behavior such that the ICx_HPPIR_EL1 and ICC_DOMHPPIR_EL3 registers return information for an HPPI only if it has Sufficient priority to be signaled.

Arm is not aware of any software component that relies on the GICv3 behavior. The GICv5 behavior helps simplify the PE interface architecture as both the logic for acknowledging an HPPI and reporting information about an HPPI rely on the same condition of Sufficient priority determination.

I_{SYCYB}

ICC_HPPIR_EL3 reports whether there is an HPPI of Sufficient priority for the EL3 Interrupt Domain.

I_{SXGQC}

In a system which only supports a single Interrupt Domain, the CPU interface never signals the FIQ interrupt.

2.9.5 Physical non-maskable interrupts

R_{CSBDX}

When there is an HPPI of Sufficient priority for the Current Physical Interrupt Domain, it is signaled to the PE with *Superpriority*[1] when all of the following are true:

- The priority of the HPPI is 0.
- One of the following is true:
 - The Current Physical Interrupt Domain is the EL3 Interrupt Domain and SCTLR_EL3.NMI is 1.
 - The Current Physical Interrupt Domain is not the EL3 Interrupt Domain and one of the following is true:
 - * Physical IRQs are routed to EL2, and SCTLR_EL2.NMI is 1.
 - * Physical IRQs are routed to EL1, and SCTLR_EL1.NMI is 1.

Otherwise, the HPPI for the Current Physical Interrupt Domain is not signaled with Superpriority.

R_{VHGPR}

When there is a Preemptive interrupt of Sufficient priority, it is IMPLEMENTATION DEFINED whether the interrupt is signaled to the PE with Superpriority, if all the following are true:

- The priority of the HPPI is 0.
- The Preemptive interrupt belongs to the EL3 Interrupt Domain.
- SCTLR_EL3.NMI is 1.

All other Preemptive interrupts are not signaled with Superpriority.

D_{SDQJW}

A physical interrupt that is signaled with Superpriority is referred to as a physical *NMI*.

I_{RPSFK}

If the priority of the HPPI is 0, but is not signaled with Superpriority due to not meeting the other conditions, then the HPPI is not an NMI.

I_{DLGVX}

An NMI can never be acknowledged by executing the GICR CDIA instruction.

Similarly, an interrupt that is not an NMI can never be acknowledged by executing the GICR CDNMI instruction. This means that when SCTLR_ELx.NMI is 0 for the current Exception level, the GICR CDNMI instruction cannot acknowledge an interrupt.

R_{SPWCF}

When FEAT_GCIE is implemented, following a Context synchronization event as the result of taking or returning from an exception, changes to whether the CPU interface has asserted or de-asserted the IRQ and FIQ signal for the Interrupt Domain have taken effect and can be observed by a read of ISR_EL1.

See also:

- [Chapter C1 Operational model](#)

2.9.6 Doorbell PPIs

R_{QCVFJ}

Doorbell PPIs are level-sensitive PPIs.

R_{PNYKZ}

There is a separate Doorbell PPI for every Non-EL3 Interrupt Domain as described in [Table 2.12](#).

Table 2.12: Doorbell PPI definitions

Non-EL3 Interrupt Domain	Doorbell PPI name
Secure	S_DB_PPI
Realm	RL_DB_PPI
Non-secure	NS_DB_PPI

R_{GSHXH}

The level of the NS_DB_PPI is asserted when all of the following are true:

- The Current Physical Interrupt Domain is one of the following:
 - Secure Interrupt Domain.
 - Realm Interrupt Domain.
- The NS_DB_PPI is assigned to the Current Physical Interrupt Domain.
- There is a Non-preemptive physical HPPI of Sufficient priority for the Non-secure Interrupt Domain.

Otherwise, the level of the NS_DB_PPI is not asserted.

The level of the S_DB_PPI is asserted when all of the following are true;

- The Current Physical Interrupt Domain is one of the following:
 - Non-secure Interrupt Domain.

- Realm Interrupt Domain.
- The S_DB_PPI is assigned to the Current Physical Interrupt Domain.
- There is a Non-preemptive physical HPPI of Sufficient priority for the Secure Interrupt Domain.

Otherwise, the level of the S_DB_PPI is not asserted.

The level of the RL_DB_PPI is asserted when all of the following are true;

- The Current Physical Interrupt Domain is one of the following:
 - Non-secure Interrupt Domain.
 - Secure Interrupt Domain.
- The RL_DB_PPI is assigned to the Current Physical Interrupt Domain.
- There is a Non-preemptive physical HPPI of Sufficient priority for the Realm Interrupt Domain.

Otherwise, the level of the RL_DB_PPI is not asserted.

R_{LWLK} When the level of a Doorbell PPI is asserted, it is asserted in finite time.

R_{YSCLK} [Table 2.13](#) lists the conditions for when a Doorbell PPI is implemented.

Table 2.13: Doorbell PPI implementation

Doorbell PPI name	Condition
S_DB_PPI	EL3 and Secure state are implemented
RL_DB_PPI	FEAT_RME is implemented
NS_DB_PPI	EL3 is implemented

I_{WGXR} When the Current Physical Interrupt Domain is a Non-EL3 Interrupt Domain, [Table 2.14](#) describes when the level of a Doorbell PPI is asserted.

Table 2.14: Doorbell PPI usage

Current Physical Interrupt Domain	Doorbell PPI	Asserted when
Non-secure	S_DB_PPI	There is a Secure Non-preemptive HPPI of Sufficient priority.
Non-secure	RL_DB_PPI	There is a Realm Non-preemptive HPPI of Sufficient priority.
Secure	NS_DB_PPI	There is a Non-secure Non-preemptive HPPI of Sufficient priority.
Secure	RL_DB_PPI	There is a Realm Non-preemptive HPPI of Sufficient priority.
Realm	S_DB_PPI	There is a Secure Non-preemptive HPPI of Sufficient priority.
Realm	NS_DB_PPI	There is a Non-secure Non-preemptive HPPI of Sufficient priority.

Multiple Doorbell PPIs can be asserted at the same time. For example, if the Current Physical Interrupt Domain is the Non-secure Interrupt Domain, and there is a Non-preemptive HPPI of Sufficient priority for both the Secure and Realm Interrupt Domains, then the levels of both S_DB_PPI and RL_DB_PPI are asserted.

S_FKVPW

When the Current Physical Interrupt Domain is a Non-EL3 Interrupt Domain, if a Doorbell PPI for another Non-EL3 Interrupt Domain is implemented, it must be assigned to the Current Physical Interrupt Domain so that it can be signaled and handled in the Current Physical Interrupt Domain. This is done by software executing at EL3 as follows:

- Prior to a switch to the Non-secure Interrupt Domain from the EL3 Interrupt Domain, the S_DB_PPI and RL_DB_PPI are assigned to the Non-secure Interrupt Domain.
- Prior to a switch to the Secure Interrupt Domain from the EL3 Interrupt Domain, the NS_DB_PPI and RL_DB_PPI are assigned to the Secure Interrupt Domain.
- Prior to a switch to the Realm Interrupt Domain from the EL3 Interrupt Domain, the S_DB_PPI and NS_DB_PPI are assigned to the Realm Interrupt Domain.

S_QFPBH

Software can manage Doorbell PPIs the same way as managing any other PPIs, including individually enabling and disabling them, and configuring their priority. Doorbell PPIs can be used to yield control to another Security State at an opportune time according to the scheduling policy in the current Security state.

R_TCBST

On taking an exception from Exception levels below EL3 to EL3, the CPU interface completes the following sequence:

1. The levels of all Doorbell PPIs are de-asserted.
2. The HPPI for each Physical Interrupt Domain is re-evaluated.

See also:

- [2.9.3 Preemptive interrupts](#).
- [2.9.4 Physical interrupt signaling](#).

2.10 The virtual CPU interface

- I_{XVDWC}** When HCR_EL2.IMO is 1 and ICH_VCTLR_EL2.V3 is 0, the GICv5 virtual CPU interface is used and all of the following are true:
- When executing at EL1, all of the following are true:
 - GICv5 System instruction that specify the Current Interrupt Domain as the domain parameter operate in the Virtual Interrupt Domain.
 - Accesses to GICv5 System registers that share an encoding between the ICC_* and ICV_* registers access the ICV_* registers.
 - Accesses to GICv3.3 Group 0, Group 1, and Common System registers are UNDEFINED.
 - Accesses to ICH_VMCR_EL2 use the GICv5 version of the register.
- See [Chapter 8 System instructions](#) and [Chapter 9 System registers](#) for more information.
- S_{KFGFJ}** The Hypervisor at EL2 configures HCR_EL2.IMO to control whether the virtual or physical CPU interface is accessed at EL1.
- I_{FCTBZ}** The value of HCR_EL2.FMO has no effect on the operation of the GICv5 CPU interface when Legacy operation is disabled because of the following reasons:
- Since SCR_EL3.FIQ is RES1 if EL3 is implemented, physical FIQ interrupts are taken to EL3.
 - When Legacy operation is disabled, the GICv5 virtual CPU interface does not signal the virtual FIQ interrupt.

2.10.1 Virtual PPIs

- G_{WWRGJ}** The architecture supports signaling virtual PPIs to software running at Exception levels below EL2. Virtual PPIs may be signaled by hypervisor software emulating a PPI source or directly injected using the corresponding physical PPI.
- R_{GZRRG}** For each implemented physical PPI, if EL2 is implemented, the corresponding virtual PPI with the same ID is implemented.
- I_{QFGQN}** For each implemented PPI, the physical and virtual PPI with the same INTID are separate interrupts.
- I_{RQBBW}** If the Virtual CPU interface is supported, virtual PPIs are configured using the following registers:

Table 2.15: Virtual PPI system registers

Register	Name
ICV_PPI_SACTIVER<n>_EL1	Set interrupt Active
ICV_PPI_CACTIVER<n>_EL1	Clear interrupt Active
ICV_PPI_ENABLER<n>_EL1	Interrupt Enable and Disable
ICV_PPI_HMR<n>_EL1	Interrupt Handling mode (Level or Edge)
ICV_PPI_SPENDR<n>_EL1	Set interrupt Pending
ICV_PPI_CPENDR<n>_EL1	Clear interrupt Pending
ICV_PPI_PRIORITYR<n>_EL1	Interrupt Priority

- I_{CBBCV}** To enable context switching, the configuration of virtual PPIs is also accessible using the following registers:

Table 2.16: Hypervisor configuration PPI system registers

Register	Name
ICH_PPI_ACTIVER<n>_EL2	Interrupt Active, accesses same state as ICV_PPI_SACTIVER<n>_EL1 and ICV_PPI_CACTIVER<n>_EL1.
ICH_PPI_DVIR<n>_EL2	Controls whether Pending physical PPIs are directly injected as virtual PPIs to the Virtual Interrupt Domain.
ICH_PPI_ENABLER<n>_EL2	Alias of ICV_PPI_ENABLER<n>_EL1.
ICH_PPI_PENDR<n>_EL2	Interrupt Pending, accesses same state as ICV_PPI_SPENDR<n>_EL1 and ICV_PPI_CPENDR<n>_EL1.
ICH_PPI_PRIORITYR<n>_EL2	Alias of ICV_PPI_PRIORITYR<n>_EL1.

R_{PQDYK} If the Virtual CPU interface is supported, a virtual PPI is permitted to be a candidate HPPI if and only if all of the following are true:

- ICV_PPI_ENABLER<n>_EL1.EN is 1.
- ICV_PPI_SPENDR<n>_EL1.PEND is 1.
- ICV_PPI_SACTIVER<n>_EL1.ACTIVE is 0.

R_{JMKJM} The CPU interface selects the virtual candidate HPPI among the virtual PPIs by selecting the highest priority PPI that meets the virtual candidate HPPI selection criteria.

R_{ZXWQD} If more than one virtual PPI meet the virtual candidate HPPI selection criteria, it is IMPLEMENTATION DEFINED which of those PPIs becomes the virtual candidate HPPI.

I_{SXKHV} A direct read or write to a PPI System register could cause an indirect write to its alias or a different PPI System register. The indirect write is only guaranteed to be visible to subsequent reads or writes if a Context synchronization event takes place after the direct read or write and before the subsequent reads or writes. See *Arm® Architecture Reference Manual, for A-profile architecture*[1] for more information.

S_{KQQTJX} When software at EL1 performs a direct write to an ICV_ register, this causes an indirect write to the corresponding ICH_ register. The value returned by a read at EL2 from the ICH_ register is guaranteed to observe the indirect write, only if there is a context synchronization event between the indirect write and the direct read. If FEAT_ExS is implemented and SCTLRL_EL2.EIS is 0, an exception taken to EL2 is not a context synchronization event, and software at EL2 must issue an ISB before reading the ICH_ registers to observe the latest value written by a guest VM.

For example, ICH_PPI_PRIORITYR<n>_EL2 is an alias of ICV_PPI_PRIORITYR<n>_EL1. A direct write of a priority value for a PPI ID to ICV_PPI_PRIORITYR<n>_EL1 causes an indirect write of the priority value to ICH_PPI_PRIORITYR<n>_EL2 for the same PPI ID. A Context Synchronization Event is required after the write to ICV_PPI_PRIORITYR<n>_EL1 to ensure that a subsequent read of ICH_PPI_PRIORITYR<n>_EL2 returns the latest priority value.

2.10.1.1 Direct injection of virtual PPIs

D_{QNZMN} The architecture supports *direct injection* of the physical PPI Pending state to the Pending state of a virtual PPI.

I_{VRZWD} When ICH_PPI_DVIR<n>_EL2.DVI is 1, the Pending state of the corresponding physical PPI INTID is directly injected to the Pending state of a virtual PPI.

R_{BMZFB} Unless a timer PPI is redirected under nested virtualization, when the Pending state of a physical PPI is directly injected to the Pending state of a virtual PPI, the INTID of the virtual PPI is the same as the INTID of the physical PPI.

See [2.10.1.2 PPI redirection under nested virtualization](#) for more information about redirection of timer PPIs under nested virtualization.

I _{PGRWM}	<p>The Pending state of a virtual PPI is one of the following:</p> <ul style="list-style-type: none"> • If the corresponding physical PPI is not directly injected, the Pending state of the virtual PPI is determined by ICH_PPI_PENDR<n>.PEND<x>. • If the corresponding physical PPI is directly injected, the Pending state of the virtual PPI is determined by ICC_PPI_xPENDR<n>.PEND<x>. <p>When the Pending state of a physical PPI is directly injected as the Pending state of a virtual PPI, the field in ICV_PPI_xPENDR<n>.PEND<x> corresponding to the virtual PPI becomes an alias of the field in ICC_PPI_xPENDR<n>.PEND<x> corresponding to the physical PPI, and the value of ICH_PPI_PENDR<n>.PEND<x> is IGNORED.</p>
I _{XXSHB}	<p>When a physical PPI is directly injected and becomes Pending, both the virtual and physical PPI are Pending and are permitted to be acknowledged in either the corresponding Physical Interrupt Domain or in the Virtual Interrupt Domain.</p> <p>Arm recommends that the physical PPI is disabled when it is directly injected.</p>
I _{KKCLT}	<p>When the Pending state of a physical PPI is directly injected as the Pending state of a virtual PPI, it is the configuration of the virtual PPI that determines when the PPI is signaled and how the PPI is handled.</p>
I _{BKZQC}	<p>For example, if the Pending state of physical PPI 23 is directly injected as the Pending state of a virtual PPI, all of the following are true:</p> <ul style="list-style-type: none"> • The Pending state of physical PPI 23 is directly injected as the Pending state of virtual PPI 23. • Virtual PPI 23 may be considered as the candidate HPPI in the Virtual Interrupt Domain when all of the following are true: <ul style="list-style-type: none"> – Physical PPI 23 is Pending. – ICV_PPI_ENABLER0_EL1.EN23 is 1. – ICV_PPI_xACTIVER0_EL1.ACTIVE23 is 0. • ICV_PPI_PRIORITYR2_EL1.PRIORITY7 determines the Priority of virtual PPI 23.
R _{JBKGB}	<p>If EL2 is not implemented, the Effective value of ICH_PPI_DVIR<n>_EL2.DVI is 0 for all implemented PPIs.</p>
I _{CCLDF}	<p>When the Pending state of a physical PPI is directly injected as the Pending state of a virtual PPI, all of the following are true:</p> <ul style="list-style-type: none"> • Writes to the corresponding field in ICH_PPI_PENDR<n> update the value of that field. • Reads from the corresponding field in ICH_PPI_PENDR<n> return the value of that field. • The value of the corresponding field in ICH_PPI_PENDR<n> does not affect the value of the field corresponding to the virtual PPI in ICV_PPI_xPENDR<n>.
S _{SVFBL}	<p>Hypervisor software may save and restore the Pending state of a directly injected Edge PPI by accessing to the fields corresponding to the physical PPI INTID in ICC_PPI_xPENDR<n>.</p>
I _{NVWGX}	<p>The Effective value of ICH_PPI_DVIR<n>_EL2.DVI is 0 if PPI <x> is not assigned to the Current Physical Interrupt Domain.</p>

2.10.1.2 PPI redirection under nested virtualization

R _{GYGNW}	<p>The outputs of the EL1 physical timer and EL1 virtual timer signal different PPI IDs when all of the following are true:</p> <ul style="list-style-type: none"> • FEAT_GCIE is implemented. • HCR_EL2.{NV,NV1} is {1,0}. • EL2 is implemented and enabled in the Security state selected by SCR_EL3.{NSE,NS}. <p>In this case, all of the following are true:</p> <ul style="list-style-type: none"> • When SCR_EL3.NS is 1, all of the following are true: <ul style="list-style-type: none"> – If ICH_PPI_DVIR0_EL2.DVI30 is 1, all of the following are true:
--------------------	---

- * Physical PPI 30 is directly injected as virtual PPI 26.
- * ICV_PPI_xPENDR0_EL2.PEND26 is an alias of ICC_PPI_PENDR_EL2.PEND30.
- * ICV_PPI_xPENDR0_EL2.PEND30 accesses the Pending state of virtual PPI 30.
- * ICH_PPI_PENDR0_EL2.PEND26 is IGNORED.
- * ICH_PPI_DVIR0_EL2.DVI26 is treated as 0 for all other purposes than a direct read of the field.
- If ICH_PPI_DVIR0_EL2.DVI27 is 1, all of the following are true:
 - * Physical PPI 27 is directly injected as virtual PPI 28.
 - * ICV_PPI_xPENDR0_EL2.PEND28 is an alias of ICC_PPI_PENDR_EL2.PEND27.
 - * ICV_PPI_xPENDR0_EL2.PEND27 accesses the Pending state of virtual PPI 27.
 - * ICH_PPI_PENDR0_EL2.PEND28 is IGNORED.
 - * ICH_PPI_DVIR0_EL2.DVI28 is treated as 0 for all other purposes than a direct read of the field.
- When SCR_EL3.NS is 0, all of the following are true:
 - If ICH_PPI_DVIR0_EL2.DVI30 is 1, all of the following are true:
 - * Physical PPI 30 is directly injected as virtual PPI 20.
 - * ICV_PPI_xPENDR0_EL2.PEND20 is an alias of ICC_PPI_PENDR_EL2.PEND30.
 - * ICV_PPI_xPENDR0_EL2.PEND30 accesses the Pending state of virtual PPI 30.
 - * ICH_PPI_PENDR0_EL2.PEND20 is IGNORED.
 - * ICH_PPI_DVIR0_EL2.DVI20 is treated as 0 for all other purposes than a direct read of the field.
 - If ICH_PPI_DVIR0_EL2.DVI27 is 1, all of the following are true:
 - * Physical PPI 27 is directly injected as virtual PPI 19.
 - * ICV_PPI_xPENDR0_EL2.PEND19 is an alias of ICC_PPI_PENDR_EL2.PEND27.
 - * ICV_PPI_xPENDR0_EL2.PEND27 accesses the Pending state of virtual PPI 27.
 - * ICH_PPI_PENDR0_EL2.PEND19 is IGNORED.
 - * ICH_PPI_DVIR0_EL2.DVI19 is treated as 0 for all other purposes than a direct read of the field.

I_QKGM

All fields corresponding to implemented PPIs in ICH_PPI_PENDR0_EL2 access the Pending state of the corresponding virtual PPI, regardless of whether a directly injected PPI is being redirected to a different virtual INTID due to nested virtualization. This means that all of the following are true:

- When Physical PPI 30 is directly injected as virtual PPI 26, all of the following are true:
 - Whether virtual PPI 26 is considered Pending is determined by the Pending state of physical PPI 30.
 - An access to ICH_PPI_PENDR0_EL2.PEND26 accesses the Pending state of virtual PPI 26.
 - The value of ICH_PPI_PENDR0_EL2.PEND26 is IGNORED, meaning that it is not used in determining whether virtual PPI 26 is considered Pending.
- When Physical PPI 27 is directly injected as virtual PPI 28, all of the following are true:
 - Whether virtual PPI 28 is considered Pending is determined by the Pending state of physical PPI 27.
 - An access to ICH_PPI_PENDR0_EL2.PEND28 accesses the Pending state of virtual PPI 28.
 - The value of ICH_PPI_PENDR0_EL2.PEND28 is IGNORED, meaning that it is not used in determining whether virtual PPI 28 is considered Pending.
- When Physical PPI 30 is directly injected as virtual PPI 20, all of the following are true:
 - Whether virtual PPI 20 is considered Pending is determined by the Pending state of physical PPI 30.
 - An access to ICH_PPI_PENDR0_EL2.PEND20 accesses the Pending state of virtual PPI 20.
 - The value of ICH_PPI_PENDR0_EL2.PEND20 is IGNORED, meaning that it is not used in determining whether virtual PPI 20 is considered Pending.
- When Physical PPI 27 is directly injected as virtual PPI 19, all of the following are true:
 - Whether virtual PPI 19 is considered Pending is determined by the Pending state of physical PPI 27.
 - An access to ICH_PPI_PENDR0_EL2.PEND19 accesses the Pending state of virtual PPI 19.
 - The value of ICH_PPI_PENDR0_EL2.PEND19 is IGNORED, meaning that it is not used in determining whether virtual PPI 19 is considered Pending.

2.10.2 Virtual priority masking

D_BMLRF

The *Virtual Priority Mask* is the value stored in ICV_PCR_EL1.PRIORITY.

I_DDQBK

See [9.1 Synchronization requirements for GICv5 System registers](#) for more information about synchronization requirements for updates to the Virtual Priority Mask.

D _{CZRGJ}	The <i>virtual running priority</i> is defined as the highest active priority stored in ICV_APR_EL1.
I _{DLHDR}	The virtual running priority is reported in ICV_HAPR_EL1.PRIORITY.
D _{WDGVW}	A virtual interrupt has <i>Sufficient priority</i> to be signaled when all of the following are true: <ul style="list-style-type: none"> • The Priority of the interrupt is higher than the <i>virtual running priority</i>. • The Priority of the interrupt is equal to or higher than the Virtual Priority Mask.
I _{GLZCB}	ICV_HPPIR_EL1 reports if there is a virtual HPPI of Sufficient priority for the Virtual Interrupt Domain. If there is a virtual HPPI of Sufficient priority, it also reports the INTID of that HPPI.
I _{MDDZR}	The behavior of ICV_HPPIR_EL1 in GICv5 is different from the ICV_HPPIRx_EL1 registers in GICv3. In GICv3, ICV_HPPIRx_EL1 return the highest priority interrupt for the corresponding interrupt group irrespective of whether the interrupt has sufficient priority to be signaled to the PE. GICv5 changes this behavior such that ICV_HPPIR_EL1 returns the highest priority interrupt only if it has Sufficient priority to be signaled.

2.10.3 Virtual interrupt signaling

I _{VDZVT}	<p>When the virtual CPU interface is not configured to use legacy operation, the GICv5 CPU interface determines the virtual HPPI for the Virtual Interrupt Domain from the virtual candidate HPPIs from the IRI and the virtual PPIs and performs the following actions:</p> <ol style="list-style-type: none"> 1. Applies priority masking based on software configured virtual priority masks and the <i>virtual running priority</i>. 2. Determines if the virtual HPPI signals the virtual IRQ interrupt to the PE. 3. When the virtual CPU interface signals the virtual IRQ interrupt, it further determines whether the virtual IRQ interrupt is signaled with Superpriority. <p>When Legacy operation is enabled, the Legacy virtual CPU interface determines the virtual HPPI from the list registers and performs the following actions:</p> <ol style="list-style-type: none"> 1. Applies interrupt grouping and priority masking based on software configured virtual priority masks and the <i>virtual running priority</i>. 2. Determines if the virtual HPPI signals the virtual IRQ interrupt or the virtual FIQ interrupt to the PE. 3. When the virtual CPU interface signals the virtual IRQ interrupt, it further determines whether the virtual IRQ interrupt should be signaled with Superpriority.
--------------------	---

This process is illustrated in [Figure 2.5](#).

See [2.10.7 Legacy virtual CPU interface](#) for more information about the Legacy virtual CPU interface:

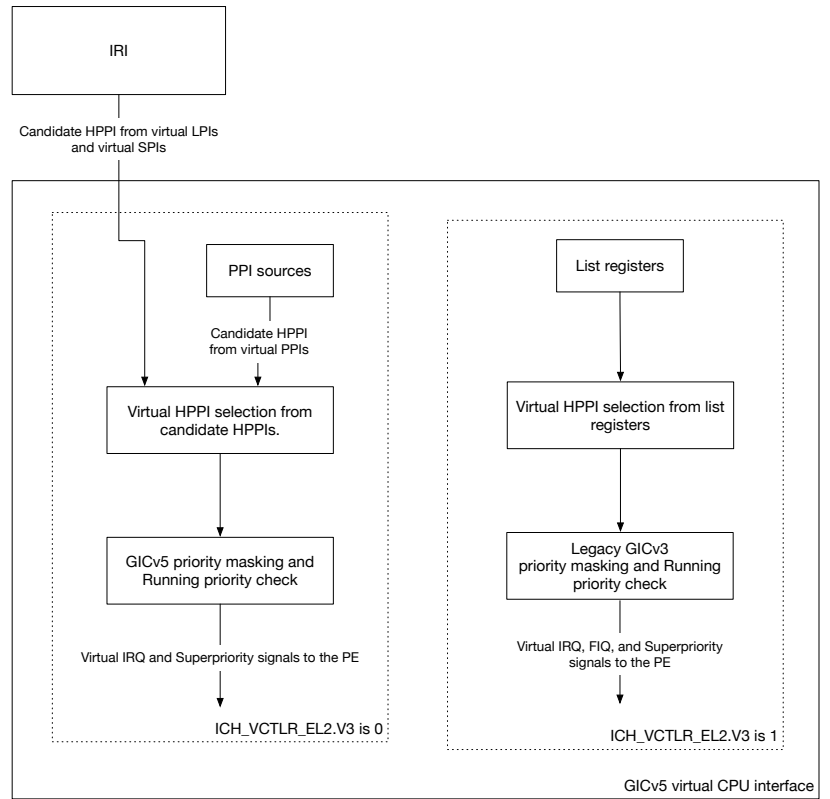


Figure 2.5: Virtual interrupt determination

R_{ZKPVN}

When EL2 is implemented and is enabled in the Security state identified by SCR_EL3.{NSE, NS} and Legacy operation is disabled, the GICv5 CPU interface determines whether there is an HPPI for the Virtual Interrupt Domain when all of the following are true:

- ICV_CR0_EL1.EN is 1.
- ICH_VCTLR_EL2.EN is 1.

The HPPI for the Virtual Interrupt Domain is determined by selecting the interrupt with the higher priority from the following:

- The candidate HPPI **presented** by the IRI for the **resident VPE**.
- The candidate HPPI selected among the virtual PPIs.

There is no HPPI for the Virtual Interrupt Domain if any of the following are true:

- ICV_CR0_EL1.EN is 0.
- ICH_VCTLR_EL2.EN is 0.
- There are no candidate HPPIs for the Virtual Interrupt Domain.

Note

The candidate HPPI **presented** by the IRI for the **resident VPE** is not considered when selecting the HPPI for the Virtual Interrupt Domain when ICH_CONTEXTR_EL2.IRICHPPIDIS is 1.

See 2.10.5 *Selecting the resident VPE* for more information about managing the **resident VPE**.

R_{JKZLR}

If EL2 is not implemented, the Effective value of ICH_VCTLR_EL2.EN is 0.

R _{XCFJD}	If the candidate HPPI presented by the IRI for the resident VPE has the same priority as the candidate HPPI selected among the virtual PPIs, it is IMPLEMENTATION DEFINED which of the two candidate HPPIs is selected as the HPPI for the Virtual Interrupt Domain.
R _{CGYDN}	When there is at least one candidate HPPI for the Virtual Interrupt Domain, the CPU interface determines the HPPI for the Virtual Interrupt Domain in finite time. When there is a change to the candidate HPPIs due to updates to the configuration and state of the virtual PPIs, or due to the IRI presenting a new candidate HPPI for the resident VPE , the CPU interface redetermines the HPPI for the Virtual Interrupt Domain in finite time.
R _{NGHYT}	The GICv5 CPU interface signals the vIRQ interrupt when there is an HPPI of Sufficient priority for the Virtual Interrupt Domain.
I _{NJPKJ}	When Legacy operation is disabled, the GICv5 virtual CPU interface does not signal the vFIQ interrupt. See 2.10.7 Legacy virtual CPU interface for more information.
I _{BVMRK}	The behavior of the Legacy virtual CPU interface including when virtual interrupts are signaled is described in [3].

2.10.4 Virtual non-maskable interrupts

R _{XKMQC}	When there is an HPPI of Sufficient priority for the Virtual Interrupt Domain, it is signaled to the PE with <i>Superpriority</i> [1] when all of the following are true: <ul style="list-style-type: none"> • The priority of the virtual HPPI is 0. • The current Exception level is EL1 or lower. • SCTLR_EL1.NMI is 1. Otherwise, the virtual HPPI is not signaled with Superpriority.
D _{QPMFG}	A virtual interrupt that is signaled with Superpriority is referred to as a <i>virtual NMI</i> .
I _{MQHTC}	If the priority of the virtual HPPI is 0, but is not signaled with Superpriority due to not meeting the other conditions, then the virtual HPPI is not a virtual NMI.

2.10.5 Selecting the resident VPE

G _{CQBBT}	Hypervisor software manages multiple VMs and VPEs and selects different VPEs to be resident on a PE over time to allow multiplexing VPEs on a single PE.
G _{PBWDL}	Separate VPEs from the same VM can be resident on separate PEs at the same time to support VMs with multiple VPEs running concurrently.
D _{JRKYV}	The resident VPE is specified by writing the corresponding VPE ID and VM ID to ICH_CONTEXTR_EL2. Each VM has a separate VPE ID namespace which means that two VPEs with the same VPE ID for different VM IDs specify different VPEs.
I _{MLMSJ}	At any given time for each PE in the system, there is either one VPE resident or no VPE resident . It is not possible for multiple VPEs to be resident on the same PE at the same time, even if those VPEs belong to different Security states.
I _{YSVWR}	The VM ID used in this specification is not directly related to the VMSA VMID[1].
I _{BJSFF}	When ICH_CONTEXTR_EL2.V is 0, no VPE is resident and ICH_CONTEXTR_EL2.{VM,VPE} are IGNORED.
R _{XVRCV}	If EL2 is not implemented, the Effective value of ICH_CONTEXTR_EL2.V is 0.
I _{WCHLY}	ICH_CONTEXTR_EL2.IRICHPPIDIS allows software to specify whether the candidate HPPIs presented by the IRI is considered when selecting the HPPI for the Virtual Interrupt Domain. GIC system instructions that operate in the virtual interrupt domain access the configuration and state of virtual LPis and SPIs belonging to the resident VM , irrespective of the value of ICH_CONTEXTR_EL2.IRICHPPIDIS.

S _{BYBKL}	Making a VPE resident with ICH_CONTEXTR_EL2.IRICHPPIDIS set to 1 allows software to emulate an environment where the virtual IRI is disabled by setting ICH_CONTEXTR_EL2.IRICHPPIDIS to 1 for all the VPEs that belong to the same disabled virtual IRI.
R _{LHFxQ}	On a write that sets ICH_CONTEXTR_EL2.V to 1, the VM and VPE are resolved in the Physical Interrupt Domain identified by SCR_EL3.{NSE,NS}.
R _{CDHXC}	If SCR_EL3.{NSE,NS} is set to a Reserved value, a write that sets ICH_CONTEXTR_EL2.V to 1 has the following CONSTRAINED UNPREDICTABLE behaviors: <ul style="list-style-type: none"> • No VPE is treated as being resident. • SCR_EL3.{NSE,NS} is treated as an UNKNOWN non-reserved value.
I _{FDPDR}	An access by a GIC System instruction to the state or configuration of a virtual LPI or SPI, when the Current Interrupt Domain is the Virtual Interrupt Domain, applies to the resident VM of that PE. Apart from GIC_VDPEND, where the target VM is specified as an argument to the instruction.
R _{RXMNN}	Changing SCR_EL3.{NSE,NS} when there is a resident VPE is CONSTRAINED UNPREDICTABLE with a choice of the following: <ul style="list-style-type: none"> • ICH_CONTEXTR_EL2.V is cleared to 0. • The CPU interface behaves as if no VPE is resident, but the value of ICH_CONTEXTR_EL2 is unchanged. <ul style="list-style-type: none"> – If ICH_CONTEXTR_EL2.V is subsequently written to 0, ICH_CONTEXTR_EL2.DB is RES0.
S _{XCDsv}	When switching Security state, Arm expects that software at EL3 writes to ICH_CONTEXTR_EL2 to either make no VPE resident or make a VPE of the new Security state resident .
I _{KWFRF}	Following a write that sets ICH_CONTEXTR_EL2.V to 1, one of the following is true: <ul style="list-style-type: none"> • The VPE becomes resident on the PE and ICH_CONTEXTR_EL2.F is set to 0. • The VPE is not resident on the PE and ICH_CONTEXTR_EL2.F is set to 1.
R _{NHCGW}	While ICH_CONTEXTR_EL2.V is 1, if ICH_CONTEXTR_EL2.VM and ICH_CONTEXTR_EL2.VPE select an invalid VPE, all of the following are true: <ul style="list-style-type: none"> • The IRI behaves as if no VPE is resident on the PE. • When a write updates ICH_CONTEXTR_EL2.V from 1 to 0, any doorbell request is ignored. • It is CONSTRAINED UNPREDICTABLE whether a read of ICH_CONTEXTR_EL2.V returns 1 or 0. <p>See also 2.10.6 Requesting VPE doorbells.</p>
R _{TXVLW}	When ICH_CONTEXTR_EL2.V is 1, if the VPE selected by ICH_CONTEXTR_EL2.{VM,VPE} becomes invalid and a new VPE selected by the same values subsequently becomes valid, it is CONSTRAINED UNPREDICTABLE whether the IRI continues to behave as if no VPE is resident on the PE or as if the new VPE had been made resident .
R _{LLMWL}	The effect of setting ICH_CONTEXTR_EL2.V to 1 is not guaranteed to be visible until after a Context synchronization event. After the Context synchronization event, if any virtual interrupt from the IRI is visible by reading the CPU interface registers, it is for the new resident VPE .
R _{LJVNC}	The Context synchronization event after setting ICH_CONTEXTR_EL2.{V,IRICHPPIDIS} to {1,0} will not complete until the IRI has presented the candidate HPPI for the new resident VPE , or the IRI has confirmed there is currently no candidate HPPI for the new resident VPE .
I _{LXZBK}	When using GICv5 Stream Protocol, when the IRS sends the SetResident_Ack() packet it indicates that it has either forwarded the first virtual interrupt for the new resident VPE or that there is no candidate HPPI for the resident VPE .
R _{WDKXF}	ICH_HPPIR_EL2 returns the HPPI for the Virtual Interrupt Domain. If there is no resident VPE , or ICH_CONTEXTR_EL2.IRICHPPIDIS is 1, the value returned represents whether there is a candidate HPPI among the virtual PPIs.

If there is a [resident VPE](#), the value returned represents whether there is a candidate HPPI among the virtual PPIs and the virtual candidate HPPI [presented](#) by the IRI for the [resident VPE](#).

R_{JSHCW} The effect of clearing ICH_CONTEXTR_EL2.V to 0, or setting ICH_CONTEXTR_EL2.IRICHPPIDIS to 1, is not guaranteed to be visible until after a Context synchronization event. After the Context synchronization event, no virtual interrupt from the IRI is visible via the CPU interface registers.

R_{VFKHN} Changing the value of ICH_CONTEXTR_EL2.VM or ICH_CONTEXTR_EL2.VPE when ICH_CONTEXTR_EL2.V is 1 is CONSTRAINED UNPREDICTABLE with the following permitted behaviors:

- The write is ignored, for all other purposes than a direct read of the register.
- The previously selected VPE is made not [resident](#) and the VPE identified by the new ICH_CONTEXTR_EL2.VM and ICH_CONTEXTR_EL2.VPE values is made [resident](#).

S_{QPFVS} To make a VPE no longer [resident](#), software performs the following sequence:

1. Write ICH_CONTEXTR_EL2, clearing ICH_CONTEXTR_EL2.V from 1 to 0. ICH_CONTEXTR_EL2.DB controls whether a doorbell interrupt is requested for the previously [resident VPE](#).
2. Issue an `ISB`.

To make a VPE [resident](#), software performs the following sequence:

1. Write ICH_CONTEXTR_EL2, setting ICH_CONTEXTR_EL2.V from 0 to 1.
 1. The ICH_CONTEXTR_EL2.VM and ICH_CONTEXTR_EL2.VPE fields specify which VPE is being made [resident](#).
 2. The Domain for the VPE is taken from SCR_EL3.{NSE,NS}.
2. Issue an `ISB`.
3. Optionally read ICH_CONTEXTR_EL2.F to check if the operation succeeded.

To change the [resident VPE](#), software must first ensure no VPE is [resident](#), by performing the following sequence:

1. Write ICH_CONTEXTR_EL2, clearing ICH_CONTEXTR_EL2.V from 1 to 0.
2. Write ICH_CONTEXTR_EL2, setting ICH_CONTEXTR_EL2.V from 0 to 1 and selecting the new [resident VPE](#).
3. Issue an `ISB`.
4. Optionally read ICH_CONTEXTR_EL2.F to check if the operation succeeded.

2.10.6 Requesting VPE doorbells

D_{GFNFV} When a [resident VPE](#) is made non-resident, software can optionally request a *VPE doorbell*.

The VPE doorbell is a request for a physical interrupt to signal that there is a candidate HPPI for the VPE when the VPE is not resident on any PE.

I_{QYLFV} When a VPE is made non-resident by writing 0 to ICH_CONTEXTR_EL2.V, software programs whether a VPE doorbell is requested by setting ICH_CONTEXTR_EL2.DB to the appropriate value.

I_{RWSJW} When a write that sets ICH_CONTEXTR_EL2.V to 0 also sets ICH_CONTEXTR_EL2.DB to 0, any previously requested doorbells are no longer requested.

I_{GVGWJ} ICH_CONTEXTR_EL2.DBPM allows software to specify a *VPE Doorbell priority mask*. If a VPE doorbell is requested, the corresponding physical interrupt is only signaled if the candidate HPPI for the VPE is greater than or equal to the value written to ICH_CONTEXTR_EL2.DBPM.

R_{VNZRM} The mechanism to configure the interrupt used for the VPE doorbell is IMPLEMENTATION DEFINED.

When the GICv5 system architecture is used, the VPE doorbell configuration mechanism is described in [4.10.7 VPE doorbells](#).

2.10.7 Legacy virtual CPU interface

G _{VXDMK}	GICv5 enables the use of unmodified GICv3.3 compatible virtual machines.
I _{BSYYP}	GICv5 provides the optional FEAT_GCIE_LEGACY extension to support GICv3.3 VMs.
D _{QHKVW}	When ICH_VCTLR_EL2.V3 is 1, <i>Legacy operation</i> is enabled.
R _{BFNZT}	When FEAT_GCIE_LEGACY is not implemented, the Effective value of ICH_VCTLR_EL2.V3 is 0.
D _{TPCQT}	When Legacy operation is enabled, the virtual CPU interface functionality is referred to as the <i>Legacy virtual CPU interface</i> .
I _{HKXPV}	When FEAT_GCIE_LEGACY is implemented and Legacy operation is enabled, the Legacy virtual CPU interface is managed through the AArch64 virtualization control system registers described in <i>Arm® Generic Interrupt Controller Architecture Specification, GIC architecture version 3 and version 4</i> [3].
I _{JBZMP}	When Legacy operation is disabled, the Effective value of ICH_HCR_EL2.En is 0 and the Legacy virtual CPU interface is disabled. When Legacy operation is enabled, ICH_HCR_EL2.En determines whether the Legacy virtual CPU interface is disabled. The effects of enabling and disabling the GICv3 Virtual CPU interface using ICH_HCR_EL2.En is described in [3].
I _{TQKBH}	When HCR_EL2.IMO is 1 and Legacy operation is enabled, the Legacy virtual CPU interface is used and all of the following are true: <ul style="list-style-type: none"> • When executing at EL1, all of the following are true: <ul style="list-style-type: none"> – All GICv5 System instructions are UNDEFINED. – Accesses to GICv5 System registers are UNDEFINED. – Accesses to GICv3.3 Group 0, Group 1, and Common System registers are to Legacy virtual CPU interface ICV_* registers. • Accesses to ICH_VMCR_EL2 use the GICv3.3 version of the register. <p>See Chapter 8 System instructions and Chapter 9 System registers for more information.</p>
R _{YZMJH}	When executing at Exception levels below EL2 and Legacy operation is enabled, the Effective value of HCR_EL2.FMO is determined by HCR_EL2.IMO: <ul style="list-style-type: none"> • When HCR_EL2.IMO is 0, the Effective value of HCR_EL2.FMO is 0. • When HCR_EL2.IMO is 1, the Effective value of HCR_EL2.FMO is 1. <p>This means that when Legacy operation is enabled and executing at EL1, software either interacts entirely with the GICv5 physical CPU interface or entirely with the Legacy virtual CPU interface.</p>
I _{PJWMZ}	When Legacy operation is enabled, HCR_EL2.FMO enables GICv3.3 virtual Group 0 interrupts to be signaled as virtual FIQ interrupts.
I _{NFJYV}	FEAT_GCIE_LEGACY does not provide Interrupt bypass support as described in <i>Arm® Generic Interrupt Controller Architecture Specification, GIC architecture version 3 and version 4</i> [3]. Therefore, when Legacy operation is enabled, the ICC_SRE_EL1.{DIB,DFB} field values are RAO/WI.
I _{LDWTK}	FEAT_GCIE_LEGACY only provides a system register interface to the Legacy virtual CPU interface. Therefore, when Legacy operation is enabled, the ICC_SRE_EL1.SRE field value is RAO/WI.
R _{XXCLD}	When Legacy operation is enabled, the ICV_CTLR_EL1.{ExtRange,RSS} field values are RAZ/WI.
I _{DQPYT}	When FEAT_GCIE_LEGACY is implemented, if ICH_LR<n>_EL2.HW is set to 1, ICH_LR<n>_EL2.pINTID identifies a PPI in the Current Physical Interrupt Domain.
I _{MTYHD}	When FEAT_GCIE_LEGACY is implemented, ICH_VTR_EL2.SEIS and ICH_HCR_EL2.TSEI are RES0, meaning that local generation of system errors and trapping them to EL2 is not supported.
I _{SHBPG}	When FEAT_GCIE_LEGACY is implemented, ICH_VTR_EL2.PRIbits and ICH_VTR_EL2.PREbits are 0b100, meaning that 5 bits of priority and preemption are available to GICv3 VMs.

I_{FPLXQ}

GICv3 recorded the active priorities for Group 0 in ICH_AP0R<n>_EL2 and active priorities for Group 1 in ICH_AP1R<n>_EL2. In GICv5, the Legacy virtual CPU interface implements the same split, however the ICH_AP1R0_EL2.P<n> fields access the same state as the corresponding fields in ICH_APR_EL2.

See also:

- [9.2 CPU interface registers](#)
- [9.6 Hypervisor control registers](#)
- [9.7 Legacy hypervisor control registers](#)

2.11 GIC synchronous exception priorities

I_{FTLLT}	The relative priority of synchronous exceptions is described in R_ZFGJP of the L.a version of <i>Arm® Architecture Reference Manual, for A-profile architecture</i> [1] .
R_{MJLRP}	Exceptions that occur as a result of attempting to execute an instruction that is UNDEFINED due to the configuration of ICH_VCTLR_EL2.V3 has a priority of 16.
R_{KNQBS}	<p>For an exception taken to EL2 as the result of a configuration control in one of the following registers, the exception priority is 22:</p> <ul style="list-style-type: none">• ICH_HCR_EL2.• ICH_HFGITR_EL2.• ICH_HFGRTR_EL2.• ICH_HFGWTR_EL2.

This is the same priority as exceptions due to the other fine-grained trap registers.

2.12 Interrupt ordering model and synchronization requirements

This section extends the Definitions of the Arm memory model in *Arm® Architecture Reference Manual, for A-profile architecture*[1].

Interrupt Location

An Interrupt Location comprises all state and configuration values relative to an interrupt. See [2.4 Interrupt types and identifiers](#) for more information about interrupt state and configuration values. An Interrupt Location is uniquely identified by an INTID.

Interrupt Effect

GIC and GICR System instructions might generate Interrupt Effects. An Interrupt Effect is an Interrupt Read Effect or an Interrupt Write Effect. Interrupt Effects can be either Explicit or Implicit. For example, a GIC System instruction with the RCFG command generates an Explicit Interrupt Read Effect, a GICR System instruction generates Implicit Interrupt Read Effects and might also generate an Implicit Interrupt Write Effect.

All Interrupt Write Effects are required to complete in finite time.

Coherence order There is a per-location *Coherence order* relation that provides a total order over all Interrupt Write Effects to that Location, starting with a notional Interrupt Write Effect of the initial value. The Coherence order of an Interrupt Location represents the order in which Interrupt Write Effects to the Interrupt Location arrive at their destination.

Reads-from The *Reads-from* relation couples Interrupt Read and Write Effects to the same Interrupt Location so that each Interrupt Read Effect is paired with exactly one Interrupt Write Effect in the execution of a program. An Interrupt Read Effect E_2 Reads-from an Interrupt Write Effect E_1 , if and only if E_1 and E_2 are to the same Interrupt Location and E_2 takes its data from E_1 .

For two Effects E_1 and E_2 if all of the following apply:

- E_1 is an Explicit Interrupt Read Effect generated by a GIC instruction with the RCFG command.
- E_1 appears in program order before E_2 .
- E_1 and E_2 are to the same Interrupt Location.
- E_2 is an Interrupt Write Effect.

then it is architecturally forbidden that E_1 Reads-from E_2 .

Coherence-before, Coherence-after An Interrupt Write Effect E_1 is *Coherence-before* an Interrupt Write Effect E_2 to the same Interrupt Location if E_1 is sequenced before E_2 in the Coherence order for the Interrupt Location.

An Interrupt Read Effect E_1 is *Coherence-before* an Interrupt Write Effect E_2 to the same Interrupt Location if E_1 Reads-from an Interrupt Write Effect E_3 and E_3 is Coherence-before E_2 .

An Effect E_2 is *Coherence-after* an Effect E_1 if E_1 is *Coherence-before* E_2 .

For two Effects E_1 and E_2 if all of the following apply:

- E_1 is an Implicit Interrupt Write Effect generated by a GICR system instruction.
- E_1 appears in program order before E_2 .
- E_1 and E_2 are to the same Interrupt Location.
- E_2 is an Interrupt Read Effect.

then it is architecturally forbidden that E_2 is Coherence-before E_1 .

For two Effects E_1 and E_2 if all of the following apply:

- E_1 is an Implicit Interrupt Write Effect generated by a GIC system instruction with the `DI` command.
- E_1 appears in program order before E_2 .
- E_1 and E_2 are to the same Interrupt Location.
- One of the following applies:
 - E_2 is an Explicit Interrupt Read Effect generated by a GIC system instruction with the `RCFG` command.
 - E_2 is an Implicit Interrupt Read Effect generated by a GICR system instruction.

then it is architecturally forbidden that E_2 is Coherence-before E_1 .

For two Effects E_1 and E_2 if all of the following apply:

- E_1 is an Explicit Interrupt Write Effect.
- E_2 is an Explicit Interrupt Read Effect.
- E_1 appears in program order before E_2 .
- E_1 and E_2 are to the same Interrupt Location.

then it is architecturally forbidden that E_2 is Coherence-before E_1 .

Successful Read-Modify-Write pair of Interrupt Effects

Two Effects E_1 and E_2 form a *Successful Read-Modify-Write pair* if and only if all of the following apply:

- E_1 is an Interrupt Read Effect.
- E_2 is an Interrupt Write Effect.
- E_1 and E_2 are to the same Interrupt Location.
- E_1 and E_2 are generated by the same instruction.

All Read-Modify-Write pair of Interrupt Effects generated by GIC and GICR System instructions are Successful Read-Modify-Write pairs of Interrupt Effects.

2.12.1 GIC and GICR ordering semantics

I_{SYSPH}

GIC System instructions generate Interrupt Effects as required by the instruction's `command`. The `domain` together with the `INTID` in `<Xt>` specify the Interrupt Location of the interrupt Effect.

A GIC System instruction with the `EOI` command does not generate any Interrupt Read Effects or Interrupt Write Effects.

A GIC System instruction with the `RCFG` command generates an Interrupt Read Effect but no Interrupt Write Effects.

All other GIC System instructions generate a Successful Read-Modify-Write pair of Interrupt Effects.

See [2.6 GIC System instructions](#) for more information about the A64 assembly language syntax for the GIC System instructions.

R_{ZWSQL}

- A GIC system instruction with the `RCFG` command generates the following Effects:
 - A Register Read Effect E_1 of the `<Xt>` register.
 - An Explicit Interrupt Read Effect E_2 to the Interrupt Location determined by the `domain` and `<Xt>`.
 - An Indirect Register Write Effect E_3 to the System Register `ICC_ICSR_EL1`.

In addition, all of the following apply:

- There is an Intrinsic Data Dependency from E_1 to E_2 .

- There is an Intrinsic Data Dependency from E_2 to E_3 .

 R_{MLYCN}

- A GIC system instruction with the `DI` command generates the following Effects:
 - A Register Read Effect E_1 of the $\langle Xt \rangle$ register.
 - An Implicit Interrupt Read Effect E_2 to the Interrupt Location determined by the domain and $\langle Xt \rangle$.
 - An Implicit Interrupt Write Effect E_3 to the Interrupt Location determined by the domain and $\langle Xt \rangle$.

In addition, all of the following apply:

- There is an Intrinsic Data Dependency from E_1 to E_2 .
- There is an Intrinsic Data Dependency from E_2 to E_3 .
- E_2 and E_3 form a Successful Read-Modify-Write pair.

 R_{XHKDN}

- A GIC system instruction with a command other than `RFCG`, `DI` or `EOI` generates the following Effects:
 - A Register Read Effect E_1 of the $\langle Xt \rangle$ register.
 - An Explicit Interrupt Read Effect E_2 to the Interrupt Location determined by the domain and $\langle Xt \rangle$.
 - An Explicit Interrupt Write Effect E_3 to the Interrupt Location determined by the domain and $\langle Xt \rangle$.

In addition, all of the following apply:

- There is an Intrinsic Data Dependency from E_1 to E_2 .
- There is an Intrinsic Data Dependency from E_2 to E_3 .
- E_2 and E_3 form a Successful Read-Modify-Write pair.

 R_{FDSPZ}

- A GICR system instruction with the `IA` or `NMIA` command generates the following Effects:
 - An Indirect System Register Read Effect E_1 to `ICC_HAPR_ELx`.
 - Implicit Interrupt Read Effect E_2, E_3, \dots, E_n to all n Interrupt Locations for the Current Interrupt Domain.
 - A Branching Effect E_{n+1} that determines the HPPI.
 - If there is an HPPI and its priority is higher than the current running priority:
 - * An Implicit Interrupt Write Effect E_{n+2} to the Interrupt Location determined as the HPPI.
 - * An Indirect System Register Write Effect E_{n+3} to `ICC_HAPR_ELx`.
 - A Register Write Effect E_{n+4} of the $\langle Xt \rangle$ register.

In addition, all of the following apply:

- There is an Intrinsic Data Dependency from E_2 to E_{n+1} .
- There is an Intrinsic Data Dependency from E_3 to E_{n+1} .
- ...
- There is an Intrinsic Data Dependency from E_n to E_{n+1} .
- There is an Intrinsic Control Dependency from E_{n+1} to E_{n+4} .
- If there is an HPPI and E_k is the Interrupt Read Effect from its Interrupt Location:
 - * There is an Intrinsic Data Dependency from E_k to E_{n+2} .
 - * There is an Intrinsic Control Dependency from E_{n+1} to E_{n+2} .
 - * There is an Intrinsic Control Dependency from E_{n+1} to E_{n+3} .
 - * E_k and E_{n+2} form a successful Read-Modify-Write pair.

2.12.2 GSB instruction semantics

D_{KZFML} The A64 assembly language syntax for GIC synchronization barrier instructions is one of the following:

- GSB SYS
- GSB ACK

D_{DLHHZ} A GSB instruction generates an Effect named after the instruction.

- A GSB SYS instruction generates a GSB SYS Effect.
- A GSB ACK instruction generates a GSB ACK Effect.

2.12.3 GIC Ordering Model

2.12.3.1 GIC Ordering Relations

GIC-hazard-ordered-before

An Effect E_1 is GIC-hazard-ordered-before an Effect E_2 if and only if all the following apply:

- E_1 is an Explicit Interrupt Read Effect generated by a GIC system instruction with the RCFG command.
- E_1 appears in program order before E_3 .
- E_1 and E_3 are to the same Interrupt Location.
- E_3 is an Interrupt Read Effect.
- E_3 is Coherence-before E_2 .
- E_2 is an Explicit Interrupt Write Effect.

If an Effect E_1 is GIC-hazard-ordered-before an Effect E_2 then E_1 is Hazard-ordered-before E_2 .

GSB-ordered-before An Effect E_1 is GSB-ordered-before an Effect E_2 if and only if any of the following apply:

- All of the following apply:
 - E_1 is an Implicit Interrupt Effect generated by a GICR system instruction.
 - E_3 is a GSB ACK Effect.
 - E_1 appears in program order before E_3 .
 - E_3 appears in program order before E_2 .
 - It is not the case that E_2 is an Implicit Instruction Memory Read Effect.
- All of the following apply:
 - E_1 is an Interrupt Effect.
 - E_3 is a GSB SYS Effect.
 - E_1 appears in program order before E_3 .
 - E_3 appears in program order before E_2 .
 - It is not the case that E_2 is an Implicit Instruction Memory Read Effect.

If an Effect E_1 is GSB-ordered-before an Effect E_2 then E_1 is Locally-ordered-before E_2 .

If an Effect E_1 is GSB-ordered-before an Effect E_2 , E_2 is an Instruction Fetch Barrier Effect and E_2 is in program-order before E_3 then E_1 is Instruction-fetch-barrier-ordered-before E_3 .

2.12.3.2 Adaptations to existing Ordering relations

The following relations are defined in the chapter “Ordering requirements defined by the formal concurrency model” in *Arm® Architecture Reference Manual, for A-profile architecture*[1] and adapted to account for the Interrupt Read or Write Effects.

Basic Dependency There is a Basic dependency from an Effect E_1 to an Effect E_2 if and only if all the following apply:

- Any of the following applies:
 - E_1 is an Explicit Memory Read Effect.
 - E_1 is an Interrupt Read Effect.
 - E_1 is a Register Read Effect.
- Any of the following applies:
 - There is a Dependency through registers and memory from E_1 to E_2 .
 - E_1 and E_2 are the same Effect.

Data dependency There is a Data dependency from an Effect E_1 to an Effect E_2 if and only if all the following apply:

- Any of the following applies:
 - E_1 is an Explicit Memory Read Effect.
 - E_1 is an Interrupt Effect.
- There is a Basic dependency from E_1 to E_3 .
- E_3 affects the data value written by E_2 .
- There exists a chain of Intrinsic Data Dependency from E_3 to E_2 .
- E_2 is an Explicit Memory Write Effect.
- It is not the case that E_1 and E_2 are generated by the same instruction.

Address dependency

There is an *Address dependency* from an Effect E_1 to an Effect E_2 if and only if all the following apply:

- Any of the following applies:
 - E_1 is an Explicit Memory Read Effect.
 - E_1 is an Interrupt Read Effect.
- There is a Basic dependency from E_1 to E_3 .
- E_3 affects the address of the Location accessed by E_2 .
- There exists a chain of Intrinsic data dependency from E_3 to E_2 .
- Any of the following applies:
 - E_2 is an Explicit Memory Effect.
 - E_2 is an Implicit Tag Memory Read Effect.
 - E_2 is an Implicit TTD Memory Read Effect.
 - E_2 is a Hardware Update Effect.
 - E_2 is a TLBI Effect.
 - E_2 is a DC CVAU Effect.

- E_2 is an IC IVAU Effect.
- E_2 is an Interrupt Effect
- It is not the case that E_1 and E_2 are generated by the same instruction.

Control dependency There is a Control dependency from an Effect E_1 to an Effect E_2 if and only if all the following apply:

- Any of the following applies:
 - E_1 is an Explicit Memory Read Effect.
 - E_1 is an Interrupt Effect.
- There is a Basic dependency from E_1 to E_3 .
- E_3 is a Conditional Branching Effect.
- E_3 appears in program order before E_2 .

Pick Basic dependency There is a *Pick Basic dependency* from an Effect E_1 to an Effect E_2 if and only if all the following apply:

- Any of the following applies:
 - E_1 is an Explicit Memory Read Effect.
 - E_1 is a Register Read Effect.
 - E_1 is an Interrupt Effect.
- Any of the following applies:
 - There is a Pick dependency through registers and memory from E_1 to E_2 .
 - E_1 and E_2 are the same Effect.

Pick Data dependency There is a Pick Data dependency from an Effect E_1 to an Effect E_2 if and only if all the following apply:

- Any of the following applies:
 - E_1 is an Explicit Memory Read Effect.
 - E_1 is an Interrupt Read Effect.
- There is a Pick Basic dependency from E_1 to E_3 .
- E_3 affects the data value written by E_2 .
- Any of the following applies:
 - There is an Intrinsic Data Dependency from E_3 to E_2 .
 - There is an Intrinsic Control Dependency from E_3 to E_2 .
 - There exists a chain of Intrinsic Data Dependency or Intrinsic Control Dependency from E_3 to E_2 .
- E_2 is an Explicit Memory Write Effect.
- It is not the case that E_1 and E_2 are generated by the same instruction.

Pick Address dependency There is a Pick Address dependency from an Effect E_1 to an Effect E_2 if and only if all the following apply:

- Any of the following applies:
 - E_1 is an Explicit Memory Read Effect.
 - E_1 is an Interrupt Read Effect.

- There is a Pick Basic dependency from E_1 to E_3 .
- E_3 affects the address of the Location accessed by E_2 .
- There exists a chain of Intrinsic Data Dependency from E_3 to E_2 .
- Any of the following applies:
 - E_2 is an Explicit Memory Effect.
 - E_2 is an Implicit Tag Memory Read Effect.
 - E_2 is an Implicit TTD Memory Read Effect.
 - E_2 is a Hardware Update Effect.
 - E_2 is a TLBI Effect.
 - E_2 is a DC CVAU Effect.
 - E_2 is an IC IVAU Effect.
 - E_2 is an Interrupt Read Effect.
- It is not the case that E_1 and E_2 are generated by the same instruction.

Pick Control dependency There is a Pick Control dependency from an Effect E_1 to an Effect E_2 if and only if all the following apply:

- Any of the following applies:
 - E_1 is an Explicit Memory Read Effect.
 - E_1 is an Interrupt Read Effect.
- There is a Pick Basic dependency from E_1 to E_3 .
- E_3 is a Conditional Branching Effect.
- E_3 appears in program order before E_2 .

2.12.3.3 GIC Observation Relations

GIC-observed-by An Effect E_1 is GIC-observed-by an Effect E_2 if and only if any of the following apply:

- All of the following apply:
 - E_1 is an Interrupt Write Effect.
 - E_2 Reads-from E_1 .
 - E_2 is an Interrupt Read Effect.
- All of the following apply:
 - E_1 is an Explicit Interrupt Effect.
 - E_1 is Coherence-before E_2 .
 - E_1 and E_2 are from different Processing Elements.
 - E_2 is an Explicit Interrupt Write Effect.

If an Effect E_1 is GIC-observed-by an Effect E_2 then E_1 is Observed-by E_2 .

2.12.3.4 Ordering Requirements

Atomicity of Successful Read-Modify-Write pair of Interrupt Effects

For two Interrupt Effects E_1 and E_2 if and only if all the following apply:

- E_1 and E_2 form a successful Read-Modify-Write pair.
- There is an Interrupt Write Effect E_3 .

then it is not the case that E_1 is Coherence-before E_3 and E_3 is Coherence-before E_2 .

An architecturally well-formed execution must not exhibit a cycle in the Ordered-before relation as defined in the *Arm[®] Architecture Reference Manual, for A-profile architecture*[\[1\]](#) and extended by the definitions in this section [GIC Ordering Model](#).

2.13 Effects on the Transactional Memory Extension

I_{MCXXC}	The Transactional Memory Extension (TME)[4] defines a set of features to support transactional memory on Armv9-A PEs. This section describes changes made by GICv5 to TME.
R_{PFDPV}	<p>The following registers are added to the subset of AArch64 state included in a transaction checkpoint:</p> <ul style="list-style-type: none"> • If FEAT_GCIE is implemented, ICC_PCR_EL1, ICV_PCR_EL1, and ICC_PCR_EL3. • If FEAT_GCIE_LEGACY is implemented, and ICV_PMR_EL1.
D_{NYKMG}	<p><i>GIC Priority Mask</i> is defined to mean the register containing the priority mask for the current interrupt regime.</p> <p>In GICv3 and GICv4, or in GICv5 when using legacy mode, the GIC Priority Mask is:</p> <ul style="list-style-type: none"> • ICC_PMR_EL1 if in EL2 or EL3. • ICC_PMR_EL1 or ICV_PMR_EL1 if in EL0 or EL1. <p>In GICv5, the GIC Priority Mask is:</p> <ul style="list-style-type: none"> • ICC_PCR_EL3 if in EL3. • ICC_PCR_EL2 if in EL2. • ICC_PCR_EL1 or ICV_PCR_EL1 if in EL0 or EL1.
I_{DNGVS}	Transactional code with sufficient privileges can change the value of DAIF or the GIC Priority Mask to mask or unmask interrupts.
R_{PVHGH}	<p>A transaction fails with IMP cause and INT set if both of the following happen:</p> <ul style="list-style-type: none"> • An unmasked interrupt delivered to a PE leads to the currently executing transaction on the PE to fail. • Upon restoring DAIF and the GIC Priority Mask, the interrupt becomes masked again and will not be taken.
R_{RNKXD}	<p>If FEAT_GCIE is implemented, the following system registers behave the same in Transactional state as Non-transactional state:</p> <ul style="list-style-type: none"> • ICC_HAPR_EL1 • ICC_HPPIR_EL1 • ICC_HPPIR_EL3 • ICC_PCR_EL1 • ICC_PCR_EL2 • ICC_PCR_EL3 • ICV_PCR_EL1 <p>If FEAT_GICv3 or FEAT_GCIE_LEGACY is implemented, the following system registers behave the same in Transactional state as Non-transactional state:</p> <ul style="list-style-type: none"> • ICC_HPPIR0_EL1 • ICC_HPPIR1_EL1 • ICC_PMR_EL1 • ICC_RPR_EL1 • ICV_PMR_EL1
I_{XVVBP}	Attempting to read or write any other GIC system register, including the PPI registers, in Transactional state fails the transaction with ERR cause.
R_{GQWJM}	While in Transactional state, the effect of GICv5 System instructions is:

Mnemonic	Instruction	Behavior
GICR	GICv5 read operation	Transaction fails with ERR cause
GIC	GICv5 write operation	Transaction fails with ERR cause

GSB	GICv5 synchronization barrier	Same effects as in non-transaction state.
-----	-------------------------------	---

 I_{NCLRB}

The restriction on use of `GIC` and `GICR` in Transactional state means that an interrupt cannot be acknowledged, deactivated or re-configured inside a transaction. Additionally, software cannot perform a Priority drop inside a transaction. It also means that the configuration and state of an interrupt cannot be queried within a transaction.

Chapter 3

GICv5 system architecture

	This section defines the GICv5 system architecture, including its components and interrupt handling behavior.
G _B CP _{XN}	The IRI manages LPIs and SPIs. If the PEs in the system implement EL2, the IRI tracks which VMs and VPEs are defined and maintains their virtual interrupt state.
I _{CKLJJ}	This section describes an IRI implemented using the GICv5 system architecture.
I _{YXYKQ}	The GICv5 system architecture defines the following classes of system component: <ol style="list-style-type: none">1. The <i>Interrupt Routing Service</i> (IRS).2. The <i>Interrupt Translation Service</i> (ITS).3. The <i>Interrupt Wire Bridge</i> (IWB).

[Figure 3.1](#) shows an overview of the GICv5 system architecture.

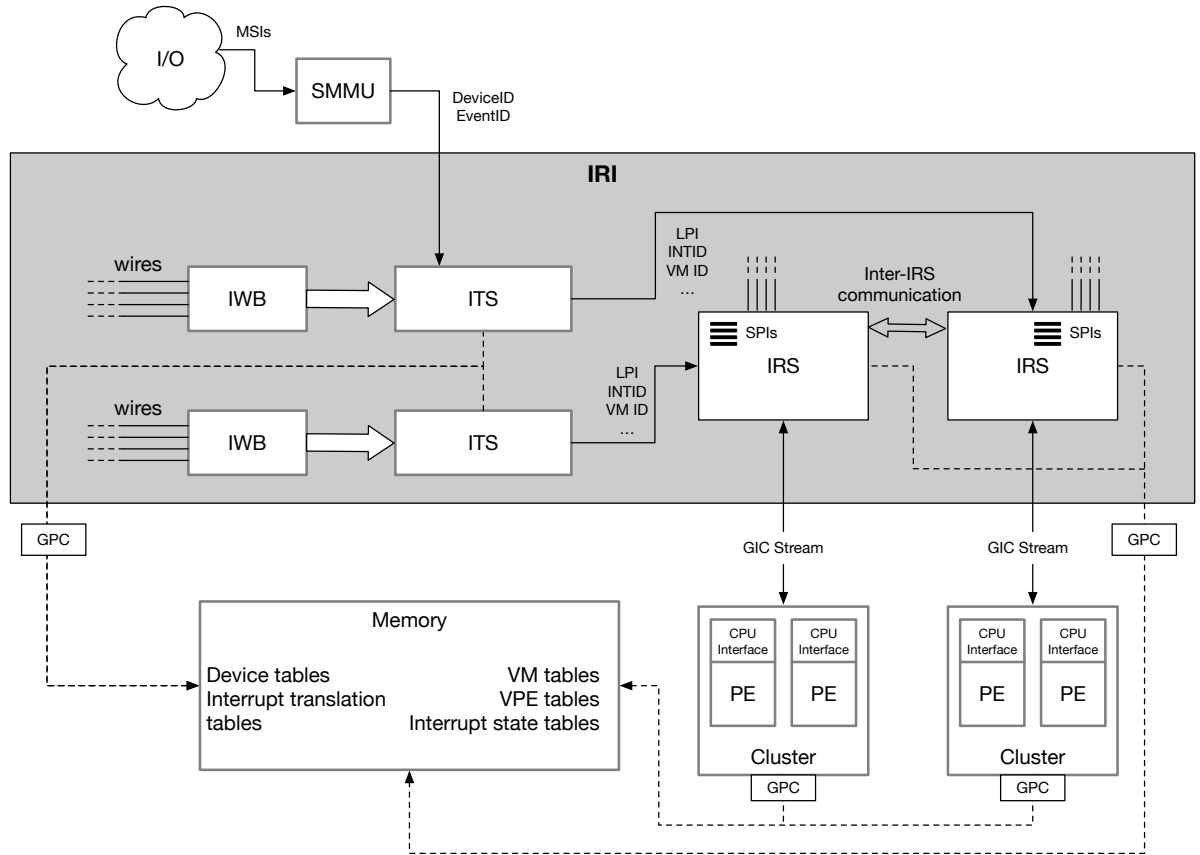


Figure 3.1: System overview

D_{MZRCX}

A GICv5 system using the GICv5 system architecture comprises the following:

- One or more IRSs.
- Zero or more ITSs.
- Zero or more IWBs.

I_{BFXPB}

A GICv5 system has optional support for GIC PMUs. See [Chapter 7 GIC Performance Monitoring Unit \(PMU\)](#) for more information.

I_{ZWPSS}

IRSs and ITSs rely on data structures stored in memory to manage the routing and translation of interrupts.

Software allocates memory for the data structures used by the IRSs and ITSs and configures these components using memory-mapped I/O interfaces.

D_{QJPLJ}

An IRS supports the following types of incoming interrupt events from peripherals:

Shared Peripheral Interrupts (SPIs) SPIs are used for interrupt signals that are directly connected to the IRS, for example, using directly wired connections or via an IMPLEMENTATION DEFINED mechanism. SPIs do not rely on memory for storage and may be used for interrupts that must meet certain early boot, reliability, or real-time guarantees. Whether SPIs provide any guarantees of fixed latency is a property of a given implementation. The number of supported SPIs is a fixed property of a given implementation.

Logical Peripheral Interrupts (LPIs) LPIs are used for interrupt signals that are translated by an ITS and forwarded to the IRS as well as for VPE doorbells. LPIs are managed by the IRSs in the system and the state and configuration of the LPIs are stored in memory allocated by system software.

Software may also signal interrupts managed by the IRS, for example to support IPIs and software emulated peripherals. See [2.5 Inter-Processor Interrupts](#) and [2.6 GIC System instructions](#) for more information.

I_{DWZTM}

When an IRS detects or receives an interrupt event, it updates the Pending state of the LPI or SPI specified by the interrupt event INTID. An interrupt event can either set or clear the Pending state depending on the interrupt event type.

The IRS supports IPIs by allowing CPU Interfaces to send commands to the IRS to make an interrupt Pending.

The IRSs keep track of all Pending, Inactive, and Enabled interrupts and route them to the target PEs. An IRS forwards an interrupt to a PE if it meets the following criteria:

- The interrupt targets that PE, either directly or through 1ofN routing.
- The interrupt is Pending, Inactive, and Enabled.
- The interrupt has the highest priority among all interrupts that are Pending, Inactive, and Enabled for that PE.

A system can include one or more IRSs and an implementation manages shared interrupt state across IRSs. Each PE is connected to exactly one IRS and an IRS can be connected to multiple PEs. The GICv5 architecture does not define how multiple IRSs communicate with each other, for example in a multi-chip system. A system with multiple IRSs supports changing the affinity of an interrupt from a PE on one IRS to a PE on another IRS, without losing interrupt state or configuration.

The IRS also manages VM state, supports virtual interrupts, and manages VPE doorbells. The IRS uses data structures stored in memory to manage the state of interrupts and configuration of VMs and VPES.

 I_{HNNWD}

An ITS provides translation of incoming interrupt events into LPI INTIDs and forwards the translation results as interrupt events for an IRS. Interrupt events can be generated from wired interrupt signals connected to an IWB or from Message Signaled Interrupts (MSIs) of other subsystems. For example, an ITS can translate MSIs generated by PCIe endpoints routed through an SMMU. It can also translate MSIs generated by other system-specific subsystems. ITS translations specify whether an event should be signaled as a physical interrupt or as a virtual interrupt that is injected directly into a VM. The ITS uses data structures stored in memory to perform translation of events.

 I_{HXBQP}

An IWB generates interrupt events from wired interrupt signals and forwards the events to an ITS. An IWB supports both edge-triggered and level-sensitive wires, and associates each wired interrupt with a Security state. The GICv5 architecture supports systems with zero to any number of IWBs.

 I_{WWTGL}

The GICv5 system architecture maintains the Arm CPU architecture security guarantees. As in the PE architecture, the IRSs and ITSs can support multiple Security states and access to memory respect that each location in memory belongs to a separate PAS. For a system that implements the Realm Management Extension (RME), this means all accesses are subject to Granule Protection Checks (GPCs). The GPC mechanism can be implemented in various ways and are outside the scope of this document. Even though the GPC is shown in [Figure 3.1](#) as occurring downstream from the ITS and IRS, it is possible to leverage the SMMU GPC mechanism for this purpose. See *Arm® System Memory Management Unit Architecture Specification, SMMU architecture version 3* [5] for more information.

Memory-mapped accesses to the IWB, ITS, and IRS registers are also subject to PAS filtering [6]. PAS filtering for the IWB occurs at the level of individual registers. PAS filtering for the ITS and IRS is performed at page granularity, either internally or externally to the component.

3.1 Interrupt Domains

- I_{LTWNQ}** In a GICv5 system, the supported Interrupt Domains depend on the Security states implemented by the PEs. See [2.3 Interrupt Domains](#) for more information.
- I_{ZGKPP}** Each GICv5 system component may implement support for a subset of the Interrupt Domains supported by the PEs, except from the IRS, which implements support for the same set Interrupt Domains as the connected PEs. See [4.3 IRS Domains](#) for more information.

Note

Implementing different Security state configurations across PEs in a system is not supported. This restriction exists for reasons outside the scope of this specification See [2] for more information.

- R_{VDTVW}** A Physical Interrupt Domain is associated with a Security state and a PAS:

Physical Interrupt Domain	Security state	PAS
Non-secure	Non-secure	Non-secure
Secure	Secure	Secure
Realm	Realm	Realm

When the EL3 Interrupt Domain is supported, it is either associated with the Secure or Root Security state and PAS as follows:

Supported Security states	EL3 Interrupt Domain Security state	EL3 Interrupt Domain PAS
Non-secure, Secure	Secure	Secure
Non-secure, Realm, Root	Root	Root
Non-secure, Realm, Secure, Root	Root	Root

- I_{KGBSN}** The PAS associated with a Physical Interrupt Domain is used in the following scenarios:
- When memory-mapped registers of the IWB, ITS, and IRS are accessed, the PAS defines the Physical Interrupt Domain of the registers being accessed.
 - When the ITS and IRS access data structures stored in memory, the access is always associated with a Physical Interrupt Domain and uses the PAS associated with that Physical Interrupt Domain.

If the GIC performs a memory access which fails a GPC, it experiences an external abort. The behavior for each GIC component is described in [Chapter 4 Interrupt routing service \(IRS\)](#) and [Chapter 5 Interrupt translation service \(ITS\)](#).

- D_{VJVNZ}** The GICv5 architecture defines the *Most Privileged Security State* (MPSS) and *Most Privileged PAS* (MPPAS) as follows:

Supported Security states	MPSS	MPPAS
Non-secure	Non-secure	Non-secure
Secure	Secure	Secure

Supported Security states	MPSS	MPPAS
Non-secure, Secure	Secure	Secure
Non-secure, Realm, Root	Root	Root
Non-secure, Realm, Secure, Root	Root	Root

See also:

- [2.3 Interrupt Domains](#)
- [Chapter 4 Interrupt routing service \(IRS\)](#)
- [Chapter 5 Interrupt translation service \(ITS\)](#)
- [Chapter 6 Interrupt Wire Bridge \(IWB\)](#)

3.2 Communication between GIC system components

R_{ZFCFX} The mechanisms used for communication between GICv5 system components are IMPLEMENTATION DEFINED. For example, the mechanism used by the ITS to send an interrupt message to the IRS is IMPLEMENTATION DEFINED.

D_{ZFVZM} The architecture supports the following flows of events between GICv5 system components:

- An IWB can send events to an ITS.
- An ITS can send events to an IRS.

Each communication path involves a source and a destination.

When a source sends an event to the destination, the event is *Accepted* when the destination has sent a reply to the source confirming that it has received the event.

An Accepted event is not guaranteed to have been handled by the destination. This means that all of the following are true:

- If the destination is an ITS, an Accepted event is not guaranteed to have been translated by the ITS.
- If the destination is an IRS, an Accepted event is not guaranteed to have been processed by the IRS.

See also [Chapter 3 GICv5 system architecture](#)

I_{VFJSX} An event is always Accepted by the destination regardless of whether the destination is enabled or disabled.

I_{MNGPV} An example interconnect of two GICv5 system components is the *AMBA[®] AXI Protocol Specification*[7].

In this example, the source is a manager and the destination is a subordinate.

The subordinate has Accepted an event when it has sent a Write Response to the manager.

Further, in this example, an ITS would send a Write Response to the IWB when the ITS has Accepted an event from the IWB.

In this example, it is assumed that transactions are not Bufferable. This means that AxCACHE[0] is deasserted and therefore no intermediate send a Read or Write Response to the manager.

I_{MZDPH} The GICv5 Stream Protocol defines an optional standard interface between an IRS and a CPU interface.

The GICv5 Stream Protocol defines a two-way communication mechanism that does not rely on the concept of Accepted as defined in this specification.

See [Chapter A1 GICv5 Stream Protocol overview](#) for more information.

3.3 Coherency considerations for GIC data structures

D _{NCTBH}	PEs and GIC system components communicate using <i>shared data structures</i> .
I _{JPHKW}	The GIC system components may be fully coherent with PE caches, IO-coherent, or not coherent. Software must take actions to ensure that writes from the PE are visible to the GIC, and vice versa.
S _{MXSNW}	Arm expects that firmware data structures describe to the OS or hypervisor whether the GIC is cache-coherent with PEs and is in the same Inner Shareability domain.
I _{QGRVW}	Arm recommends that all GIC system components are in the same Inner Shareability domain as the PEs.

Chapter 4

Interrupt routing service (IRS)

└_{BRKFW}

An IRS implements support for LPIs and SPIs. An IRS is either connected to one or more PEs, or it is stand-alone with no direct PE connections. A system can contain multiple IRSs.

The GICv5 architecture does not define the topological relationship between PEs and IRSs. Each PE is connected to exactly one IRS and an IRS can be connected to multiple PEs. An IRS can also be standalone and not be connected to any PE. This can be useful to support system topologies where an IRS and its associated ITS are located separately from PEs, allowing a more flexible routing of interrupts and power management of components close to the PE.

The mechanisms by which multiple IRSs communicate are IMPLEMENTATION DEFINED. The IRS stores interrupt state and configuration in memory provisioned to the IRS by software. In a system with multiple IRSs, the Affinity of an interrupt in an Interrupt Domain can be changed from a PE connected to an IRS to another one connected to another IRS, without losing interrupt state or configuration, if both IRSs can access the provisioned memory for that Interrupt Domain.

To support interrupt signaling in scenarios where memory is unavailable, such as during early boot, power transitions, or RAS handling, the IRS supports SPIs.

SPIs do not require a data structure allocated by software in order to be signaled, or to be managed and handled by PEs. Every SPI is assigned to an Interrupt Domain and the IRS exposes a programming interface to configure which Interrupt Domain an SPI is assigned to.

The IRS also provides support for virtual interrupts and manages VMs and VPEs. The IRS supports storing interrupt state and configuration in memory-backed data structures.

An IRS is responsible for:

- Generating SET_EDGE, SET_LEVEL, and CLEAR interrupt events for SPIs as a result of interrupts connected to the IRS being signaled.

- Receiving SET_EDGE, SET_LEVEL, and CLEAR interrupt events from ITSs for LPIs.
- Generating SET_EDGE LPI interrupt events as a result of writes to the SETLPI register, if implemented.
- Selecting the physical candidate HPPIs among LPIs and SPIs for each PE connected to the IRS.
- Selecting the virtual candidate HPPIs among virtual LPIs and virtual SPIs for each resident VPE on PEs connected to the IRS.
- Delivering virtual and physical HPPIs to the CPU interface.
- [Processing](#) Interrupt Effects from GIC System instructions executed on PEs.
- Managing residency state of VPEs on PEs, including:
 - Generating doorbell interrupts for non-resident VPEs.

I_{NZXZQ} In this section, the term PE(s) refers to PE(s) connected to the IRS.

I_{FZLBL} Arm strongly recommends that all PEs connected to a GIC are application PEs running a common system software stack.

The architecture permits the PEs connected to an IRS to be a combination of application PEs and non-application PEs, as long as all PEs implement a GICv5 compliant CPU interface.

However, the GIC architecture does not provide features to provide isolation or prevent interference between equally privileged host software running in the same Interrupt Domain. Therefore, if the PEs connected to an IRS are running separate system software stacks, a mechanism to manage the shared IRS and other GIC components must be supported by system software.

I_{QBJVX} The IRSs replace parts of the functionality of the Redistributors and Distributors in GICv3.

I_{GCWBM} The IRS manages LPIs and SPIs.

R_{FMTFZ} When a system includes more than one IRS, each IRS supports the same number of interrupt ID bits for each interrupt type.

R_{DRRHH} Each PE is associated with exactly one IRS. Zero or more PEs can be connected to a single IRS.

D_{LJQ SX} The IRS where an interrupt is signaled is referred to as the *local IRS* for that interrupt.

Similarly, the IRS that a PE is connected to is referred to as the *local IRS* for that PE.

D_{CHPRY} A system with more than one IRS is referred to as a *multi-IRS system*.

4.1 Communication between the IRS and the CPU interface

R_{HWCBT}	The mechanism for communication between a CPU interface in a PE and the local IRS is IMPLEMENTATION DEFINED.
I_{NTXCC}	<p>Arm recommends that an implementation uses the GICv5 Stream Protocol for communication between PEs and IRSs.</p> <p>See Chapter A1 GICv5 Stream Protocol overview for more information.</p>
I_{LTXQF}	Access to virtual interrupt configuration and state is made in the context of the resident VPE on the PE. When the PE switches the Current Physical Interrupt Domain, the PE and the IRS must ensure that software cannot access configuration and state for interrupts in the context of a VPE from the previous Physical Interrupt Domain. The GICv5 Stream Protocol avoids this by tagging all commands with the Physical Interrupt Domain.
I_{SBMVW}	The IRS manages the state and configuration of the LPis and SPIs, and selects a candidate HPPI from among LPis and SPIs for each Interrupt Domain for each connected PE. The PE selects the HPPI for each Interrupt Domain from the candidate HPPI and the PPIs managed in the PE.

4.2 Signaling interrupts

D _{YDND}	An interrupt event is <i>processed</i> when the Interrupt Effects of the event are Ordered-before an Interrupt Read Effect to the Interrupt Location.
D _{KLYXV}	<p>An interrupt is <i>signaled</i> to an IRS when one of the following events occurs:</p> <ul style="list-style-type: none"> • An interrupt connected to the IRS as an SPI is asserted or de-asserted and generates an interrupt event to change the Pending state of the SPI. • An ITS associated with the IRS generates an interrupt event to change the Pending state of an LPI. • A write to the IRS_SETLPI register occurs to set the Pending state of an LPI. • The IRS <i>processes</i> Interrupt Effects of a GIC System instruction that sets the Pending state of an interrupt.
R _{JVVS}	Updates to the Pending state of interrupts occur only in response to defined signaling events. Speculative updates are not permitted.
R _{QHGBY}	The mechanism for signaling SPI interrupt events to the IRS is IMPLEMENTATION DEFINED.
R _{MJNXL}	The mechanism for communication between an ITS and the IRS is IMPLEMENTATION DEFINED.
R _{FNGRT}	When multiple interrupt events of different types are generated for the same Interrupt Location, their Interrupt Effects are ordered in the same order in which the events are generated.
R _{VGJBD}	<p>An interrupt event received from an ITS contains the following information:</p> <ul style="list-style-type: none"> • Interrupt Domain. • Physical or virtual interrupt. • VM ID for a virtual interrupt. • LPI INTID. • Event type: SET_EDGE, SET_LEVEL, or CLEAR.
R _{ZHJMZ}	<p>If more than one mapping for an LPI exists, either in a single ITS or across multiple ITSs, such that both mappings are capable of generating events for the same LPI to one or more IRSs concurrently, all of the following are true for events with multiple mappings:</p> <ul style="list-style-type: none"> • The order in which the events are received by the IRSs is UNKNOWN. • The order in which the events are <i>processed</i> by the IRSs is UNKNOWN. • It is CONSTRAINED UNPREDICTABLE whether the IRS <i>processes</i> all the events or ignores some of the events.
R _{DJDSY}	<p>The following sequence avoids the CONSTRAINED UNPREDICTABLE behavior caused by multiple concurrent ITS mappings to the same LPI:</p> <ol style="list-style-type: none"> 1. For each ITS where a mapping to the LPI exists, the following sequence is performed: <ol style="list-style-type: none"> 1. The mapping is removed. 2. A synchronization request is issued. 3. A synchronization request is issued to the connected IRS. 2. A new mapping can be created on an ITS to the LPI. <p>See also 5.2.4 ITS synchronization requests and 4.5 IRS synchronization requests.</p>
I _{JNQGW}	<p>An IRS supports the following interrupt events generated for SPIs or received from ITSs to update the state and configuration of an interrupt:</p> <ul style="list-style-type: none"> • SET_EDGE. • SET_LEVEL. • CLEAR. <p>An IRS supports the following interrupt events when <i>processing</i> Interrupt Effects of a GIC System instruction that sets the Pending state of an interrupt:</p> <ul style="list-style-type: none"> • SET. • CLEAR.

When a write to IRS_SETLPI occurs, the IRS generates a SET_EDGE event locally for the specified LPI.

R_{HHKMN}

When the IRS [processes](#) an interrupt event, there is an Interrupt Write Effect to the interrupt state as follows:

- For a SET event, the interrupt Pending state is set to Pending.
- For a SET_EDGE event, the interrupt Handling mode is set to Edge and the Pending state is set to Pending.
- For a SET_LEVEL event, the interrupt Handling mode is set to Level and the interrupt Pending state is set to Pending.
- For a CLEAR event, the interrupt Pending state is set to Idle.

I_{DRNMD}

As interrupts are asynchronous, the Interrupt Write Effects of an operation by a PE might be overwritten by a subsequent event received from the interrupt source.

For SPIs where the Trigger mode is level-sensitive and where an interrupt signal is physically connected to the SPI, updates to the Pending state from GIC System instruction executed on a PE are permitted to be ignored by the IRS, as the Pending state may be determined by the current state of signal.

R_{QKYQM}

When a mapping for an LPI exists and the mapping is for events generated by an IWB, and the events are for a wire that is configured as level-sensitive in the IWB, the IRS is permitted to generate a CLEAR event for the LPI when the ITS mapping is updated or when the wire is assigned to a different Interrupt Domain in the IWB.

R_{DQTDV}

The SET_EDGE and SET_LEVEL interrupt events generate the following Interrupt Write Effects:

- An Interrupt Write Effect E₁ that updates the Handling mode.
- An Interrupt Write Effect E₂ that updates the Pending state.

For these events, E₁ is Ordered-before E₂.

I_{NFMBM}

The IRS_SETLPI register can be used in a system that does not support virtualization to support MSIs without the use of an ITS.

A hypervisor can also emulate this architected functionality to allow a VM to program a device assigned to the VM to generate an MSI without requiring the hypervisor to emulate an ITS.

R_{LCSNJ}

If an interrupt event specifies an unreachable INTID at the IRS where it is signaled, the event has no effect on any interrupt.

I_{XZDTX}

If an interrupt event specifies an unreachable INTID, and software error reporting is supported, this is reported with an appropriate error code in IRS_SWERR_STATUSR.

See also:

- [2.2 The GICv5 CPU interface](#)
- [Chapter 5 Interrupt translation service \(ITS\)](#)

4.3 IRS Domains

- R_{TCWZB}

The IRSs support interrupt routing and delivery with separate configurations for each supported Interrupt Domain.
- D_{LDYmZ}

The GICv5 architecture defines an *IRS Domain* that provides interrupt routing and management services for an Interrupt Domain. An IRS Domain comprises register state and interrupt configuration structures across all IRSs in the system. Each IRS Domain is configured independently from other IRS Domains. The IST and the VM table are shared between all IRSs for the same IRS Domain, but are not shared across IRS Domains. The data structures are accessed with the PAS associated with the IRS Domain.
- I_{FGLKF}

Figure 4.1 shows an overview of the IRS Domains and the register frames for each IRS Domain for a single-IRS system.

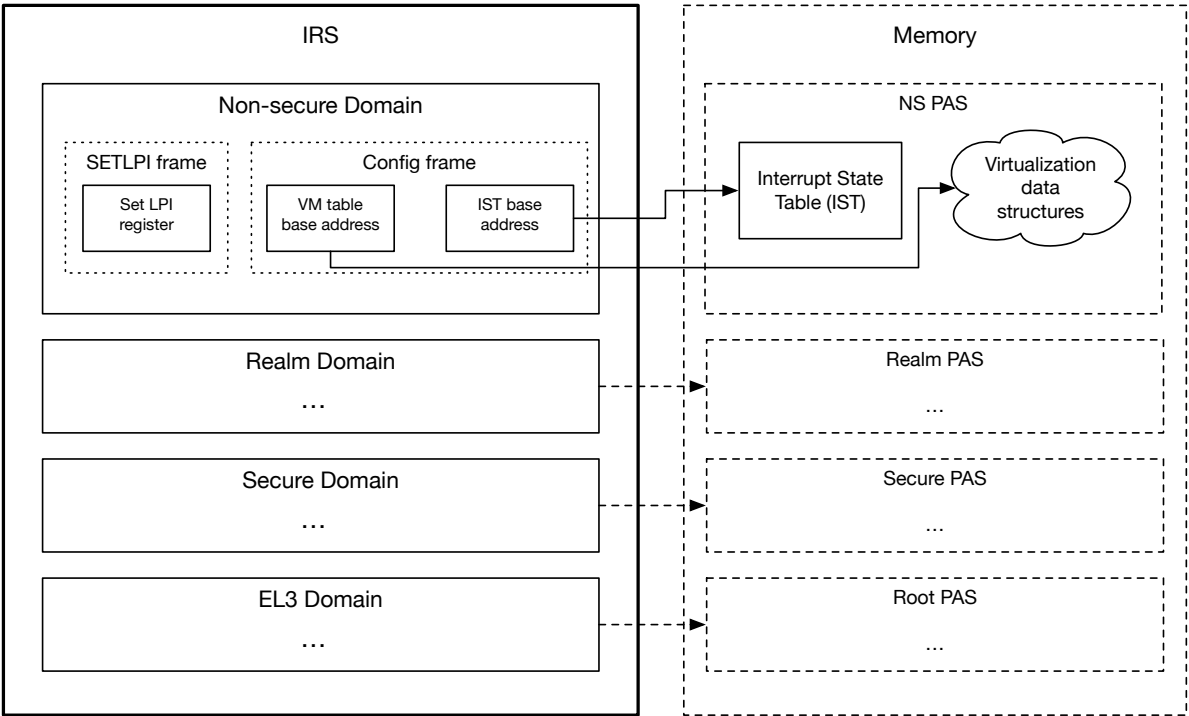


Figure 4.1: IRS Domains in a single-IRS system.

- R_{GSNPJ}

Each IRS implements a separate IRS Domain for each Security state supported by the PEs in the system. Every PE connected to an IRS implements the same set of Security states.
- I_{KZGKR}

An IRS Domain is associated with the PAS of its corresponding Interrupt Domain.
- R_{XXWDT}

An IRS exposes separate IRS configuration register frames for each IRS Domain.
- I_{CKHSF}

The IRS Domain for an IRS configuration register frame is reported by IRS_IDR0.INT_DOM.
- I_{TCTJK}

A VM is associated with exactly one IRS Domain which defines the Security state of the VM and its virtual interrupts.
- Similarly, the VPEs in the VM are associated with the IRS Domain and Security state of the VM.
- For example, a VPE in a Non-secure VM receives Non-secure virtual interrupts, and a VPE in a Realm VM receives Realm virtual interrupts.

R_{NKVFJ}

Each non-EL3 Physical Interrupt Domain has a separate namespace for Virtual Machines (VMs).

R_{MHKVK}

VMs are supported for the Secure, Realm, or Non-secure Interrupt Domains. VMs are not supported in the EL3 Interrupt Domain.

See also:

- [3.1 Interrupt Domains](#)

4.4 IRS Configuration

R_{JJPMF}

For each Interrupt Domain, each IRS exposes the following register frames:

- **IRS_CONFIG_FRAME**: An IRS configuration frame for settings that are common to all PEs connected to the IRS.
- **IRS_SETLPI_FRAME**: If **IRS_IDR0.SETLPI** is 1, a SETLPI register frame.

For each Interrupt Domain for each IRS, the IRS configuration register frame occupies a 64KB region in the system address space.

For each Interrupt Domain for each IRS, if **IRS_IDR0.SETLPI** is 1, the IRS SETLPI register frame occupies a 64KB region in the system address space.

Each register frame is 64KB-aligned.

See [10.2 IRS register frames](#) for more information.

I_{GVFTM}

There is at most one SETLPI register frame for each IRS for an Interrupt Domain.

I_{MZDWR}

The SETLPI register frame is not implemented when support for physical LPIs is not implemented.

S_{MTWNL}

The base addresses of the IRS configuration frame and the IRS SETLPI register frame, if implemented, for each IRS and Interrupt Domain, are provided to software by firmware.

R_{QHDWM}

The value of **IRS_IDR0.SETLPI** is the same for all IRSs and all Interrupt Domains in the system.

I_{NCXKF}

Each IRS has a unique **IRSID** which is reported in **IRS_IDR0.IRSID**. The **IRSID** is unique for each IRS in the system. The **IRSID** reported in **IRS_IDR0.IRSID** is the same for each **IRS_CONFIG_FRAME** that belongs to the same IRS.

I_{NRPRB}

The number of PEs connected to an IRS is reported by **IRS_IDR1.PE_CNT**.

R_{FRRGD}

Every PE is assigned a unique PE interrupt Affinity ID across all IRSs in the system.

R_{GSKQZ}

The PE interrupt Affinity ID is the same across all implemented Interrupt Domains.

R_{YBQML}

If a system is reset with the same configuration, each PE has the same interrupt Affinity as before the reset.

I_{XDFPV}

If a system is reset and reconfigured in a way that presents a different system, for example by physically partitioning IRSs and PEs in a different topology after the system reset, each PE may be assigned a new interrupt Affinity ID in the new topology.

I_{CWCZR}

The PE interrupt Affinity ID is reported by the PE in **ICC_IAFFIDR_EL1**.

I_{KZCQC}

The PE interrupt Affinity ID is used to configure the target PE for Targeted interrupts.

It can also be used to discover which IRS a PE is connected to and vice versa.

I_{LXNBN}

Information about each PE connected to an IRS may be accessed by writing to **IRS_PE_SELR** to select a PE and accessing the following registers:

- **IRS_PE_STATUSR**.
- **IRS_PE_CR0**.

IRS_PE_STATUSR.IDLE reports whether the effects of a write to any of the above registers, including **IRS_PE_SELR**, have completed.

IRS_PE_STATUSR.V reports whether the value written to **IRS_PE_SELR** selects a valid PE connected to the IRS.

IRS_PE_CR0 provides access to configuration settings for the selected PE.

S_{HHMP}

To configure a PE connected to an IRS, software performs the following sequence:

1. Software selects the PE by writing to **IRS_PE_SELR**.
2. Software polls **IRS_PE_STATUSR.IDLE** until it reads as 1 and ensures that **IRS_PE_STATUSR.V** is 1.
3. Software can access the configuration of a PE by accessing **IRS_PE_CR0**.

4. If software wrote to the PE configuration register, software polls IRS_PE_STATUSR.IDLE until it reads as 1 to ensure the effects of the write are complete.

R_{KCZYL} When an IRS is reset, each PE connected to the IRS is reset as follows:

- If the IRS supports 1ofN interrupt routing, 1ofN PE selection is enabled for the PE.

I_{PPXSP} The configuration of each PE at IRS reset is consistent with a read of IRS_PE_CR0.DPS returning 0 when the PE is selected in IRS_PE_SEL and there was no write to IRS_PE_CR0 since the reset.

I_{MCXJN} Arm expects that firmware describes the association of PEs to an IRS to system software using firmware data structures.

The architecture also provides a mechanism for software to discover the association of PEs to an IRS which may be useful for system bring-up and debugging.

S_{BCBQP} Software can discover the association of IRSs and PEs by following this sequence:

1. When a PE is powered on, software reads ICC_IAFFIDR_EL1 to obtain the Interrupt Affinity ID for that PE.
2. For each IRS, software writes the Interrupt Affinity ID for that PE to IRS_PE_SEL and performs the following actions:
 1. Poll IRS_PE_STATUSR.IDLE until it reads as 1.
 2. Read IRS_PE_STATUSR.V. If the value returned is 1, the IAFFID specifies a valid PE connected to that IRS.
 3. Otherwise, continue to the next IRS.

I_{HTRN} The number of supported IAFFID bits for physical interrupts is reported in IRS_IDR1.IAFFID_BITS.

R_{GCGDD} The number of supported IAFFID bits for physical interrupts is the same across all IRSs and Interrupt Domains.

I_{ZFSNT} The number of supported IAFFID bits is sufficient to uniquely identify all physical PEs in the system.

I_{KVXQW} The number of supported VPE ID bits for virtual interrupts is reported in IRS_IDR4.VPE_ID_BITS.

R_{NPTWH} When [processing](#) an Interrupt Write Effect generated by a GIC System instruction that updates the Affinity of an interrupt, all of the following are true:

- If the interrupt is a physical interrupt, unimplemented upper bits of IAFFID above IRS_IDR1.IAFFID_BITS are RES0.
- If the interrupt is a virtual interrupt, unimplemented upper bits of IAFFID above IRS_IDR4.VPE_ID_BITS are RES0.

See [4.9.2 The VPE table](#) for more information about the relationship between the IAFFID of a virtual interrupt and the VPE ID.

R_{QQCXG} When IRS_CR0.IRSEN is 1 and IRS_CR0.IDLE is 1, an IRS is *enabled* for the Interrupt Domain and the IRS selects a candidate HPPI for each connected PE for the Interrupt Domain, following the rules in this chapter.

Otherwise, the IRS is *disabled* for the Interrupt Domain and the IRS does not select any candidate HPPI for the connected PEs for the Interrupt Domain.

I_{ZYYLL} An IRS can access data structures to which it has a valid pointer, irrespective of whether the IRS is enabled for the corresponding Interrupt Domain.

S_{SMLQY} To quiesce IRS accesses to IRS data structure, software makes the physical IST and VM table invalid, and polls the relevant status bits to ensure that all IRS memory accesses to these data structures have completed.

R_{GQZNP} An IRS [processes](#) incoming interrupt events irrespective of the value of IRS_CR0.IRSEN.

R_{LKCZP} An IRS [processes](#) Interrupt Effects generated by GIC System instructions, irrespective of the value of IRS_CR0.IRSEN.

I_{TQMVQ}

Interrupts may need to be delivered to an IRS before the IRS has been enabled.

For example, an edge-triggered SPI may cause wake-up of an IRS that is in a low power mode and the signal should result in the Pending state of the SPI being updated as soon as the IRS is powered on.

See [4.11 IRS power management](#) for more information.

See also:

- [2.6.1 LPI and SPI configuration](#)
- [4.4.1 Enabling and disabling the IRS](#)
- [4.6 Interrupt configuration and state](#)
- [4.7 The interrupt state table \(IST\)](#)
- [4.8 Physical interrupts](#)
- [4.9 Virtualization data structures](#)

4.4.1 Enabling and disabling the IRS

R_{JHYGJ}

When a write to IRS_CR0.IRSEN changes the value from 0 to 1, the IRS begins a transition from disabled to enabled for the Interrupt Domain.

The transition from disabled to enabled is complete when IRS_CR0.IDLE is 1.

The transition from disabled to enabled completes in finite time.

Enabling an IRS does not affect the state and configuration of LPIs.

R_{SBXYL}

When a write to IRS_CR0.IRSEN changes the value from 1 to 0, the IRS begins a transition from enabled to disabled for the Interrupt Domain.

The transition from enabled to disabled is complete when IRS_CR0.IDLE is 1.

The transition from enabled to disabled completes in finite time.

When the transition from enabled to disabled is complete, for each candidate HPPIs that were previously selected by the IRS for connected PEs for the Interrupt Domain while the IRS was enabled, one of the following is true:

- The candidate HPPI is acknowledged by the PE before the IRS is disabled.
- The candidate HPPI is not acknowledged and is no longer selected as the candidate HPPI for the PE.

I_{ZMTMZ}

The architecture allows enabling an IRS for an Interrupt Domain either before or after software provisions a valid physical IST. See [4.8.1 Physical LPIs](#) for more information.

See also:

- [4.7 The interrupt state table \(IST\)](#)
- [4.9 Virtualization data structures](#)

4.5 IRS synchronization requests

R_{HDTJ}

Software can request synchronization of interrupt events for the IRS Domain by writing 1 to IRS_SYNCR.SYNC.

Writing 1 to IRS_SYNCR.SYNC requests synchronization of interrupt events and ensures that the following events for the IRS Domain are [processed](#):

- All interrupt events generated as a result of an SPI being asserted or de-asserted before the write to IRS_SYNCR.
- All interrupt events Accepted by the IRS before the write to IRS_SYNCR.
- All interrupt events generated as a result of a write to IRS_SETLPI before the write to IRS_SYNCR.
- All VPE doorbell events generated before the write to IRS_SYNCR.

Following a write to IRS_SYNCR.SYNC that requests synchronization, when IRS_SYNC_STATUSR.IDLE is 1, all of the following are true:

- Any Interrupt Write Effect E_1 to an Interrupt Location from a processed event is Ordered-before an Interrupt Read Effect E_2 , when all of the following are true:
 - There is a Memory Read Effect E_3 to IRS_SYNC_STATUSR that reads IRS_SYNC_STATUSR.IDLE is 0.
 - E_1 is Ordered-before E_3 .
 - E_3 is Ordered-before E_2 .

See [2.12 Interrupt ordering model and synchronization requirements](#) for more information about observability of Interrupt Effects by PEs.

Note

An interrupt that meets the candidate HPPI conditions after an interrupt event has been processed is not required to be acknowledged by an interrupt acknowledge instruction executed on a PE following an IRS synchronization request. This is because selecting the candidate HPPI is only required to happen in finite time. See [4.8.4 Physical interrupt signaling](#) and [4.10.4 Virtual interrupt signaling](#) for more information.

I_{MJPSN}

An interrupt event generates the following Interrupt Write Effects:

- If the interrupt event type is SET_EDGE or SET_LEVEL, an Interrupt Write Effect E_1 that updates the Pending state and an Interrupt Write Effect E_2 that updates the Handling mode.
- If the interrupt event type is SET or CLEAR, an Interrupt Write Effect that updates the Pending state.

S_{BBHRQ}

Software can use synchronization of events to ensure that interrupt events have been [processed](#) before repurposing an interrupt.

For example, if software disables a PCIe function, it can ensure that no additional Interrupt Effects originate from that function by performing the following sequence:

1. Ensure that all MSIs have completed from the PCIe function.
2. Issue a synchronization request to the ITS and poll for completion.
3. Issue a synchronization request to the IRS and poll for completion.

After this sequence, software can remove any old translations in the ITS and create new translations to the same INTID, knowing that Interrupt Effects to the INTID can only come from the new translations.

I_{MXYLQ}

In a multi-IRS system, an IRS synchronization request applies to all IRSs in the system for the IRS Domain.

R_{ZPNFL}

If a write to IRS_SYNCR occurs when IRS_SYNC_STATUSR.IDLE is 0 on any other IRS for the same IRS Domain, it is CONSTRAINED UNPREDICTABLE whether the write synchronizes any interrupt events.

See also:

- [3.2 Communication between GIC system components](#)

Chapter 4. Interrupt routing service (IRS)

4.5. IRS synchronization requests

- [*Chapter 5 Interrupt translation service \(ITS\)*](#)
- [*5.2.4 ITS synchronization requests*](#)

4.6 Interrupt configuration and state

I _{JVVTZ}	<p>The IRS manages the following states and configurations for each LPI and SPI:</p> <ul style="list-style-type: none"> • Pending. • Active. • Enabled. • Handling mode. • Priority. • Routing mode. • Affinity.
I _{BWPPP}	<p>The IRS further manages the following for each SPI:</p> <ul style="list-style-type: none"> • Trigger mode. • Interrupt Domain assignment. • VM assignment.
D _{HJCLV}	<p>An INTID specifying an LPI or SPI is <i>reachable</i> if the INTID.ID is within the configured range for the interrupt type and if the IRS can access its configuration and state. Otherwise, the INTID is <i>unreachable</i>.</p>
R _{PGDXN}	<p>When the PE executes any of the following GIC System instructions to request the configuration of an unreachable INTID, the IRS indicates a failure in the response to the PE:</p> <ul style="list-style-type: none"> • GIC CDRCFG. • GIC VDRCFG. • GIC LDRCFG. <p>When the PE executes any other GIC System instruction that accesses the state and configuration of an unreachable INTID, there are no Interrupt Effects generated as a result of executing the instruction.</p>
I _{SRPZD}	<p>The maximum number of INTID.ID bits supported by an IRS is reported by IRS_IDR2.ID_BITS.</p>
R _{SVDYQ}	<p>The value of IRS_IDR2.ID_BITS is the same for all IRSs for all Interrupt Domains in the system.</p>
I _{YKSNC}	<p>Arm recommends that an IRS supports a number of INTID.ID bits greater than or equal to the number of INTID.ID bits supported by the PEs connected to the IRS.</p>
R _{BBZWX}	<p>If an IRS supports fewer INTID.ID bits than a connected PE and the PE that specifies an INTID.ID beyond the INTID.ID range supported by the IRS, the INTID specifies an unreachable interrupt.</p> <p>See also 2.6.1 LPI and SPI configuration.</p>
I _{GTVLX}	<p>Support for physical LPIs is optional. Whether the IRS supports physical LPIs is reported in IRS_IDR2.LPI.</p>
I _{BNDTT}	<p>Arm recommends that IRSs in a system designed to run standard operating systems implements support for physical LPIs. If physical LPIs are not implemented, sufficient physical SPIs that are not connected to any interrupt sources must be implemented to be used for IPIs.</p>
I _{YKQDJ}	<p>Arm expects that support for LPIs is required for GICv5 systems in a future version of the <i>Arm® Base System Architecture 1.0C</i>[2].</p>
R _{MGQGH}	<p>When the IRSs do not implement support for LPIs, there are no ITSs or IWBs in the system.</p>
R _{RCGPT}	<p>The value of IRS_IDR2.LPI is the same for all IRSs for all Interrupt Domains in the system.</p>
I _{ZKLJR}	<p>The minimum number of LPI INTID.ID bits supported by an IRS is reported in IRS_IDR2.MIN_LPI_ID_BITS.</p> <p>Arm expects that most hardware implementations will not require a minimum number of LPI INTID.ID bits. However, specifying a minimum number of LPI INTID.ID bits is useful for virtualization.</p>
R _{YHQLW}	<p>The value of IRS_IDR2.MIN_LPI_ID_BITS is the same for all IRSs for all Interrupt Domains in the system.</p>

<code>I_PQBVVF</code>	<p>The definition of reachable and unreachable interrupts applies to both physical and virtual interrupts.</p> <p>The conditions for when a physical interrupt is reachable differ from when a virtual interrupt is reachable.</p> <p>See 4.8.1 Physical LPIs and 4.8.2 Physical SPIs for more information about when physical interrupts are reachable.</p> <p>See 4.10 Virtual interrupts for more information about when virtual interrupts are reachable.</p>
<code>R_ZMRVR</code>	<p>The Handling mode property of an interrupt controls whether the interrupt should be handled using edge-triggered or level-sensitive semantics as follows:</p> <ul style="list-style-type: none"> • When the interrupt Handling mode is Edge, the IRS clears the Pending state when the interrupt is acknowledged by the PE. • When the interrupt Handling mode is Level, the IRS does not clear the Pending state when the interrupt is acknowledged by the PE.
<code>I_XBVHV</code>	The Priority property stores the Priority of an Interrupt.
<code>I_XWHLD</code>	<p>The number of implemented Priority bits in the IRS is reported by <code>IRS_IDR1.PRI_BITS</code>, and the minimum is 1 bit.</p> <p>The number of Priority bits supported may be different across Interrupt Domains.</p>
<code>I_FRYMQ</code>	Some operating systems have requirements for a minimum number of interrupt priority levels. These requirements will be captured as part of the <i>Arm® Base System Architecture 1.0C</i> [2] .
<code>R_TZYHS</code>	For each Interrupt Domain, the number of supported Priority bits for the Interrupt Domain is the same across all IRSs in the system.

4.7 The interrupt state table (IST)

I _{LQMKQ}	An <i>Interrupt State Table</i> (IST) is used by the IRS to store interrupt state and configuration.
D _{PDQBB}	<p>An IST uses either a linear or 2-level structure.</p> <p>For a linear IST, all of the following are true:</p> <ul style="list-style-type: none">• The level 1 IST is omitted.• The IST consists of a single level 2 IST.
D _{XTNLJ}	<p>An IST is used to manage the state and configuration of interrupts in each of the following <i>IST contexts</i>:</p> <ul style="list-style-type: none">• Physical LPI IST context: Used to manage physical LPIs for each Interrupt Domain. See 4.8.1 Physical LPIs for more information.• Virtual LPI IST context: Used to manage virtual LPIs for VMs. See 4.10.1 Virtual LPIs for more information.• Virtual SPI IST context: Used to manage virtual SPIs for VMs. See 4.10.2 Virtual SPIs for more information.
D _{VBJJZ}	<p>For a linear IST, the base address of the IST specifies the location of the level 2 IST.</p> <p>For a 2-level IST, the base address of the IST specifies the location of the level 1 IST.</p> <p>Figure 4.2 shows an overview of the linear and 2-level structure.</p>

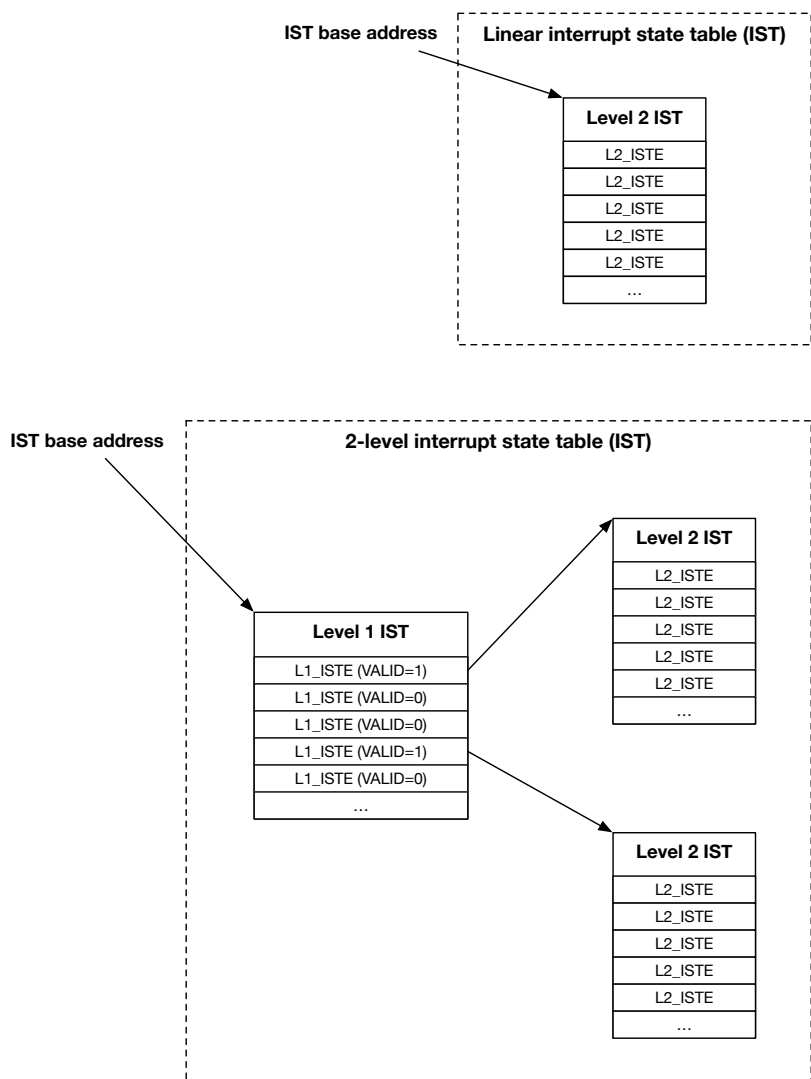


Figure 4.2: IST structure overview.

<code>S_{NWLKP}</code>	Software is responsible for allocating the IST from memory in the PAS associated with the Interrupt Domain where the IST is used.
<code>I_{MWLZJ}</code>	An IST is indexed using the <code>INTID.ID</code> for an interrupt within the configured range for the IST.
<code>D_{BXYVH}</code>	An IST is either considered <i>valid</i> or <i>invalid</i> . The mechanism that determines whether the IST is valid or invalid depends on the IST context used. When the IST transitions from being invalid to being valid, the IST <i>becomes valid</i> . When the IST transitions from being valid to being invalid, the IST <i>becomes invalid</i> .
<code>I_{VVKXF}</code>	An IST is either valid or invalid across all IRSs.
<code>R_{SCLJY}</code>	When an IST is valid, the IRS Domain is permitted to access the IST for any reason, including speculative reads. When an IST is invalid, the IRS Domain does not access the IST.
<code>D_{XQVBD}</code>	An IST stores a number of INTIDs for a single <code>INTID.Type</code> . The mechanism that determines the number of INTIDs stored in an IST depends on the IST context used.

I _{XTPFF}	An IST is used to manage the state and configuration for each INTID in the configured interrupt range for the IST context used.
I _{NQBVH}	<p>The configuration of an IST is controlled by the following parameters:</p> <p>STRUCTURE Selects if the IST uses a linear or 2-level structure.</p> <p>ID_BITS Selects how many INTIDs are stored in the IST.</p> <p>The number of INTIDs stored in the IST determines the size of the single level 2 IST when using a linear table and the size of the level 1 IST when using a 2-level table.</p> <p>The number of INTIDs stored in the IST also limits the supported size of each level 2 IST entry when the IRS stores metadata in the IST.</p> <p>L2SZ For a 2-level IST, selects the size of each level 2 IST. The number of INTIDs stored in each level 2 IST depends on the size of a level 2 IST entry.</p> <p>Arm recommends a linear IST is used when the total number of INTIDs can be stored in a single level 2 IST.</p> <p>Level 2 IST entry size Selects the size of each level 2 IST entry.</p> <p>The supported entry sizes depend on the number of INTIDs stored in the IST as well as the IRS requirements for storing metadata in the IST. See 4.7.4 IST metadata for more information.</p> <p>These parameters are programmed in a way that is specific to the IST context used. For the physical IST context, each parameter is programmed in the physical IST base and configuration registers. For the virtual IST context, each parameter is programmed in the VM table entry.</p> <p>See 4.8.1 Physical LPIs, 4.10.1 Virtual LPIs, and 4.10.2 Virtual SPIs for more information.</p>
D _{TWFKY}	<p>The size of the IST is calculated based on the number of INTIDs and the size of each entry, including any metadata if implemented, as follows:</p> <ul style="list-style-type: none"> For a linear IST, the size of the level 2 IST is (Number of INTIDs) * (Level 2 IST entry size). For a 2-level IST, the size of the level 1 IST is ((Number of INTIDs) / (Level 2 IST size / Level 2 IST entry size)) * (Level 1 IST entry size).
I _{QBVSN}	IRS_IDR2.IST_LEVELS reports whether 2-level IST support is implemented.
R _{ZGTCN}	The value of IRS_IDR2.IST_LEVELS is the same for all IRSs for all Interrupt Domains in the system.
I _{NCSJH}	For a 2-level IST, an implementation may only implement support for a subset of the possible L2SZ values. The supported L2SZ values are reported by IRS_IDR2.IST_L2SZ.
R _{HVL TJ}	The value of IRS_IDR2.IST_L2SZ is the same for all IRSs for all Interrupt Domains in the system.
I _{PSTQD}	<p>Arm strongly recommends that an implementation implements the L2SZ values corresponding to the stage 1 and stage 2 translation granule sizes in the PEs in the same system.</p> <p>See also:</p> <ul style="list-style-type: none"> 4.8.1 Physical LPIs 4.10.1 Virtual LPIs 4.10.2 Virtual SPIs

4.7.1 Level 2 IST management

D _{HMZP}	<p>A level 2 IST is considered <i>valid</i>, if the IST is valid and one of the following is true:</p> <ul style="list-style-type: none"> The IST uses a linear structure. The IST uses a 2-level structure and the corresponding L1_ISTE.VALID is 1. <p>Otherwise, the level 2 IST is considered <i>invalid</i>.</p>
-------------------	---

D _{LWRCC}	<p>A level 2 IST <i>becomes valid</i> in any of the following scenarios:</p> <ul style="list-style-type: none"> • The IST uses a linear structure and the IST becomes valid. • The IST uses a 2-level structure and any of the following occur: <ul style="list-style-type: none"> – A write to the level 2 IST map register for the IST context makes the level 2 IST valid. – The IST becomes valid and the corresponding L1_ISTE.VALID is 1.
I _{XLXLR}	<p>The level 2 IST map register to use depends on the IST context used.</p> <p>See 4.8.1 Physical LPIs, 4.10.1 Virtual LPIs, and 4.10.2 Virtual SPIs for more information.</p>
I _{ZDJBW}	<p>When a 2-level IST becomes valid, if L1_ISTE.VALID is 1 for more than one level 1 IST entry, more than one level 2 IST may become valid at the same time.</p>
R _{CWRMM}	<p>When the IST is valid and uses a 2-level structure, if software writes to an invalid level 1 IST entry and sets L1_ISTE.VALID to 1, the behavior of the IRS Domain is UNPREDICTABLE.</p> <p>The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the Interrupt Domain or in an access to memory that does not follow the behaviors described in 4.12 IRS memory access rules.</p>
R _{LFYMV}	<p>If software performs a write to any field in a valid level 1 IST entry, the IRS Domain behavior is UNPREDICTABLE.</p> <p>The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the Interrupt Domain or in an access to memory that does not follow the behaviors described in 4.12 IRS memory access rules.</p>
R _{DDPXH}	<p>To recover from UNPREDICTABLE behavior resulting from an incorrect write to a level 1 IST entry, the IST is made invalid.</p>
I _{FSYTK}	<p>The architecture does not support making an individual level 2 IST invalid when using a 2-level structure for an IST.</p>
R _{FKYJV}	<p>When a write to the level 2 IST map register for the IST context makes the corresponding level 2 IST valid, all of the following are true:</p> <ul style="list-style-type: none"> • The IRS generates the following Effects: <ul style="list-style-type: none"> – A Register Write Effect E₁ to the status register for the IST context which sets the IDLE field to 0. – A Memory Write Effect E₂ to the Location corresponding to the level 1 IST entry which updates L1_ISTE.VALID from 0 to 1. – A Register Write Effect E₃ to the status register for the IST context which sets the IDLE field to 1. • There is an Intrinsic order dependency from the Register Write Effect E₁ to the Memory Write Effect E₂. • There is an Intrinsic order dependency from the Memory Write Effect E₂ to the Register Write Effect E₃. <p>This means that if there is a Memory Read Effect E₄ which Reads-from the Register Write Effect E₃, then E₂ is Ordered-before E₄.</p>
I _{LNDTY}	<p>When the IST is valid and uses a 2-level structure, a level 2 IST is made valid by the following sequence:</p> <ol style="list-style-type: none"> 1. Software writes the address of the new level 2 IST to L1_ISTE.L2_ADDR, with L1_ISTE.VALID remaining to be 0. 2. Software writes to the level 2 IST map register for the IST context. 3. The IRS performs a write to the level 1 IST entry that updates L1_ISTE.VALID from 0 to 1. 4. The IRS reports that the effects of the write to the level 2 IST map register for the IST context are complete. <p>See 3.3 Coherency considerations for GIC data structures for more information about ensuring that memory accesses are coherent between the IRS and a PE.</p>
R _{SDVFP}	<p>When the IST is valid and uses a 2-level structure, if software writes to the level 1 IST entry when the status register IDLE field is 0 for the IST context, it is CONSTRAINED UNPREDICTABLE whether the IRS writes back the entry using the updated or old level 1 IST entry value.</p> <p>The CONSTRAINED UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the Interrupt Domain or in an access to memory that does not follow the behaviors described in 4.12 IRS memory access rules.</p>

D_{XCRJV} A level 2 IST *becomes invalid* when the IST becomes invalid.

4.7.2 Initialization of level 2 IST entries

R_{LXJHZ} When a level 2 IST becomes valid, the behavior is UNPREDICTABLE if all of the following are true:

- IRS_IDR2.ISTMD is 1.
- The metadata area of one or more level 2 IST entries are non-zero. See [4.7.4 IST metadata](#) for more information about the metadata area.

The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the Interrupt Domain or in an access to memory that does not follow the behaviors described in [4.12 IRS memory access rules](#).

R_{XJJYH} To recover from UNPREDICTABLE behavior resulting from a level 2 IST containing non-zero metadata becoming valid, the IST is made invalid.

I_{BNPSN} A level 2 IST with non-zero metadata may become valid as part of an IMPLEMENTATION DEFINED IRS restore mechanism. See [4.11 IRS power management](#) and [4.7.4 IST metadata](#) for more information.

R_{ZFBKH} When a level 2 IST becomes valid and L2_ISTE.HWU is not 0b00 for one or more level 2 IST entries, the behavior is UNPREDICTABLE.

The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the Interrupt Domain or in an access to memory that does not follow the behaviors described in [4.12 IRS memory access rules](#).

R_{XZWXJX} To recover from UNPREDICTABLE behavior resulting from a level 2 IST containing entries with a non-zero HWU field becoming valid, the IST is made invalid.

I_{GXNWX} A level 2 IST with non-zero HWU values may become valid as part of an IMPLEMENTATION DEFINED IRS restore mechanism. See [4.11 IRS power management](#) for more information.

R_{GXLJR} When a level 2 IST becomes valid and L2_ISTE.Pending is 1 for one or more level 2 IST entries, it is CONSTRAINED UNPREDICTABLE whether the INTID is selected as the candidate HPPI even though it meets the candidate HPPI conditions.

I_{CLSGT} If a level 2 IST containing entries with a non-zero Pending field becomes valid, once the corresponding interrupts are made Pending as a result of an interrupt being signaled, the IRS considers the interrupt when selecting a candidate HPPI.

I_{XBHGH} See [4.8.4 Physical interrupt signaling](#) and [4.10.4 Virtual interrupt signaling](#) for more information about the candidate HPPI conditions.

I_{HMFBB} A level 2 IST containing Pending interrupts may become valid as part of an IMPLEMENTATION DEFINED IRS restore mechanism. See [4.11 IRS power management](#) for more information.

R_{PFQHM} When a level 2 IST becomes valid, the behavior is UNPREDICTABLE if all of the following are true for one or more level 2 IST entries:

- L2_ISTE.IRM is 1.
- L2_ISTE.IAFFID is not 0.

The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the Interrupt Domain or in an access to memory that does not follow the behaviors described in [4.12 IRS memory access rules](#).

R_{VZCQM} To recover from UNPREDICTABLE behavior resulting from a level 2 IST containing entries with a non-zero 1ofN interrupt hint becoming valid, the IST is made invalid.

I_{KJLYK} A level 2 IST containing non-zero 1ofN interrupt hints may become valid as part of an IMPLEMENTATION DEFINED IRS restore mechanism. See [4.11 IRS power management](#) for more information.

4.7.3 INTID state and configuration

D _{PTGRK}	For each INTID in the configured range for an IST, there is a corresponding level 2 IST entry, if the IST is valid and one of the following is true: <ul style="list-style-type: none"> • The IST uses a linear structure. • The IST uses a 2-level structure and the level 2 IST is valid. Otherwise, there is no level 2 IST entry corresponding to the INTID.
R _{RFGXL}	When a level 2 IST becomes valid, the configuration of each INTID is consistent with the values stored in the corresponding IST entry.
R _{RFXND}	When the IST is valid, the state and configuration of each INTID in the configured range for the IST are coherent across all IRSs.
R _{NQXGQ}	After a level 2 IST becomes valid, the values stored in the level 2 IST entries are UNKNOWN. The values stored in the level 2 IST entries remain UNKNOWN after the IST becomes invalid.
R _{BTYMB}	If software performs a write to a level 2 IST entry when the IST is valid, the behavior is UNPREDICTABLE. The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the Interrupt Domain or in an access to memory that does not follow the behaviors described in 4.12 IRS memory access rules .
R _{XYLQ}	To recover from UNPREDICTABLE behavior resulting from an incorrect write to a level 2 IST entry, the IST is made invalid.

4.7.4 IST metadata

I _{RCFTT}	IRS_IDR2.ISTMD reports whether an IRS stores metadata in the level 2 IST entries.
D _{JKSZH}	If an IRS stores metadata in the level 2 IST entries, the bytes containing the metadata is called the <i>metadata area</i> .
R _{WDNFJ}	The value of IRS_IDR2.ISTMD is the same for all IRSs for each Interrupt Domain in the system.
I _{HPJWB}	The size of a level 2 IST entry is 4 bytes, 8 bytes, or 16 bytes. If the size of a level 2 IST entry is more than 4 bytes, the higher address bytes are RES0. The layout of the 4 lowest address bytes is defined by the architecture in 11.2.4 L2_ISTE, Level 2 interrupt state table entry .
I _{CMFVK}	The metadata area is stored in the highest address bytes of the IST entry.
I _{QRFDQ}	IRS_IDR2.ISTMD_SZ describes the minimum number of INTID.ID bits which requires a level 2 IST entry size of 16 bytes to store metadata.
R _{THDPW}	The value of IRS_IDR2.ISTMD_SZ is the same for all IRSs for all Interrupt Domains in the system.
I _{FFWQC}	The minimum size of the level 2 IST entries used for the physical IST depends on the whether the IRS stores metadata in the level 2 IST entries and on the number of LPis stored in the IST. IRS_IDR2.ISTMD reports whether the IRS stores metadata in the level 2 IST entries. The minimum size of the level 2 IST entries is 4 bytes, 8 bytes, or 16 bytes as follows: <ul style="list-style-type: none"> • If the IRS does not store metadata in the IST, the minimum size of the level 2 IST entries is 4 bytes. • If the IRS stores metadata in the IST and the number of INTID.ID bits is less than the value in IRS_IDR2.ISTMD_SZ, the minimum size of the level 2 IST entries is 8 bytes. • If the IRS stores metadata in the IST and the number of INTID.ID bits is greater than or equal to IRS_IDR2.ISTMD_SZ, the minimum size of the level 2 IST entries is 16 bytes.

4.7.5 Example IST structures

I_{SSCKN} The size of a level 1 IST entry is 8 bytes and the size of a level 2 IST entry may scale with the number of INTIDs stored in the IST.

[Table 4.1](#) shows some example IST structures for an implementation that does not require storing metadata in the IST and therefore the size of the level 2 IST entry does not scale with the configured number of INTIDs.

[Table 4.2](#) shows some example IST structures for an implementation that requires storing metadata in the IST. The examples assume that the size of a level 2 IST entry is 8 bytes up to and including 16 bits of LPI ID, and 16 bytes for more than 16 bits of LPI ID.

S_{BRMDG} The L1 size and L2 size columns in [Table 4.1](#) and [Table 4.2](#) indicate the required maximum physically contiguous allocation by software for the configuration.

Table 4.1: Example IST structures without metadata storage

STRUCTURE	LPI_ID_BITS	L2_ISTE size	L2SZ	L1 size	L2 size	Maximum number of LPIs
Linear	10	4 bytes	-	-	4 KB	1024
Linear	14	4 bytes	-	-	64 KB	16,384
Linear	16	4 bytes	-	-	256 KB	65,536
2-level	19	4 bytes	0b00, 10 bits	4 KB	4 KB	524,288
2-level	24	4 bytes	0b00, 10 bits	128 KB	4 KB	16,777,216
2-level	16	4 bytes	0b01, 12 bits	128 B	16 KB	65,536
2-level	24	4 bytes	0b01, 12 bits	32 KB	16 KB	16,777,216
2-level	16	4 bytes	0b10, 14 bits	32 B	64 KB	65,536
2-level	24	4 bytes	0b10, 14 bits	8 KB	64 KB	16,777,216

Table 4.2: Example IST structures with metadata storage

STRUCTURE	LPI_ID_BITS	L2_ISTE size	L2SZ	L1 size	L2 size	Maximum number of LPIs
Linear	9	8 bytes	-	-	4 KB	512
Linear	13	8 bytes	-	-	64 KB	8,192
Linear	16	8 bytes	-	-	512 KB	65,536
2-level	16	8 bytes	0b00, 9 bits	1 KB	4 KB	65,536
2-level	17	16 bytes	0b00, 8 bits	4 KB	4 KB	131,072
2-level	24	16 bytes	0b00, 8 bits	512 KB	4 KB	16,777,216
2-level	16	8 bytes	0b01, 11 bits	256 B	16 KB	65,536
2-level	24	16 bytes	0b01, 10 bits	128 KB	16 KB	16,777,216
2-level	16	8 bytes	0b10, 13 bits	64 B	64 KB	65,536
2-level	24	16 bytes	0b10, 12 bits	32 KB	64 KB	16,777,216

I_{LWNEZ} For a 2-level IST, an implementation may only implement support for a subset of the possible L2SZ values. The supported L2SZ values are reported by IRS_IDR2.IST_L2SZ.

R_{FVVBH} The value of IRS_IDR2.IST_L2SZ is the same for all IRSs for all Interrupt Domains in the system.

I_{MCFNW} IRS_IDR2.{MIN_LPI_ID_BITS,ID_BITS} report the range of valid values for IRS_IST_CFGR.LPI_ID_BITS.

See also:

- [11.2 IRS Data Structures](#)

4.8 Physical interrupts

D _{RVLFH}	Physical interrupts are interrupts that belong to a Physical Interrupt Domain.
R _{ZQGC}	A physical interrupt is only permitted to be selected as the candidate HPPI for the corresponding Physical Interrupt Domain.
I _{NYGDM}	For a description of how the IRS supports virtualization, see 4.10 Virtual interrupts .

4.8.1 Physical LPis

R _{VWVB}	The physical IST contains the configuration and state for the physical LPis in an IRS Domain.
I _{BHYV}	When support for physical LPis is not implemented, there is no physical IST. IRS_IDR2.LPI reports whether support for physical LPis is implemented. See 4.6 Interrupt configuration and state for more information.
I _{XZPJ}	The number of LPis in an IRS Domain is configured in IRS_IST_CFGR.LPI_ID_BITS. Software allocates the physical IST for an IRS Domain by writing a valid address to IRS_IST_BASER.ADDR and writing 1 to IRS_IST_BASER.VALID to make the IST valid.
R _{QXDB}	The physical IST is indexed using an INTID.ID value in the physical LPI range, from 0 to $(2^{\text{IRS_IST_CFGR.LPI_ID_BITS}} - 1)$.
R _{TXDP}	When an IRS accesses the physical IST, the IRS does not access any memory location derived from an INTID which is outside the configured LPI range. This includes situations where the INTID is programmed in the parameters of a GIC System instruction executed on a PE or the INTID is stored in the IST metadata areas.
R _{LFZVT}	An INTID specifying an LPI in the Physical Interrupt Domain is reachable, if all of the following are true: <ul style="list-style-type: none"> • Support for physical LPis is implemented. • The physical IST is valid. • There is a level 2 IST entry corresponding to the INTID. • The INTID is in the configured physical LPI range for the IRS Domain. Otherwise, the INTID is unreachable.
I _{MZRS}	In the physical LPI IST context , the base address of the physical IST for an IRS Domain is stored in IRS_IST_BASER.ADDR.
I _{XCVJT}	In the physical LPI IST context , IRS_IST_BASER.VALID determines whether the physical IST is valid.
I _{KYYP}	A write to IRS_IST_BASER completes when IRS_IST_STATUSR.IDLE is 1.
S _{KMPNW}	The architecture relies on making a physical IST valid or invalid to support the following software sequences: Initial configuration after system reset Software is expected to make the IST valid as part of the boot process. Soft reboot Software can reclaim the memory used for the IST by making the IST invalid. Making an IST valid or invalid is not intended to be used in power management sequences. See 4.11 IRS power management for more information about power management of an IRS.
I _{ZRRXD}	In a multi-IRS system, the physical IST is shared across all IRSs for an IRS Domain.
D _{MRGCV}	The physical IST <i>becomes valid</i> when a write that updates IRS_IST_BASER.VALID from 0 to 1 completes. The physical IST <i>becomes invalid</i> when a write that updates IRS_IST_BASER.VALID from 1 to 0 completes. See 4.7 The interrupt state table (IST) for more information about an IST becoming valid and invalid.
R _{NSNN}	In a multi-IRS system, when a write to IRS_IST_BASER.VALID completes, a subsequent read of any of the following registers on any IRS in the IRS Domain returns the same value. <ul style="list-style-type: none"> • IRS_IST_BASER.

- IRS_IST_CFGR.

R _{KLBCH}	<p>If a write occurs to IRS_IST_BASER.VALID when IRS_IST_STATUSR.IDLE is 0 on any other IRS for the same IRS Domain, it is CONSTRAINED UNPREDICTABLE which configuration and base address is used for the IST.</p> <p>The CONSTRAINED UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the IRS Domain or in an access to memory that does not follow the behaviors described in 4.12 IRS memory access rules.</p>
R _{SSTWT}	<p>When the physical IST becomes invalid and an LPI for the IRS Domain is selected as the candidate HPPI for a PE for the Interrupt Domain, one of the following is true:</p> <ul style="list-style-type: none"> • The candidate HPPI is acknowledged by the PE before the IST becomes invalid. • The candidate HPPI is not acknowledged and is no longer selected as the candidate HPPI for the PE.
R _{LJXGD}	<p>When the physical IST becomes invalid, if there are Accepted events for physical LPIs that are not yet processed, all of the following are true:</p> <ul style="list-style-type: none"> • The events are dropped and have no effect on any interrupt. • If software error reporting is supported, the IRS is permitted to report an error.
I _{KQGD}	<p>When the physical IST is valid, all IRSs are permitted to access the physical IST for any reason, including speculative reads.</p> <p>When the physical IST is invalid, the IRSs do not access the physical IST.</p> <p>See 4.7 The interrupt state table (IST) for more information.</p>
I _{JVVM}	<p>The configuration of the physical IST is controlled using the following register fields:</p> <ul style="list-style-type: none"> • IRS_IST_CFGR.STRUCTURE. • IRS_IST_CFGR.LPI_ID_BITS. • IRS_IST_CFGR.L2SZ. • IRS_IST_CFGR.ISTSZ.
I _{MBKRT}	<p>When IRS_IDR2.IST_levels is 0, IRS_IST_CFGR.STRUCTURE is RES0.</p>
I _{CLPKT}	<p>IRS_IDR2.{MIN_LPI_ID_BITS,ID_BITS} report the range of valid values for IRS_IST_CFGR.LPI_ID_BITS.</p>
I _{TQCYG}	<p>In the physical LPI IST context, when the physical IST is valid and uses a 2-level structure, a physical level 2 IST is made valid by writing to IRS_MAP_L2_ISTR.</p> <p>The effects of the write are complete when IRS_IST_STATUSR.IDLE is 1.</p>
R _{CMYTZ}	<p>If a write occurs to IRS_MAP_L2_ISTR when IRS_IST_STATUSR.IDLE is 0 on any other IRS for the same IRS Domain, it is CONSTRAINED UNPREDICTABLE whether any effects of any of the writes complete successfully.</p> <p>See also:</p> <ul style="list-style-type: none"> • 4.7 The interrupt state table (IST) • 11.2 IRS Data Structures

4.8.2 Physical SPIs

G _{BJKB}	<p>SPIs are designed to enable implementations to meet the following requirements:</p> <ul style="list-style-type: none"> • SPIs can be used to connect interrupt signals directly to an IRS, without recourse to an ITS. • SPIs can generate interrupts to PEs without requiring memory being allocated to the IRS. • SPIs can reliably generate interrupts to PEs in the presence of memory system errors. • SPIs can be implemented to provide predictable low latency guarantees for interrupt delivery. • SPIs can be assigned to VMs to support system partitioning using virtualization techniques.
R _{TLLZS}	<p>The storage location for physical SPIs is IMPLEMENTATION DEFINED.</p>
I _{ZFRNY}	<p>The number of SPIs across all IRSs in the system across all Interrupt Domains is reported by IRS_IDR5.SPI_RANGE.</p>

R _{KXMTR}	The value reported by IRS_IDR5.SPI_RANGE is the same across all Interrupt Domains across all IRSs in a system.
R _{MHRZT}	<p>An INTID specifying an SPI in a Physical Interrupt Domain is reachable, when all of the following are true:</p> <ul style="list-style-type: none"> • The SPI is implemented. • The SPI is assigned to the Interrupt Domain. • The SPI is not assigned to a VM. • The INTID.ID is less than or equal to IRS_IDR5.SPI_RANGE - 1. <p>Otherwise, the INTID is unreachable.</p>
I _{BVXHP}	See 4.10.2.1 Assigning physical SPIs to VMs for more information about assigning an SPI to a VM.
I _{HZSKP}	The configuration and state of a reachable SPI can be accessed via the GIC instructions on the PE.
D _{KZSTV}	<p>An SPI is <i>managed</i> by a single IRS.</p> <p>This means that an SPI can be assigned to an Interrupt Domain and to a VM only on the IRS where the SPI is managed.</p>
R _{VDLHR}	Each SPI is identified by a unique INTID.
I _{XYNMZ}	<p>The range of SPIs managed by an IRS is reported in IRS_IDR6.SPI_IRS_RANGE</p> <p>If IRS_IDR6.SPI_IRS_RANGE is 0, the IRS does not manage any SPIs, and IRS_IDR7.SPI_BASE is RES0.</p> <p>The minimum INTID.ID of the SPIs managed by an individual IRS is reported by IRS_IDR7.SPI_BASE.</p> <p>The range of SPI INTIDs managed by the IRS can be calculated as IRS_IDR7.SPI_BASE + IRS_IDR6.SPI_IRS_RANGE - 1.</p>
I _{YMKMH}	<p>If IRS_IDR5.SPI_RANGE is 0, no SPIs are implemented on any IRS, and all of the following are true:</p> <ul style="list-style-type: none"> • IRS_IDR6.SPI_IRS_RANGE is RES0. • IRS_IDR7.SPI_BASE is RES0.
R _{GCVXZ}	Not all SPIs in the range reported by IRS_IDR7.SPI_BASE and IRS_IDR6.SPI_IRS_RANGE are required to be implemented.
R _{CSQHP}	<p>The number reported by IRS_IDR6.SPI_IRS_RANGE is the same across all Interrupt Domains.</p> <p>The number reported by IRS_IDR7.SPI_BASE is the same across all Interrupt Domains.</p>
R _{ZYGST}	<p>For each IRS, all of the following are true for the range reported by IRS_IDR7.SPI_BASE and IRS_IDR6.SPI_IRS_RANGE:</p> <ul style="list-style-type: none"> • The range does not overlap with the range reported by other IRSs in the system. • The minimum and maximum SPI INTID included in the range are less than IRS_IDR5.SPI_RANGE.
I _{QNBQC}	<p>It is permitted to implement SPIs that are not connected to any interrupt input signal. Such SPIs can only be made Pending as a result of receiving a SET interrupt event from a PE.</p> <p>This may be useful to support IPIs without requiring the configuration and state for IPIs to be stored in memory.</p> <p>See also 2.5 Inter-Processor Interrupts.</p>
R _{PDMTJ}	<p>When an SPI is assigned to an Interrupt Domain, all of the following are true:</p> <ul style="list-style-type: none"> • The SPI is only considered for a candidate HPPI in the Interrupt Domain that the SPI is assigned to. • The SPI is only reachable in the Interrupt Domain that the SPI is assigned to.
I _{YTLRC}	This means that an SPI is unreachable from any Interrupt Domain other than the SPI is assigned to.
D _{HWBBY}	<p>The <i>SPI IRS configuration</i> for each SPI consists of the following configurations, which are separate from the SPI INTID configuration and state:</p> <ul style="list-style-type: none"> • The assignment of an SPI to an Interrupt Domain.

- The assignment of an SPI to a VM.
- The Trigger mode of the SPI.

R_{ZYJHT} The SPI IRS configuration for an SPI can be accessed from the following IRS Domains:

- The IRS Domain corresponding to the Interrupt Domain that the SPI is assigned to.
- The EL3 IRS Domain.

I_{SZZNP} Software accesses the SPI IRS configuration state using one of the following registers after selecting an implemented SPI managed by the IRS:

- IRS_SPI_CFGR.
- IRS_SPI_DOMAINR.
- IRS_SPI_VMR.

Software selects an SPI by writing the SPI INTID to IRS_SPI_SEL.

IRS_SPI_STATUSR.IDLE reports whether the effects of a write to any of these registers is completed.

IRS_SPI_STATUSR.V reports whether the value written to IRS_SPI_SEL selects an implemented SPI that is managed by the IRS and where the configuration can be accessed in the IRS Domain.

IRS_SPI_STATUSR.V is updated on a write to any of the following registers:

- IRS_SPI_CFGR.
- IRS_SPI_SEL.
- IRS_SPI_VMR.

If an SPI is assigned to a new Interrupt Domain after being selected by IRS_SPI_SEL, then a subsequent write to any of the registers above in the old Interrupt Domain causes IRS_SPI_STATUSR.V to become 0.

I_{RSFPY} When an SPI is selected using IRS_SPI_SEL and IRS_SPI_STATUSR.{V,IDLE} is {1,1}, IRS_SPI_CFGR.TM provides access to the Trigger mode of the SPI.

The Trigger mode of an SPI determines when the IRS generates events for an SPI and defines the type of the events.

Following a write that updates IRS_SPI_CFGR.TM to 1, the effects of the write are complete when IRS_SPI_STATUSR.{V,IDLE} is {1,1}.

See [4.6 Interrupt configuration and state](#) for more information about the Handling mode.

R_{QBXV} When the Trigger mode of an SPI is edge-triggered, the IRS generates a SET_EDGE event when the interrupt signal connected to an SPI is asserted.

When the Trigger mode of an SPI is level-sensitive, the IRS generates a SET_LEVEL event when the interrupt signal connected to an SPI is asserted and a CLEAR event when the signal is de-asserted.

The IRS does not generate a CLEAR event when the Trigger mode is edge-triggered and the signal is de-asserted.

I_{BGJKL} For each SPI, it is IMPLEMENTATION DEFINED whether the Trigger mode can be programmed by software.

If the Trigger mode cannot be programmed by software, access to the corresponding field is RO.

S_{DCPVJ} Software can detect that the Trigger mode of an SPI cannot be programmed by attempting to update the Trigger mode and reading it back. If the Trigger mode cannot be programmed, the Trigger mode is not updated.

R_{KBPXL} When the Trigger mode of an SPI is updated from level-sensitive to edge-triggered and the interrupt signal is asserted, the IRS generates a CLEAR event for the SPI.

For an SPI connected to an interrupt signal, when the Trigger mode of the SPI is updated from edge-triggered to level-sensitive, the IRS generates a SET_LEVEL event if the interrupt signal is asserted, and a CLEAR event if the signal is de-asserted.

For an SPI not connected to an interrupt signal, when the Trigger mode of the SPI is updated from edge-triggered to level-sensitive, the IRS generates a CLEAR event for the SPI.

When an update of the Trigger mode of an SPI generates an Interrupt Write Effect E_1 to the SPI Interrupt Location, all of the following are true:

- There is a Register Write Effect E_2 to IRS_SPI_STATUSR which updates the IDLE field from 1 to 0.
- There is a Register Write Effect E_3 to IRS_SPI_STATUSR which updates the IDLE field from 0 to 1.
- There is an Intrinsic order dependency from the Register Write Effect E_2 to the Interrupt Write Effect E_1 .
- There is an Intrinsic order dependency from the Interrupt Write Effects E_1 to the Register Write Effect E_3 .

This means that if there is a Memory Read Effect E_4 which Reads-from the Register Write Effect E_3 , then E_1 is Ordered-before E_4 .

R_{LKBBQ}

When the Trigger mode of an SPI is level-sensitive and connected to an interrupt input signal, the IRS is permitted to IGNORE SetPending commands received from PEs specifying that SPI.

D_{XKYXL}

A write to IRS_SPI_RESAMPLER causes the IRS to *resample* the signal connected to an SPI, if all of the following are true:

- The SPI is managed on the IRS where the write occurs.
- The SPI IRS configuration can be accessed in the IRS Domain where the write occurs.

Otherwise, the write has no effect beyond reporting a software error using EC 0x2B if supported.

R_{DMTFM}

When the signal of an SPI is resampled, all of the following are true:

- If the SPI Trigger mode is level-sensitive, all of the following are true:
 - If the signal is asserted at the time of the resample, the IRS generates a SET_LEVEL event.
 - If the signal is de-asserted at the time of the resample, the IRS generates a CLEAR event.
- If the SPI Trigger mode is edge-triggered, all of the following are true:
 - If the signal is asserted at the time of the resample, the IRS generates a SET_EDGE event.
 - If the signal is de-asserted at the time of the resample, the IRS generates no events.

R_{VYMH}

When a resample of an SPI generates an Interrupt Write Effects E_1 to the SPI Interrupt Location, all of the following are true:

- There is a Register Write Effect E_2 to IRS_SPI_STATUSR which updates the IDLE field from 1 to 0.
- There is a Register Write Effect E_3 to IRS_SPI_STATUSR which updates the IDLE field from 0 to 1.
- There is an Intrinsic order dependency from the Register Write Effect E_2 to the Interrupt Write Effects E_1 .
- There is an Intrinsic order dependency from the Interrupt Write Effects E_1 to the Register Write Effect E_3 .

This means that if there is a Memory Read Effect E_4 which Reads-from the Register Write Effect E_3 , then E_1 is Ordered-before E_4 .

S_{SNRXC}

Software may need to resample the state of an SPI to ensure that the Pending state of the SPI INTID corresponds to the state of the signal connected to the SPI. For example, if software has updated the Pending state or Handling mode of the SPI using GIC system instructions, it can use the resample mechanism to recover from this scenario.

Software may also need to resample an SPI if configuring the SPI for a source with level-sensitive semantics as Edge-triggered. In this case, software can resample the SPI after handling each event to ensure that interrupt events are not lost if the signal remained asserted while handling the interrupt.

D_{YGLYC}

The *reset state of the SPI IRS configuration* is the following:

- If the SPI Trigger mode can be programmed by software, it is reset to edge-triggered.
- The SPI is not assigned to a VM.
- The SPI is assigned to the Interrupt Domain corresponding to the MPSS supported by the IRS.

D_{TVVRZ}

The *reset state of an SPI INTID* is the following:

- The interrupt is Disabled.
- The interrupt is Inactive.
- The interrupt is Idle.
- If 1ofN interrupt routing is not supported, the SPI Routing mode is Targeted.
- All other interrupt configuration and state are UNKNOWN.

- R_{QRHKE}** When an IRS is reset, all of the following are true:
- All SPI INTIDs are reset to their reset state.
 - The IRS configuration for each SPI is reset to its reset state.
- S_{HZPVX}** To configure the assignment of an SPI to a VM, software performs the following sequence:
1. Software selects the SPI by writing to IRS_SPI_SEL.
 2. Software polls IRS_SPI_STATUS.IDLE until it reads as 1 and ensures that IRS_SPI_STATUS.V is 1.
 3. Software accesses the VM assignment of the SPI by accessing IRS_SPI_VMR.
 4. Following a write to IRS_SPI_VMR, software polls IRS_SPI_STATUS.IDLE until it reads as 1, and ensures that IRS_SPI_STATUS.V is 1, to ensure the effects of the write are complete.
- S_{GQDWY}** To configure the assignment of an Interrupt Domain, software performs the following sequence using the EL3 Interrupt Domain IRS configuration frame:
1. Software selects the SPI by writing to IRS_SPI_SEL.
 2. Software polls IRS_SPI_STATUS.IDLE until it reads as 1 and ensures that IRS_SPI_STATUS.V is 1.
 3. Software accesses the Interrupt Domain assignment of the SPI by accessing IRS_SPI_DOMAINR.
 4. Following a write to IRS_SPI_DOMAINR, software polls IRS_SPI_STATUS.IDLE until it reads as 1.
- To check whether an SPI is statically assigned to an Interrupt Domain or it can be dynamically assigned to an Interrupt Domain, software can attempt to update the assignment of the SPI to an Interrupt Domain and read back the assigned domain from IRS_SPI_DOMAINR.DOMAIN when IRS_SPI_STATUSR.IDLE is 1.
- I_{DMFCZ}** An SPI may be statically assigned to an Interrupt Domain.
- I_{BCQRT}** Arm recommends only statically assigning SPIs to an Interrupt Domain when it is known that the interrupt can only be handled in that Domain. For example, RAS interrupts may be statically assigned to the EL3 Domain. Statically assigning interrupts to a Domain may restrict how that interrupt can be used. For example, static assignment may prevent the interrupt from being directly assigned to Realms.
- R_{JJHJX}** When an SPI is assigned to a new Interrupt Domain as a result of selecting the SPI by writing to IRS_SPI_SEL and writing to IRS_SPI_DOMAINR.DOMAIN, the effects of the write are complete when IRS_SPI_STATUSR.IDLE is 1.
- When the effects are complete, all of the following are true:
- If the SPI was selected as the candidate HPPI for a PE in the old Interrupt Domain, one of the following is true:
 - The SPI is acknowledged in the old Interrupt Domain before the SPI is assigned to the new Interrupt Domain.
 - The SPI is not acknowledged in the old Interrupt Domain and no longer selected as the candidate HPPI in the old Interrupt Domain.
 - If the SPI was assigned to a VM in the old Interrupt Domain, it is CONSTRAINED UNPREDICTABLE whether the configuration and state of the virtual SPI INTID is UNKNOWN or reset to its reset state.
 - The SPI is not assigned to a VM in the new Interrupt Domain.
 - The SPI INTID is reset to its reset state.
- See also:
- [3.1 Interrupt Domains](#)
 - [4.6 Interrupt configuration and state](#)
 - [4.7 The interrupt state table \(IST\)](#)
 - [4.8.4 Physical interrupt signaling](#)

4.8.3 Physical interrupt routing

- R_{LFYMS}** The Routing mode of a physical interrupt determines how the target PE is selected for the interrupt in the following ways:
- If the Routing mode is Targeted, the target PE is the PE specified by the interrupt Affinity.

	<ul style="list-style-type: none"> • If the Routing mode is 1ofN, the target PE is determined dynamically.
D _{VBQJG}	When an interrupt Routing mode is Targeted and the interrupt Affinity specifies a PE, the interrupt is said to be <i>targeted to</i> that PE.
R _{HGPQH}	IRS support for 1ofN interrupt routing is optional.
D _{VPZNM}	An interrupt whose Routing mode is <i>1ofN</i> is referred to as a <i>1ofN interrupt</i> .
I _{GHKQS}	IRS_IDR0.ONE_N reports whether 1ofN interrupt routing is supported.
R _{FGVHX}	The value of IRS_IDR0.ONE_N is the same for all IRSs for all Interrupt Domains in the system.
R _{ZMTLR}	<p>If 1ofN interrupt routing is not supported, all of the following are true:</p> <ul style="list-style-type: none"> • If a GIC System instruction generates an Interrupt Effect that updates the Routing mode of a physical interrupt, the Routing mode is set to Targeted regardless of the value provided in <Xt>. • L2_ISTE.IRM is treated as 0 when the physical IST becomes valid.
R _{GJNGF}	Each physical 1ofN interrupt can be acknowledged by at most one PE.
R _{STMZL}	The storage location for the Affinity of physical 1ofN interrupts may be used for an IMPLEMENTATION DEFINED purpose to handle routing of physical 1ofN interrupts, with the restriction that an Affinity value of 0 has no IMPLEMENTATION DEFINED meaning.
R _{QZWMR}	An IRS is permitted to update the Affinity of physical 1ofN interrupts.
R _{DJNNR}	A PE may only be selected as the target for a 1ofN interrupt, if 1ofN PE selection is enabled for the PE for the Interrupt Domain
R _{CDXNW}	<p>An IRS Domain selects a target PE for a reachable 1ofN interrupt in finite time, when all of the following are true:</p> <ul style="list-style-type: none"> • The interrupt is Pending. • The interrupt is Inactive. • The interrupt is Enabled. • At least one IRS is enabled. • At least one PE connected to an enabled IRS has 1ofN PE selection enabled.
R _{TMBZH}	An IRS Domain may select a new target PE for a physical 1ofN interrupt at any time.
I _{DBQSC}	<p>Selecting a target PE does not guarantee that the interrupt is selected as the candidate HPPI for the target PE. For example, the selected target PE may be masking the interrupt or another interrupt may be the candidate HPPI for the selected target PE.</p> <p>See 4.8.4 Physical interrupt signaling for more information about selecting the physical candidate HPPI.</p>
R _{FYRLW}	The mechanism that an IRS uses to select a target PE among several possible target PEs is IMPLEMENTATION DEFINED.
I _{LZSJY}	<p>The architecture does not require that higher priority 1ofN interrupt are delivered before lower priority 1ofN interrupts.</p> <p>For example, the following situation may occur:</p> <ul style="list-style-type: none"> • The IRS selects PE 0 for a high priority 1ofN interrupt (INTID A). • The IRS selects PE 1 for a low priority 1ofN interrupt (INTID B). • There are several interrupts with higher priority than INTID A that are Targeted to PE 0. <p>In this situation, PE 1 may handle INTID B before PE 0 handles INTID A.</p>
I _{XRMBH}	<p>Whether a PE has 1ofN PE selection enabled or disabled is configured separately for each IRS Domain.</p> <p>This configuration setting is accessed using the following sequence:</p> <ol style="list-style-type: none"> 1. Select the PE using IRS_PE_SEL on the IRS where the PE is connected. 2. Poll IRS_PE_STATUSR.IDLE until it reads as 1. 3. Access IRS_PE_CR0.DPS.

4. Poll IRS_PE_STATUSR.IDLE until it reads as 1.

R_{LQMBW} When 1ofN selection is disabled for PE for an Interrupt Domain, the corresponding PE is not selected as the target for 1ofN interrupts belonging to that Interrupt Domain.

I_{SHQTF} When IRS_PE_CR0.DPS is updated for a PE, and a physical 1ofN interrupt is selected as the physical candidate HPPI for the PE, one of the following is true:

- The PE acknowledges the interrupt before IRS_PE_STATUSR.IDLE is 1.
- The PE does not acknowledge the interrupt.

See [4.8.4 Physical interrupt signaling](#) for more information.

4.8.4 Physical interrupt signaling

R_{LXYDJ} The Pending state of a physical LPI identified by an INTID is updated when all of the following are true:

- The physical LPI as identified by the INTID is reachable.
- Any of the following occur:
 - An ITS generates an interrupt event that specifies the INTID.
 - A write to the IRS_SETLPI register specifies the INTID.
 - An IRS [processes](#) an Interrupt Effect generated by a GIC System instruction that updates the LPI Pending state.

Otherwise, the physical LPI Pending state is not updated.

I_{DZSFL} The Interrupt Domain of an LPI that is being signaled is determined by how the interrupt is signaled as follows:

- The Interrupt Domain is specified as part of the interrupt event received from the ITS.
- The Interrupt Domain as reported in IRS_IDR0.INT_DOM for the IRS where a write to IRS_SETLPI occurs.
- The Interrupt Domain that the GIC System instruction operates in.

R_{XYLLD} The Pending state of a physical SPI identified by an INTID is updated when all of the following are true:

- The physical SPI as identified by the INTID is reachable.
- Any of the following occur:
 - An interrupt connected to an IRS as an SPI is asserted or de-asserted and generates an interrupt event for the physical SPI.
 - An IRS [processes](#) an Interrupt Effect generated by a GIC System instruction that updates the SPI Pending state.

Otherwise, the physical SPI Pending state is not updated.

I_{SKNHH} The Interrupt Domain of an SPI that is being signaled is determined by the Interrupt Domain that the SPI is assigned to.

See [4.8.2 Physical SPIs](#) and [4.2 Signaling interrupts](#) for more information.

I_{SZZVX} If an SPI that is assigned to a VM is signaled, the SPI is a virtual SPI and the rules for when its Pending state is updated are different. See [4.10 Virtual interrupts](#) and [4.10.4 Virtual interrupt signaling](#) for more information.

D_{NFYQX} The IRS Domain selects a physical candidate HPPI for a PE for an Interrupt Domain if, and only if, at least one physical interrupt meets the physical *candidate HPPI conditions*:

- The interrupt is Enabled.
- The interrupt is Pending.
- The interrupt is Inactive.
- One of the following is true:
 - The interrupt Routing mode is Targeted and the interrupt Affinity specifies the PE.
 - The interrupt Routing mode is 1ofN and the PE is selected as the target PE for the interrupt.

R _{WSBLV}	<p>For a PE and an Interrupt Domain, all of the following apply to the physical candidate HPPI selection:</p> <ul style="list-style-type: none"> • If at least one interrupt satisfies the physical candidate HPPI conditions, the IRS Domain selects one of these as the physical candidate HPPI in finite time. • If no interrupts satisfy the physical candidate HPPI conditions, the IRS Domain does not select any interrupt as the physical candidate HPPI.
R _{HGTSH}	<p>If more than one interrupt satisfies the physical candidate HPPI conditions for a PE for an Interrupt Domain, the IRS Domain selects the interrupt with the highest Priority as the physical candidate HPPI in finite time.</p> <p>If there is more than one physical candidate HPPI with the same Priority, it is IMPLEMENTATION DEFINED which of those interrupts is selected as the physical candidate HPPI.</p>
I _{BGBFD}	<p>The architecture only requires that the highest Priority interrupt that satisfies the candidate HPPI conditions is selected in finite time, because the set of interrupts that satisfy the conditions changes over time as the state and configuration of interrupts are updated. An implementation may need to consider the Priority of all the interrupts that satisfy the conditions which may require processing time.</p>
R _{KFYBX}	<p>When at least one physical interrupt meets the physical <i>candidate HPPI conditions</i> for a PE for an Interrupt Domain, and the IRS is enabled for the Interrupt Domain, the IRS selects a physical candidate HPPI in finite time.</p>
I _{PZVZZ}	<p>See 4.11 IRS power management for information about the IRS behavior if it has selected a physical candidate HPPI for a PE and the PE is offline.</p>
I _{ZPJZL}	<p>An interrupt might be signaled on one IRS and be targeted to a PE connected to a different IRS. Whether that interrupt can be selected as the physical candidate HPPI for the PE depends only on the configuration of the IRS local to the PE. If the IRS on which the interrupt was originally signaled is subsequently disabled, that has no effect on whether the interrupt is selected as the physical candidate HPPI for the targeted PE.</p>
R _{XWWNP}	<p>When a physical interrupt that has been selected as the physical candidate HPPI for a PE for an Interrupt Domain is no longer selected as the physical candidate HPPI, the PE will either acknowledge the old physical candidate HPPI within finite time, or the old physical candidate HPPI will not be acknowledged.</p>
R _{CHNVV}	<p>If the IRS updates the configuration of an interrupt that has been selected as the physical candidate HPPI for an Interrupt Domain for a PE, and the interrupt configuration data was communicated to the PE, the IRS will select the same or a new physical interrupt as the physical candidate HPPI.</p> <p>The IRS then communicates the selected physical candidate HPPI to the PE, including any updated interrupt configuration data, in finite time.</p> <p>For more information about ordering of such events, see 2.12 Interrupt ordering model and synchronization requirements and Chapter B1 Interrupt ordering litmus tests.</p>

Note

The following sequence is an example of the IRS communicating updated interrupt configuration data because there was an update to the configuration data that was communicated to a PE:

1. An interrupt is selected as the physical candidate HPPI for a PE and its Priority is communicated to the PE.
 2. The PE executes a GIC System instruction that updates the Priority of the interrupt.
 3. The IRS selects the same interrupt as the physical candidate HPPI for the PE.
 4. The IRS communicates that the interrupt is still the candidate HPPI with the updated Priority value to the PE.
-

I _{TMRWF}	<p>Updates to interrupt state or configuration can lead to the IRS selecting a different candidate HPPI. For example, the current candidate HPPI may become Disabled, or a higher priority interrupt may be updated to meet the candidate HPPI conditions. In such scenarios, the IRS selects a new candidate HPPI in finite time. However, the PE may acknowledge the old candidate HPPI before the IRS communicates to the PE that a new candidate HPPI has been selected.</p>
--------------------	--

4.9 Virtualization data structures

Figure 4.3 shows an overview of the data structures used to support VMs in the IRS.

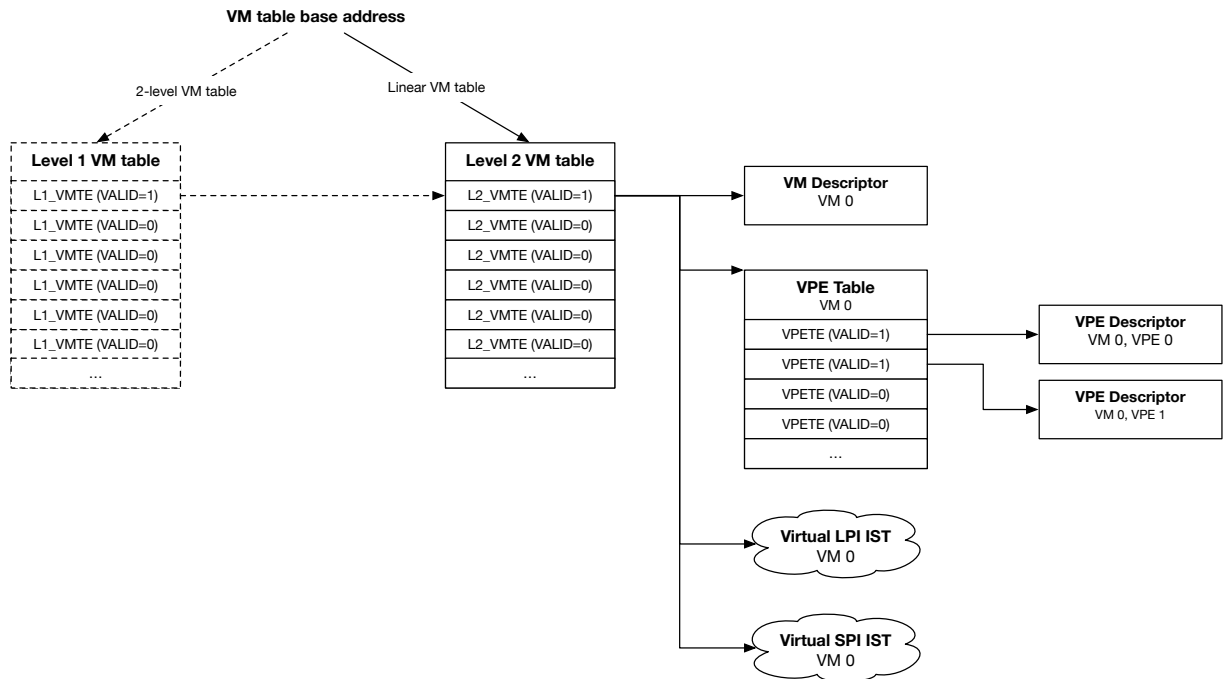


Figure 4.3: IRS virtualization structures

The GICv5 architecture defines the following *virtualization data structures* used by an IRS:

- The Virtual Machine Table (VM table).
- The Virtual PE Table (VPE table).

Each level 2 entry in the VM table describes the configuration of a VM and contains pointers to the following child data structures:

- A VM descriptor, if required.
- A VPE table.
- A virtual LPI IST.
- A virtual SPI IST.

Each entry in the VPE table describes the configuration of the VPE and contains pointers to the VPE descriptor for each VPE in the VM.

A VM supports virtual LPIs, virtual SPIs, or both.

Virtual LPIs and virtual SPIs are stored in separate virtual ISTs.

Software is responsible for allocating the virtualization data structures from memory in the PAS associated with the IRS Domain where the VM table is used.

See also:

- [4.7 The interrupt state table \(IST\)](#)
- [4.10 Virtual interrupts](#)
- [11.2 IRS Data Structures](#)

4.9.1 The VM table

<code>I_PVHPD</code>	<p>The VM table contains the configuration of the VMs defined in an IRS Domain.</p> <p>For each VM, the VM table contains pointers to other data structures associated with that VM.</p> <p>In a multi-IRS system, the VM table is shared between all IRSs for the IRS Domain.</p>
<code>D_PZXYM</code>	<p>The VM table uses either a linear or 2-level structure.</p> <p>For a linear VM table, all of the following are true:</p> <ul style="list-style-type: none"> • The level 1 VM table is omitted. • The VM table consists of a single level 2 VM table.
<code>I_DSPSD</code>	IRS_IDR3.VMT_LEVELS reports whether 2-level VM table support is implemented.
<code>R_KQYKZ</code>	<p>The levels supported for the VM table are the same across all Interrupt Domains, except for the EL3 Interrupt Domain, where VMs are not supported.</p> <p>The levels supported for the VM table are the same across all IRSs.</p>
<code>D_NRLZZ</code>	<p>For a linear VM table, the base address of the VM table specifies the location of the level 2 VM table.</p> <p>For a 2-level VM table, the base address of the VM table specifies the location of the level 1 VM table.</p>
<code>D_LSPBK</code>	When the VM table uses a 2-level table structure, the size of each level 2 VM table is 4 KB.
<code>R_MKPCF</code>	<p>The structure of the VM table is controlled using the following register fields:</p> <ul style="list-style-type: none"> • IRS_VMT_CFGR.STRUCTURE: Selects if the VM table uses a linear or 2-level structure. • IRS_VMT_CFGR.VM_ID_BITS: Selects how many VMs the IRS Domain supports. This impacts the size of the level 2 VM table when using a linear structure and the size of the level 1 VM table when using a 2-level structure.
<code>I_HCGJW</code>	The VM table is indexed using a VM ID ranging from 0 through $(2^{\text{IRS_VMT_CFGR.VM_ID_BITS}}) - 1$.
<code>R_NXJZZ</code>	<p>When an IRS accesses the VM table for an IRS Domain, the IRS does not access any memory location derived from a VM ID which is outside the configured VM ID range for the IRS Domain.</p> <p>This includes situations where the VM ID is programmed in the parameters to a GIC System instruction, the VM ID is received from an ITS, or the VM ID is stored in a physical SPI VM assignment register.</p>
<code>I_ZKGTX</code>	The maximum VM ID range supported is reported by IRS_IDR3.VM_ID_BITS.
<code>R_XYKND</code>	<p>The maximum number of VM ID bits are the same across all Interrupt Domains, except for the EL3 Interrupt Domain, where VMs are not supported.</p> <p>The maximum number of VM ID bits is the same across all IRSs.</p>

4.9.1.1 The VM table base address and configuration registers

<code>I_GMXKS</code>	The base address of the VM table is stored in IRS_VMT_BASER.ADDR.
<code>I_JLYPB</code>	IRS_VMT_BASER.VALID determines whether the VM table is valid.
<code>I_BNSSH</code>	A write to IRS_VMT_BASER completes when IRS_VMT_STATUSR.IDLE is 1.
<code>I_YJKRP</code>	In a multi-IRS system, the VM table is shared across all IRSs for an IRS Domain.
<code>D_VJFSP</code>	<p>The VM table becomes valid when a write that updates IRS_VMT_BASER.VALID from 0 to 1 completes.</p> <p>The VM table becomes invalid when a write that updates IRS_VMT_BASER.VALID from 1 to 0 completes.</p>

R_{RLYYJ} In a multi-IRS system, when a write to IRS_VMT_BASER.VALID completes, the values of the following registers on that IRS are returned on a subsequent read on any IRS in the IRS Domain:

- IRS_VMT_BASER.
- IRS_VMT_CFGR.

R_{WXTKP} If a write occurs to IRS_VMT_BASER.VALID when IRS_VMT_STATUSR.IDLE is 0 on any other IRS for the same IRS Domain, it is CONSTRAINED UNPREDICTABLE which configuration and base address is used for the VM table.

The CONSTRAINED UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the IRS Domain or in an access to memory that does not follow the behaviors described in [4.12 IRS memory access rules](#).

R_{QRSZL} When the VM table is valid, all IRSs are permitted to access the virtualization data structures for any reason, including speculative reads.

When the VM table is invalid, the IRSs do not access the virtualization data structures table.

See [4.9 Virtualization data structures](#) for the definition of virtualization data structures.

4.9.1.2 Level 2 VM table management

D_{YPLVV} A level 2 VM table is considered *valid*, if the VM table is valid and one of the following is true:

- The VM table uses a linear structure.
- The VM table uses a 2-level structure and the corresponding L1_VMTE.VALID is 1.

Otherwise, the level 2 VM table is considered *invalid*.

D_{QKXL} A level 2 VM table *becomes valid* in any of the following scenarios:

- The VM table uses a linear structure and the VM table becomes valid.
- The VM table uses a 2-level structure and any of the following occur:
 - A write to IRS_VMAP_L2_VMTR makes the level 2 VM table valid.
 - The VM table becomes valid and the corresponding L1_VMTE.VALID is 1.

I_{CSRYH} When a 2-level VM table becomes valid, if L1_VMTE.VALID is 1 for more than one level 1 VM table entry, more than one level 2 VM table may become valid at the same time.

R_{BXWBF} When the VM table is valid and uses a 2-level structure, if software writes to an invalid level 1 VM table entry and sets L1_VMTE.VALID to 1, the behavior of the IRS Domain is UNPREDICTABLE.

The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the Interrupt Domain or in an access to memory that does not follow the behaviors described in [4.12 IRS memory access rules](#).

R_{XBDXW} For a 2-level VM table, if a write updates a valid level 1 VM table entry, the IRS Domain behavior is UNPREDICTABLE.

The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the IRS Domain or in an access to memory that does not follow the behaviors described in [4.12 IRS memory access rules](#).

I_{DPCRG} The architecture does not support making an individual level 2 VM table invalid when using a 2-level structure for the VM table.

R_{RNSYR} To recover from UNPREDICTABLE behavior resulting from an incorrect write to a level 1 VM table entry, the VM table is made invalid.

R_{TZWTX} When a write to IRS_VMAP_L2_VMTR makes the corresponding level 2 VM table valid, all of the following are true:

- The IRS generates the following Effects:
 - A Register Write Effect E₁ to IRS_VMT_STATUSR which sets the IDLE field to 0.
 - A Memory Write Effect E₂ to the Location corresponding to the level 1 VM table entry which updates L1_VMTE.VALID from 0 to 1.

- A Register Write Effect E_3 to IRS_VMT_STATUSR which sets the IDLE field to 1.
- There is an Intrinsic order dependency from the Register Write Effect E_1 to the Memory Write Effect E_2 .
- There is an Intrinsic order dependency from the Memory Write Effect E_2 to the Register Write Effect E_3 .

This means that if there is a Memory Read Effect E_4 which Reads-from the Register Write Effect E_3 , then E_2 is Ordered-before E_4 .

I_{YTWVR}

When the VM table is valid and uses a 2-level structure, a level 2 VM table is made valid by the following sequence:

1. Software writes the address of the new level 2 VM table to L1_VMTE.L2_ADDR, with L1_VMTE.VALID remaining to be 0.
2. Software writes to IRS_VMAP_L2_VMTR.
3. The IRS performs a write to the level 1 VM table entry that updates L1_VMTE.VALID from 0 to 1.
4. The IRS reports that the effects of the write to IRS_VMAP_L2_VMTR are complete.

See [3.3 Coherency considerations for GIC data structures](#) for more information about ensuring that memory accesses are coherent between the IRS and a PE.

R_{KZZPQ}

When the VM table is valid and uses a 2-level structure, if software writes to the level 1 VM table entry when IRS_VMT_STATUSR.IDLE is 0, it is CONSTRAINED UNPREDICTABLE whether the IRS writes back the entry using the updated or old level 1 entry value.

The CONSTRAINED UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the Interrupt Domain or in an access to memory that does not follow the behaviors described in [4.12 IRS memory access rules](#).

R_{KGTMJ}

If a write occurs to IRS_VMAP_L2_VMTR when IRS_VMT_STATUSR.IDLE is 0 on any other IRS for the same IRS Domain, it is CONSTRAINED UNPREDICTABLE whether any effects of any of the writes complete successfully.

D_{CFNYD}

A level 2 VM table *becomes invalid* when the VM table becomes invalid.

4.9.1.3 VM management

D_{VBQTG}

A VM and the corresponding level 2 VM table entry are considered *valid* when all of the following are true:

- The level 2 VM table is valid.
- The corresponding L2_VMTE.VALID is 1.

Otherwise the VM and the corresponding level 2 VM table entry are considered *invalid*.

R_{YBMTV}

If a VM ID is outside the configured VM ID range for the IRS Domain, the VM ID is treated as invalid.

I_{JQFYM}

When a VM is valid, all of the following are true:

- Virtual interrupts in the VM can be signaled.
- VPEs can become resident.
- VPEs can access the configuration and state of reachable virtual interrupts.

When a VM is invalid, none of the above are true.

I_{JFWCL}

When the VM table is invalid, all of the following are true:

- All VM IDs specify invalid VMs.
- All VPEs are invalid.
- The IRS does not select any virtual candidate HPPIs to the PEs connected to the IRS.
- When a PE executes a GIC System instruction that operates in the Virtual Interrupt Domain, all Interrupt Effects generated by that instruction, that are Ordered-after the VM table becomes invalid, are IGNORED.

See [Chapter 2 PE architecture](#) for more information about the relationship between a resident VPE and the Virtual Interrupt Domain.

D_{QXMZT}

A VM *becomes valid* in any of the following scenarios:

- A write to IRS_VMAP_VMR makes the VM valid.
- The level 2 VM table becomes valid and the corresponding L2_VMTE.VALID is 1.

<code>R_VQKZB</code>	<p>When a level 2 VM table is valid, if software writes to an invalid level 2 VM table entry and sets <code>L2_VMTE.VALID</code> to 1, the behavior of the IRS Domain is UNPREDICTABLE.</p> <p>The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the Interrupt Domain or in an access to memory that does not follow the behaviors described in 4.12 IRS memory access rules.</p>
<code>R_HPDGY</code>	To recover from UNPREDICTABLE behavior resulting from a write to an invalid level 2 VM table entry that sets <code>L2_VMTE.VALID</code> to 1, the VM table is made invalid.
<code>R_KBKBFB</code>	<p>When a write to <code>IRS_VMAP_VMR</code> makes the corresponding VM valid, all of the following are true:</p> <ul style="list-style-type: none"> The IRS generates the following Effects: <ul style="list-style-type: none"> A Register Write Effect E_1 to <code>IRS_VMT_STATUSR</code> which sets the IDLE field to 0. A Memory Write Effect E_2 to the Location corresponding to the level 1 VM table entry which updates <code>L2_VMTE.VALID</code> from 0 to 1. A Register Write Effect E_3 to <code>IRS_VMT_STATUSR</code> which sets the IDLE field to 1. There is an Intrinsic order dependency from the Register Write Effect E_1 to the Memory Write Effect E_2. There is an Intrinsic order dependency from the Memory Write Effect E_2 to the Register Write Effect E_3. <p>This means that if there is a Memory Read Effect E_4 which Reads-from the Register Write Effect E_3, then E_2 is Ordered-before E_4.</p>
<code>I_LTQMV</code>	<p>When a level 2 VM table is valid, a VM is made valid by the following sequence:</p> <ol style="list-style-type: none"> Software initializes the corresponding level 2 VM table entry, with <code>L2_VMTE.VALID</code> remaining to be 0. Software writes to <code>IRS_VMAP_VMR</code>. The IRS performs a write to the level 2 VM table entry that updates <code>L2_VMTE.VALID</code> from 0 to 1. The IRS reports that the effects of the write to <code>IRS_VMAP_VMR</code> are complete. <p>See 3.3 Coherency considerations for GIC data structures for more information about ensuring that memory accesses are coherent between the IRS and a PE.</p>
<code>I_PYJLW</code>	<p>When a VM becomes valid, the following fields are interpreted by the IRS Domain:</p> <ul style="list-style-type: none"> <code>L2_VMTE.VMD_ADDR</code> <code>L2_VMTE.VPE_ID_BITS</code> <code>L2_VMTE.VPET_ADDR</code> <code>L2_VMTE.LPI_IST_VALID</code> <code>L2_VMTE.SPI_IST_VALID</code> <p>If the <code>L2_VMTE.LPI_IST_VALID</code> field is 1 when the VM becomes valid, the following fields are interpreted by the IRS Domain:</p> <ul style="list-style-type: none"> <code>L2_VMTE.LPI_ID_BITS</code> <code>L2_VMTE.LPI_IST_STRUCTURE</code> <code>L2_VMTE.LPI_IST_ADDR</code> <code>L2_VMTE.LPI_ISTSZ</code> <p>If the <code>L2_VMTE.SPI_IST_VALID</code> field is 1 when the VM becomes valid, the following fields are interpreted by the IRS Domain:</p> <ul style="list-style-type: none"> <code>L2_VMTE.SPI_ID_BITS</code> <code>L2_VMTE.SPI_IST_STRUCTURE</code> <code>L2_VMTE.SPI_IST_ADDR</code> <code>L2_VMTE.SPI_ISTSZ</code> <p>See 4.9.2 The VPE table for information about valid entries in the VPE table when a VM becomes valid.</p>
<code>R_BYJYM</code>	<p>For a write to a valid level 2 VM table entry, all of the following are true:</p> <ul style="list-style-type: none"> If <code>L2_VMTE.LPI_IST_VALID</code> is 0, the following fields are permitted to be updated: <ul style="list-style-type: none"> <code>L2_VMTE.LPI_ID_BITS</code>. <code>L2_VMTE.LPI_IST_STRUCTURE</code>.

- L2_VMTE.LPI_IST_ADDR.
- L2_VMTE.LPI_ISTSZ.
- If L2_VMTE.SPI_IST_VALID is 0, the following fields are permitted to be updated:
 - L2_VMTE.SPI_ID_BITS.
 - L2_VMTE.SPI_IST_STRUCTURE.
 - L2_VMTE.SPI_IST_ADDR.
 - L2_VMTE.SPI_ISTSZ.

Updating any other field in a valid level 2 VM table entry results in UNPREDICTABLE behavior.

The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the IRS Domain or in an access to memory that does not follow the behaviors described in [4.12 IRS memory access rules](#).

I_HQDFR

See [4.10.1 Virtual LPIs](#) and [4.10.2 Virtual SPIs](#) for information about how the virtual ISTs become valid and the corresponding L2_VMTE.LPI_IST_VALID and L2_VMTE.SPI_IST_VALID are updated for a valid VM.

R_LDKSX

When a level 2 VM table is valid, the IRS behavior is UNPREDICTABLE if there is a write to the level 2 VM table entry when IRS_VMT_STATUSR.IDLE is 0 because of a write to any of the following registers:

- IRS_VMAP_VMR.
- IRS_VMAP_VPER.
- IRS_VMAP_VISTR.
- IRS_VMAP_L2_VISTR.

The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the Interrupt Domain or in an access to memory that does not follow the behaviors described in [4.12 IRS memory access rules](#).

R_FFCKZ

If a write occurs to any of the following registers when IRS_VMT_STATUSR.IDLE is 0 on any other IRS for the same IRS Domain, it is CONSTRAINED UNPREDICTABLE whether any effects of any of the writes complete successfully:

- IRS_VMAP_VMR.
- IRS_VMAP_VPER.
- IRS_VMAP_VISTR.
- IRS_VMAP_L2_VISTR.

D_BTКBB

A VM becomes invalid in any of the following scenarios:

- A write to IRS_VMAP_VMR makes the VM invalid.
- The VM table becomes invalid.

R_KFTFZ

To recover from UNPREDICTABLE behavior resulting from an incorrect write to a valid level 2 VM table entry, the VM is made invalid by a write to IRS_VMAP_VMR that sets IRS_VMAP_VMR.U to 1.

R_XXVLZ

If a VM becomes invalid and a VPE in the VM is resident on a PE, the resident VPE is treated as belonging to an invalid VM and all of the following are true:

- If a virtual candidate HPPI was selected for the VPE, one of the following is true:
 - The virtual candidate HPPI is acknowledged in finite time.
 - The virtual candidate HPPI is not acknowledged.
- GIC System instructions operating in the Virtual Interrupt Domain and specify the VM are treated as specifying an invalid VM.
- Doorbell requests from the PE where the VPE is resident are IGNORED.

If the VM with the same VM ID subsequently becomes valid, then it's CONSTRAINED UNPREDICTABLE whether the resident VPE is still treated as being resident and belonging to the new VM, or is treated as belonging to an invalid VM.

S_WKTCN

Software can make a VM invalid and reclaim all resources related to the VM by performing the following sequence:

1. Make all VPEs non-resident.
2. Unmap and invalidate all ITS mappings that target the VM.
3. Unassign all SPIs assigned to the VM.

4. Make the VM invalid by writing to IRS_VMAP_VMR.

4.9.1.4 The VM descriptor

- I_{ZVXTY}** An implementation may require a VM descriptor for each VM.
Whether a VM descriptor is required is reported by IRS_IDR3.VMD.
- I_{QYYFH}** The VM descriptor is memory provided by software to the IRS Domain to manage the VM. The content and usage of the VM descriptor is IMPLEMENTATION DEFINED. The VM descriptor cannot be moved once the VM is made valid.
- I_{YBLTD}** If a VM descriptor is required, a valid L2_VMTE contains the base addresses of the VM descriptor.
- R_{BWKBG}** If a VM descriptor is required, when a VM becomes valid, the memory pointed to by L2_VMTE.VMD_ADDR is expected to contain all zeros.
- If the VM descriptor does not contain all zeros, all of the following are true:
- It is CONSTRAINED UNPREDICTABLE whether the IRS Domain selects a target VPE for virtual 1ofN interrupts for the VM.
 - The VM configuration, including the doorbell settings, is reset to UNKNOWN values.

- R_{HQPRD}** The format and content of the VM descriptor is IMPLEMENTATION DEFINED.
- I_{BBCFD}** The size of the VM descriptor is reported by IRS_IDR3.VMD_SZ.
- R_{VNVHD}** Whether a VM descriptor is required, and its size, is the same across all Interrupt Domains, except for the EL3 Interrupt Domain, where VMs are not supported.
- Whether a VM descriptor is required, and its size, is the same across all IRSs.

4.9.1.5 Example VM table structures

- I_{BSPBV}** The size of a level 1 VMTE is 8 bytes and the size of a level 2 VMTE is 32 bytes.
[Table 4.3](#) shows some example VM table structures.
- S_{VKKQW}** The L1 size and L2 size columns in [Table 4.3](#) indicate the required maximum physically contiguous allocation by software for the configuration.

Table 4.3: Example VM table structures

STRUCTURE	VM_ID_BITS	L1 size	L2 size	Maximum number of VMs
Linear	7	-	4 KB	128
Linear	11	-	64 KB	2,048
Linear	15	-	1 MB	32,768
Linear	16	-	2 MB	65,536
2-level	16	4 KB	4 KB	65,536

4.9.2 The VPE table

- I_{BGKFQ}** The VPE table contains the configuration of the VPEs defined for a VM.
For each VPE, the VPE table contains pointers to other data structures associated with that VPE.
In a multi-IRS system, the VPE table is shared between all IRSs for the IRS Domain.
The VPE table is not permitted to be shared across multiple VMs.

D _{PFHDR}	The VPE table uses a linear structure.
D _{MPMBR}	The base address of the VPE table specifies the location of the VPE table.
I _{WXRZV}	L2_VMTE.VPE_ID_BITS selects maximum number of VPEs for the VM.
S _{PLJRT}	When a VM is made valid, it is configured to support a maximum number of VPEs. This value cannot be modified once the VM is made valid. If more VPEs are required than the maximum configured in the L2_VMTE.VPE_ID_BITS field, then the VM can be made invalid and a new VM can be made valid with a larger maximum number of VPEs.

Note

L2_VMTE.VPE_ID_BITS specifies the maximum number of entries in the VPE table. For some of those entries, VPETE.VALID may be 0, and therefore the number of valid VPEs in the VM may be smaller than the value specified by L2_VMTE.VPE_ID_BITS.

I _{DRHQP}	The VPE table is indexed using a VPE ID ranging from 0 to $(2^{L2_VMTE.VPE_ID_BITS}) - 1$.
R _{WSDXY}	When an IRS accesses the VPE table for a VM, the IRS does not access any memory location derived from a VPE ID which is outside the configured VPE ID range for the VM. This includes situations where the VPE ID is programmed in the parameters to a GIC System instruction, written to a GIC System register, or the VPE ID is stored in a virtualization data structure entry.
I _{QCQBD}	The maximum supported VPE ID is reported by IRS_IDR4.VPE_ID_BITS.
R _{HBNVF}	The number of supported VPE ID bits is the same across all Interrupt Domains, except for the EL3 Interrupt Domain, where VMs are not supported. The number of supported VPE ID bits is the same across all IRSs.
R _{NQDMK}	For virtual interrupts, when an Interrupt Affinity ID is specified, it is interpreted as the VPE ID.
S _{RMDNH}	A virtual Interrupt Affinity ID is interpreted as the VPE ID for virtual interrupts and used directly to index into the VPE table. Arm expects that hypervisor software at EL2 presents a virtual IRS with an emulated virtual IRS_PE_SEL register. The value written to the emulated IRS_PE_SEL.IAFFID field should select the emulated configuration registers corresponding to the VPE with the corresponding IAFFID and VPE ID.

4.9.2.1 The VPE table base address and configuration

I _{CFPXL}	The base address of the VPE table for a VM is stored in the corresponding L2_VMTE.VPET_ADDR. See 4.9.1.3 VM management for more information.
I _{LTQZN}	A VPE table becomes valid when its VM becomes valid. When a VPE table becomes valid, the VPE table can contain all invalid entries, a combination of valid and invalid entries, or all valid entries.
R _{PSYQV}	When VPETE.VALID is 1, the IRS may speculatively access any memory location derived from the address and configuration data in the VPE table entry. See 4.12 IRS memory access rules for more information.

4.9.2.2 VPE management

D _{BKCHX}	A VPE for a VM, and the corresponding VPE table entry, are considered <i>valid</i> when all of the following are true: <ul style="list-style-type: none"> • The VM is valid. • The corresponding VPETE.VALID is 1.
--------------------	--

	Otherwise, the VPE and the corresponding VPE table entry are considered <i>invalid</i> .
R _{JKDCN}	If a VPE ID is outside the configured VPE ID range for the VM, the VPE ID is treated as specifying an invalid VPE.
G _{QNPFW}	The architecture supports multiple ways to create a VM: <ul style="list-style-type: none"> • A VM can be made valid without containing any valid VPEs, and each VPE is subsequently made valid. • A VM can be made valid with pre-configured valid VPEs to reduce the overhead of setting up a new VM.
D _{HNRJX}	A VPE <i>becomes valid</i> in any of the following scenarios: <ul style="list-style-type: none"> • A write to IRS_VMAP_VPER makes the VPE valid for an already valid VM. • A VM becomes valid and VPETE.VALID is 1 for the corresponding VPE table entry for the VM.
R _{LSYNH}	When a VPE table is valid, if software writes to an invalid VPE table entry and sets VPETE.VALID to 1, the behavior of the IRS Domain is UNPREDICTABLE. The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the Interrupt Domain or in an access to memory that does not follow the behaviors described in 4.12 IRS memory access rules .
R _{XRNSL}	When a write to IRS_VMAP_VPER makes the corresponding VPE valid, all of the following are true: <ul style="list-style-type: none"> • The IRS generates the following Effects: <ul style="list-style-type: none"> – A Register Write Effect E₁ to IRS_VMT_STATUSR which sets the IDLE field to 0. – A Memory Write Effect E₂ to the Location corresponding to the level 1 VM table entry which updates VPETE.VALID from 0 to 1. – A Register Write Effect E₃ to IRS_VMT_STATUSR which sets the IDLE field to 1. • There is an Intrinsic order dependency from the Register Write Effect E₁ to the Memory Write Effect E₂. • There is an Intrinsic order dependency from the Memory Write Effect E₂ to the Register Write Effect E₃. <p>This means that if there is a Memory Read Effect E₄ which Reads-from the Register Write Effect E₃, then E₂ is Ordered-before E₄.</p>
I _{SBCHK}	When a VPE table is valid, a VPE is made valid by the following sequence: <ol style="list-style-type: none"> 1. Software initializes the corresponding VPE table entry, with VPETE.VALID remaining to be 0. 2. Software writes to IRS_VMAP_VPER. 3. The IRS performs a write to the VPE table entry that updates VPETE.VALID from 0 to 1. 4. The IRS reports that the effects of the write to IRS_VMAP_VPER are complete. <p>See 3.3 Coherency considerations for GIC data structures for more information about ensuring that memory accesses are coherent between the IRS and a PE.</p>
I _{TTRHV}	When the VPE becomes valid, the corresponding VPETE.VPED_ADDR specifies the location of the VPE descriptor. See 4.9.2.3 The VPE descriptor for more information.
R _{YLKWF}	When a VPE table is valid, the IRS behavior is UNPREDICTABLE if there is a write to the VPE table entry when IRS_VMT_STATUSR.IDLE is 0 because of a write to IRS_VMAP_VPER. The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the Interrupt Domain or in an access to memory that does not follow the behaviors described in 4.12 IRS memory access rules .
R _{LJJVP}	If a write updates any field in a valid VPE table entry, the IRS Domain behavior is UNPREDICTABLE. The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the IRS Domain or in an access to memory that does not follow the behaviors described in 4.12 IRS memory access rules .
I _{LVSWT}	The architecture does not support making an individual VPE invalid without making the corresponding VM invalid.
R _{PKQLK}	To recover from UNPREDICTABLE behavior resulting from an incorrect write to a VPE table entry, the VM is made invalid.
D _{VDSVN}	A VPE <i>becomes invalid</i> when the corresponding VM becomes invalid.

4.9.2.3 The VPE descriptor

I _{YWLWM}	The VPE descriptor is memory provided by software to the IRSs to manage the VPE. The content and usage of the VPE descriptor is IMPLEMENTATION DEFINED. The VPE descriptor cannot be moved once the VPE is created, for as long as the VM remains valid.
R _{DLKTX}	When a VPE becomes valid, the memory pointed to by VPETE.VPED_ADDR is expected to contain all zeros. If the VPE descriptor does not contain all zeros, all of the following are true: <ul style="list-style-type: none"> It is CONSTRAINED UNPREDICTABLE whether the IRS Domain selects a virtual candidate HPPI for the VPE. The VPE configuration, including the doorbell settings, is reset to UNKNOWN values.
R _{DHCDH}	The format and content of the VPE descriptor is IMPLEMENTATION DEFINED.
I _{QPDNG}	The size of the VPE descriptor is reported by IRS_IDR4.VPED_SZ.
R _{TJMHJ}	The size of the VPE descriptor is the same across all Interrupt Domains, except for the EL3 Interrupt Domain, where VMs are not supported. The size of the VPE descriptor is the same across all IRSs.

4.9.2.4 Example VPE table structures

I _{LFWWJ}	The size of a VPE table entry is 8 bytes. Table 4.4 shows some example VPE table structures. The VP_ID_BITS column represents the value stored in the L2_VMTE.VPE_ID_BITS field for the VM containing the VPE table.
--------------------	--

Table 4.4: Example VPE table sizes

VPE_ID_BITS	VPE table size	Maximum number of VPEs
2	32 B	4
3	64 B	8
4	128 B	16
7	1 KB	128
9	4 KB	512
14	128 KB	16,384
16	512 KB	65,536

4.10 Virtual interrupts

D _{YJQXX}	Virtual interrupts are interrupts that belong to a VM.
R _{XKHNL}	A virtual interrupt is only permitted to be selected as the candidate HPPI for the Virtual Interrupt Domain.
I _{BNPQR}	The IRS supports the following types of virtual interrupts for VMs: <ul style="list-style-type: none"> • Virtual LPIs. • Virtual SPIs. <p>Both types of virtual interrupts are stored in virtual ISTs.</p>
I _{ZDNVV}	Virtual interrupts are routed to a target VPE and selected as the virtual candidate HPPI for the PE where the target VPE is resident. See 4.10.6 VPE residency and 4.10.3 Virtual interrupt routing for more information.
R _{KPKLS}	IRS support for virtualization is optional.
I _{FYJLC}	IRS_IDR0.VIRT indicates whether virtualization is supported.
R _{DGHCV}	IRS_IDR0.VIRT is 0 for the EL3 Interrupt Domain.
R _{KDSPG}	When at least one PE connected to an IRS implements EL2, IRS_IDR0.VIRT is 1 for all IRSs for all Interrupt Domains other than the EL3 Interrupt Domain. When none of the PEs in the system implement EL2, IRS_IDR0.VIRT is 0 for all IRSs for all Interrupt Domains.
R _{JJPQJ}	The value of IRS_IDR0.VIRT is the same for all IRSs for the same Interrupt Domain in the system.
R _{TQFCT}	When virtualization is not supported, the IRS does not select any virtual candidate HPPIs to the PEs and any virtual interrupts signaled to the IRS are IGNORED.
R _{TZBQY}	The number of Priority bits supported for virtual interrupts corresponds to the number of Priority bits supported in the Interrupt Domain where the VM is defined. See 4.6 Interrupt configuration and state for more information about the number of supported Priority bits.

4.10.1 Virtual LPIs

D _{JXYLG}	The virtual LPIs for a VM are stored in the <i>virtual LPI IST</i> for the VM.
I _{RMNHL}	The virtual LPI IST is allocated separately for each VM. The virtual LPI IST for a VM is shared across all IRSs in the system. A virtual LPI IST is not permitted to be shared across multiple VMs.
I _{WPHXH}	The number of virtual LPIs in a VM is configured in L2_VMTE.LPI_ID_BITS. Software allocates and provides the virtual LPI IST for a VM by writing a valid address to L2_VMTE.LPI_IST_ADDR and writing 1 to L2_VMTE.LPI_IST_VALID.
I _{WPZYD}	The virtual LPI IST is indexed using the INTID.ID in the range from 0 to $(2^{L2_VMTE.LPI_ID_BITS}) - 1$.
R _{CJYYQ}	When an IRS accesses the virtual LPI IST for a VM, the IRS does not access any memory location derived from an INTID which is outside the configured virtual LPI range for the VM. This includes situations where the INTID is programmed in the parameters to a GIC System instruction or the INTID is stored in the IST metadata areas.
R _{DJZRM}	A virtual LPI specified by an INTID for a VM specified by a VM ID is reachable when all of the following are true: <ul style="list-style-type: none"> • The VM ID is valid. • The INTID is within the configured virtual LPI range for the VM. • The virtual LPI IST for the VM is valid.

- There is a level 2 IST entry corresponding to the INTID in the virtual LPI IST.

Otherwise, the virtual LPI is unreachable.

I_{GQQSW}

In the [virtual LPI IST context](#), the base address of the virtual LPI IST for a VM is stored in L2_VMTE.LPI_IST_ADDR.

I_{JKPTM}

In the [virtual LPI IST context](#), L2_VMTE.LPI_IST_VALID determines if the virtual LPI IST is valid for a VM.

D_{WDHXY}

A virtual LPI IST for a VM *becomes valid* in one of the following scenarios:

- The VM becomes valid and L2_VMTE.LPI_IST_VALID is 1 in the corresponding level 2 VM table entry.
- The VM is valid and a write to IRS_VMAP_VISTR makes the virtual LPI IST valid.

See [4.7 The interrupt state table \(IST\)](#) for more information about an IST becoming valid and invalid.

R_{XKGC}

When a write to IRS_VMAP_VISTR makes the corresponding LPI IST for the VM valid or invalid, all of the following are true:

- The IRS generates the following Effects:
 - A Register Write Effect E₁ to IRS_VMT_STATUSR which sets the IDLE field to 0.
 - A Memory Write Effect E₂ to the Location corresponding to the level 2 VM table entry which updates L2_VMTE.LPI_IST_VALID.
 - A Register Write Effect E₃ to IRS_VMT_STATUSR which sets the IDLE field to 1.
- There is an Intrinsic order dependency from the Register Write Effect E₁ to the Memory Write Effect E₂.
- There is an Intrinsic order dependency from the Memory Write Effect E₂ to the Register Write Effect E₃.

This means that if there is a Memory Read Effect E₄ which Reads-from the Register Write Effect E₃, then E₂ is Ordered-before E₄.

I_{ZTSBX}

When a VM is valid, the virtual LPI IST is made valid by the following sequence:

1. Software configures the virtual LPI IST by writing to the corresponding level 2 VM table entry, with L2_VMTE.LPI_IST_VALID remaining to be 0.
2. Software writes to IRS_VMAP_VISTR.
3. The IRS performs a write to the level 2 VM table entry that updates L2_VMTE.LPI_IST_VALID from 0 to 1.
4. The IRS reports that the effects of the write to IRS_VMAP_VISTR are complete.

See [3.3 Coherency considerations for GIC data structures](#) for more information about ensuring that memory accesses are coherent between the IRS and a PE.

See [4.9.1.3 VM management](#) for more information about which fields are permitted to be updated when the VM is valid.

I_{BHBYF}

See [4.9.1.3 VM management](#) for more information about updates to a level 2 VM table entry and the use of the virtual map registers.

S_{KBXNL}

The architecture supports the following ways to configure virtual LPIs for a VM:

- Software can write a valid IST pointer and IST configuration data to the level 2 VM table entry after the VM is made valid and write to the map virtual IST register.
- Software can provide a valid IST pointer in the level 2 VM table entry when the VM is made valid.

Adding an IST to add virtual LPI support for a VM can, for example, be done in response to a guest operating system booting and configuring a virtual IRS emulated by software.

Providing a valid IST pointer at the time the VM is made valid can be used to support VM migration, or to support suspend and resume of VMs.

I_{SYDWQ}

When a virtual LPI IST becomes valid, the following fields are interpreted by the IRS Domain:

- L2_VMTE.LPI_ID_BITS
- L2_VMTE.LPI_IST_STRUCTURE
- L2_VMTE.LPI_IST_ADDR
- L2_VMTE.LPI_ISTSZ

I _{LTXVZ}	When IRS_IDR2.IST_levels is 0, L2_VMTE.LPI_IST_STRUCTURE is treated as 0.
I _{NMFLD}	<p>IRS_IDR2.{MIN_LPI_ID_BITS,ID_BITS} report the range of valid values for L2_VMTE.LPI_ID_BITS.</p> <p>In a multi-IRS system, when the virtual LPI IST is valid for a VM, and the VM table is valid, all IRSs are permitted to access the virtual LPI IST for any reason, including speculative reads.</p> <p>See 4.9.1.1 The VM table base address and configuration registers for more information.</p>
R _{TDTCTJ}	When the virtual LPI IST is invalid for a VM, the IRS Domain does not access the virtual LPI IST.
D _{PCVXZ}	<p>A virtual LPI IST for a VM <i>becomes invalid</i> in one of the following scenarios:</p> <ul style="list-style-type: none"> • The VM becomes invalid. • The VM is valid and a write to IRS_VMAP_VISTR makes the virtual LPI IST invalid. <p>See 4.7 The interrupt state table (IST) for more information about an IST becoming valid and invalid.</p>
I _{QFTNB}	<p>When a VM is valid, the virtual LPI IST is made invalid by the following sequence:</p> <ol style="list-style-type: none"> 1. Software writes to IRS_VMAP_VISTR. 2. The IRS performs a write to the level 2 VM table entry that updates L2_VMTE.LPI_IST_VALID from 1 to 0. 3. The IRS reports that the effects of the write to IRS_VMAP_VISTR are complete. <p>See 3.3 Coherency considerations for GIC data structures for more information about ensuring that memory accesses are coherent between the IRS and a PE.</p>
R _{CXWBT}	<p>If a virtual LPI IST for a VM becomes invalid, a VPE in the VM is resident on a PE, and a virtual LPI is the candidate HPPI for the Virtual Interrupt Domain, one of the following is true:</p> <ul style="list-style-type: none"> • The candidate HPPI is acknowledged by the PE before the IST becomes invalid. • The candidate HPPI is not acknowledged and is no longer selected as the candidate HPPI for the VPE.
R _{KYMTZ}	<p>When a virtual LPI IST becomes invalid, if there are Accepted events for virtual LPIs targeting the VM that are not yet processed, all of the following are true:</p> <ul style="list-style-type: none"> • The events are dropped and have no effect on any interrupt. • If software error reporting is supported, the IRS is permitted to report an error.
S _{BWSKR}	<p>If hypervisor software wishes to update the configuration of a virtual LPI IST, for example as the result of guest software programming emulated IRS registers, software should perform the following sequence:</p> <ol style="list-style-type: none"> 1. Make the virtual LPI IST invalid. 2. Update the relevant fields in the level 2 VM table entry. 3. Make the virtual LPI IST valid.
I _{PCNMM}	<p>In the virtual LPI IST context, when the virtual LPI IST is valid and uses a 2-level structure, a virtual level 2 IST is made valid by writing to IRS_VMAP_L2_VISTR.</p> <p>The write specifies that the map applies to the LPI IST by setting IRS_VMAP_L2_VISTR.TYPE to LPI.</p> <p>The effects of the write are complete when IRS_VMT_STATUSR.IDLE is 1.</p> <p>See 4.7.1 Level 2 IST management for more information.</p>
I _{SBDCY}	<p>See 4.9.1.3 VM management for more information about the restrictions for the use of IRS_VMAP_L2_VISTR concurrently with updates to the virtualization data structures for the VM and the use of other virtual map registers across IRSs.</p> <p>See also:</p> <ul style="list-style-type: none"> • 4.7 The interrupt state table (IST) • 4.9 Virtualization data structures

4.10.2 Virtual SPIs

D _{JSJKN}	The virtual SPIs for a VM are stored in the <i>virtual SPI IST</i> for the VM.
I _{YZPNL}	<p>The virtual SPI IST is allocated separately for each VM.</p> <p>The virtual SPI IST for a VM is shared across all IRSs in the system.</p> <p>A virtual SPI IST is not permitted to be shared across multiple VMs.</p>
I _{PVHMD}	<p>The number of virtual SPIs in a VM is configured in L2_VMTE.SPI_ID_BITS.</p> <p>Software allocates and provides the virtual SPI IST for a VM by writing a valid address to L2_VMTE.SPI_IST_ADDR and writing 1 to L2_VMTE.SPI_IST_VALID.</p>
I _{DXXHY}	The virtual SPI IST is indexed using the INTID.ID in the range from 0 through $(2^{L2_VMTE.SPI_ID_BITS}) - 1$.
R _{JQZKY}	<p>When an IRS accesses the virtual SPI IST for a VM, the IRS does not access any memory location derived from an INTID which is outside the configured virtual SPI range for the VM.</p> <p>This includes situations where the INTID is programmed in the parameters of a GIC System instruction executed on a PE or the INTID is stored in the IST metadata areas.</p>
R _{MYLWH}	<p>A virtual SPI specified by an INTID for a VM specified by a VM ID is reachable when all of the following are true:</p> <ul style="list-style-type: none"> • The VM is valid. • The virtual SPI IST for the VM is valid. • The INTID is within the configured virtual SPI range for the VM. • There is a level 2 IST entry corresponding to the INTID in the virtual SPI IST. <p>Otherwise, the virtual SPI is unreachable.</p>
I _{HRDGP}	In the virtual SPI IST context , the base address of the virtual SPI IST for a VM is stored in L2_VMTE.SPI_IST_ADDR.
I _{CSNJM}	In the virtual SPI IST context , L2_VMTE.SPI_IST_VALID determines if the virtual SPI IST is valid for a VM.
D _{BKKXD}	<p>A virtual SPI IST for a VM <i>becomes valid</i> in one of the following scenarios:</p> <ul style="list-style-type: none"> • The VM becomes valid and L2_VMTE.SPI_IST_VALID is 1 in the corresponding level 2 VM table entry. • The VM is valid and a write to IRS_VMAP_VISTR makes the virtual SPI IST valid. <p>See 4.7 The interrupt state table (IST) for more information about an IST becoming valid and invalid.</p>
R _{MLTBD}	<p>When a write to IRS_VMAP_VISTR makes the corresponding SPI IST for the VM valid or invalid, all of the following are true:</p> <ul style="list-style-type: none"> • The IRS generates the following Effects: <ul style="list-style-type: none"> – A Register Write Effect E₁ to IRS_VMT_STATUSR which sets the IDLE field to 0. – A Memory Write Effect E₂ to the Location corresponding to the level 2 VM table entry which updates L2_VMTE.SPI_IST_VALID. – A Register Write Effect E₃ to IRS_VMT_STATUSR which sets the IDLE field to 1. • There is an Intrinsic order dependency from the Register Write Effect E₁ to the Memory Write Effect E₂. • There is an Intrinsic order dependency from the Memory Write Effect E₂ to the Register Write Effect E₃. <p>This means that if there is a Memory Read Effect E₄ which Reads-from the Register Write Effect E₃, then E₂ is Ordered-before E₄.</p>
I _{CDTQL}	<p>When a VM is valid, the virtual SPI IST is made valid by the following sequence:</p> <ol style="list-style-type: none"> 1. Software configures the virtual SPI IST by writing to the corresponding level 2 VM table entry, with L2_VMTE.SPI_IST_VALID remaining to be 0. 2. Software writes to IRS_VMAP_VISTR. 3. The IRS performs a write to the level 2 VM table entry that updates L2_VMTE.SPI_IST_VALID from 0 to 1. 4. The IRS reports that the effects of the write to IRS_VMAP_VISTR are complete.

See [3.3 Coherency considerations for GIC data structures](#) for more information about ensuring that memory accesses are coherent between the IRS and a PE.

See [4.9.1.3 VM management](#) for more information about which fields are permitted to be updated when the VM is valid.

S_{LHQJY} Arm expects that hypervisor software provides a valid virtual SPI when a VM is made valid with a fixed number of SPIs for the lifetime of the VM. This is analogous to how SPIs work on a physical system.

I_{ZZGXF} When a virtual SPI IST becomes valid, the following fields are interpreted by the IRS Domain:

- L2_VMTE.SPI_ID_BITS
- L2_VMTE.SPI_IST_STRUCTURE
- L2_VMTE.SPI_IST_ADDR
- L2_VMTE.SPI_ISTSZ

I_{ZKTZX} When IRS_IDR2.IST_levels is 0, L2_VMTE.SPI_IST_STRUCTURE is treated as 0.

I_{RGNRJ} IRS_IDR2.ID_BITS reports the maximum value for L2_VMTE.SPI_ID_BITS.

I_{TWRCJ} See [4.9.1.3 VM management](#) for more information about updates to a level 2 VM table entry and the use of the virtual map registers.

In a multi-IRS system, when the virtual SPI IST is valid for a VM, and the VM table is valid, all IRSs are permitted to access the virtual SPI IST for any reason, including speculative reads.

See [4.9.1.1 The VM table base address and configuration registers](#) for more information.

R_{KYYR} When the virtual SPI IST is invalid for a VM, the IRS Domain does not access the virtual SPI IST.

D_{FYKVT} A virtual SPI IST for a VM *becomes invalid* in one of the following scenarios:

- The VM becomes invalid.
- The VM is valid and a write to IRS_VMAP_VISTR makes the virtual SPI IST invalid.

See [4.7 The interrupt state table \(IST\)](#) for more information about an IST becoming valid and invalid.

I_{PZQWR} When a VM is valid, the virtual SPI IST is made invalid by the following sequence:

1. Software writes to IRS_VMAP_VISTR.
2. The IRS performs a write to the level 2 VM table entry that updates L2_VMTE.SPI_IST_VALID from 1 to 0.
3. The IRS reports that the effects of the write to IRS_VMAP_VISTR are complete.

See [3.3 Coherency considerations for GIC data structures](#) for more information about ensuring that memory accesses are coherent between the IRS and a PE.

R_{JZQCG} If a virtual SPI IST for a VM becomes invalid, a VPE in the VM is resident on a PE, and a virtual SPI is the candidate HPPI for the Virtual Interrupt Domain, one of the following is true:

- The candidate HPPI is acknowledged by the PE before the IST becomes invalid.
- The candidate HPPI is not acknowledged and is no longer selected as the candidate HPPI for the VPE.

R_{SGLWK} When a virtual SPI IST becomes invalid, if there are Accepted events for virtual SPIs targeting the VM that are not yet **processed**, all of the following are true:

- The events are dropped and have no effect on any interrupt.
- If software error reporting is supported, the IRS is permitted to report an error.

I_{GVQDH} In the **virtual SPI IST context**, when the virtual SPI IST is valid and uses a 2-level structure, a virtual level 2 IST is made valid by writing to IRS_VMAP_L2_VISTR.

The write specifies that the map applies to the SPI IST by setting IRS_VMAP_L2_VISTR.TYPE to SPI.

The effects of the write are complete when IRS_VMT_STATUSR.IDLE is 1.

See [4.7.1 Level 2 IST management](#) for more information.

See also:

- [4.7 The interrupt state table \(IST\)](#)
- [4.9 Virtualization data structures](#)

4.10.2.1 Assigning physical SPIs to VMs

D_{WGVHQ}

A physical SPI is *assigned to a VM* when all of the following are true:

- A write sets IRS_SPI_VMR.VIRT to 1.
- The write specifies a valid VM.
- The write completes successfully as indicated by IRS_SPI_STATUSR.{V,IDLE} being {1,1}.

I_{KJJNL}

Assigning SPIs to a VM is only supported when IRS_IDR0.VIRT is 1.

When IRS_IDR0.VIRT is 0, IRS_SPI_VMR is RAZ/WI.

I_{BSYVR}

When an SPI is statically assigned to the EL3 Interrupt Domain, it is not possible to assign the SPI to a VM.

Access to IRS_SPI_VMR is RAZ/WI when the SPI is assigned to the EL3 Interrupt Domain.

R_{PYSML}

For each SPI that is not statically assigned to the EL3 Interrupt Domain, it is IMPLEMENTATION DEFINED whether the SPI can be assigned to a VM.

When support for assigning an SPI to a VM is implemented, the SPI is assigned to a VM in the Interrupt Domain where the SPI is assigned.

When support for assigning an SPI to a VM is not implemented, access to IRS_SPI_VMR, when that SPI is selected, is RAZ/WI.

I_{LKYQR}

Arm recommends that when IRS_IDR0.VIRT is 1, for all SPIs that are connected to an interrupt input signal, that support for assigning the SPI to a VM is implemented.

R_{NZFQQ}

When an SPI is assigned to a VM, all of the following are true for the physical SPI:

- The SPI is not considered for a physical candidate HPPI.
- The SPI is unreachable.

The effects on virtual SPIs to which physical SPIs are assigned are described in [4.10.4 Virtual interrupt signaling](#).

R_{NSTHN}

If a physical SPI is assigned to a VM, GIC System instructions that update the Pending state of the physical SPI are IGNORED.

R_{VYNBP}

When a physical SPI is assigned to a virtual SPI, the virtual SPI INTID is reset to its reset state.

If virtual 1ofN interrupt routing is not supported, the reset state of Routing mode for a virtual SPI is Targeted.

When a physical SPI is unassigned from a virtual SPI, all of the following are true:

- The physical SPI INTID is reset to its reset state.
- The virtual SPI INTID is reset to its reset state.

I_{BGSSV}

The Trigger mode of an SPI determines the type of the events that the IRS generates for the SPI, regardless of whether the SPI is assigned to a VM.

R_{HHBSB}

If IRS_SPI_VMR.VM_ID does not specify a valid VM for an SPI written to IRS_SPI_SEL, it is CONSTRAINED UNPREDICTABLE whether the SPI is treated as not being assigned to any VM, or assigned to an invalid VM.

If the SPI is treated as being assigned to an invalid VM, all of the following are true:

- The SPI is not considered for the physical or virtual candidate HPPI for any PE.
- The SPI is unreachable.

4.10.3 Virtual interrupt routing

D _{PMQVK}	<p>The Routing mode of a virtual interrupt determines how the target VPE is selected for that interrupt in the following ways:</p> <ul style="list-style-type: none"> • If the Routing mode is Targeted, the target VPE is the VPE specified by the interrupt Affinity. • If the Routing mode is 1ofN, the target VPE is determined dynamically.
D _{SZXHC}	<p>When an interrupt Routing mode is Targeted and the interrupt Affinity specifies a VPE, the interrupt is said to be <i>targeted to</i> that VPE.</p>
I _{VSDPM}	<p>If the Routing mode of a virtual interrupt is Targeted, the interrupt Affinity specifies the target VPE. If the Routing mode of a virtual interrupt is 1ofN, the interrupt Affinity is IMPLEMENTATION DEFINED.</p>
R _{BFGYB}	<p>IRS support for virtual 1ofN interrupt routing is optional.</p>
I _{JQMTB}	<p>IRS_IDR0.VIRT_ONE_N reports whether virtual 1ofN interrupt routing is supported.</p>
R _{XLGVV}	<p>The value of IRS_IDR0.VIRT_ONE_N is the same across all Interrupt Domains, except for the EL3 Interrupt Domain, where VMs are not supported.</p> <p>The value of IRS_IDR0.VIRT_ONE_N is the same across all IRSs.</p>
R _{JCSBY}	<p>If virtual 1ofN interrupt routing is not supported, all of the following are true:</p> <ul style="list-style-type: none"> • If a GIC System instruction generates an Interrupt Effect that sets the Routing mode of a virtual interrupt, the Routing mode remains set to Targeted regardless of the value provided in <Xt>. • L2_ISTE.IRM is treated as 0 when a virtual IST becomes valid.
S _{QQDND}	<p>When virtual 1ofN interrupt routing is not supported, the IRM field in the value encoded in the <Xt> parameter of the GIC <code>xDAFF</code> System instruction is RES0 and GIC <code>xDRCFG</code> instructions always return 0 in the IRM field in ICC_ICSR_EL1 when requesting the configuration of a virtual interrupt.</p> <p>This means that virtual 1ofN interrupts are treated as Targeted interrupts and are always delivered to the VPE specified by the interrupt Affinity.</p> <p>Hypervisor software can support virtual 1ofN interrupt routing in systems where virtual 1ofN interrupt routing is not implemented by emulating the 1ofN behavior.</p> <p>Hypervisor software can emulate virtual 1ofN by trapping guest access to the configuration of virtual interrupts and maintain a shadow configuration of each interrupt used by the physical IRS and updating the target VPE of the shadow 1ofN interrupts to emulate a 1ofN selection algorithm.</p>
R _{TBBTW}	<p>Each virtual 1ofN interrupt can be acknowledged by at most one VPE.</p>
R _{FHDFL}	<p>The storage location for the Affinity of virtual 1ofN interrupts may be used for an IMPLEMENTATION DEFINED purpose to handle routing of virtual 1ofN interrupts, with the restriction that an Affinity value of 0 has no IMPLEMENTATION DEFINED meaning.</p>
R _{FRMXL}	<p>An IRS is permitted to update the Affinity of virtual 1ofN interrupts.</p>
I _{BYXVC}	<p>When virtual 1ofN is supported, the VPE can be configured to have 1ofN PE selection enabled or disabled by accessing IRS_VPE_CR0.DPS.</p>
R _{SCWSJ}	<p>When virtual 1ofN is supported, an IRS does not choose a VPE which has 1ofN PE selection disabled as the target VPE for a 1ofN interrupt.</p>
R _{XMNZS}	<p>When virtual 1ofN is supported, the IRS Domain selects a target VPE for a reachable virtual 1ofN interrupt belonging to a valid VM in finite time, when all of the following are true:</p> <ul style="list-style-type: none"> • The interrupt is Enabled. • The interrupt is Pending. • The interrupt is Inactive. • At least one IRS is enabled. • At least one VPE in the VM has 1ofN PE selection enabled.

The target VPE is selected among the VPEs in the VM that the virtual 1ofN interrupt belongs to.

Note

The target VPE selection behavior does not guarantee that the interrupt is selected as the virtual candidate HPPI for a PE. For example, the selected target VPE may be masking the interrupt, the VPE may not be resident, another virtual interrupt may be the virtual candidate HPPI for the target VPE.

R _{SKVBZ}	An IRS Domain may select a new target VPE for a virtual 1ofN interrupt at any time.
R _{QYKSH}	When virtual 1ofN is supported, the mechanism that an IRS uses to select a target VPE among several possible target VPEs is IMPLEMENTATION DEFINED.
S _{RBXBB}	When virtual 1ofN is supported, the VPE 1ofN PE selection enable control may be used to support virtualization of IRS_PE_CR0.DPS. Arm expects that hypervisor software at EL2 traps the guest programming of the virtual IRS_PE_CR0.DPS field and configures the VPE accordingly.
I _{DWNMY}	When IRS_VPE_CR0.DPS is updated for a VPE that is resident on a PE, and a virtual 1ofN interrupt is selected as the virtual candidate HPPI for the VPE, one of the following is true: <ul style="list-style-type: none"> • The PE acknowledges the interrupt before IRS_VPE_STATUSR.IDLE is 1. • The PE does not acknowledge the interrupt. <p>See 4.10.4 Virtual interrupt signaling for more information.</p>

4.10.4 Virtual interrupt signaling

R _{GMMHX}	The Pending state of a virtual LPI identified by an INTID and a VM ID is updated when all of the following are true: <ul style="list-style-type: none"> • The virtual LPI as identified by the INTID is reachable. • The VM as identified by the VM ID is valid. • Any of the following occur: <ul style="list-style-type: none"> – An ITS generates an interrupt event that specifies the VM ID and INTID. – A GIC System instruction generates an Interrupt Write Effect that updates the Pending state of the virtual LPI. <p>Otherwise, the virtual LPI Pending state is not updated.</p>
I _{YDCMB}	The IRS SETLPI register does not support signaling virtual interrupts.
R _{CQTDI}	The Pending state of a virtual SPI identified by an INTID and a VM ID is updated when all of the following are true: <ul style="list-style-type: none"> • The virtual SPI as identified by the INTID is reachable. • The VM as identified by the VM ID is valid. • Any of the following occur: <ul style="list-style-type: none"> – An interrupt connected to an IRS as an SPI is asserted or de-asserted and generates an interrupt event for the virtual SPI because the SPI is assigned to the VM. – A GIC System instruction generates an Interrupt Write Effect that updates the Pending state of the virtual SPI. <p>Otherwise, the virtual SPI Pending state is not updated.</p>
I _{BZFTG}	When a virtual interrupt is signaled to the IRS, the IRS updates the Pending state of the virtual interrupt regardless of whether the target VPE is resident or non-resident.

D _{HCZPD}	<p>The IRS selects a virtual candidate HPPI for a VPE for an Interrupt Domain, when at least one virtual interrupt meets the virtual <i>candidate HPPI conditions</i>:</p> <ul style="list-style-type: none"> • The interrupt is Enabled. • The interrupt is Pending. • The interrupt is Inactive. • One of the following is true: <ul style="list-style-type: none"> – The interrupt Routing mode is Targeted and the interrupt Affinity specifies the VPE. – The interrupt Routing mode is 1ofN and the VPE is selected as the target VPE for the interrupt.
R _{BFYMP}	<p>For a virtual Targeted interrupt whose interrupt Affinity specifies an invalid VPE that is updated from invalid to valid, whether the interrupt is considered for the virtual candidate HPPI selection for the VPE depends on whether there is an update to the following states and configurations for the virtual interrupt:</p> <ul style="list-style-type: none"> • Pending. • Active. • Enabled. • Routing mode. • Affinity. <p>If there is no update to any of the mentioned states and configurations of the virtual interrupt after the VPE becomes valid, it is CONSTRAINED UNPREDICTABLE whether the interrupt is considered for the virtual candidate HPPI selection for the VPE.</p> <p>If there is an update to any of the mentioned states and configurations for the virtual interrupt after the VPE becomes valid, the virtual interrupt is considered for the virtual candidate HPPI selection for the VPE:</p>
R _{CPCKK}	<p>For a VPE in a VM, all of the following apply to the virtual candidate HPPI selection:</p> <ul style="list-style-type: none"> • If at least one interrupt satisfies the virtual candidate HPPI conditions, the IRS selects one of these as the virtual candidate HPPI in finite time. • If no interrupts satisfy the virtual candidate HPPI conditions, the IRS does not select any interrupt as the virtual candidate HPPI.
R _{ZDTYM}	<p>If more than one interrupt satisfies the virtual candidate HPPI conditions for a VPE in a VM, the IRS selects the interrupt with the highest Priority as the virtual candidate HPPI in finite time.</p> <p>If there is more than one virtual candidate HPPI with the same Priority, it is IMPLEMENTATION DEFINED which of those interrupts is selected as the virtual candidate HPPI.</p>
R _{VBPZM}	<p>An IRS selects a virtual candidate HPPI for a connected PE in finite time, when all of the following are true:</p> <ul style="list-style-type: none"> • There is a resident VPE on the PE. • The IRS is enabled for the Physical Interrupt Domain corresponding to the resident VPE. • There is at least one virtual interrupt that meets the virtual <i>candidate HPPI conditions</i> for the VPE.
I _{WZNJR}	<p>If the IRS Domain has selected a virtual candidate HPPI for a non-resident VPE, it might generate a VPE doorbell. See 4.10.7 VPE doorbells for more information.</p>
R _{XBNGN}	<p>When a virtual candidate HPPI that has been selected for a VPE is no longer the virtual candidate HPPI for the VPE, and if the VPE is resident on a PE, one of the following is true:</p> <ul style="list-style-type: none"> • The PE acknowledges the old virtual candidate HPPI within finite time. • The old virtual candidate HPPI is not acknowledged.
R _{ZFVRC}	<p>If the IRS updates the configuration of an interrupt that has been selected as the virtual candidate HPPI for the Virtual Interrupt Domain for a PE, and the interrupt configuration data was communicated to the PE, the IRS will select the same or a new virtual interrupt as the virtual candidate HPPI.</p> <p>The IRS then communicates the selected virtual candidate HPPI to the PE, including any updated interrupt configuration data, in finite time.</p>

For more information about ordering of such events, see [2.12 Interrupt ordering model and synchronization requirements](#) and [Chapter B1 Interrupt ordering litmus tests](#).

See also:

- [4.8.4 Physical interrupt signaling](#)

4.10.5 VPE selection and configuration

I_{BNNMF}

Software accesses the configuration and samples the HPPI of a VPE by selecting the VPE and accessing one of the following registers:

- IRS_VPE_CR0.
- IRS_VPE_DBR.
- IRS_VPE_HPPIR.

Software selects a VPE by writing the VM ID and VPE ID to IRS_VPE_SELR.{VM_ID,VPE_ID} and writing 1 to IRS_VPE_SELR.S.

IRS_VPE_STATUSR.IDLE reports whether the effects of the write IRS_VPE_SELR.S are complete.

IRS_VPE_STATUSR.V reports whether the value written to IRS_VPE_SELR selected a valid VPE.

IRS_VPE_STATUSR.V is updated on a write to any of the following registers:

- IRS_VPE_CR0.
- IRS_VPE_DBR.
- IRS_VPE_HPPIR.
- IRS_VPE_SELR.

R_{MVDZK}

If a VPE is selected on more than one IRS in a multi-IRS system, and a write occurs to any of the VPE configuration registers on one IRS while IRS_VPE_STATUSR.IDLE is 0 on any other IRS for the same IRS Domain due to a write to the configuration registers, it is CONSTRAINED UNPREDICTABLE which one of the following is true:

- The write does not successfully update the VPE and IRS_VPE_STATUSR.V is 0 on one of the IRSs where a write occurs.
- None of the writes updates the VPE configuration.
- Both of the writes successfully update the VPE configuration in UNKNOWN order.

R_{RLHSL}

When a VPE becomes valid, the VPE configuration is reset to the following values:

- If virtual 1ofN is supported, the VPE can receive both Targeted and 1ofN interrupts.
- If virtual 1ofN is not supported, the VPE can only receive Targeted interrupts.
- There are no doorbell settings for the VPE.

I_{MFSZW}

The configuration of a VPE when it becomes valid is consistent with a read of IRS_VPE_CR0.DPS returning 0, assuming the VPE is selected in IRS_VPE_SELR and there was no write to IRS_VPE_CR0 since the VPE became valid.

I_{HHWBZ}

A write to any of the VPE configuration registers is only guaranteed to have successfully completed when IRS_VPE_STATUSR.{V,IDLE} is {1,1}.

If the VPE becomes invalid after the VPE is selected, a write to any of the VPE configuration registers results in IRS_VPE_STATUSR.V being 0.

I_{RXTXP}

A read of IRS_VPE_HPPIR returns whether there is a candidate HPPI for the selected VPE.

If there is a candidate HPPI for the selected VPE, IRS_VPE_HPPIR.{TYPE,ID} returns the INTID of the candidate HPPI.

I_{PNRVY}

IRS_VPE_HPPIR reports the virtual candidate HPPI for a VPE. When virtual 1ofN interrupt routing is not supported, the reported candidate HPPI is a targeted interrupt. When virtual 1ofN interrupt routing is supported, the reported candidate HPPI may be either a targeted or 1ofN virtual interrupt.

4.10.6 VPE residency

D _Y TZXF	The IRS tracks, for each VPE, whether the VPE is resident VPE , and if it is resident, which PE it is resident on. If a VPE is not resident on any PE, the VPE is <i>non-resident</i> .
I _L GVMO	A VPE is made resident or non-resident on a PE as a result of a write to ICH_CONTEXTR_EL2.
R _Z FXNN	If a write to ICH_CONTEXTR_EL2 to make a VPE resident specifies an invalid VPE, the write does not cause a VPE to be considered resident by the IRS.
R _J JKJV	If a PE requests that a VPE is made resident on a PE when the VPE is already resident on another PE, it is CONSTRAINED UNPREDICTABLE which one of the following applies: <ul style="list-style-type: none"> • The request is IGNORED. • The request completes and all of the following are true: <ul style="list-style-type: none"> – It is CONSTRAINED UNPREDICTABLE which of the PEs are allowed to acknowledge and configure interrupts from the VPE context. – It is CONSTRAINED UNPREDICTABLE whether making a VPE non-resident from some or all of the PEs results in the VPE being resident on any PEs or non-resident.

To recover from the CONSTRAINED UNPREDICTABLE behavior, software makes the VPE non-resident on all PEs where it was previously made resident, and the VPE can subsequently be made resident on a single PE.

See also:

- [4.9 Virtualization data structures](#)

4.10.7 VPE doorbells

G _M GLGB	The architecture provides a mechanism to signal a physical interrupt when a virtual interrupt is available for a non-resident VPE. This allows a hypervisor to make scheduling decisions about when to run a VPE.
D _J NQYP	A <i>VPE doorbell</i> is the mechanism that generates a physical interrupt event when the VPE doorbell conditions are met.
R _C WZMW	A VPE doorbell event is generated for a VPE, when all of the following <i>VPE doorbell conditions</i> are true: <ul style="list-style-type: none"> • The VPE is not resident on any PE. • The VPE doorbell settings are valid for the VPE. • A VPE doorbell is requested for the VPE. • Any of the following are true: <ul style="list-style-type: none"> – There is at least one Targeted virtual interrupt for the VPE that is Pending, Inactive, and Enabled. – The 1ofN doorbell conditions are met for the VM and the VPE is the 1ofN VPE doorbell target. See 4.10.8 1ofN doorbells for more information. • The interrupt causing the doorbell event to be generated has a priority greater than or equal to the doorbell priority mask.

Note

The VPE doorbell priority mask is checked both when there is a Targeted virtual interrupt causing the VPE doorbell to be generated, and when the VPE doorbell is generated because the 1ofN doorbell conditions are met.

This means that when the 1ofN doorbell conditions are met, they only cause a VPE doorbell event to be generated when the virtual 1ofN interrupt Priority is greater than or equal to the doorbell priority mask.

R _P CXCN	When a VPE is created, the doorbell settings for the VPE are not valid. The doorbell settings become valid when the VPE is selected using IRS_VPE_SEL and a write to IRS_VPE_DBR occurs.
---------------------	--

R _{KDQNS}	When support for LPIs is implemented, VPE doorbells generate SET_EDGE targeting physical LPIs in the same Interrupt Domain as where the VPE is defined. Otherwise, the mechanism to signal a VPE doorbell is IMPLEMENTATION DEFINED.
I _{GBLPY}	An LPI used for a VPE doorbell can be configured the same way as any other LPI including configuring its Priority, Enabled value, Routing mode, and Affinity when relevant.
I _{KTNWK}	The doorbell settings of a VPE can be accessed by selecting the VPE by writing its VM ID and VPE ID to IRS_VPE_SEL and accessing IRS_VPE_DBR.
I _{MZXXV}	The interrupt used for the VPE doorbell is programmed by writing to IRS_VPE_DBR.INTID.
D _{MGQPJ}	A VPE doorbell is <i>requested</i> when IRS_VPE_DBR.REQ_DB is 1.
I _{CVJVS}	A VPE doorbell is requested in any of the following ways: <ul style="list-style-type: none"> When software makes a VPE non-resident, it can program whether a VPE doorbell is requested. Software can access IRS_VPE_DBR.REQ_DB using the VPE selection and configuration interface on an IRS.
R _{WTZVF}	Once a VPE doorbell event is generated, the doorbell is no longer requested for the VPE.

Note

The above rules mean that if a VPE doorbell is **requested** and a corresponding event is generated, an event is only generated once, until it is **requested** again, because the doorbell stops being **requested** when an event is generated.

S _{RGZBY}	Once a VPE doorbell event has been generated, no further VPE doorbell events will be generated for that VPE until the VPE doorbell is requested again. The VPE doorbell may be requested again by hypervisor software after making the VPE resident and subsequently making it non-resident.
I _{SRMLF}	IRS_VPE_DBR.REQ_DB reflects the value of ICH_CONTEXTR_EL2.DB when the VPE was last made non-resident, or the last write to IRS_VPE_DBR.REQ_DB, whichever happened last. Once a VPE doorbell event is generated, this bit is cleared.
I _{BWQVY}	A doorbell priority mask is set when the PE makes a VPE non-resident and requests a VPE doorbell. The IRS Domain only generates a VPE doorbell if there is a virtual interrupt for the VPE with a priority greater than or equal to the minimum priority programmed in the doorbell priority mask, accessible via IRS_VPE_DBR.DBPM.
S _{WWFWZ}	When making a VPE non-resident, ICH_CONTEXTR_EL2.DBPM sets the minimum priority for a virtual interrupt to trigger the VPE doorbell. Software can use this field to ensure that doorbells are only generated for virtual interrupts which would be handled by the VPE if it were resident.
R _{DXSYV}	When the VPE doorbell conditions are met, the VPE doorbell interrupt event is generated in finite time. If the VPE doorbell conditions are met for a short amount of time and are then no longer met, for example due to a reconfiguration of an interrupt, and the VPE doorbell event was not yet generated, it is IMPLEMENTATION DEFINED whether a VPE doorbell event is generated.
R _{RFWPN}	If a VPE doorbell event is generated for an INTID that is unreachable, the event is treated as any other interrupt event for an unreachable INTID and the event does not make any interrupt Pending.
I _{XJQLG}	If a VPE doorbell event is generated for an INTID that is unreachable, the doorbell is no longer requested for the VPE.
I _{DTTDM}	If a VPE doorbell event is generated for an INTID that is unreachable, and software error reporting is supported, this error is reported in IRS_SWERR_STATUSR.

R_{DHHRM} If a VPE doorbell uses an LPI where any of the following are true, it is **CONSTRAINED UNPREDICTABLE** whether the IRS makes the LPI Pending when the VPE doorbell event is generated, and whether any IRS synchronization requests affect the LPI:

- There is an ITS mapping to the LPI.
- There is another VPE doorbell that uses the LPI.

See also:

- [2.10.6 Requesting VPE doorbells](#)

4.10.8 1ofN doorbells

G_{THLBD} The architecture supports generating a VPE doorbell event when there is a virtual 1ofN interrupt for a VM and there are no resident VPEs with 1ofN PE selection enabled for the VM. This allows a hypervisor to make scheduling decisions about when to run a VPE to handle a 1ofN interrupt.

D_{FJTLC} A *1ofN doorbell* is the mechanism that generates a VPE doorbell when the 1ofN doorbell conditions are met. The VPE doorbell is generated for the *1ofN VPE doorbell target*.

I_{JRYKN} The 1ofN doorbell configuration for a VM is managed by selecting the VM and accessing **IRS_VM_DBR**. Software selects a VM by writing the VM ID to **IRS_VM_SELR.VM_ID**.

IRS_VM_STATUSR.IDLE reports whether the effects of the write **IRS_VM_SELR** are complete.

IRS_VM_STATUSR.V reports whether the value written to **IRS_VM_SELR** selects a valid VM.

IRS_VM_STATUSR.V is updated on a write to any of the following registers:

- **IRS_VM_SELR**.
- **IRS_VM_DBR**.

I_{NWBYP} Whether 1ofN doorbells are enabled or disabled for the VM is programmed in **IRS_VM_DBR.EN**.

I_{FKZMP} The 1ofN VPE doorbell target is programmed in **IRS_VM_DBR.VPE_ID**.

D_{QTNVT} The *1ofN doorbell conditions* are met when all of the following are true:

- Virtual 1ofN interrupt routing is supported.
- There is no VPE belonging to the VM where all of the following are true:
 - The VPE is resident.
 - The VPE has 1ofN PE selection enabled.
- 1ofN doorbells are enabled for the VM.
- There is at least one virtual 1ofN interrupt for the VM that is Pending, Inactive, and Enabled.

See [4.10.7 VPE doorbells](#) for more information about additional conditions for generating a VPE doorbell event.

R_{VVLVP} When a VM is made valid, 1ofN doorbells are disabled for the VM.

4.10.9 Save and restore of virtual interrupts

G_{NSYHS} The architecture is designed to support saving the state and configuration of all virtual interrupts in a VM to memory, and restoring the saved state and configuration from memory.

This is to support common virtualization use cases such as VM live migration across physical hosts, snapshotting of VM state, and provisioning of VMs from a pre-booted state.

D_{CZQCL} This chapter refers to saving the state of a *source* VM and restoring it to a *destination* VM.

This is regardless of whether the source and destination VMs are running on separate physical machines, are separate VMs with distinct VM IDs on the same physical machine, or are separate instances of a VM executing at different times with the same VM ID on the same physical machine.

D_{MGVXX}

A VM is *Quiesced* if all of the following are true:

- No VPEs in the VM are resident.
- There is no change in the Pending state of any virtual interrupt for the VM.
- There is no change to the number of valid PEs in the VM.
- The VM does not become invalid.

I_{G_NLWX}

To Quiesce a VM, software should ensure that all of the following are true:

- The Pending states of LPis are not updated by events from an ITS.
- The Pending states of SPIs are not updated by events generated by the IRS.
- The Pending states of LPis or SPIs are not updated as a result of executing the GIC \times DPEND instruction on a PE.
- No VPE is resident.
- The virtual LPI IST and virtual SPI IST do not become valid or invalid.
- No VPE in the VM is made valid.
- The VM is not made invalid.

S_{JVVJM}

Software can prevent updates to virtual interrupts belonging to a VM when saving the state of virtual interrupts by using any of the following mechanisms:

- Preventing devices from generating interrupts for the VM.
- Unmapping ITS events that are mapped to virtual interrupts belonging to the VM.
- Issuing a synchronization request on every ITS that had events mapped to virtual interrupts belonging to the VM.
- Unassigning SPIs that are assigned to virtual SPIs in the VM.
- Issuing a synchronization request on all IRSs in the system.

S_{JKLZX}

To support saving the state of a source VM and restoring it to a destination VM, software is expected to perform the following steps:

1. Quiesce the source VM.
2. Save the state and configuration of all virtual interrupts to the virtual ISTs.
3. Save a copy of the virtual IST from the source VM.
4. Scan the copy of the virtual IST and record which virtual interrupts are Pending.
5. Update all entries in the copy of the virtual IST as follows:
 1. Set the L2_ISTE.Pending field to 0.
 2. Set IAFFID to 0 if IRM is 1 and IRS_IDR0.ONE_N is 1.
 3. Set the metadata area to 0 if IRS_IDR2.ISTMD is 1.
6. Configure the destination VM to use the copy of the virtual IST as the new IST.
7. Make the new virtual IST valid.
8. For each virtual interrupt that was recorded as Pending, make the virtual interrupt Pending using the GIC VDPEND System instruction.

The destination VM must be configured with valid VPEs before making the virtual interrupts Pending to ensure that the interrupts are considered for the virtual candidate HPPI for the VPEs.

I_{ZKFGL}

The state and configuration of virtual interrupts is written to the virtual ISTs by writing the VM ID to IRS_SAVE_VMR.VM_ID and writing 1 to IRS_SAVE_VMR.S when the VM is Quiesced.

The operation is complete when IRS_SAVE_VM_STATUSR.IDLE is 1.

Following a write that sets IRS_SAVE_VMR.S or IRS_SAVE_VMR.Q to 1, IRS_SAVE_VM_STATUSR.Q reports whether the VM is Quiesced since the last write that set IRS_SAVE_VMR.S to 1.

Note

Saving the state and configuration of virtual interrupts applies to the ISTs that are valid for the VM. For example,

if a VM does not have a virtual SPI IST, but does have a valid virtual LPI IST, writing 1 to IRS_SAVE_VMR.S saves the state of the virtual LPIs to the virtual LPI IST.

R _{XPMBB}	Following a write that sets IRS_SAVE_VMR.S to 1, for as long as the specified VM is Quiesced, the values stored in the virtual ISTs are consistent with the state and configuration of the virtual interrupts.
I _{DBSGD}	Saving the state and configuration of virtual interrupts is requested on any IRS and saves the state of all virtual interrupts for the VM to the virtual IST. In a multi-IRS system, an implementation coordinates this operation across all IRSs.
R _{CFDQJ}	If a write occurs to IRS_SAVE_VMR when IRS_SAVE_VM_STATUSR.IDLE is 0 on any other IRS for the same IRS Domain, it is CONSTRAINED UNPREDICTABLE whether the state and configuration of virtual interrupts are saved to the virtual ISTs and the value returned in IRS_SAVE_VM_STATUSR.Q is UNKNOWN on both IRSs.
I _{SYNGG}	Following a write that sets IRS_SAVE_VMR.S to 1, the state and configuration of virtual interrupts are stored in the virtual ISTs using the format defined in 11.2.4 L2_ISTE, Level 2 interrupt state table entry .

4.11 IRS power management

G _{CRTH}	<p>The architecture supports implementations that enable the following power management operations:</p> <ul style="list-style-type: none"> • System suspend and resume: All interrupt state and configuration, and register state of all IRSs in the system for all Interrupt Domains, is stored to memory on suspend and restored from memory on resume. • Opportunistic power management: The system may decide to power down one or more IRSs while preserving the illusion to software that the IRSs are fully operational.
R _{FYSXL}	<p>Support for system suspend and resume is supported using an IMPLEMENTATION DEFINED sequence.</p> <p>The IMPLEMENTATION DEFINED sequence must support saving all of the following to memory for all Interrupt Domains:</p> <ul style="list-style-type: none"> • Physical LPI state and configuration. • Physical SPI state and configuration. • If virtualization is supported in the Interrupt Domain, GIC-managed state related to the VM table, including all of the following: <ul style="list-style-type: none"> – State and configuration of virtual interrupts stored in the virtual ISTs. – State stored in the VM descriptors, including the configuration of VMs. – State stored in the VPE descriptors, including the configuration of VPEs. <p>The IMPLEMENTATION DEFINED sequence must support restoring all of the information from memory to the values saved during suspend.</p>
I _{XHGLK}	Data structures used by mechanisms for IRS save and restore for system suspend are IMPLEMENTATION DEFINED.
S _{NWXQM}	Arm expects that system suspend and resume is managed by system-specific software.
R _{HJFWL}	Support for opportunistic power management is IMPLEMENTATION DEFINED. The opportunistic power management sequence must preserve the illusion to software that the IRSs remain powered on.
D _{BGYZW}	<p>Each PE connected to an IRS is either <i>offline</i> or <i>online</i>.</p> <p>Arm expects that a PE is offline when the PE is physically powered off and is unable to execute GIC System instructions or write to GIC System registers.</p>
I _{TCRZH}	Whether a PE is online or offline is reported in IRS_PE_STATUSR.ONLINE for the PE selected in IRS_PE_SEL.
R _{PMSHG}	The mechanism to detect whether a PE is offline or online is IMPLEMENTATION DEFINED.
D _{XGKCK}	When a PE is offline, the IRS generates a <i>Wake Request</i> to a PE to bring it online.
R _{CLDWD}	The mechanism through which an IRS generates a Wake Request to a PE is IMPLEMENTATION DEFINED.
R _{HPCRY}	<p>When the IRS generates a Wake Request to a PE, the request is generated in finite time.</p> <p>The Wake Request does not guarantee that the PE becomes online. The Wake Request is treated as a hint to the power control subsystem.</p>
I _{FYYJF}	<p>When a PE becomes offline, and the PE is not expected to become online when receiving a Wake Request, Arm strongly recommends that the system is configured to avoid sending Wake Requests to the PE.</p> <p>The system is configured not to generate Wake Requests to the PE when all of the following are true for all IRS Domains:</p> <ul style="list-style-type: none"> • No Targeted interrupts are programmed to target the PE. • 1ofN PE selection is disabled for the PE. <p>An example of when a PE is not expected to become online when receiving a Wake Request is PSCI_CPU_OFF. See <i>Arm® Power State Coordination Interface</i>[8] for more information about PSCI.</p>
R _{YPBHL}	<p>When a PE is offline, all of the following are true:</p> <ul style="list-style-type: none"> • The PE does not execute any instructions that generate Interrupt Effects.

- If the IRS has selected a physical candidate HPPI for a PE for any Interrupt Domain, and the IRS is enabled for that Interrupt Domain, the IRS generates a Wake Request to the PE.

When a PE is online, the IRS does not generate any Wake Requests to the PE.

R_{FVFKC}

When a PE becomes offline and the IRS has selected any candidate HPPIs for the PE, each candidate HPPI is either acknowledged before the PE becomes offline or is not acknowledged while the PE is offline.

When the PE subsequently becomes online, any selected candidate HPPIs for the PE for enabled Interrupt Domains, can be acknowledged by the PE.

4.12 IRS memory access rules

R _{VHRSJ}	All IRS data structures are little-endian.
R _{FRXXQ}	The IRS accesses all its data structures using Normal memory types. See [1] for more information about memory access types.
S _{XFKQH}	Software allocates memory for the data structures and programs the physical address of the allocated data structures in IRS registers. Arm expects that such allocations will typically be done from Conventional memory.
I _{DPJZK}	<p>IRS_CR1 specifies the Shareability and Cacheability attributes for memory accesses performed by the IRS.</p> <p>Changes to IRS_CR1 can only be made when the IRS does not access memory, specifically when no data structures are valid on the IRS.</p> <p>The most recent values written to IRS_CR1 are used when the IRS accesses memory following a write to make any data structure valid.</p> <p>Arm recommends that IRS_CR1 is initialized with the same Shareability and Cacheability attributes on all IRSs in the system for an IRS Domain before making any data structures valid.</p>
R _{XKMVR}	<p>An IRS does not access any memory location which is not derived from address and configuration data stored in registers and structures belonging to the IRS Domain.</p> <p>When a memory address is stored in an address field in a register or data structure entry, and that register or entry includes a field indicating that the address is not valid, the IRS does not derive any memory location from that address field.</p> <p>This includes scenarios in which the IRS behavior is CONSTRAINED UNPREDICTABLE, UNPREDICTABLE, or IMPLEMENTATION DEFINED.</p>
I _{GTTMT}	<p>When a memory address is stored in a data structure, and that entry includes a field indicating whether the address is valid, the address only becomes invalid as a result of a write to an IRS map register.</p> <p>In this case, IRS Domain is guaranteed to not derive an address from the entry once the effects of the write to the map register are complete.</p>
R _{CPXND}	<p>When the IRS accesses an IRS data structure, and the contents of that data structure are not defined as UNKNOWN at the time of the access, the access must be a 64-bit aligned, single-copy atomic access of at least 64 bits in size.</p> <p>For data structures where the entries are smaller than 64 bits, accesses are permitted to be smaller than 64 bits, provided that each access is:</p> <ul style="list-style-type: none"> • Single-copy atomic, • exactly the size of the entry, and • aligned to the size of the entry.
R _{LRRFV}	<p>If an IRS access to any IRS data structure is performed into any PCIe address space, then the IRS is permitted to return an UNKNOWN value or terminate the access and report the error.</p> <p>In this case, the IRS is permitted to stop the operation as follows:</p> <ul style="list-style-type: none"> • If the operation was due to an incoming interrupt event, the IRS is permitted to drop the interrupt event. • If the operation was requested as the result of executing a GIC System instruction or writing to a GIC System register, the IRS is permitted to indicate a failure back to the PE or IGNORE the operation.
R _{RFYLF}	<p>When an IRS accesses a memory location in a PAS, it relies on information stored in registers and data structures structures that are accessible only within the same PAS to validate that the access is permitted.</p> <p>This includes scenarios in which the IRS behavior is CONSTRAINED UNPREDICTABLE, UNPREDICTABLE, or IMPLEMENTATION DEFINED.</p>

For example, an IRS access to a data structure is permitted only if the memory location lies within a region defined by a base address and size held in a register or another data structure entry, where that defining information resides in the same PAS as the access itself.

This prevents software running in a Security state from directing the IRS to access any memory location using another PAS than the one associated with the software's Security state.

I _{FJRHT}	<p>If the IRS experiences an external abort during a memory access to an IRS data structure, the IRS stops the operation.</p> <p>If software error reporting is supported, the error is reported with IRS_SWERR_STATUSR.EC.</p>
R _{LVNHK}	If the IRS experiences an external abort during a memory access to an IRS data structure, it is IMPLEMENTATION DEFINED whether the IRS reports a RAS error using an IMPLEMENTATION DEFINED mechanism.
R _{GZJNW}	If the IRS experiences an external abort as part of an operation that occurs as a result of a GIC System instruction or write to a GIC System register on a PE, the effects of the instruction are UNKNOWN and the IRS may optionally report an error.
I _{LQCPD}	<p>For example, if the IRS experiences an external abort when trying to access the IST as part of handling a GIC \rightarrowxDDIS instruction, the instruction is permitted to execute as a NOP and the IRS is permitted to report an error using the software error reporting mechanism and by raising IMPLEMENTATION DEFINED RAS errors if the abort is caused by a memory system error.</p>
I _{DPBMS}	<p>As another example, if the IRS experiences an external abort when trying to access the IST as part of handling a GIC \timesDRCFG instruction, the IRS is permitted to indicate a failure to the PE that results in ICC_ICSR_EL1.F being set to 1, and the IRS may optionally report an error.</p>
R _{KMZHR}	If the IRS experiences an external abort as part of processing an interrupt that is being signaled, the effects on the state and configuration of the signaled interrupt is UNKNOWN and the IRS is permitted to drop the event. The IRS may optionally report an error using the software error reporting mechanism and by raising IMPLEMENTATION DEFINED RAS errors if the abort is caused by a memory system error.
R _{ZRWQP}	<p>If an IRS data structure overlaps with any other IRS or ITS data structure, the IRS behavior is UNPREDICTABLE. This applies to both IRS data structures for physical and virtual interrupts.</p> <p>The UNPREDICTABLE behavior may result in loss of interrupt configuration and state, but must not result in access to memory outside the PAS associated with the IRS Domain or in access to any memory location which is not derived from address and configuration data in valid IRS registers.</p>

Software can recover from this situation by performing the following sequence:

1. Disabling the ITS.
2. Making the corresponding data structures invalid in the IRS.
3. Reconfigure the base addresses of the data structure to avoid any overlap.

See also [5.3 Translation structures](#).

I _{QTYKR}	<p>The following changes to system state when an interrupt event is processed by an IRS are not required to result in a write to memory:</p> <ul style="list-style-type: none"> • Interrupt Write effects generated by interrupt events. • Interrupt Write effects generated by GIC System instructions executed on a PE. • Configuration changes generated by GIC System instructions executed on a PE. • Configuration changes generated by writes to IRS registers.
--------------------	---

In these cases, the change to system state does not correspond directly to any Memory Write effect to IRS data structures.

However, the change may cause an update to architecturally invisible IRS caches and Memory Write effects to write back cached data to IRS data structures where the content of the data structure is described as UNKNOWN.

R_{FZBVM}

An IRS is permitted to implement caching of interrupt state, VMs, and VPE configurations using architecturally invisible IRS-specific caches.

The IRS caches are not considered data or system caches belonging to the memory system, meaning they are not managed by any cache maintenance instructions issued by a PE.

When an IRS data structure is invalid, the IRS does not cache previously accessed data in IRS-specific caches.

4.13 IRS support for MPAM

I_{LKQJP}

The Memory System Resource Partitioning and Monitoring Memory, MPAM, architecture defines per-transaction attributes that affect system behavior or the behavior of components that the transactions pass through or Completers that satisfy a transaction [1].

The additional attributes are two identifiers:

- Partition ID, or PARTID.
- Performance Monitoring Group, or PMG.

The PARTID and PMG are both interpreted within a PARTID space. The PARTID space used depends on the Security state associated with an IRS Domain, and the PARTID and PMG values used may be programmed independently for each IRS Domain. See [1] for more information about MPAM information bundles.

I_{WXHXM}

In the IRS, support for MPAM, is indicated in IRS_IDR0.MPAM.

If MPAM is supported, the supported PARTID and PMG width is indicated in IRS_MPAM_IDR.{PARTID_MAX,PMG_MAX}.

R_{RMZLS}

If MPAM is supported, IRS accesses to memory are associated with the PARTID and PMG programmed in IRS_MPAM_PARTID_R.{PARTID,PMG}.

I_{NSQVT}

In systems without support for RME, the PARTID space used by the IRS is determined by the memory system attribute MPAM_NS.

In systems with support for RME, the PARTID space used by the IRS is determined by the memory system attribute MPAM_SP.

I_{HHRJR}

The IRS architecture has optional support for MPAM PARTID space selection indicated by IRS_MPAM_IDR.HAS_MPAM_SP.

If IRS_MPAM_IDR.HAS_MPAM_SP is 0, each IRS for each Interrupt Domain uses a default MPAM PARTID space.

The following table shows the MPAM PARTID space used for accesses made by the IRSs for each Interrupt Domain if IRS_MPAM_IDR.HAS_MPAM_SP is 0 and the system does not support RME:

Interrupt Domain	MPAM PARTID space
Secure	Secure PARTID space
Non-Secure	Non-secure PARTID space
EL3	Secure PARTID space

The following table shows the MPAM PARTID space used for accesses made by the IRSs for each Interrupt Domain if IRS_MPAM_IDR.HAS_MPAM_SP is 0 and the system supports RME:

Interrupt Domain	MPAM PARTID space
Secure	Secure PARTID space
Non-Secure	Non-secure PARTID space
EL3	Root PARTID space.
Realm	Realm PARTID space.

If IRS_MPAM_IDR.HAS_MPAM_SP is 1, the IRS uses the MPAM PARTID specified by IRS_MPAM_PARTID_R.MPAM_SP.

R_{NNWLJQ}

If an IRS without support for MPAM is integrated in a system that supports MPAM, the PARTID and PMG used for each IRS for each Interrupt Domain is IMPLEMENTATION DEFINED.

4.14 IRS support for Memory Encryption Contexts

I _{FPYMT}	The Memory Encryption Contexts feature, FEAT_MEC, provides finer-grained memory encryption contexts, within the Realm physical address space, to be assigned to Realms, with policy controlled by Realm EL2 [1].
G _{HJFY}	The GICv5 architecture ensures that an IRS can function correctly in systems with support for Memory Encryption Contexts (MEC).
I _{HJNXC}	<p>In a system with support for MEC, data may be shared between the PEs and an IRS in the following scenarios:</p> <ul style="list-style-type: none"> • A PE allocates and initializes data structures used by an IRS. • An IRS writes the contents of virtual ISTs to memory to support VM migration. <p>In such a system, all IRS Realm data structures are managed and configured by Realm EL2, and virtualized software running at EL1 is not expected to have direct access to the IRS configuration registers. Therefore, the IRS only supports configuration of a single MECID used for all IRS memory accesses to the Realm PAS. This allows software running at Realm EL2 to initialize data structures that are read by the IRS and allows the IRS to write the content of the virtual ISTs that are read by software running at Realm EL2.</p>
I _{YQFXX}	<p>Support for Memory Encryption Contexts (MEC) for an IRS is indicated in IRS_IDR0.MEC for the Realm Interrupt Domain. If the MEC feature is supported, the supported MECID width is indicated in IRS_MEC_IDR.MECIDSIZE for the Realm Interrupt Domain.</p> <p>Arm strongly recommends that the MECID bit width supported by the IRS matches or exceeds the width supported by the PEs in the system.</p>
R _{YGVPR}	If the MEC feature is supported, IRS accesses to memory are associated with a MECID that identifies the Memory Encryption Context of the access.
R _{QKQOW}	Accesses made by the IRS for Secure, Non-secure, and Root PA spaces are issued with the default MECID of zero.
I _{TKDGT}	Accesses made by the IRS to Realm PA space are associated with the global Realm PAS IRS MECID programmed in IRS_MEC_MECID_R.MECID.
R _{KCCQN}	If an IRS without support for the Realm Interrupt Domain is integrated in a system that supports MEC, all IRS accesses for that IRS are treated as having the default MECID of zero.
R _{GMPJJ}	<p>In a multi-IRS system, if all of the following are true, the IRS Domain behavior is UNPREDICTABLE:</p> <ul style="list-style-type: none"> • One of the following is true: <ul style="list-style-type: none"> – IRS_IST_BASER.VALID is 1. – IRS_VMT_BASER.VALID is 1. • An IRS uses a different MECID than another IRS, for the same PA space. <p>The UNPREDICTABLE behavior must not result in access to memory outside the PAS associated with the Interrupt Domain or in an access to memory that does not follow the behaviors described in 4.12 IRS memory access rules.</p> <p>To recover from the UNPREDICTABLE behavior, the IRS data structures are made invalid and the same MECID is configured across all IRSs for the same PA space.</p>

4.15 IRS support for software error reporting

I _{GZRSS}	The IRS specifies a mechanism to report errors because of incorrect programming.
R _{QWVQB}	<p>The IRS detects software errors when it performs any of the following operations:</p> <ul style="list-style-type: none"> • The IRS processes an interrupt event in response to a signaled interrupt. • The IRS accesses state and configuration of interrupts in response to any of the following events: <ul style="list-style-type: none"> – A system register access in the PE interface. – Execution of a GIC system instruction in the PE interface. – An MMIO access via the IRS configuration register frame.
I _{TQQTs}	<p>Updates to the state and configuration of interrupts via the PE interface or the MMIO interface can be cached in the IRS caches. The IRS could experience an external abort, when it commits the cached information to the IRS data structures in memory at a specified address. The IRS uses the software error reporting mechanism to indicate such errors.</p> <p>For example, an update to the configuration of a physical LPI is cached in the IRS caches. The address of the physical IST in IRS_IST_BASER.ADDR is invalid, and the IRS experiences an external abort when it commits the cached information to the L2_ISTE in memory. The IRS sets IRS_SWERR_STATUSR.EC to report this error.</p> <p>Due to the effects of caching, this error could be reported after the operation to update the configuration of the physical LPI has completed from software's perspective. This means that two reads of IRS_SWERR_STATUSR can report errors without any interrupts or GIC System instructions executed on a PE having been processed by the IRS between the two reads.</p>
R _{YHKJZ}	IRS support for software error reporting is optional.
I _{NMHQJ}	IRS_IDR0.SWE reports whether software error reporting is supported.
R _{HWQTK}	The value of IRS_IDR0.SWE is the same for all Interrupt Domains in an IRS in the system.
I _{ZTBKN}	When IRS_IDR0.SWE is 1, the IRS uses the IRS_SWERR_STATUSR.IMP_EC to report any IMPLEMENTATION DEFINED errors detected by the IRS.
I _{NLLSH}	IRS_SWERR_STATUSR.EC is 0 when an IMPLEMENTATION DEFINED error is reported by the IRS.
I _{YMQZW}	<p>Software error information can be read on each IRS for each Interrupt Domain via the following registers:</p> <ul style="list-style-type: none"> • IRS_SWERR_STATUSR. • IRS_SWERR_SYNDROMER0. • IRS_SWERR_SYNDROMER1.
I _{FCYCT}	When a software error is reported, the value of IRS_SWERR_STATUSR.V is 1. Otherwise, no software error is reported and fields in this register are UNKNOWN.
I _{ZLMMJ}	<p>When the value of IRS_SWERR_STATUSR.V is 1, all of the following are true:</p> <ul style="list-style-type: none"> • IRS_SWERR_STATUSR.EC specifies the fault that caused the software error. • IRS_SWERR_STATUSR.S0V indicates whether IRS_SWERR_SYNDROMER0 contains valid error syndrome information. • IRS_SWERR_STATUSR.S1V indicates whether IRS_SWERR_SYNDROMER1 contains valid error syndrome information. • IRS_SWERR_STATUSR.OF indicates whether multiple software errors were detected.
I _{DMCGM}	<p>When IRS_SWERR_STATUSR.S0V is 1, IRS_SWERR_SYNDROMER0 reports the following information for the software error:</p> <ul style="list-style-type: none"> • The type and ID of the interrupt that resulted in the software error. • Whether the interrupt is virtual or physical. • If the interrupt is virtual, the VM ID of the VM that the interrupt is assigned to.

<code>I_HQSCK</code>	When <code>IRS_SWERR_STATUSR.S1V</code> is 1, <code>IRS_SWERR_SYNDROMER1</code> reports the address of the IRS data structure associated with the software error.
<code>R_XBYRF</code>	For each reported error, the values of <code>IRS_SWERR_STATUSR.{S0V,S1V}</code> are IMPLEMENTATION DEFINED and set independently.
<code>S_MTHCH</code>	<p>Arm recommends that software performs the following sequence to either clear the last error, or detect whether new errors were reported:</p> <ol style="list-style-type: none"> 1. Read <code>IRS_SWERR_STATUSR</code> and determine which fields need to be cleared to zero. 2. In a single-copy atomic write to <code>IRS_SWERR_STATUSR</code>: <ol style="list-style-type: none"> 1. Write ones to all the W1C fields that are nonzero in the read value. 2. Write zero to all the W1C fields that are zero in the read value. 3. Write zero to all the RW fields. 3. Read back <code>IRS_SWERR_STATUSR</code> after the write. If the value read back is the same as the value that was written, all W1C fields that were non-zero are cleared, and no new errors were reported. Otherwise, one or more new errors were reported.

See also:

- [10.2.1.32 `IRS_SWERR_STATUSR`](#).
- [10.2.1.33 `IRS_SWERR_SYNDROMER0`](#).
- [10.2.1.34 `IRS_SWERR_SYNDROMER1`](#).

Chapter 5

Interrupt translation service (ITS)

D_{MBMKY}

The Interrupt Translation Service (ITS) is responsible for generating *ITS events* and translating them into *interrupt events* for delivery to the associated Interrupt Routing Service (IRS) in each Interrupt Domain.

Each translation specifies whether the resulting interrupt is physical or virtual, the associated Interrupt Domain, the interrupt's INTID, and, if the interrupt is virtual, the VMID of the target VM.

The ITS uses translation structures in the form of a Device Table (DT) and Interrupt Translation Tables (ITT). These translation structures reside in memory and may be partially cached by the ITS. For more information, see [5.3 Translation structures](#).

[Figure 5.1](#) illustrates the overall ITS translation flow.

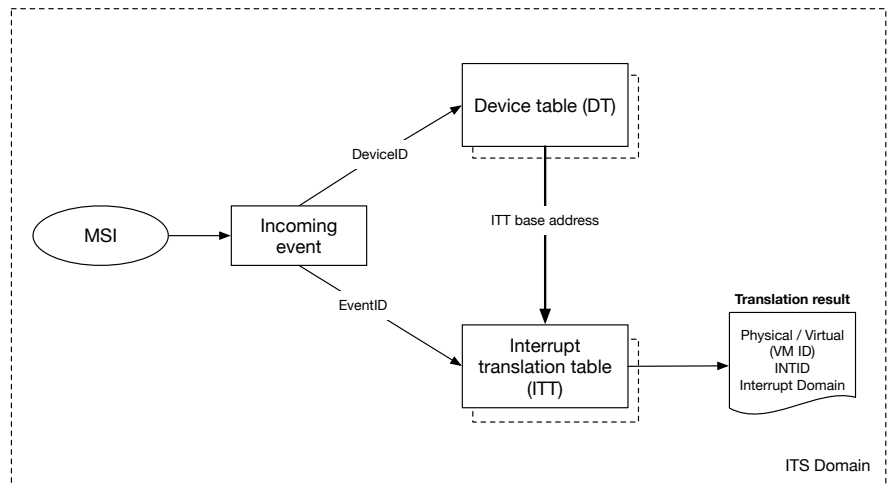


Figure 5.1: GICv5 ITS translation flow

I_{VHDBC} Translation of an ITS event may involve internal cache lookups and memory accesses to retrieve translation data.

I_{QQNWW} An example of the ITS translating an event into an interrupt occurs when a Message Signaled Interrupt (MSI) is generated by a device. In this case, the MSI write targets an ITS translation register. The ITS captures the write and generates an ITS event. This event contains the DeviceID of the originator of the write and the EventID value provided in the write data. The ITS uses the programmed entries in the Device Table (DT) and Interrupt Translation Table (ITT) to translate the ITS event into an interrupt. The resulting interrupt is then forwarded to the associated IRS for delivery to a PE.

R_{ZZQCC} An ITS event is generated by the ITS when any of the following occurs:

- A memory write to ITS_TRANSLATER or ITS_RL_TRANSLATER.
- A write to ITS_GEN_EVENTR.
- An IWB forwards an event using an IMPLEMENTATION DEFINED communication mechanism.
- A system peripheral forwards an event using an IMPLEMENTATION DEFINED communication mechanism.

D_{XCCZK} An ITS event contains all of the following information:

- DeviceID.
- EventID.
- [Interrupt event type](#).
- Originating Interrupt Domain.

D_{ZVRSC} The *originating Interrupt Domain* of an ITS event is determined as follows:

- If the event is generated by a write to ITS_TRANSLATER, ITS_RL_TRANSLATER, or ITS_GEN_EVENTR, the originating Interrupt Domain is the ITS Domain associated with the PAS of the register.
- If the event is generated by an IWB, the originating Interrupt Domain is the Interrupt Domain that the corresponding wire is assigned to.
- If the event is generated by a system peripheral using an IMPLEMENTATION DEFINED mechanism, the originating Interrupt Domain is determined using an IMPLEMENTATION DEFINED mechanism.

Note

The originating Interrupt Domain may not correspond to the ITS Domain that translates the event. See [5.1 ITS Domains](#) for more information.

R_{RTVBF} When multiple SET_LEVEL and CLEAR events are received for the same DeviceID and EventID, they must be translated and forwarded to the IRS in the order in which they are received.

R _{BZTJM}	When the ITS receives a memory write transaction to ITS_TRANSLATER, the ITS generates an ITS event that is translated by the ITS in finite time.
I _{NBTNK}	ITS events can be Message Signaled Interrupts (MSIs), for example generated by PCIe endpoints, or can be events generated by IWBs or other system components. Arm expects that MSIs are typically generated as a result of an MSI source performing a write to ITS_TRANSLATER.
S _{HFCND}	Software ensures that a PCIe posted write to an ITS_TRANSLATER_FRAME is Accepted by the ITS, by ensuring that the PCIe subsystem issues a Completion TLP following the posted MSI write. See <i>PCI Express® Base Specification Revision 6.0</i> [9] for more information on the PCIe TLP ordering rules.
R _{CHRYR}	For an incoming memory write transaction to ITS_TRANSLATER Accepted by the ITS, the ITS handles the incoming write independently from the state of the source.
I _{CNGJZ}	The ITS Accepts an incoming memory write transaction to ITS_TRANSLATER without waiting for the IRS to Accept a corresponding interrupt event and without waiting for the IRS to present the corresponding interrupt to a PE. The IRS separates the process of making an interrupt Pending from signaling the interrupt to the PE. See Chapter 4 Interrupt routing service (IRS) for more information.
I _{YDBCQ}	Interrupt events sent by an ITS to the associated IRS contain all of the following: <ul style="list-style-type: none"> • LPI ID. • Physical or virtual interrupt. • VM ID if virtual. • Interrupt event type. • Interrupt Domain.
R _{NYJRF}	An Interrupt event is never sent to the associated IRS speculatively. It is only sent when an ITS event is generated and translated by the ITS.
I _{PTZXC}	The mechanism used by the ITS to communicate events to the IRS is IMPLEMENTATION DEFINED. See 3.2 Communication between GIC system components for more information.
R _{HPDKY}	An ITS is associated with a single IRS and all interrupt events generated by an ITS are sent to that IRS.
R _{DXDFC}	The DeviceID in an ITS event is a unique identifier within the DeviceID namespace of each ITS. The DeviceID is unique for each device that can generate events to the ITS. For example, Arm expects that the 16-bit Requester ID from a PCIe Root Complex is used to derive the DeviceID.

Note

Each ITS has a separate namespace for the DeviceID which is shared across the Interrupt Domains.

R _{DJZVX}	The number of DeviceID bits supported by an ITS is in the range from 0 to 32 inclusively, and is reported by ITS_IDR1.DEVICEID_BITS.
I _{VCJFT}	An ITS can support 0 DeviceID bits, which means that the ITS only supports a single device.
R _{WXGDP}	The DeviceID namespace is unique for each ITS. For example, DeviceID 0 on ITS A identifies a device distinct from the device identified by DeviceID 0 on ITS B.
R _{CMZTW}	The number of EventID bits supported by an ITS is in the range from 0 to 16 inclusively, and is reported by ITS_IDR2.EVENTID_BITS.

R _{WCDFX}	<p>The EventID namespace is separate for each ITS DeviceID.</p> <p>For example, EventID 0 for DeviceID 0 on ITS A identifies an event distinct from the event identified by EventID 0 for DeviceID 1 on ITS A.</p>
I _{TNTVS}	<p>The number of EventID bits supported for a DeviceID is programmed in L2_DTE.EVENTID_BITS. See 5.3.2 The Interrupt Translation Table (ITT) for more information.</p>
R _{RYJGM}	<p>The mechanism by which a DeviceID is communicated to the ITS is IMPLEMENTATION DEFINED.</p> <p>However, a unique DeviceID is provided for each requesting device, and the DeviceID is presented to the ITS when a write to the ITS translation registers occurs in a manner that cannot be spoofed by any agent capable of performing writes.</p>
G _{TTLHL}	<p>The ITS translation mechanism supports <i>interrupt isolation</i>.</p> <p>This allows an untrusted software agent such as a virtual machine to directly control the EventID used by an interrupt source without being able to generate interrupts with INTIDs used for other interrupt sources.</p>
I _{RWJXV}	<p>Interrupt isolation can be achieved through the use of an ITS by only allowing the hypervisor to control the ITS. The hypervisor therefore controls the translation of each DeviceID and EventID value pair into an INTID for the Physical Interrupt Domain or a specified VM. In this scenario, the ITS only translates EventIDs programmed by the hypervisor for a given DeviceID, and since the DeviceID is not under control of the VM, the hypervisor fully controls which INTIDs can be signaled irrespective of the EventID programmed by the VM.</p>
R _{QWVTC}	<p>For the interrupt isolation mechanism to work, the DeviceID used to identify a device context assigned to untrusted software must be static and not under direct or indirect control of the untrusted software.</p>
I _{WRXRW}	<p>Each ITS has a unique ITSID which is reported in ITS_IDR0.ITSID. The ITSID is unique for each ITS in the system.</p>

5.1 ITS Domains

G _{QYXKS}	An ITS operates independently for each Interrupt Domain.
D _{WNGLR}	The GICv5 architecture defines an <i>ITS Domain</i> which provides translation services for an Interrupt Domain. An <i>ITS Domain</i> comprises register state and translation structures independently configured from other ITS Domains on the same ITS. An ITS Domain is associated with an Interrupt Domain and its corresponding PAS.
I _{PXYHF}	Each ITS Domain implements a separate ITS_CONFIG_FRAME. The ITSID reported in ITS_IDR0.ITSID is the same for each ITS_CONFIG_FRAME that belongs to the same ITS.
R _{PMQOB}	The ITS Domain that translates an ITS event is determined as follows: <ul style="list-style-type: none"> • If the ITS event is generated by a write to ITS_TRANSLATER, the ITS Domain is the originating Interrupt Domain. • If the ITS event is generated by a write to ITS_GEN_EVENTR, the ITS Domain is the ITS Domain specified by ITS_GEN_EVENTR.TARGET_DOMAIN. • If the ITS event is generated by a write to ITS_RL_TRANSLATER, the ITS Domain is the Realm ITS Domain. • If the ITS event is generated by an IWB, the ITS Domain is the originating Interrupt Domain. • If the event is generated by a system peripheral using an IMPLEMENTATION DEFINED mechanism, the ITS Domain is the ITS Domain that corresponds to the originating Interrupt Domain.
R _{YJNYY}	If an ITS is associated with an IWB, the ITS provides an ITS Domain for each Interrupt Domain implemented by the IWB.
R _{DZYMW}	An ITS does not support any Interrupt Domain not supported by the IRSs.
I _{XCBRD}	The IRS associated with the ITS implements support for all the ITS Domains provided by the ITS, because an IRS implements support for all the Interrupt Domains supported by the PEs in the system.
I _{KFVMT}	An ITS exposes separate control register frames for each ITS Domain.
I _{KKYFF}	An ITS exposes separate translation register frames for each ITS Domain.
I _{NJHMG}	An ITS implements at least one ITS translate register frame, ITS_TRANSLATE_FRAME, for each ITS Domain. If an ITS Domain implements more than one translate register frame, the DeviceID space is shared across all the translate register frames belonging to the same ITS.
R _{FGHHH}	Each ITS Domain can have a maximum of 256 ITS translate register frames.
R _{TZMYP}	Each ITS translate register frame has an 8 bit identifier that is unique within an ITS Domain.
I _{CXPWJ}	The ITS translation register frame identifier is used by software when issuing ITS synchronization requests. See 5.2.4 ITS synchronization requests for more information about synchronization requests.
S _{HPMNL}	Arm expects that the ITS translation register frame identifier, as well as the relationship between ITS translate register frames and an ITS Domain and its configuration register frame, is communicated to system software via firmware data structures.
I _{TNYVZ}	The ITS does not provide a mechanism to identify via registers which ITS translation register frames are associated with which ITS configuration register frames.
I _{GCJFM}	Different from GICv3, the GICv5 architecture does not require the control register frame and translation register frame to be contiguous with respect to each other.
I _{YCFZC}	The base address of each register frame for each ITS Domain should be provided to software by firmware.
I _{XFTTF}	The ITS Domain for a register frame is reported in ITS_IDR0.INT_DOM.

- R_{PSXLG}** All memory accesses performed by the ITS are subject to GPT checks.
An access is performed using the PAS associated with the ITS Domain where the translation is requested.
- I_{DRTDZ}** [Figure 5.2](#) shows an overview of the ITS Domain registers and translation structures.

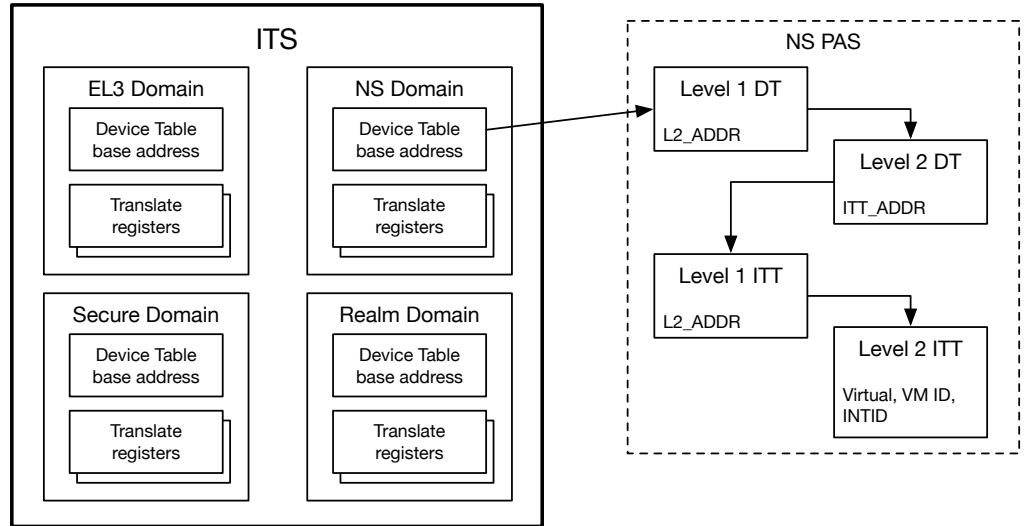


Figure 5.2: GICv5 ITS Domains

Translation structures are only shown for the Non-secure domain but exist separately for each ITS Domain.

- I_{SRZCQ}** Arm expects an interrupt source to be configured to use the translation register frame of an ITS Domain corresponding to the Security state and PAS that manages the device. This allows end-to-end isolation of interrupt events and interrupt configuration within a Security state. For example, a device managed by Software in the Secure state can generate MSIs to the Secure ITS Domain in the Secure PAS, and only software with access to the Secure PAS can configure this interrupt, and only that specific device can generate a Secure interrupt with the configured INTID.

- R_{SXLJQ}** The following ITS Domains support virtual interrupts:
- Non-secure.
 - Realm.
 - Secure.

The EL3 ITS Domain does not support virtual interrupts.

- R_{NTTVJ}** An event can only be translated to generate interrupt events for a VM in the Interrupt Domain corresponding to the ITS Domain. This is a consequence of the VM ID namespace being separate in each Interrupt Domain.

See also:

- [3.1 Interrupt Domains](#)

5.1.1 Supporting Realm interrupts from Non-secure writes

- G_{YLHMP}** The GICv5 architecture supports direct injection of MSIs generated by Non-secure devices to Realms.
- I_{BGWGH}** An ITS event is translated in the Realm ITS Domain, even though the event is generated as a result of a write to a register in the Non-secure PAS, if any of the following are true:
- The event is generated in response to a write to ITS_RL_TRANSLATER.
 - The event is generated in response to a write to ITS_GEN_EVENTR in the Non-secure PAS, and the write sets ITS_GEN_EVENTR.TARGET_DOMAIN to 0b01.

<code>I_JKXQV</code>	<p>Software in the Realm Security state controls whether an ITS event generated by a write to an ITS register in the Non-secure PAS is permitted to be translated in the Realm ITS Domain by programming <code>L2_ITTE.DAC</code> to allow such a translation.</p> <p>See 5.3.2 The Interrupt Translation Table (ITT) for more information about the DAC field.</p>
<code>S_GYSSD</code>	<p>When an untrusted PCIe function is assigned to a VM, its MSI-X vectors are programmed to generate MSI writes to <code>ITS_TRANSLATER</code>. The MSI writes use the Non-secure PAS.</p> <p>The ITS events are generated as a result of the write to a register in the Non-secure Interrupt Domain and translated in the Non-secure ITS Domain. If the translation is successful, LPIs are generated in the Non-secure Interrupt Domain.</p> <p>When the function is assigned to Realm, its MSI-X vectors can be programmed to generate MSI writes to the <code>ITS_RL_TRANSLATER</code>. The MSI writes still use the Non-secure PAS, however, they are translated in the Realm ITS Domain. If the translation is successful, LPIs are generated in the Realm Interrupt Domain.</p>
<code>S_SZVBR</code>	<p>On a system that implements the Realm Management Extension, for an untrusted device that is assigned to Realm, there are scenarios where the host hypervisor must be able to inject a virtual LPI into Realm. The host hypervisor and the target VPE of Realm could be running on different PEs.</p> <p>For example, an untrusted <code>VirtIO[10]</code> device is assigned to Realm. The device is emulated in the host hypervisor. The host hypervisor must be able to inject virtual LPIs into Realm.</p> <p>The host hypervisor can write to <code>ITS_GEN_EVENTR</code> with <code>TARGET_DOMAIN</code> set to <code>0b01</code> in the Non-secure PAS to generate these ITS events translated by the Realm ITS Domain.</p> <p>See also:</p> <ul style="list-style-type: none"> • 5.3.2 The Interrupt Translation Table (ITT). • 10.3.2.2 ITS_RL_TRANSLATER. • 11.1.4 L2_ITTE, Level 2 interrupt translation table entry.

5.2 Operation

R_{SXWCM}

When ITS_CR0.ITSEN is 1, the ITS Domain is *enabled* and all of the following are true:

- The ITS translates ITS events and may generate corresponding interrupt messages to the associated IRS.
- The ITS Domain may perform memory accesses to the ITS Domain's translation structures as part of processing ITS events or as a result of handling ITS register accesses.

When ITS_CR0.ITSEN is 0 and ITS_CR0.IDLE is 1, the ITS Domain is *disabled* and all of the following are true:

- The ITS does not process ITS events; ITS events are dropped and have no effects on the ITS.
- The ITS Domain does not generate outgoing interrupt events to the associated IRS.
- The ITS Domain does not perform memory accesses to the ITS Domain's translation structures.

R_{RNQH}

An ITS event is ignored if it cannot be translated by the ITS.

I_{XFFRW}

If software error reporting is supported, the ITS reports if an ITS event cannot be translated using the appropriate error code in ITS_SWERR_STATUSR.EC.

I_{HYKFL}

A change to ITS_CR0.ITSEN is not guaranteed to be observed by the ITS until ITS_CR0.IDLE is 1.

See also:

- [5.2.1 Enabling and disabling the ITS](#)

5.2.1 Enabling and disabling the ITS

R_{FXDRJ}

When a write to ITS_CR0.ITSEN changes the value from 0 to 1, the ITS Domain begins a transition from disabled to enabled, and all of the following are true:

- The transition completes in finite time.
- The transition is complete when ITS_CR0.IDLE is 1.
- The ITS will process all events that are Accepted after the transition is complete.

Note

While the ITS is transitioning from disabled to enabled, transactions arriving at the ITS will observe the ITS being enabled or disabled. The architecture only guarantees transactions arriving after the transition from disabled to enabled is complete to observe that the ITS is enabled.

R_{YSGFF}

When a write to ITS_CR0.ITSEN changes the value from 1 to 0, the ITS Domain begins a transition from enabled to disabled, and all of the following are true:

- The transition completes in finite time.
- The transition is complete when ITS_CR0.IDLE is 1.
- All ITS events, Accepted before the transition is complete, are processed.
- All outgoing interrupt events generated for the associated IRS before the transition is complete, are Accepted by the IRS.

Note

While the ITS is transitioning from enabled to disabled, transactions arriving at the ITS will observe the ITS as being enabled or disabled. The architecture only guarantees that transactions arriving after the transition from enabled to disabled is complete to observe that the ITS is disabled.

- S_{DNGJG}** When ITS_CR0.ITSEN is 1, the ITS accesses the translation structures described in [5.3 Translation structures](#). Software is expected to initialize the base address and configuration registers for the structures before enabling the ITS.
- For example, software can provide a device table where all entries are invalid, or it can be a device table with valid entries containing the addresses of one or more valid ITTs, and initialize the ITS_DT_BASER and ITS_DT_CFGR to point to the device table.
- S_{ZFWSP}** When the ITS Domain is disabled, the ITS Domain no longer accesses translation structures and software can reclaim the memory used to hold the translation structures.
- See also:
- [5.3 Translation structures](#)

5.2.2 Interrupt event types

- D_{CBDTG}** The *type* of an interrupt event generated by the ITS for an IRS is one of the following:
- SET_EDGE: The IRS should make the interrupt Pending and set its Handling mode to Edge.
 - SET_LEVEL: The IRS should make the interrupt Pending and set its Handling mode to Level.
 - CLEAR: The IRS should make the interrupt Idle.
- R_{JFGLF}** A write to ITS_TRANSLATER generates a SET_EDGE event to the ITS.
- I_{LKLYH}** The ITS only supports MSI sources that generate interrupts with edge-triggered semantics. MSI sources that are designed to signal interrupts with level-sensitive semantics using message based writes are not supported by the ITS.
- R_{PQCKZ}** SET_LEVEL events are only supported for events generated by an IWB.
- R_{TDXLH}** The interrupt event type generated to the IRS is the same as the ITS event type.
- I_{LQZSR}** When a write updates an ITS translation structure entry, the ITS does not generate a CLEAR event as a result of the update to the translation structures. For EventIDs that represent Level interrupts and that had a valid translation before the update, if those EventIDs do not have a valid translation after the update, the corresponding interrupts may still be in the Pending state in the IRS. See [5.3 Translation structures](#) for more information.
- S_{KSFMZ}** Software must disable or explicitly clear the Pending state of a Level interrupt in the IRS after unmapping the event in the ITS to avoid that the interrupt is signaled to a PE. An Edge interrupt may also remain Pending after being unmapped in the ITS, but its Pending state would be cleared when the interrupt is acknowledged.

5.2.3 Software generated ITS events

- I_{QTNPX}** Software may generate an ITS event for an EventID and DeviceID in an ITS Domain via the following ITS registers:
- ITS_GEN_EVENT_DIDR. A write to this register selects the DeviceID of the event.
 - ITS_GEN_EVENT_EIDR. A write to this register selects the EventID of the event.
 - ITS_GEN_EVENTR. The TARGET_DOMAIN field in this register selects the ITS Domain where the ITS event is translated.
 - A write of 1 to the R field requests the ITS to generate an ITS event of the selected type for the specified EventID and DeviceID.
 - ITS_GEN_EVENT_STATUSR. A read of this register reports when generation of an ITS event is complete.
- S_{VKLMK}** An event generated by writing to ITS_GEN_EVENTR may be synchronized once the write is complete to ensure that the corresponding interrupt event has been Accepted by the IRS by writing 1 to ITS_SYNCR.SYNC and polling ITS_SYNC_STATUSR.IDLE until it is 1.

5.2.4 ITS synchronization requests

I_{TWLSC}

Software requests synchronization of ITS events for the ITS by writing 1 to ITS_SYNCR.SYNC.

ITS_SYNCR.SYNCALL specifies whether the synchronization request should apply to all Accepted ITS events or is only required to apply to those specified by ITS_SYNCR.DEVICE_ID.

See [10.3.1.28 ITS_SYNCR](#) for more information about the effects of a write to ITS_SYNCR.SYNC.

S_{KDJBY}

Software can use synchronization of events to ensure that the Pending state of Level interrupts is cleared by the IRS.

This is, for example, useful when changing the Interrupt Domain association of a wire in the IWB.

See also:

- [3.2 Communication between GIC system components](#)
- [4.5 IRS synchronization requests](#)
- [6.2 IWB support for multiple Interrupt Domains](#)

5.3 Translation structures

<code>I_KBTLK</code>	<p>The GICv5 architecture defines two translation structures which can be used by an ITS:</p> <ul style="list-style-type: none"> • The <i>Device Table</i> (DT). • The <i>Interrupt Translation Table</i> (ITT).
<code>R_SDBRM</code>	All ITS translation structures are little-endian.
<code>R_SVKHY</code>	The ITS accesses its translation structures using Normal memory types. See [1] for more information about memory access types.
<code>S_KZJVX</code>	Software allocates memory for the translation structures and programs the physical address of the allocated translation structures in ITS registers. Arm expects that such allocations will typically be done from Conventional memory [1].
<code>I_XTRLs</code>	<p>ITS_CR1 specifies the Shareability and Cacheability attributes for memory accesses performed by the ITS.</p> <p>ITS_CR1 is read-only when the ITS is enabled.</p> <p>The most recent values written to ITS_CR1 are used when the ITS accesses memory following a write that enables the ITS.</p>
<code>D_ZQBDN</code>	<p>The GICv5 architecture defines two types of translation structures entries:</p> <ul style="list-style-type: none"> • An entry in a DT, a <i>device table entry</i> (DTE). • An entry in an ITT, an <i>interrupt translation table entry</i> (ITTE).
<code>D_YMZNc</code>	<p>The VALID field of an ITS translation structure entry defines the validity of the entry:</p> <ul style="list-style-type: none"> • When the VALID field is 1, the entry is a <i>valid</i> entry. • When the VALID field is 0, the entry is an <i>invalid</i> entry.
<code>R_MJPHW</code>	The DT translates a DeviceID into an ITT base address. The DeviceID is used as an index into the DT.
<code>R_XHKFY</code>	The ITT translates a per-device EventID into an LPI ID, a physical/virtual qualifier, and a VM ID for virtual interrupts. The EventID is used as an index into the ITT.
<code>R_TLTPV</code>	<p>An ITS Domain has a valid translation for the DeviceID and EventID of an ITS event, if all of the following are true:</p> <ul style="list-style-type: none"> • One of the following is true: <ul style="list-style-type: none"> – The DT does not use a 2-level table structure. – The level 1 DTE for the DeviceID is valid. • The level 2 DTE for the DeviceID is valid. • One of the following is true: <ul style="list-style-type: none"> – The ITT does not use a 2-level table structure. – The level 1 ITTE for the EventID is valid. • The level 2 ITTE for the EventID is valid. • One of the following is true: <ul style="list-style-type: none"> – The event is generated by a write to a register in the PAS native to the ITS Domain. – The event is generated by a write to a register in the Non-secure PAS, and all of the following are true: <ul style="list-style-type: none"> * The event is translated in the Realm ITS Domain. * L2_ITTE.DAC is 0b01 in the L2_ITTE for the EventID. <p>Otherwise, if any of the above conditions are false, the ITS Domain does not have a valid translation for the ITS event.</p>
<code>R_PGcPL</code>	<p>If any of the following is true, it is CONSTRAINED UNPREDICTABLE whether any events are sent to the IRS:</p> <ul style="list-style-type: none"> • The ITS Domain has more than one valid translation to the same physical INTID. • The ITS Domain has more than one valid translation to the same virtual INTID with the same VM ID.

See also:

- 11.1 ITS Data Structures

5.3.1 The Device Table (DT)

D _{XJQMP}	The DT structure is either linear or 2-level.
I _{GRMKZ}	ITS_IDR1.DT_LEVELS reports whether 2-level DT support is implemented.
R _{WKPKD}	The value of ITS_IDR1.DT_LEVELS is the same for all ITS Domains for an ITS in the system.
I _{VZYWW}	Arm strongly recommends that an implementation that supports more than 9 bits of DeviceID implements 2-level device tables.
I _{SWCZW}	The GICv5 architecture defines the formats of the level 1 DTE and level 2 DTE.
I _{ZMRVW}	The base address of the DT is stored in ITS_DT_BASER.ADDR.
I _{VZVJN}	When ITS_CR0.ITSEN is 1, accesses to fields in ITS_DT_BASER and ITS_DT_CFGR are RO .
I _{WBLVK}	When the DT uses a 2-level table structure, the address stored in ITS_DT_BASER.ADDR is the base address of the level 1 table. When the DT uses a linear table structure, the address stored in ITS_DT_BASER.ADDR is the base address of the array of level 2 DTEs.
S _{LDVFG}	Software is responsible for allocating the DT from memory in the PAS associated with the ITS Domain where the DT is used.
I _{KRKBH}	The structure of the DT is controlled using the following register and translation structure fields: <ul style="list-style-type: none"> ITS_DT_CFGR.DEVICEID_BITS: Selects how many bits of DeviceID the DT can translate. This impacts the size of the level 2 DT when using a linear structure and the size of the level 1 DT when using 2-level structure. ITS_DT_CFGR.STRUCTURE: Selects if the DT uses a linear or 2-level structure. ITS_DT_CFGR.L2SZ: When using a 2-level DT, this field configures the number of DeviceID bits resolved by each level 2 DT. L1_DTE.SPAN: When using a 2-level DT, for the level 1 entry corresponding to a range of DeviceIDs, this field configures the number of entries in the level 2 DT.
I _{PZYZM}	The supported values for ITS_DT_CFGR.L2SZ are reported in ITS_IDR1.L2SZ.
I _{NTYVG}	The size of a level 1 DTE and a level 2 DTE is 8 bytes.

Table 5.1 shows some example DT structures:

Table 5.1: Example DT structures.

STRUCTURE	DEVICEID_BITS	L2SZ	L1 size	L2 size	Maximum number of devices
Linear	4	-	-	128 B	16
Linear	9	-	-	4 KB	512
Linear	13	-	-	64 KB	8,192
Linear	18	-	-	2 MB	262,144
2-level	18	0b00, 9 bits	4 KB	4 KB	262,144

STRUCTURE	DEVICEID_BITS	L2SZ	L1 size	L2 size	Maximum number of devices
2-level	27	0b00, 9 bits	2 MB	4 KB	134,217,728
2-level	22	0b01, 11 bits	16 KB	16 KB	4,194,304
2-level	26	0b10, 13 bits	64 KB	64 KB	67,108,864
2-level	32	0b10, 13 bits	4 MB	64 KB	4,290,000,000

In the table above, the L2SZ column shows the value of ITS_DT_CFGR.L2SZ along with its interpretation in bits of DeviceID resolved by each level.

See [11.1 ITS Data Structures](#) and [10.3.1.6 ITS_DT_CFGR](#) for more information.

SYMMZX

The L1 size and L2 size columns in [Table 5.1](#) indicate the required maximum physically contiguous allocation by software for the configuration. Arm expects that for large systems with many devices, a 2-level structure with a large level 1 table can be allocated by software on system boot.

SFGKBW

L1_DTE.SPAN field is used to limit the amount of storage required for the level 2 DT. For example, if the L2SZ is configured at 9 bits, and only two devices are differentiated using bit 0 of the DeviceID for the DeviceID space defined by bits N:9, then a SPAN of 1 can be used to only require that two level 2 DTEs are allocated. See [Figure 5.3](#) for an illustration where the L2SZ is configured at 9 bits and where two level 1 DTEs use different SPAN values.

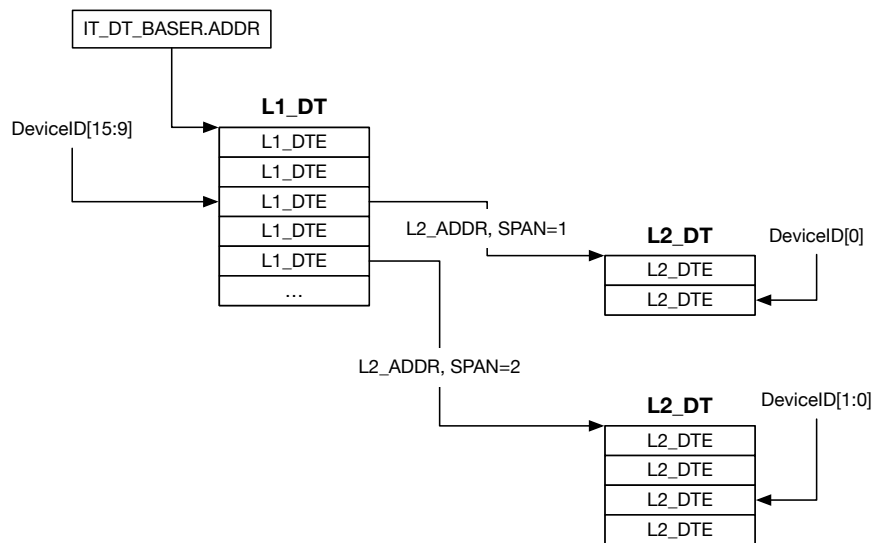


Figure 5.3: Example device table using L1_DTE.SPAN

SCQRVC

ITS_DT_CFGR.DEVICEID_BITS determines the number of level 2 DTEs the DT contains. When using a 2-level structure, the number of DeviceID bits resolved by each level 2 table is decided by ITS_DT_CFGR.L2SZ, and software must allocate enough level 1 DTEs to cover the full DeviceID space for the DT.

DQFTBH

A DeviceID is considered *invalid* if any of the following is true:

- The DeviceID is larger than $(2^{\text{ITS_DT_CFGR.DEVICEID_BITS}}) - 1$.
- DeviceIDs which do not have a corresponding level 2 DTE because of how the L1_DTE.SPAN field is configured.

RHGKJK

Events signaled with an invalid DeviceID are ignored and are not translated by the ITS.

<code>I_XSFXB</code>	<p>A level 1 DT is aligned to the size of the level 1 table.</p> <p>A level 2 DTE array is aligned to the size of the array.</p> <p>See 10.3.1.5 ITS_DT_BASER for more information.</p>
<code>I_CFRYB</code>	When limiting a level 2 DT size by using the level 1 SPAN field, the level 2 DT address is aligned to the size of the level 2 DTE array.
<code>R_XNJYG</code>	<p>When an ITS Domain is enabled, a write to any field in a valid level 2 DTE other than when setting the <code>L2_DTE.VALID</code> field to 0, results in <code>CONSTRAINED UNPREDICTABLE</code> behavior with a choice of:</p> <ul style="list-style-type: none"> • The old values are used for the device. • The new values are used for the device. • A combination of the old and new values is used for the device.
<code>I_DHHJN</code>	<p>When the ITS Domain is enabled, changing the address of an ITT for a valid level 2 DTE, requires performing the following series of actions to avoid <code>UNPREDICTABLE</code> behavior:</p> <ol style="list-style-type: none"> 1. Performing a write that sets the <code>L2_DTE.VALID</code> field to 0. 2. Ensuring that the write is visible to the ITS Domain by performing the necessary cache invalidation. 3. Writing the address and configuration fields of the new ITT to the level 2 DTE. 4. Performing a write that sets the <code>L2_DTE.VALID</code> field to 1. <p>See 5.4.2 ITS cache management for DeviceIDs for more information about ITS device cache management.</p>

5.3.2 The Interrupt Translation Table (ITT)

<code>D_PPZDD</code>	The ITT structure is either linear or 2-level.
<code>I_RPQVS</code>	<code>ITS_IDR1.ITT_LEVELS</code> reports whether 2-level ITT support is implemented.
<code>R_TJCMV</code>	The value of <code>ITS_IDR1.ITT_LEVELS</code> is the same for all ITS Domains in an ITS in the system.
<code>I_PLTFB</code>	The GICv5 architecture defines the formats of the level 1 ITTE and level 2 ITTE.
<code>I_QJTXY</code>	The base address of the ITT is stored in <code>L2_DTE.ITT_ADDR</code> .
<code>I_FMQQL</code>	<p>When the ITT uses a 2-level table structure, the address stored in <code>L2_DTE.ITT_ADDR</code> is the base address of the level 1 table.</p> <p>When the ITT uses a linear table structure, the address stored in <code>L2_DTE.ITT_ADDR</code> is the base address of the array of level 2 ITTEs.</p>
<code>D_CFDKB</code>	<p>A level 2 ITT <i>overlaps</i> with another level 2 ITT, if the same address for a level 2 ITTE can be reached via separate DTs or ITTs.</p> <p>For example, a level 2 ITT overlaps with another level 2 ITT, if any of the following are true:</p> <ul style="list-style-type: none"> • Two separate level 2 DTEs point to the same level 1 ITT. • Two separate level 2 DTEs point to the same level 2 ITT. • Two separate level 1 ITTEs point to the same level 2 ITT. • The base address of a level 2 ITT is between the base and end addresses of another level 2 ITT.
<code>R_XGPSP</code>	<p>If any part of either a level 1 or level 2 ITT overlaps with another ITT, the ITS behavior is <code>CONSTRAINED UNPREDICTABLE</code> with any of the following:</p> <ul style="list-style-type: none"> • Any ITS cache invalidation by DeviceID or EventID may not work as expected. • ITS events may not be translated. • The ITS may translate an ITS event resulting in sending an interrupt event to the IRS with <code>UNKNOWN</code> data. <p>Software can recover from this situation by disabling the ITS and configuring the translation structures without any overlap before re-enabling the ITS.</p> <p>See 5.5 ITS memory access rules for more information.</p>

<code>S_{YDTKM}</code>	Software is responsible for allocating the ITT from memory in the PAS associated with the ITS Domain where the ITT is used.
<code>I_{SRXXT}</code>	<p>The structure of the ITT is controlled using the following fields in the ITS translation structures:</p> <ul style="list-style-type: none"> • <code>L2_DTE.EVENTID_BITS</code>: Selects how many bits of EventID the ITT can translate. This impacts the size of the level 2 ITT when using a linear structure and the size of the level 1 ITT when using 2-level structure. • <code>L2_DTE.ITT_STRUCTURE</code>: Selects if the ITT uses a linear or 2-level structure. • <code>L2_DTE.ITT_L2SZ</code>: When using a 2-level ITT, this field configures the number of EventID bits resolved by each level 2 ITT. • <code>L1_ITTE.SPAN</code>: When using a 2-level ITT, for the level 1 entry corresponding to a range of EventIDs, this field configures the number of entries in the level 2 ITT.
<code>I_{HPWVL}</code>	The supported values for <code>L2_DTE.ITT_L2SZ</code> are reported in <code>ITS_IDR1.L2SZ</code> .
<code>I_{NJHFT}</code>	The size of a level 1 ITTE and a level 2 ITTE is 8 bytes.

[Table 5.2](#) shows some example ITT structures:

Table 5.2: Example ITT structures.

STRUCTURE	EVENTID_BITS	L2SZ	L1 size	L2 size	Maximum number of events
Linear	4	-	-	128 B	16
Linear	9	-	-	4 KB	512
Linear	11	-	-	16 KB	2,048
Linear	13	-	-	64 KB	8,192
Linear	16	-	-	512 KB	65,536
2-level	11	0b00, 9 bits	32 B	4 KB	2,048
2-level	16	0b00, 9 bits	1 KB	4 KB	65,536

In the table above, the L2SZ column shows the value of `L2_DTE.ITT_L2SZ` field along with its interpretation in bits of EventID resolved by each level.

See [11.1 ITS Data Structures](#) for more information.

<code>S_{FYKKL}</code>	The L1 size and L2 size columns in Table 5.2 indicate the required maximum physically contiguous allocation by software for the configuration. Arm expects that 4KB allocations will satisfy the capabilities of most devices. For example, PCIe MSI-X supports a maximum of 2,048 events for a device. 2,048 events can either be supported with a linear 16KB table or using a 2-level structure with 4 level 1 entries each pointing to 4KB level 2 tables.
<code>S_{YSMWP}</code>	The <code>L1_ITTE.SPAN</code> field can be used to limit the amount of storage required for the level 2 ITT. For example, if a 2-level ITT contains entries for 544 events and the L2SZ is configured at 9 bits, then the last level 1 ITTE can configure a SPAN of 5 such that the second level 2 ITT only contains 32 entries.
<code>S_{ZFWKT}</code>	<code>L2_DTE.EVENTID_BITS</code> determines the number of level 2 ITTEs that the ITT contains. When using a 2-level structure, the number of EventID bits resolved by each level 2 table is decided by <code>L2_DTE.ITT_L2SZ</code> , and software must allocate enough level 1 ITTEs to cover the full EventID space for the ITT.

D _{XVTYL}	An EventID is considered <i>invalid</i> if any of the following is true: <ul style="list-style-type: none"> • The EventID is larger than $(2^{\text{L2_DTE.EVENTID_BITS}}) - 1$. • EventIDs which are not described by level 2 ITTEs because of how the L1_ITTE.SPAN field is configured.
R _{XJGVV}	Events signaled with an invalid EventID are ignored and are not translated by the ITS.
I _{XRKXZ}	If software error reporting is supported, the ITS reports if an ITS event cannot be translated due to an invalid EventID using the appropriate error code in ITS_SWERR_STATUSR.EC.
R _{FVJTG}	The ITTs are always aligned to the size of the table.
I _{WBPQX}	When limiting a level 2 ITT size by using the level 1 SPAN field, the level 2 ITT address is aligned to the size of the level 2 ITTE array.
I _{CKRKW}	Software may configure the ITS to translate an ITS event, generated by a write to a register in the Non-secure PAS, using the Realm ITS Domain. Such an event does not violate the separation between Interrupt Domains, because the event is only translated successfully in the Realm ITS Domain if L2_ITTE.DAC is 0b01 for the event.

See also:

- [5.1.1 Supporting Realm interrupts from Non-secure writes.](#)

I _{TSBSN}	Software may read translation information for an EventID and DeviceID in an ITS Domain via the following ITS registers: <ul style="list-style-type: none"> • ITS_DIDR. A write to this register selects the DeviceID of the event. • ITS_EIDR. A write to this register selects the EventID of the event. • ITS_READ_EVENTR. A write of 1 to the R field requests that the ITS returns translation information for the selected event. • ITS_STATUSR. A read of this register reports when the translation information for the event is available. • ITS_READ_EVENT_DATAR. A read of this register provides the translation information for the event once it is available.
--------------------	--

See also:

- [10.3.1.4 ITS_DIDR.](#)
- [10.3.1.7 ITS_EIDR.](#)
- [10.3.1.8 ITS_GEN_EVENTR.](#)
- [10.3.1.9 ITS_GEN_EVENT_STATUSR.](#)
- [10.3.1.10 ITS_GEN_EVENT_EIDR.](#)
- [10.3.1.11 ITS_GEN_EVENT_DIDR.](#)
- [10.3.1.22 ITS_READ_EVENTR.](#)
- [10.3.1.23 ITS_READ_EVENT_DATAR.](#)
- [10.3.1.24 ITS_STATUSR.](#)

5.4 ITS cache management

R_{FXXDG}	An ITS is permitted to cache any part of the ITS translation structures at any time when they are accessible by the ITS.
I_{KKDTR}	The ITS caches are not considered data or system caches belonging to the memory system, meaning they are not managed by any cache maintenance instructions issued by a PE.

Note

This means that information from all of level 1 DTEs, level 2 DTEs, level 1 ITTEs, and level 2 ITTEs may be cached by the ITS.

I_{NRPBV}	An ITS is permitted to cache both valid and invalid entries of ITS translation structures.
R_{XVTHB}	Entries held in ITS caches are associated with the ITS Domain that contains the ITS translation structures used to derive the entry.

Note

This ensures isolation between ITS Domains and ensures that modifying translation structures and performing cache maintenance for one ITS Domain does not affect the behavior of another ITS Domain.

R_{MKQFF}	When an ITS Domain is disabled, the ITS caches contain no data from ITS translation structures associated with that ITS Domain.
R_{JTCBT}	Cached information from an ITS translation structure entry is not guaranteed to remain in the ITS caches.
I_{QSGHD}	Cached information from an ITS translation structure entry is not guaranteed to remain in the ITS caches. This means that when the ITS caches an ITS translation structure entry with different values than what is stored in memory, the ITS may naturally evict the cached entry and read the updated translation structure entry from memory, without any ITS cache invalidation operation having been performed.
R_{HQFDL}	An update to an ITS translation structure entry is not guaranteed to be visible to the ITS until an appropriate cache invalidation operation has completed.
I_{GMVTV}	When software updates the fields in an ITS translation structure entry, the ITS is permitted to cache both the old and the new entry. For as long as both entries exist in the ITS cache, each translation of a DeviceID and EventID affected by the updated entry is permitted to use either the old or the new entry.
S_{PJDOF}	To avoid the ITS using the old translation entries, software should perform the following steps: <ol style="list-style-type: none"> 1. Ensure that no events are generated for that DeviceID. 2. Synchronize events Accepted by the ITS that use the old translation structures. See also 5.2.4 ITS synchronization requests for more information. 3. Update the fields in the ITS translation structure entry and perform the appropriate cache invalidation operation. See also 5.4.1 ITS cache management for EventIDs and 5.4.2 ITS cache management for DeviceIDs for more information. 4. Enable events for that DeviceID.
I_{WBWHP}	ITS cache invalidation is done for a device or an event. A device is identified by its DeviceID. An event is identified by its DeviceID and EventID.
I_{BPDHC}	The ITS provides cache invalidation operations that apply to any cached information from ITS translation structure entries for the specified device or event. An invalidation operation applies to the cached information from ITS translation structures irrespective of the values stored in the ITS translation structures in memory.

For example, for an event, the state of the corresponding level 2 ITTE could be as follows:

- There is cached information from a valid level 2 ITTE for that event.
- There is no cached information from the level 2 DTE for the DeviceID of the event.
- The level 2 DTE for the DeviceID of the event is invalid.

The ITS cache invalidation operation for that event invalidates the cached information from the valid level 2 ITTE even though the level 2 DTE is invalid.

<code>I_LNLHV</code>	ITS cache invalidation operations invalidate cached information regardless of the value of the <code>VALID</code> field in the cached translation structure entry.
<code>I_WZQON</code>	The effects of ITS cache invalidation operations are complete when <code>ITS_STATUSR.IDLE</code> is 1.
<code>R_ZXXJF</code>	When <code>ITS_STATUSR.IDLE</code> transitions from 0 to 1, any ITS event that used information from ITS translation structure entries prior to the invalidate have been Accepted by the IRS.

5.4.1 ITS cache management for EventIDs

<code>I_RPBTZ</code>	<p>An ITS cache invalidation operation for a single EventID is performed as follows:</p> <ol style="list-style-type: none"> 1. The EventID is written to <code>ITS_EIDR.EVENT_ID</code> and the DeviceID to <code>ITS_DIDR.DEVICE_ID</code>. 2. A write to <code>ITS_INV_EVENTR</code> is performed with the fields set to the following: <ul style="list-style-type: none"> • <code>ITS_INV_EVENTR.L1</code> is set to 0. • <code>ITS_INV_EVENTR.I</code> is set to 1.
<code>R_KYPHL</code>	An ITS cache invalidation operation for an EventID invalidates all cached information from the L2 ITTE corresponding to the EventID.
<code>I_GLSYM</code>	<p>The ITS supports invalidating a range of EventIDs corresponding to the number of EventIDs described by a single level 2 ITT. The size of the range is selected using <code>ITS_INV_EVENTR.ITT_L2SZ</code>.</p> <p>An ITS cache invalidation operation for a range of EventIDs is performed as follows:</p> <ol style="list-style-type: none"> 1. Any EventID within the range is written to <code>ITS_EIDR.EVENT_ID</code> and the DeviceID is written to <code>ITS_DIDR.DEVICE_ID</code>. 2. A write to <code>ITS_INV_EVENTR</code> is performed with the fields set to the following: <ul style="list-style-type: none"> • <code>ITS_INV_EVENTR.L1</code> is set to 1. • <code>ITS_INV_EVENTR.ITT_L2SZ</code> is set to the size of a level 2 ITT corresponding to the range of EventIDs that are invalidated. • <code>ITS_INV_EVENTR.I</code> is set to 1.
<code>R_PBJYK</code>	An ITS cache invalidation operation for a range of EventIDs invalidates all cached information from the L1 ITTE and L2 ITTEs for the DeviceID corresponding to the invalidate operation's DeviceID and for the EventID bits[15:N] corresponding to the invalidate operation's EventID's bits[15:N], where N equals $(9 + 2 * \text{ITS_INV_EVENTR.ITT_L2SZ})$.

Note

An ITS cache invalidation for a range of EventIDs applies to cached information from that range of EventIDs, regardless of whether a linear or 2-level structure is used for the ITT when the information is cached.

<code>I_JNKRH</code>	In an invalidation operation for a range of EventIDs, for each EventID in the range, the ITS invalidates the same cached information as when an invalidate operation is performed for a single EventID. See rule KYPHL for more information.
<code>I_FLZSK</code>	The effects of a write to <code>ITS_INV_EVENTR</code> are complete when <code>ITS_STATUSR.IDLE</code> is 1.

See also:

- [10.3.1.4 ITS_DIDR](#).
- [10.3.1.7 ITS_EIDR](#).

- [10.3.1.17 ITS_INV_EVENTNR.](#)
- [10.3.1.24 ITS_STATUSR.](#)

5.4.2 ITS cache management for DeviceIDs

I _{TKFBL}	<p>An ITS cache invalidation operation for a single DeviceID is performed as follows:</p> <ol style="list-style-type: none"> 1. The DeviceID is written to ITS_DIDR.DEVICE_ID. 2. A write to ITS_INV_DEVICER is performed with the fields set to the following: <ul style="list-style-type: none"> • ITS_INV_DEVICER.L1 is set to 0. • ITS_INV_DEVICER.EVENTID_BITS is set to the number of EventID bits to which the invalidation operation applies. • ITS_INV_DEVICER.I is set to 1.
R _{NWHYF}	<p>An ITS cache invalidation operation for a DeviceID invalidates all of the following:</p> <ul style="list-style-type: none"> • All cached information from the L2 DTE that corresponds to the DeviceID. • All cached information from all L1 ITTEs that correspond to the DeviceID. • All cached information from L2 ITTEs for the DeviceID where the EventID is less than $2^{\text{ITS_INV_DEVICER.EVENTID_BITS}}$.
I _{DWLGT}	<p>The ITS supports invalidating a range of DeviceIDs corresponding to the number of DeviceIDs described by a single level 2 DT.</p> <p>An ITS cache invalidation operation for a range of DeviceIDs is performed as follows:</p> <ol style="list-style-type: none"> 1. A DeviceID within the range is written to ITS_DIDR.DEVICE_ID. 2. A write to ITS_INV_DEVICER is performed with the fields set to the following: <ul style="list-style-type: none"> • ITS_INV_DEVICER.L1 is set to 1. • ITS_INV_DEVICER.I is set to 1.
R _{FWYGP}	<p>An ITS cache invalidation operation for a range of DeviceIDs invalidates all cached information from the L1 DTE, L2 DTEs, L1 ITTEs, and L2 ITTEs for DeviceID bits[31:N] corresponding to the invalidate operation's DeviceID bits[31:N], where N equals $(9 + 2 * \text{ITS_DT_CFGR.L2SZ})$ and the effective value of EVENTID_BITS is the value reported in ITS_IDR2.EVENTID_BITS.</p>
I _{NTHYD}	<p>When a cache invalidation operation for a range of DeviceIDs is performed, the Effective value of EVENTID_BITS used to invalidate each DeviceID is the maximum implemented number of EventID bits reported in ITS_IDR2.EVENTID_BITS.</p>
I _{XBVDN}	<p>In an invalidation operation for a range of DeviceIDs, for each DeviceID in the range, the ITS invalidates the same cached information as when an invalidate operation is performed for a single DeviceID.</p> <p>This means that cached information from the following translation structures is invalidated for each DeviceID when invalidating a range of DeviceIDs:</p> <ul style="list-style-type: none"> • The L2 DTE that corresponds to the DeviceID. • All L1 ITTEs that correspond to the DeviceID. • All L2 ITTEs that correspond to the DeviceID. <p>See rule NWHYF for more information.</p>
I _{WWQDF}	<p>The effects of a write to ITS_INV_DEVICER are complete when ITS_STATUSR.IDLE is 1.</p>
I _{QNNQG}	<p>When a write to a level 1 DTE sets VALID to 0, and there is cached information from a valid level 2 DTE for a DeviceID in the range covered by that level 1 DTE, the ITS may also contain cached information of the previously valid level 1 DTE.</p> <p>This means that the ITS is permitted to treat the DeviceID either as valid or invalid, and to ensure the ITS treats the DeviceID as invalid, the level 1 DTE cache entries for the DeviceID range must be invalidated.</p> <p>See also:</p>

- [10.3.1.6 ITS_DT_CFGR.](#)
- [10.3.1.16 ITS_INV_DEVICER.](#)
- [10.3.1.24 ITS_STATUSR.](#)

5.5 ITS memory access rules

R _{BXBXJ}	<p>If an ITS access to any ITS translation structure occurs in PCIe address space, then the ITS is permitted to return an UNKNOWN value or terminate the access and report the error.</p> <p>In this case, the ITS is permitted to drop the translation and not send any outgoing interrupt events to the IRS.</p>
R _{GRXYB}	<p>When an ITS accesses a memory location in a PAS, it relies on information stored in registers and translation structures that are accessible only within the same PAS to validate that the access is permitted.</p> <p>This includes scenarios in which the ITS behavior is CONSTRAINED UNPREDICTABLE, UNPREDICTABLE, or IMPLEMENTATION DEFINED.</p> <p>For example, an ITS access to a translation structure is permitted only if the memory location lies within a region defined by a base address and size held in a register or another translation structure entry, where that defining information resides in the same PAS as the access itself.</p> <p>This prevents software running in a Security state from directing the ITS to access any memory location using another PAS than the one associated with the software's Security state.</p>
R _{BLYPY}	<p>An ITS does not access any memory location which is not derived from address and configuration data in the ITS registers and translation structures.</p> <p>When a memory address is stored in an address field in a register or translation structure entry, and that register or entry contains a field indicating that the address is not valid, the ITS does not derive any memory location from that address field.</p> <p>This includes scenarios in which the ITS behavior is CONSTRAINED UNPREDICTABLE, UNPREDICTABLE, or IMPLEMENTATION DEFINED.</p>
R _{LTFTQ}	<p>When the ITS accesses an ITS translation structure, the access must be a 64-bit aligned, single-copy atomic access of at least 64 bits in size.</p>
I _{JXSPX}	<p>If the ITS experiences an external abort during a memory access to an ITS translation structure, the ITS stops the operation and reports the error.</p> <p>If software error reporting is supported, the error is reported with ITS_SWERR_STATUSR.EC in the range from 0x00 to 0x04.</p>
R _{LSCDT}	<p>If the ITS experiences an external abort as part of translating an event, the translation is dropped and no outgoing interrupt event is sent to the IRS.</p>
R _{PGRJQ}	<p>If an ITS translation structure overlaps with any other ITS or IRS translation structure, the ITS behavior is UNPREDICTABLE.</p> <p>The UNPREDICTABLE behavior may result in loss of interrupt configuration and state, but must not result in access to memory outside the PAS associated with the ITS Domain or in access to any memory location which is not derived from address and configuration of the DT or ITT.</p> <p>For more information about the IRS translation structures, see 4.7 The interrupt state table (IST) and 4.9 Virtualization data structures.</p> <p>Software can recover from this situation by performing the following sequence:</p> <ol style="list-style-type: none">1. Disabling the ITS.2. Making the corresponding translation structures invalid in the IRS.3. Reconfigure the base addresses of the translation structure to avoid overlap.

5.6 ITS support for MPAM

I_{XZGHC} The Memory System Resource Partitioning and Monitoring Memory, MPAM, architecture defines per-transaction attributes that affect system behavior or the behavior of components that the transactions pass through or Completers that satisfy a transaction [1].

The ITS supports MPAM with the following additional attributes:

- Partition ID, or PARTID.
- Performance Monitoring Group, or PMG.

The PARTID and PMG are both interpreted within a PARTID space. The PARTID space used depends on the Security state associated with an ITS Domain, and the PARTID and PMG values used may be programmed independently for each ITS Domain. See [1] for more information about MPAM information bundles.

I_{F BPMW} In the ITS, support for MPAM is indicated by ITS_IDR0.MPAM.

If MPAM is supported, the supported PARTID and PMG width is indicated in ITS_MPAM_IDR.{PARTID_MAX,PMG_MAX}.

R_{FCMHZ} If MPAM is supported, ITS accesses to memory are associated with the PARTID and PMG programmed in ITS_MPAM_PARTID_R.{PARTID,PMG}.

I_{THJMY} In systems without support for RME, the PARTID space used by the ITS is determined by the memory system attribute MPAM_NS.

In systems with support for RME, the PARTID space used by the ITS is determined by the memory system attribute MPAM_SP.

I_{VCSKZ} The ITS architecture has optional support for MPAM PARTID space selection indicated by ITS_MPAM_IDR.HAS_MPAM_SP.

If ITS_MPAM_IDR.HAS_MPAM_SP is 0, each ITS Domain uses a default MPAM PARTID space.

The following table shows the MPAM PARTID space used for accesses made by the ITS Domains if ITS_MPAM_IDR.HAS_MPAM_SP is 0 and the system does not support RME:

ITS Domain	MPAM PARTID space
Secure	Secure PARTID space
Non-Secure	Non-secure PARTID space
EL3	Secure PARTID space

The following table shows the MPAM PARTID space used for accesses made by the ITS Domains if ITS_MPAM_IDR.HAS_MPAM_SP is 1 and the system supports RME:

ITS Domain	MPAM PARTID space
Secure	Secure PARTID space
Non-Secure	Non-secure PARTID space
EL3	Root PARTID space.
Realm	Realm PARTID space.

If ITS_MPAM_IDR.HAS_MPAM_SP is 1, the ITS uses the MPAM PARTID specified by ITS_MPAM_PARTID_R.MPAM_SP.

R_{BJQFG} If an ITS without support for MPAM is integrated in a system that supports MPAM, the PARTID and PMG used for each ITS Domain is IMPLEMENTATION DEFINED.

5.7 ITS support for Memory Encryption Contexts

<code>I_{PRBCL}</code>	The Memory Encryption Contexts feature, FEAT_MEC, provides finer-grained memory encryption contexts, within the Realm physical address space, to be assigned to Realms, with policy controlled by Realm EL2 [1].
<code>G_{ZHPQC}</code>	The GICv5 architecture ensures that an ITS can function correctly in systems with support for Memory Encryption Contexts (MEC).
<code>I_{YQGHF}</code>	<p>In a system with support for MEC, the PEs write ITS translation structures to memory and these are read by the ITS.</p> <p>In such a system, all ITS Realm translation structures are managed and configured by Realm EL2, and virtualized software running at EL1 is not expected to have direct access to the ITS configuration registers or translation structures. Therefore, the ITS only supports configuration of a single MECID used for all ITS memory accesses to the Realm PAS. This allows software running at Realm EL2 to write to the ITS translation structures and allows the ITS to read the translations.</p>
<code>I_{VRVQV}</code>	<p>In the ITS, support for Memory Encryption Contexts (MEC) is indicated in ITS_IDR0.MEC for the Realm ITS Domain. If the MEC feature is supported, the supported MECID width is indicated in ITS_MEC_IDR.MECIDSIZE for the Realm ITS Domain.</p> <p>Arm strongly recommends that the MECID bit width supported by the ITS matches or exceeds the width supported by the PEs in the system.</p>
<code>R_{XKTWP}</code>	If the MEC feature is supported, ITS accesses to memory are associated with a MECID that identifies the Memory Encryption Context of the access.
<code>R_{XTVVQ}</code>	Accesses made by the ITS for Secure, Non-secure, and Root PA spaces are issued with the default MECID of zero.
<code>I_{FPSSV}</code>	Accesses made by the ITS to Realm PA space are associated with the global Realm PAS ITS MECID programmed in ITS_MEC_MECID_R.MECID.
<code>R_{DGJXQ}</code>	If an ITS without support for the Realm ITS Domain is integrated in a system that supports MEC, all ITS accesses for that ITS are treated as having the default MECID of zero.

5.8 ITS support for software error reporting

I _{FWKKS}	The ITS specifies a mechanism to report errors because of incorrect programming.
R _{FBVQN}	The ITS detects software errors while processing ITS events. Otherwise, the ITS does not detect software errors.
R _{RLHPC}	Software error codes that refer to incoming events apply to incoming events generated by all supported mechanisms. See Chapter 5 Interrupt translation service (ITS) for more information about the supported mechanisms for incoming events.
R _{RLJNX}	ITS support for software error reporting is optional.
I _{HKRKL}	ITS_IDR0.SWE reports whether software error reporting is supported.
R _{QRYDD}	The value of ITS_IDR0.SWE is the same for all ITS Domains in an ITS in the system.
I _{FCPZL}	When ITS_IDR0.SWE is 1, the ITS uses the ITS_SWERR_STATUSR.IMP_EC to report any IMPLEMENTATION DEFINED errors detected by the ITS.
I _{ZGMTB}	ITS_SWERR_STATUSR.EC is 0 when an IMPLEMENTATION DEFINED error is reported by the ITS.
I _{ZNSVM}	When L2_DTE.DSWE is 0, the ITS does not report errors for the following error codes in ITS_SWERR_STATUSR.EC: <ul style="list-style-type: none"> • 0x03: Failed lookup of L1_ITTE due to an external abort. • 0x04: Failed lookup of L2_ITTE due to an external abort. • 0x06: An incoming event could not be translated because L2_DTE.VALID is 0. • 0x07: An incoming event could not be translated because L1_ITTE.VALID is 0. • 0x08: An incoming event could not be translated because L2_ITTE.VALID is 0. • 0x0A: An incoming event could not be translated because the EventID > (2 ^ L2_DTE.EVENTID_BITS) - 1. • 0x0C: An incoming event could not be translated because the EventID exceeds the L1_ITTE.SPAN. • 0x0D: An incoming event could not be translated because the event is associated with the Non-secure Interrupt Domain and L2_ITTE.DAC = 0.
S _{FPVFM}	When a device is assigned to a VM, the Hypervisor can set L2_DTE.DSWE to 0 in the L2 DTE of the device to mask reporting of software errors caused by misconfiguration of the device by VM.
I _{ZKJKN}	Software error information can be read in an ITS Domain via the following ITS registers: <ul style="list-style-type: none"> • ITS_SWERR_STATUSR. • ITS_SWERR_SYNDROMER0. • ITS_SWERR_SYNDROMER1.
I _{YSSCL}	When a software error is reported in the ITS Domain, the value of ITS_SWERR_STATUSR.V is 1. Otherwise, no software error is reported in the ITS Domain and fields in this register are UNKNOWN.
I _{PVQYF}	When the value of ITS_SWERR_STATUSR.V is 1, all of the following are true in that ITS Domain: <ul style="list-style-type: none"> • ITS_SWERR_STATUSR.EC specifies the fault that caused the software error. • ITS_SWERR_STATUSR.S0V indicates whether ITS_SWERR_SYNDROMER0 contains valid error syndrome information. • ITS_SWERR_STATUSR.S1V indicates whether ITS_SWERR_SYNDROMER1 contains valid error syndrome information. • ITS_SWERR_STATUSR.OF indicates whether multiple software errors were detected.
I _{ZBHQQ}	When ITS_SWERR_STATUSR.S0V is 1, ITS_SWERR_SYNDROMER0 reports the following information for the software error: <ul style="list-style-type: none"> • The DeviceID and EventID of the incoming event that resulted in the software error. • The Interrupt Domain the incoming event is associated with.
I _{BBZNG}	When ITS_SWERR_STATUSR.S1V is 1, ITS_SWERR_SYNDROMER1 reports the address of the ITS translation structure associated with the software error.

<code>R_LKLOG</code>	For each reported error, the values of <code>ITS_SWERR_STATUSR.{S0V,S1V}</code> are IMPLEMENTATION DEFINED and set independently.
<code>S_NZTWH</code>	<p>Arm recommends that software performs the following sequence of operations to either clear the last error, or detect if new errors were reported:</p> <ol style="list-style-type: none">1. Read <code>ITS_SWERR_STATUSR</code> and determine which fields need to be cleared to zero.2. In a single-copy atomic write to <code>ITS_SWERR_STATUSR</code>:<ol style="list-style-type: none">1. Write ones to all the W1C fields that are nonzero in the read value.2. Write zero to all the W1C fields that are zero in the read value.3. Write zero to all the RW fields.3. Read back <code>ITS_SWERR_STATUSR</code> after the write. If the value read back is the same as the value that was written, all W1C fields that were non-zero are cleared, and no new errors were reported. Otherwise, one or more new errors were reported. <p>See also:</p> <ul style="list-style-type: none">• 10.3.1.25 ITS_SWERR_STATUSR.• 10.3.1.26 ITS_SWERR_SYNDROMER0.• 10.3.1.27 ITS_SWERR_SYNDROMER1.

Chapter 6

Interrupt Wire Bridge (IWB)

I _{FJFTJ}	This section describes the architecture of an IWB. The IWB detects changes to the state of input wires and signals an ITS. The ITS generates ITS events from the signals and translates them into interrupt events that are forwarded to an IRS.
R _{LMVTM}	A system is permitted to have zero, one, or many IWB instances.
R _{FYRGC}	Each IWB is associated with a single ITS in the system.
D _{YQYMC}	Every input wire connected to the IWB is <i>asserted</i> or <i>de-asserted</i> .
R _{GXTZZ}	The mechanism for making an input wire asserted or de-asserted is IMPLEMENTATION DEFINED.
I _{HBSZB}	<p>For example, a wire connecting a device that signals <i>interrupt events</i> to the IWB may assert the input wire whenever the device signals an event and de-assert the input wire at any time after signaling the event. Only the assertion (and not the de-assertion) of the input wire matters for such a device.</p> <p>As another example, a wire connecting a device signals <i>interrupt state</i> to the IWB may keep the input wire asserted to indicate a state in the device and keep the wire de-asserted to indicate a different state in the device. The input wire would be asserted and de-asserted when the device transitions between such states.</p>
I _{DXLSJ}	The number of implemented wire control registers IWB_WTMR<n>, IWB_WDOMAINR<n>, and IWB_WENABLER<n> is reported in IWB_IDR0.IW_RANGE.
I _{LTGCB}	The maximum number of implemented wires for a single IWB instance is 65536.
R _{MMWGD}	The number of implemented wires by an IWB is less than or equal to (IWB_IDR0.IW_RANGE + 1) * 32.
D _{RLQKQ}	Each input wire to an IWB can be uniquely identified by its <i>input wire index</i> which is a number from 0 through (IWB_IDR0.IW_RANGE + 1) * 32 - 1.
I _{ZQSGD}	Accesses to wire control fields for wires that are not implemented in the wire control registers are RAZ/WI.

R _{VKMRC}	The IWB is identified with a DeviceID that is unique in the DeviceID namespace of the associated ITS. All wire events signaled to the ITS from the same IWB are signaled using the same DeviceID.
R _{RMRTF}	An input wire uses the input wire index as the EventID when signaling an event to the ITS.
I _{MBFPR}	Each input wire is signaled with an EventID that is unique across all input wires implemented by the IWB.
I _{QHCJD}	Arm strongly recommends implementing wires contiguously to avoid the need to allocate a sparse ITT for the ITS.
S _{HYHGF}	The relationship between wires and the connected interrupt source is available to system software in firmware tables.
R _{RVNRP}	The mechanism used by the IWB to communicate events to the ITS is IMPLEMENTATION DEFINED.
I _{TVRKT}	IWB_CR0.IWBEN controls whether the IWB is enabled or disabled.
I _{RQJMN}	The Trigger mode of each wire connected to an IWB is configured to be either edge-triggered or level-sensitive.
R _{WLHDQ}	When the IWB is enabled, events are generated according to the input wire settings: <ul style="list-style-type: none"> • If an input wire is enabled and configured as edge-triggered, the IWB generates a SET_EDGE event when the input wire becomes asserted. • If an input wire is enabled and configured as level-sensitive, the IWB generates a SET_LEVEL event when the input wire becomes asserted and a CLEAR event when the level becomes de-asserted. • If an input wire is disabled, the IWB does not generate any events as a result of the wire becoming asserted or de-asserted. <p>When the IWB is disabled, it generates no events to the ITS.</p>
R _{MXGNY}	When the IWB generates an event as a result of a wire being asserted or de-asserted, the event is generated in finite time.
R _{PKSVB}	For an edge-triggered input wire, if the wire is asserted multiple times without the IWB having sent an event to the ITS, the IWB is permitted to only generate a single SET_EDGE event and send it to the ITS. Once an event is generated and sent to the ITS, when the wire is subsequently asserted, the IWB generates a new SET_EDGE event in finite time.
I _{HTKHT}	For example, when a wire that is configured as edge-triggered goes through the following transition, the IWB is permitted to generate only a single SET_EDGE event: <ol style="list-style-type: none"> 1. The wire is asserted 2. The wire is de-asserted 3. The wire is asserted
R _{PHYZS}	If an input wire is configured as level-sensitive, the IWB is permitted to discard the SET_LEVEL and CLEAR events when the wire is asserted and de-asserted before the IWB sends any event to the ITS.
I _{HTBRD}	For example, for a wire that is configured as level-sensitive and goes through the following transition, the IWB is permitted not to generate any events to the ITS: <ol style="list-style-type: none"> 1. The wire is asserted 2. The wire is de-asserted
R _{GJJZH}	For an edge-triggered input wire, when the IWB or the input wire is disabled, the IWB is permitted to detect that the input wire is asserted and generate a SET_EDGE event later when the IWB and the input wire are enabled. If the input wire's Trigger mode is updated from edge-triggered to level-sensitive when the IWB or input wire is disabled, the IWB does not generate any SET_EDGE events when the IWB and the input wire are subsequently enabled.
I _{BGSBG}	Changing the value of IWB_CR0.IWBEN does not affect the values in other IWB registers.

R _{VWTHK}	<p>When a write to IWB_CR0.IWBEN changes the value from 0 to 1, the IWB begins a transition from disabled to enabled.</p> <p>The transition from disabled to enabled is complete when IWB_CR0.IDLE is 1.</p> <p>When the transition from disabled to enabled is complete, the IWB processes the following events:</p> <p>For wires where IWB_WTMR<n>.TM<x> is 1 (level-sensitive) and IWB_WENABLER<n>.WEN<x> is 1:</p> <ul style="list-style-type: none"> • If the wire is asserted when the transition is complete, the IWB generates a SET_LEVEL event. • If the wire is de-asserted when the transition is complete, the IWB generates a CLEAR event. <p>For wires where IWB_WTMR<n>.TM<x> is 0 (edge-triggered):</p> <ul style="list-style-type: none"> • The IWB does not generate any events for edge-triggered input wires when the IWB becomes enabled.
I _{XFJJK}	<p>If a level-sensitive input wire becomes asserted or de-asserted during the IWB transitions from disabled to enabled, the IWB either generates a single event corresponding to the final state of the input wire or generates multiple events for the input wire where the last event corresponds to the final state of the input wire.</p>
R _{DRQYB}	<p>When a write to IWB_CR0.IWBEN changes the value from 1 to 0, the IWB begins a transition from enabled to disabled.</p> <p>When the IWB is transitioning from enabled to disabled, the IWB is not guaranteed to detect when input wires are asserted or de-asserted.</p> <p>The transition from enabled to disabled is complete when IWB_CR0.IDLE is 1.</p> <p>When the transition from enabled to disabled is complete, the IWB has generated a CLEAR event for every level-sensitive wire where it had sent a SET_LEVEL event and no corresponding CLEAR. These events are Accepted by the ITS when the transition is complete.</p>
	<hr/> <p>Note</p> <p>If the assignment of an input wire to an Interrupt Domain is changed when IWB_CR0.IDLE is 0, and before a corresponding CLEAR event is Accepted by the ITS, the IWB does not guarantee that a CLEAR event is sent or Accepted by the ITS when IWB_CR0.IDLE is 1.</p> <p>See 6.2 IWB support for multiple Interrupt Domains for more information about assigning input wires to Interrupt Domains.</p> <hr/>
R _{XCTCW}	<p>If the IWB is enabled or disabled when IWB_WENABLE_STATUSR.IDLE is 0 for any PAS, it is CONSTRAINED UNPREDICTABLE whether the IWB generates events for level-sensitive wires affected by the write to ongoing write to IWB_WENABLER<n>.</p> <p>See 6.1 IWB wire control registers for more information about individually enabling and disabling wires.</p>
R _{XDGGG}	<p>Following a write to IWB_CR0.IWBEN, when IWB_CR0.IDLE is 1, all events generated as a result of the write are Accepted by the ITS.</p>
S _{XWPYQ}	<p>When the IWB is reset, software is expected to configure wired interrupts using a sequence similar to the following:</p> <ul style="list-style-type: none"> • Software executing in a Security state that can access the MPPAS of the IWB performs the following steps: <ol style="list-style-type: none"> 1. Individually disable all input wires to the IWB. 2. Initialize the Interrupt Domain association for the input wires in the IWB. 3. Enable the IWB. • Software executing in any Security state performs the following steps to configure a wired interrupt in an Interrupt Domain that the input wire is assigned to: <ol style="list-style-type: none"> 1. Enable and configure the IRS, including providing the necessary data structures in memory. 2. Enable and configure the associated ITS, including providing the necessary data structures in memory.

3. Configure a valid translation in the ITS for the DeviceID and EventID associated with the IWB and input wire.
4. Configure the input wire in the IWB as edge-triggered or level-sensitive.
5. Individually enable the wire in the IWB.

See also:

- [Chapter 4 *Interrupt routing service \(IRS\)*](#)
- [4.8.2 *Physical SPIs*](#)
- [Chapter 5 *Interrupt translation service \(ITS\)*](#)

6.1 IWB wire control registers

D _{FKXGT}	<p>The IWB allows software to control the behavior of the IWB for each input wire through the following <i>wire control registers</i>:</p> <ul style="list-style-type: none"> IWB_WENABLER<n>: Enables and disables individual wires. IWB_WTMR<n>: Configures if individual wires are edge-triggered or level-sensitive. IWB_WDOMAINR<n>: Selects the Interrupt Domain that the individual wires are assigned to.
D _{WLWFZ}	<p>The following terms are used to indicate programming of a single wire as opposed to a configuration setting which applies to all of the IWB:</p> <ul style="list-style-type: none"> <i>Individually enabling</i> a wire refers to a write to IWB_WENABLER<n>.WEN<x> that updates the value from 0 to 1. <i>Individually disabling</i> a wire refers to a write to IWB_WENABLER<n>.WEN<x> that updates the value from 1 to 0.
R _{SDJBH}	<p>When individually enabling one or more wires, the IWB behaves as follows:</p> <ul style="list-style-type: none"> For wires where IWB_WTMR<n>.TM<x> is 1 (level-sensitive): <ul style="list-style-type: none"> If the wire is asserted at the time of the write to IWB_WENABLER<n>, the IWB generates a SET_LEVEL event. If the wire is de-asserted at the time of the write to IWB_WENABLER<n>, the IWB generates a CLEAR event. For wires where IWB_WTMR<n>.TM<x> is 0 (edge-triggered): <ul style="list-style-type: none"> The IWB does not generate any events for the wire. <p>When individually disabling one or more wires, the IWB behaves as follows:</p> <ul style="list-style-type: none"> For wires where IWB_WTMR<n>.TM<x> is 1 (level-sensitive): <ul style="list-style-type: none"> If a SET_LEVEL event was previously generated for that wire without a corresponding CLEAR event, the IWB generates a CLEAR event. For wires where IWB_WTMR<n>.TM<x> is 0 (edge-triggered): <ul style="list-style-type: none"> The IWB does not generate any events for the wire. <p>Writing to IWB_WENABLER<n> when IWB_CR0.IWBEN is 0 has no effect other than changing the values stored in IWB_WENABLER<n>.</p>
I _{WLHND}	<p>The effects of a write to IWB_WENABLER<n> are complete when a read using the same PAS as the write of IWB_WENABLE_STATUSR.IDLE returns 1.</p>
R _{HJXNH}	<p>When the effects of a write to IWB_WENABLER<n> are complete, all of the following are true:</p> <ul style="list-style-type: none"> Each input wire affected by the write becomes either disabled or enabled. All CLEAR events generated as a result of individually disabling wires are Accepted by the ITS.

Note

The IWB is required to generate a CLEAR event as a result of individually disabling a level-sensitive wire. The effects of the write to IWB_WENABLER<n> are not complete until the CLEAR events are accepted by the ITS.

The IWB is not required to generate any events and ensure that they are Accepted by the ITS before IWB_WENABLE_STATUSR.IDLE is 1 when individually enabling a wire.

S _{DDQDS}	<p>The IWB is required to generate CLEAR events when individually disabling a wire to support the following use case:</p> <ul style="list-style-type: none"> To re-sample the level of a level-sensitive interrupt, software can rely on a CLEAR being sent when the effects of individually disabling a wire are complete, and subsequent SET_EDGE or SET_LEVEL events
--------------------	--

being sent in finite time.

- Software may remove a mapping for an event in the ITS and ensure that the ITS does not report failed translations for that wire due to events being sent by the IWB after the wire has been disabled.

Note

If the assignment of an input wire to an Interrupt Domain is changed when IWB_WENABLE_STATUSR.IDLE is 0 due to a write to IWB_WENABLER<n> that disables that wire, and before a corresponding CLEAR event is Accepted by the ITS, the IWB does not guarantee that a CLEAR event is sent or Accepted by the ITS when IWB_WENABLE_STATUSR.IDLE is 1.

See [6.2 IWB support for multiple Interrupt Domains](#) for more information about assigning input wires to Interrupt Domains.

S_{TFNWQ}

Software can poll IWB_WENABLE_STATUSR.IDLE until it is 1, using the same PAS as was used to individually disable the wire, to ensure that all events for a wire that has been disabled are Accepted by the ITS and that no further events will be generated for that wire.

Software can rely on this functionality combined with an ITS synchronization request to ensure that all events are translated by the ITS before unmapping the event to avoid spurious translation errors being logged by the ITS.

I_{VXFZX}

Arm strongly recommends that a translation exists in the ITS for the wire before individually enabling the wire, and that the translation is not changed or made invalid in the ITS until the wire is individually disabled.

Further, Arm strongly recommends that a wire is individually disabled before changing the Interrupt Domain that the wire is assigned to. See [6.2 IWB support for multiple Interrupt Domains](#) for more information.

R_{HWHSV}

It is IMPLEMENTATION DEFINED whether software is permitted to program the Trigger mode of individual wires.

If software is not permitted to program the Trigger mode, any access to the Trigger mode of individual wires is **RO**.

I_{YQYHL}

The permission to program the Trigger mode is expressed by the pseudocode function `IsWireConfigRO()`.

R_{MQFZD}

Writing to IWB_WTMR<n>.TM<x> when any of the following are true results in a CONSTRAINED UNPREDICTABLE behavior:

- The wire is enabled.
- IWB_WENABLE_STATUSR.IDLE is 0 due to a write to IWB_WENABLER<n> affecting that wire.

The CONSTRAINED UNPREDICTABLE behavior is one of the following:

- The write is IGNORED and the configuration is not updated.
- The configuration is updated and no events are generated as a result of changing the configuration.
- The configuration is updated and the IWB generates a SET_EDGE, SET_LEVEL, or CLEAR event to the ITS.

S_{TQWV}

Software is expected to configure the wire to be edge-triggered or level-sensitive when the wire is individually disabled.

I_{BNWSF}

Software may *resample* a wire by writing the input wire index to IWB_WRESAMPLER.

R_{QSYNQ}

A wire can only be resampled by writing the input wire index to IWB_WRESAMPLER using one of the following PA spaces:

- The MPPAS of the IWB.
- The PAS associated with the Interrupt Domain that the wire is assigned to.

Writing the input wire index to IWB_WRESAMPLER using any other PAS has no effects.

R_{HFBMBM}

Resampling an enabled wire has the following effects:

- For wires where IWB_WTMR<n>.TM<x> is 1 (level-sensitive), all of the following are true:
 - If the wire is asserted, the IWB generates a SET_LEVEL event.
 - If the wire is de-asserted, the IWB generates a CLEAR event.
- For wires where IWB_WTMR<n>.TM<x> is 0 (edge-triggered), all of the following are true:
 - If the wire is asserted, the IWB generates a SET_EDGE event.
 - If the wire is de-asserted, the IWB generates no events.

Resampling a disabled wire has no effects and generates no events.

R_{RKLLJP}

Events generated as a result of resampling a wire are generated in finite time.

R_{MQVNP}

If a wire is being individually enabled or disabled at the same time that a wire is resampled, it is CONSTRAINED UNPREDICTABLE whether the resample has any effects.

See also:

- [4.5 IRS synchronization requests](#)
- [5.2.4 ITS synchronization requests](#)
- [10.4 IWB register frames](#)

6.2 IWB support for multiple Interrupt Domains

R _{ZDJNJ}	An IWB only implements support for Interrupt Domains supported by PEs in the system. It is possible for an IWB to implement a subset of the Interrupt Domains that the PEs support.
I _{WYBFB}	For a system with more than one IWB, it is possible for each IWB to implement a different set of Interrupt Domains. The Interrupt Domains implemented by an IWB are reported in IWB_IDR0.INT_DOMS.
R _{ZLWKJ}	The MPPAS of an IWB is determined from the subset of Interrupt Domains implemented by the IWB as follows:

Implemented Interrupt Domains	IWB MPPAS
Non-secure	Non-secure
Secure	Secure
Non-secure, Secure, EL3	Secure
Non-secure, Realm, EL3	Root
Non-secure, Realm, Secure, EL3	Root

I _{VITYPK}	For example, if an IWB implements the Non-secure, Secure, and EL3 Interrupt Domains, the MPPAS of the IWB is the Secure PAS. If the PEs in the system implement FEAT_RME, the MPPAS of the PEs is the Root PAS. Software executing in either Root state or Secure state can access the MPPAS of the IWB.
R _{ZGJQC}	The ITS associated with the IWB provides an ITS Domain for each Interrupt Domain implemented by the IWB.
D _{TDVQF}	Each input wire connected to the IWB is <i>assigned</i> to an Interrupt Domain.
R _{QWBDZ}	The assignment of a wire to an Interrupt Domain is configured in IWB_WDOMAINR<n>. For each wire, the assignment of a wire to an Interrupt Domain is either fixed or software programmable.
S _{MMLFL}	The assignment of a wire to an Interrupt Domain is either fixed or software programmable. Software checks the configuration by attempting to assign a new Interrupt Domain for the wire in IWB_WDOMAINR<n> using the IWB MPPAS and reading back the value once IWB_WDOMAIN_STATUSR.IDLE is 1. If the value returned has not been updated, the wire assignment to the Interrupt Domain is fixed. Otherwise, the wire assignment is software programmable.
I _{NZYYK}	The Interrupt Domains supported by the IWB are reported in IWB_IDR0.INT_DOMS.
R _{GFGYY}	When the IWB communicates an event related to a wire to the ITS, the Interrupt Domain of the event is the Interrupt Domain that the wire is assigned to.
I _{VGSJN}	To provide isolation between Interrupt Domains, the configuration of an input wire is restricted based on the PAS used to write to IWB wire registers.
I _{PVVMN}	IWB_WDOMAINR<n> is only accessible in the MPPAS of the IWB. For any other PAS, the registers are RAZ/WI.
I _{VFDHJ}	The fields in IWB_WENABLER<n> and IWB_WTMR<n> are modified only through one of the following PA spaces: <ul style="list-style-type: none"> • The MPPAS of the IWB. • The PAS associated with the Interrupt Domain that the wire is assigned to.
I _{HTBXZ}	For example, in an implementation supporting all four Interrupt Domains, an input wire configured to be assigned to the Secure Interrupt Domain can only be configured and enabled when writing to the MMIO registers using either the Secure or Root PAS.

Chapter 6. Interrupt Wire Bridge (IWB)
6.2. IWB support for multiple Interrupt Domains

I _{FTMVV}	The effects of a write to IWB_WDOMAINR<n> are complete when IWB_WDOMAIN_STATUSR.IDLE is 1.
R _{RXCQB}	<p>When a write to IWB_WDOMAINR<n> is complete, for each wire that was affected by the write, the wire is assigned to the new Interrupt Domain.</p> <p>Events generated after the write to IWB_WDOMAINR<n>, but before the write completes, for wires affected by the write, are either signaled in the old Interrupt Domain or the new Interrupt Domain of the wires.</p>
I _{YZCLL}	The IWB does not generate any events as a result of a write to IWB_WDOMAINR<n>.
R _{BMZMZ}	<p>If the IWB is being enabled or disabled at the same time that a wire is being assigned to a new Interrupt Domain, the IWB is not required to generate any events for that wire as a result of enabling or disabling the IWB.</p> <p>See Chapter 6 Interrupt Wire Bridge (IWB) for more information about enabling and disabling the IWB.</p>
R _{TXFJS}	<p>If a wire is being individually enabled or disabled at the same time that a wire is being assigned to a new Interrupt Domain, the IWB is not required to generate any events for that wire as a result of individually enabling or disabling the wire.</p> <p>See 6.1 IWB wire control registers for more information about individually enabling and disabling wires.</p>
S _{XWKVY}	<p>Software can use the following sequence to change the assignment of a wire from an old Interrupt Domain to a new Interrupt Domain:</p> <ol style="list-style-type: none"> 1. Software in the old Interrupt Domain performs the following actions: <ol style="list-style-type: none"> 1. Disable the interrupt in the IRS where the wire is signaled. 2. Individually disable the wire in the IWB. 3. Wait until IWB_WENABLE_STATUSR.IDLE is 1 to ensure that no further events will be generated for the wire. 4. Perform a synchronization request in the ITS and wait until it has completed. 5. Unmap the event from the ITS and invalidate the required ITS caches. 6. Request that software running in a Security state that can access the MPPAS of the IWB, assigns the input wire to the new Interrupt Domain. 2. Software in the Security state that can access the MPPAS of the IWB performs the following actions: <ol style="list-style-type: none"> 1. Writes to IWB_WDOMAINR<n>.WDDOM<x> to assign the wire to the new Interrupt Domain. 2. Waits until IWB_WDOMAIN_STATUSR.IDLE is 1. 3. Software in the Security state that can access the MPPAS of the IWB notifies software in the new Interrupt Domain that the wire has been re-assigned. 4. Software in the new Interrupt Domain performs the following actions: <ol style="list-style-type: none"> 1. Allocate and configure an interrupt in the IRS. 2. Map the wire event in the ITS for the new Interrupt Domain. 3. Individually enable the wire in the IWB.

See also:

- [10.2 IRS register frames](#)
- [10.4 IWB register frames](#)

Chapter 7

GIC Performance Monitoring Unit (PMU)

D _{WQRHY}	A GICv5 system may implement one or more <i>GIC Performance Monitoring Units (PMU)</i> that counts events for one or more GICv5 system components.
R _{GSJJJ}	A GIC PMU is an implementation of the <i>Arm® CoreSight™ Architecture Performance Monitoring Unit Architecture</i> [11].
I _{CXHZH}	The GIC PMU registers for each GIC PMU are accessible in a separate 64KB aligned GIC_PMU_FRAME.
R _{QZHLT}	The base address of the GIC_PMU_FRAME is IMPLEMENTATION DEFINED.
S _{BPXJG}	Arm expects that firmware describes the base address of ITS PMU register frames to system software using firmware data structures.
R _{VGHTV}	For a GIC PMU, monitor 0 is 64 bits. The sizes of all other monitors are IMPLEMENTATION DEFINED. This means that PMCFGR.SIZE is 0b111111.
S _{HTCFV}	Arm expects that monitor 0 is used to count GIC cycles.
D _{LTMWH}	A GIC PMU counts events from GIC components. This is referred to as the <i>agent being monitored</i> in the <i>Arm® CoreSight™ Architecture Performance Monitoring Unit Architecture</i> [11]. The agent being monitored is one of the following: <ul style="list-style-type: none">• The GIC PMU counts events for a single ITS.• The GIC PMU counts events for a single IRS.• The GIC PMU counts events for a single IRS and all of its associated ITSs.
R _{FKJCW}	When a GIC PMU counts events for a system component that supports multiple Interrupt Domains, the GIC PMU counts events for all the supported Interrupt Domains.

Figure 7.1 shows different implementation options for GIC PMUs in a GICv5 system.

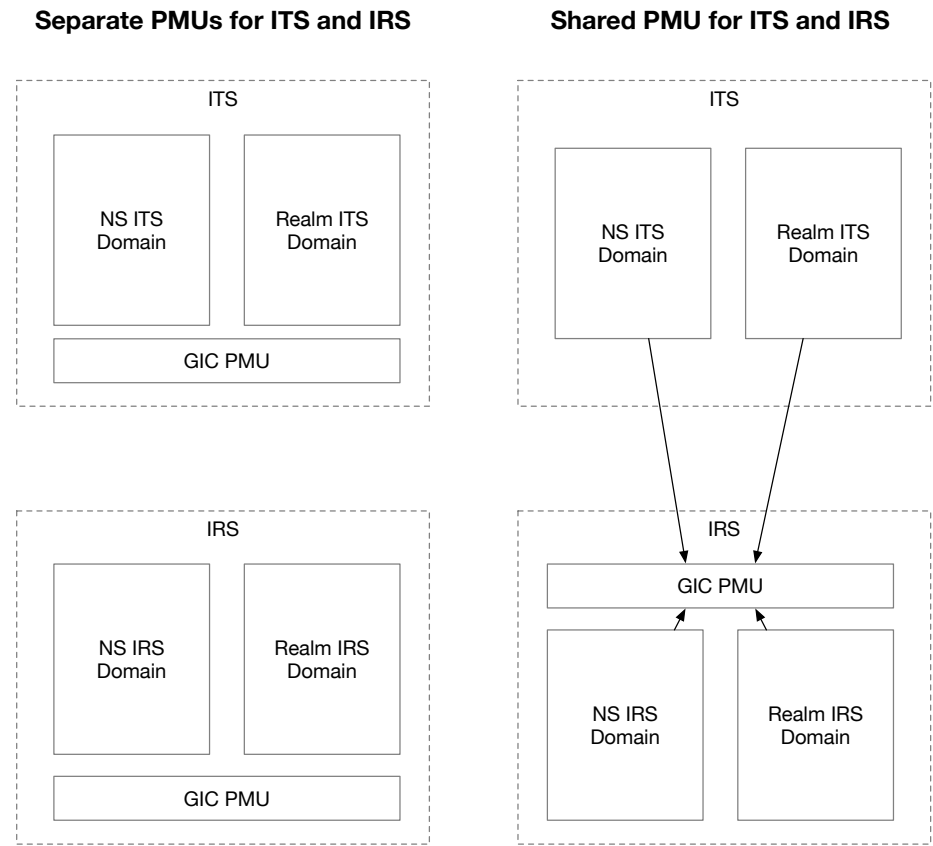


Figure 7.1: GIC PMU implementation options

Note

The GIC PMU implementation options are shown using only the Non-secure and Realm Interrupt Domains for illustrative purposes. The GIC PMU is also supported for the Secure and EL3 Interrupt Domains.

I_{HLZWW}

GIC_PMDR0.IRS_PMU reports whether the GIC PMU counts events for an IRS.

GIC_PMDR0.ITS_PMU reports whether the GIC PMU counts events for one or more ITSs.

R_{BMCHW}

GIC_PMDR0.{IRS_PMU, ITS_PMU} is not permitted to be {0, 0}.

See also:

- [10.5.1 GIC_PMU_FRAME, GIC PMU register frame](#)

7.1 CoreSight PMU extensions

R _{FVCHK}	A GIC PMU implements the 64-bit programmers' model extension defined in [11].
I _{SCHGW}	A GIC PMU implements the direct access to enables and overflows extension because it is always implemented when the 64-bit programmers' model extension is implemented.
R _{TGDSP}	<p>A GIC PMU implements the observability and access control extension defined in [11] if, for any of the components in the agent being monitored, all of the following are true:</p> <ul style="list-style-type: none"> • Multiple Interrupt Domains are supported. • Any of the following are true: <ul style="list-style-type: none"> – The agent supports the Secure Interrupt Domain. – The agent supports the EL3 Interrupt Domain. <p>Otherwise, it is IMPLEMENTATION DEFINED whether the GIC PMU implements the observability and access control extension.</p>
R _{SKSTY}	If a GIC PMU does not implement the observability and access control extension, it is IMPLEMENTATION DEFINED which physical address space is used to access the GIC PMU registers.
I _{JMZMR}	GIC_PMDR0.OACE reports whether the observability and access control extension is implemented.
R _{QZCCM}	<p>When the Observability and access control extension is implemented, all of the following are true:</p> <ul style="list-style-type: none"> • It is IMPLEMENTATION DEFINED whether PMSCR and PMROOTCR are implemented in the GIC PMU register frame or at an IMPLEMENTATION DEFINED location. • PMSCR is only accessible in the MPPAS for the system.
R _{PMHYN}	<p>When the Observability and access control extension is implemented, and the agent being monitored supports the Realm interrupt domain, all of the following are true:</p> <ul style="list-style-type: none"> • The configuration of PMSCR.NSRA has no impact on the access controls to the GIC PMU configuration registers. • The configuration of PMSCR.NSMSI has no impact on physical address space used for message-signaled interrupts from the GIC PMU. • The following fields are added to PMROOTCR, Root and Realm Control Register:

RA, bits [6:4] Register Access.

This field determines physical address space for which register access is enabled for the PMU.

RA	Description
0b000	Root register access is enabled. Access from other address spaces is disabled, meaning accesses to all PMU registers are RAZ/WI.
0b001	Root and Realm register access is enabled. Access from other address spaces is disabled, meaning accesses to all PMU registers are RAZ/WI.
0b010	Root and Secure register access is enabled. Access from other address spaces is disabled, meaning accesses to all PMU registers are RAZ/WI.
0b011	All access is enabled.

Other values are reserved.

For the CoreSight management registers, 0xFA8 to 0xFFC, it is IMPLEMENTATION DEFINED whether these registers are RO or RAZ/WI when register access is disabled by this field.

The reset value of this field is IMPLEMENTATION DEFINED, and depends on the security policy of the component implementing this register.

RMSI, bits [11:10] Root control for MSI PA space

When PMCFGR.MSI is 1, this field determines the physical address space used for MSIs.

RA	Description
0b000	MSIs are generated using the Root PAS. If PMROOTCR.RA is not 0b000, this value results in the Non-secure PAS being used for MSIs.
0b001	MSIs are generated using the Realm PAS. If PMROOTCR.RA is 0b010, this value results in the Non-secure PAS being used for MSIs.
0b010	MSIs are generated using the Secure PAS. If PMROOTCR.RA is 0b0x1, this value results in the Non-secure PAS being used for MSIs.
0b011	MSIs are generated using the Non-secure PAS.

Other values are reserved.

The reset value of this field is IMPLEMENTATION DEFINED, and depends on the security policy of the component implementing this register.

R _{TSZTF}	It is IMPLEMENTATION DEFINED whether GIC PMU implements the Freeze on overflow extension defined in [11].
R _{KLJSW}	A GIC PMU does not implement the Halt-on-debug extension defined in [11]. For GIC PMU, PMCFGR.HDBG is 0.
R _{XHQVD}	A GIC PMU does not implement the Fixed-function cycle counter extension defined in [11].
R _{VQYFY}	A GIC PMU does not implement the monitor group extension defined in [11]. For GIC PMU, PMCFGR.NCG is 0.
R _{CZMLC}	A GIC PMU does not implement the Counter chaining extension defined in [11].
R _{DSXHC}	A GIC PMU does not implement the event counter threshold and edge detection extensions defined in [11].
I _{SKNKR}	The <i>Arm® CoreSight™ Architecture Performance Monitoring Unit Architecture</i> [11] defines Reusable event filter definitions. When a GIC PMU implements the observability and access control extension, the GIC PMU implements the Security operating state filtering. Otherwise, the GIC PMU does not use the Security operating state filtering. See 7.4 Event filtering for more information. No other Reusable event filter definitions are used.
R _{JZTKK}	It is IMPLEMENTATION DEFINED whether GIC PMU implements support for the Snapshot extension defined in [11].
I _{ZRJLK}	If the Snapshot extension is implemented, PMCFGR.SS is 1.
R _{XTRZB}	It is IMPLEMENTATION DEFINED whether the GIC PMU does implements the Trace generation extension defined in [11].
R _{RRYKG}	A GIC PMU does not implement the Export extension defined in [11]. For GIC PMU, PMCFG.EX is 0.
R _{GYDWW}	A GIC PMU does not implement the Dual-page extension defined in [11].

7.2 GIC PMU Overflow interrupt

R_{HVTMV}

A GIC PMU implements an overflow interrupt.

The permitted overflow interrupt types depend on which type of agent is being monitored as follows:

- If the GIC PMU counts events for an IRS, or for an IRS and its associated ITSs, the overflow interrupt is an SPI connected to that IRS.
- If the GIC PMU counts events for a single ITS, the overflow interrupt is one of the following:
 - An interrupt signal connected to an SPI on the associated IRS.
 - An interrupt wire signal connected to an IWB.
 - An ITS event generated from the ITS PMU as a system peripheral and translated by the ITS.

R_{DRWTL}

If the overflow interrupt is an ITS event generated from the GIC PMU as a system peripheral and translated by the ITS, all of the following are true:

- The GIC PMU uses the message-signaled interrupt functionality to configure the overflow interrupt as defined in [11].
- PMCFGR.MSI is 1.
- PMIRQCR0 is IGNORED and access to the register is permitted to be RES0 or RAZ/WI.
- The DeviceID for the ITS event is IMPLEMENTATION DEFINED.
- The EventID for the ITS event is programmed using PMIRQCR1.DATA.
- The originating Interrupt Domain is the Interrupt Domain corresponding to the physical address space configured for the message-signaled interrupt.

I_{LLFNW}

The *Arm® CoreSight™ Architecture Performance Monitoring Unit Architecture* [11] defines rules and controls that specify the physical address space used for message-signaled interrupts.

These rules are extended to also support the Root and Realm physical address space for message-signaled interrupts. See [7.1 CoreSight PMU extensions](#) for more information.

R_{PTFKG}

When Non-secure register access is enabled, message-signaled interrupts are always Non-secure.

Note

This replaces an incorrect statement in [11] that states the opposite rule for Non-secure register access and the MSI physical address space.

7.3 GIC PMU event types

D_{QTVKY}

GIC PMU events are identified by a 4-bit PMEVTTYPE value specifying an *event type*, and a 12-bit PMEVTID value specifying the *event* for the event type.

The PMEVTTYPE space is interpreted as follows:

- 0b0000: Architected IRS event.
- 0b0001: IMPLEMENTATION DEFINED IRS event.
- 0b0010: Architected ITS event.
- 0b0011: IMPLEMENTATION DEFINED ITS event.

All PMEVTTYPE values not defined above are reserved.

See [7.5 IRS PMU events](#) and [7.6 ITS PMU events](#) for information about the architected events.

I_{RWFQW}

An implementation may implement support for fewer than 16 bits of PMEVTID. In this case, unimplemented upper bits are RES0.

I_{VSWDJ}

When an event is selected using GIC_PMEVTYPER<n>.{PMEVTTYPE, PMEVTID}, GIC_PMEVTYPER<n>.V reports whether an implemented event is selected.

7.4 Event filtering

I _{JLHGK}	The GIC PMU supports filtering of events using various filters. The available filters depend on the event type and the event. See 7.5 IRS PMU events and 7.6 ITS PMU events for information about which filters are supported for the event types and architected events.
I _{NYZYT}	When an event is selected using GIC_PMEVTYPER<n>.{PMEVTYPE, PMEVTID}, GIC_PMEVTYPER<n>.FS reports whether filtering is supported for the selected event.
R _{MZLWX}	When more than one filter for an event is enabled, the event is counted if and only if it matches on all enabled filters.
I _{RFZHB}	<p>When a GIC PMU implements the observability and access control extension, the GIC PMU supports filtering of events based on their association with the Security state using the Security operating state filtering defined as part of the Reusable event filter definitions defined in <i>Arm® CoreSight™ Architecture Performance Monitoring Unit Architecture</i>[11].</p> <p>The Security operating state filtering State match controls are implemented using GIC_PMEVTYPER<n>.{RL,EL3,NS,S}.</p> <p>By filtering events using the Security state, the GIC PMU filters events matching the corresponding Interrupt Domains.</p>
D _{DGSKD}	<p>On some filters, the GIC PMU supports <i>filtering on a range</i> of values instead of matching on an exact value.</p> <p>When filtering on a range, the range is specified by splitting the value into X most significant bits that must match, and Y least significant bits that are allowed to differ as follows:</p> <ul style="list-style-type: none"> • Bit[Y-1] is 0. That is, the most significant bit in the group of bits that do not need to match is zero. • Bits[(Y-2:0)] are 1. That is, the remaining bits that are allowed to differ are all programmed to 1. • The remainder of bits (the X most-significant bits, from bit Y upwards) contain a value that must match the corresponding bits for the value associated with monitored event.
I _{GSLNM}	GIC_PMEVTYPER<n>.FSPAN reports whether filtering by a range of values is supported for the selected event.
R _{MZTYJ}	For each event that may support filtering on a range of values, it is IMPLEMENTATION DEFINED whether filtering on a range of values is supported or filtering is only supported on an exact value.
S _{LKPXH}	Arm expects that when a virtual GIC PMU is presented to a VM, the virtual GIC PMU does not support filtering on a range of values for identifiers where ranges in the host system do not correspond to similar ranges in the guest system.
I _{VMFDS}	<p>The following are examples (in binary) of specifying ranges for GIC PMU filtering:</p> <ul style="list-style-type: none"> • 0000:0000:0001:1011:1111:0111:1111:0111 matches 0000:0000:0001:1011:1111:0111:1111:xxxx • 0000:0000:0001:1011:1111:0111:1111:0110 matches 0000:0000:0001:1011:1111:0111:1111:011x • 0000:0000:0001:1011:1111:0101:1111:1111 matches 0000:0000:0001:1011:1111:01xx:xxxx:xxxx

7.5 IRS PMU events

R_{JPY}XX The architected IRS PMU events are listed in [Table 7.3](#).

Table 7.3: Architected IRS PMU events

PMEVTYPE	PMEVTID	Name	Description
0	0	PMIRSCC	Counts IRS cycles.
0	1	PMIRSPENDSETIN	Counts incoming interrupt events to make an interrupt Pending. The events are counted when received from an SPI, an ITS, or another IRS.
0	2	PMIRSPENDCLEARIN	Counts incoming interrupt events to make an interrupt Idle. The events are counted when received from an SPI, an ITS, or another IRS.
0	3	PMIRSPENDSETOUT	Counts incoming interrupt events to make an interrupt Pending that are forwarded to another IRS.
0	4	PMIRSPENDCLRROUT	Counts incoming interrupt events to make an interrupt Idle that are forwarded to another IRS.
0	5	PMIRSPENDSETPRO	Counts interrupt events to make an interrupt Pending that are processed at this IRS.
0	6	PMIRSPENDCLRPRO	Counts interrupt events to make an interrupt Idle that are processed at this IRS.
0	7	PMIRSPECMDPEND	Counts commands received from a PE to update an interrupt Pending state.
0	8	PMIRSPECMDHANDLE	Counts commands received from a PE related to handling of an interrupt. Includes only Activate and Deactivate commands.
0	9	PMIRSPECMDCFG	Counts commands received from a PE to update or request the configuration and state of an interrupt.
0	10	PMIRSVPEMIGR	Counts when a VPE is made resident on an IRS and was not previously resident on that IRS.
0	11	PMIRSISTMISS	IST cache miss. Counts for each interrupt effect requiring an access to the IST.
0	12	PMIRSVMTMISS	VM table cache miss. Counts for each interrupt effect to a virtual interrupt requiring an access to the VM table.
0	13	PMIRSVPETMISS	VPE table cache miss. Counts for each interrupt effect to a virtual interrupt requiring an access to the VPE table.
0	14	PMIRSVMDMISS	VM descriptor miss. Counts for each interrupt effect to a virtual interrupt requiring an access to the VM descriptor.
0	15	PMIRSVPEDMISS	VPE descriptor miss. Counts for each interrupt effect to a virtual interrupt requiring an access to the VPE descriptor.
0	16	PMIRSL1ISTLU	Level 1 IST table access. Counts for each memory access to a level 1 IST.
0	17	PMIRSL2ISTLU	Level 2 IST table access. Counts for each memory access to a level 2 IST.

PMEVTYPE	PMEVTID	Name	Description
0	18	PMIRSL1VMTLU	Level 1 VM table access. Counts for each memory access to a level 1 VM table.
0	19	PMIRSL2VMTLU	Level 2 VM table access. Counts for each memory access to a level 2 VM table.
0	20	PMIRSVPETLU	VPE table access. Counts for each memory access to a VPE table.
0	21	PMIRSVMDLU	VM descriptor access. Counts for each memory access to a VM descriptor.
0	22	PMIRSVPEDLU	VPE descriptor access. Counts for each memory access to a VPE descriptor.
0	23	PMIRSVPEDB	Counts when a VPE doorbell event is generated.
0	24	PMIRS1NPESEL	Counts when a new PE or VPE selected for a 1ofN interrupt.
0	25	PMIRS1NDB	Counts when the 1ofN doorbell configuration for a VM causes a VPE doorbell to be generated.

All other PMEVTID values for the architected IRS event type are reserved.

R _{GFFRF}	In a multi-IRS system, the events are counted only for the IRS being monitored by the GIC PMU.
S _{ZKMFB}	Software is expected to combine events from multiple GIC PMUs to gather counts of events globally for the system.
R _{CTPKW}	When the IRS cycle counter event PMIRSCC is monitored, the event always counts IRS cycles, unless the IRS is in a low power state.
R _{HBFMK}	When an incoming interrupt event is received at the same IRS that processed the event, the event is counted both as incoming and as processing for the same interrupt event.
R _{KKJVF}	Table 7.4 lists the GIC System instructions that are processed by the IRS and recorded by IRS PMU events when the PE indicates to the IRS that the instruction is being executed. See 2.6 GIC System instructions for more information.

Table 7.4: IRS PMU events counting processing of GIC System instructions

PMIRSPECMDPEND	PMIRSPECMDHANDLE	PMIRSPECMDCFG
GIC xDPEND	GICR CDIA	GIC xDDIS
	GICR CDNMA	GIC xDEN
	GIC xDDI	GIC xDRCFG
		GIC xDPRI
		GIC xDAFF
		GIC xDHM

R _{XJKZC}	When monitoring any of the following events, the event is counted when the IRS performs any access to IRS data structures in memory as part of processing an Interrupt Effect or when making a VPE resident: <ul style="list-style-type: none"> • 11: PMIRSISTMISS. • 12: PMIRSVMTMISS.
--------------------	---

- 13: PMIRSVPETMISS.
- 14: PMIRSVMDMISS.
- 15: PMIRSVPEDMISS.

R_{JJNGJ}

The IRS data structure access events are counted for an access to the corresponding IRS data structure in memory for any of the following reasons:

- There was an access to the data structure related to an Interrupt Effect.
- There was an access to the data structure related to making a VPE resident.
- There was a speculative access performed by the IRS.

7.5.1 IRS PMU events filtering

R_{CSDPD}

The GIC PMU supports the following filter types for the IRS events:

- Filtering on INTID.
- Filtering on physical or virtual interrupts, and a specific VM ID for virtual interrupts (VM and VM ID).
- Filtering on whether an access to an IRS data structure was a read or a write (RW).
- Filtering on the which PE sent a PE command to the IRS (Source PE).

Table 7.5 shows which filters are supported for the architected IRS events.

Table 7.5: Supported filters for IRS PMU events

PMEVTID	Name	INTID filter	VM and VM ID filter	RW filter	Source PE filter
0	PMIRSCC	No	No	No	No
1	PMIRSPENDSETIN	Yes	Yes	No	No
2	PMIRSPENDCLEARIN	Yes	Yes	No	No
3	PMIRSPENDSETOUT	Yes	Yes	No	No
4	PMIRSPENDCLROUT	Yes	Yes	No	No
5	PMIRSPENDSETPRO	Yes	Yes	No	No
6	PMIRSPENDCLRPRO	Yes	Yes	No	No
7	PMIRSPECMDPEND	Yes	Yes	No	Yes
8	PMIRSPECMDHANDLE	Yes	Yes	No	Yes
9	PMIRSPECMDCFG	Yes	Yes	No	Yes
10	PMIRSVPEMIGR	No	Yes ^a	No	Yes
11	PMIRSISTMISS	Yes	Yes	Yes	No
12	PMIRSVMTMISS	No	Yes	Yes	No
13	PMIRSVPETMISS	No	Yes	Yes	No
14	PMIRSVMDMISS	No	Yes	Yes	No
15	PMIRSVPEDMISS	No	Yes	Yes	No
16	PMIRSL1ISTLU	Yes	Yes	Yes	No
17	PMIRSL2ISTLU	Yes	Yes	Yes	No
18	PMIRSL1VMTLU	No	Yes	Yes	No
19	PMIRSL2VMTLU	No	Yes	Yes	No

PMEVTID	Name	INTID filter	VM and VM ID filter	RW filter	Source PE filter
20	PMIRSVPETLU	No	Yes	Yes	No
21	PMIRSVMDLU	No	Yes	Yes	No
22	PMIRSVPEDLU	No	Yes	Yes	No
23	PMIRSVPEDB	No	Yes	No	No
24	PMIRS1NPESEL	Yes	Yes	No	No
25	PMIRS1NDB	No	Yes	No	No

- a. The event is not counted when filtering on physical events only.

All other PMEVTID values for the architected IRS event type are reserved.

7.6 ITS PMU events

R_{GCNFI} The architected ITS PMU events are listed in [Table 7.6](#).

Table 7.6: Architected ITS PMU events

PMEVTYPE	PMEVTID	Name	Description
2	0	PMITSCC	Counts ITS cycles.
2	1	PMITSTRQ	Counts for each translation request performed by the ITS.
2	2	PMITSDIDMISS	DeviceID cache miss. Counts for each translation of a DeviceID that could not be satisfied from cache.
2	3	PMITSEIDMISS	EventID cache miss. Counts for each translation of an EventID that could not be satisfied from cache.
2	4	PMITSL1DTLU	L1 device table lookup. Counts every time a level 1 DTE is looked up in memory. Does not count when the entry is read from cache.
2	5	PMITSL2DTLU	L2 device table lookup. Counts every time a level 2 DTE is looked up in memory. Does not count when the entry is read from cache.
2	6	PMITSL1ITTLU	L1 interrupt translation table lookup. Counts every time a level 1 ITTE is looked up in memory. Does not count when the entry is read from cache.
2	7	PMITSL2ITTLU	L2 interrupt translation table lookup. Counts every time a level 2 ITTE is looked up in memory. Does not count when the entry is read from cache.

All other PMEVTID values for the architected ITS event type are reserved.

R_{MSFLT}	When the ITS cycle counter event PMITSCC is monitored, it counts regardless of the event being translated by the ITS, unless the ITS is in a low power state.
R_{YFQW}	The ITS DeviceID and EventID cache miss events, PMITSDIDMISS and PMITSEIDMISS, are counted when the translation causes any access to the ITS translation structures in memory to perform the translation.
R_{ZJTPS}	The ITS table lookup events are counted for each access to an ITS translation structure in memory caused by performing a translation or by speculative accesses.
R_{GKBVK}	ITS events are counted regardless of whether the translation was successful or not.

7.6.1 ITS PMU events filtering

R_{HXGVQ}	A GIC PMU supports filtering of ITS events by any of the following combinations of DeviceID and EventID: <ul style="list-style-type: none"> • No filtering on DeviceID or EventID. • DeviceID. • DeviceID and EventID.
R_{LNVGM}	Filtering on DeviceID and EventID is supported for all other ITS events than ITS event 0 PMITSCC.
I_{ZJXHT}	When a GIC PMU counts events for an IRS and multiple associated ITSs, a monitor always counts ITS events for a single ITS specified by GIC_PMEVFILT2R<n>.ITSID.

I_{XGBDQ}

When GIC_PMEVFILT2R<n>.FILTER_DID is 0, the monitor counts the selected PMU event for all translation requests.

When GIC_PMEVFILT2R<n>.FILTER_DID is 1 and GIC_PMEVFILT2R<n>.FILTER_DID_SPAN is 0, the monitor counts the selected PMU event for translation requests where the DeviceID matches the DeviceID specified in GIC_PMEVFILTR<n>.DEVICE_ID.

When GIC_PMEVFILT2R<n>.FILTER_DID is 1 and GIC_PMEVFILT2R<n>.FILTER_DID_SPAN is 1, the monitor counts the selected PMU event for translation requests where the DeviceID matches a range of DeviceIDs specified in GIC_PMEVFILTR<n>.DEVICE_ID.

If the ITS being monitored supports less than 2 bits of DeviceID, GIC_PMEVFILT2R<n>.FILTER_DID_SPAN is RES0 and matching a range of DeviceIDs is not supported.

I_{NTCKY}

When GIC_PMEVFILT2R<n>.FILTER_DID is 0, the monitor counts the selected PMU event for all translation requests, and filtering on EventID is not supported.

When GIC_PMEVFILT2R<n>.FILTER_DID is 1 and GIC_PMEVFILT2R<n>.FILTER_EID is 0, the monitor counts the selected PMU event for translation requests where the DeviceID matches the DeviceID filter value, for all EventIDs.

When GIC_PMEVFILT2R<n>.FILTER_DID is 1, GIC_PMEVFILT2R<n>.FILTER_EID is 1, and GIC_PMEVFILT2R<n>.FILTER_EID_SPAN is 0, the monitor counts the selected PMU event for translation requests where the DeviceID matches the DeviceID filter value, and the EventID matches the EventID specified in GIC_PMEVTILFTR<n>.EVENT_ID.

When GIC_PMEVFILT2R<n>.FILTER_DID is 1, GIC_PMEVFILT2R<n>.FILTER_EID is 1, and GIC_PMEVFILT2R<n>.FILTER_EID_SPAN is 1, the monitor counts the selected PMU event for translation requests where the DeviceID matches the DeviceID filter value, and the EventID matches the EventID matches a range of EventIDs specified in GIC_PMEVTILFTR<n>.EVENT_ID.

If the ITS supports less than 2 bits of EventID, GIC_PMEVFILT2R<n>.FILTER_EID_SPAN is RES0 and matching a range of EventIDs is not supported.

Chapter 8

System instructions

This chapter describes the GICv5 system instructions.

8.1 System instructions for the Current Interrupt Domain

8.1.1 GIC CDAFF, Interrupt Set Target in the Current Interrupt Domain

The GIC CDAFF characteristics are:

Purpose

Sets the routing mode and target PE for the specified INTID in the Current Interrupt Domain.

If the Current Interrupt Domain is the Virtual Interrupt Domain, this operation applies to the virtual INTID in the VM identified by the resident VPE.

Configuration


This instruction is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to GIC CDAFF are UNDEFINED.

Attributes

GIC CDAFF is a 64-bit System instruction.

Field descriptions

The GIC CDAFF bit assignments are:



Bits [63:48]

Reserved, RES0.

IAFFID, bits [47:32]

The interrupt Affinity value.

The IRS may support fewer than 16 bits of IAFFID. Unimplemented upper bits are treated as RES0 by the IRS.

When IRM is 1, this field provides an IMPLEMENTATION DEFINED hint to 1ofN selection algorithm. A value of 0 means that no hint is provided.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

IRM, bit [28]

Controls how the interrupt is routed.

If the GIC IRS does not support 1ofN distribution of the specified interrupt Type, this configuration is treated as RES0.

IRM	Meaning
0b0	The interrupt Routing mode is Targeted
0b1	The interrupt Routing mode is 1ofN

Bits [27:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC CDAFF

This instruction has no effect if the INTID is unreachable.

This instruction is treated as a NOP if the operation applies to the Virtual Interrupt Domain and there is no resident VPE.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC CDAFF, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b000	0b1100	0b0001	0b011

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGITR_EL2.GICCDAFF == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_AFF);
    else
        GIC(X[t, 64], GICInstrDomain_CD, GICInstr_AFF);
elseif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_AFF);
elseif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_AFF);

```

8.1.2 GIC CDDI, Interrupt Deactivate in the Current Interrupt Domain

The GIC CDDI characteristics are:

Purpose

Clears the Active state of the specified INTID at the CPU interface in the Current Interrupt Domain.

This instruction applies to all interrupt types, including PPIs.

If the Current Interrupt Domain is the Virtual Interrupt Domain and the interrupt type is a PPI, this operation applies to the virtual PPI identified by the specified INTID.

If the Current Interrupt Domain is the Virtual Interrupt Domain and the interrupt type is not a PPI, this operation applies to the virtual INTID in the VM identified by the resident VPE.

Configuration


This instruction is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to GIC CDDI are UNDEFINED.

Attributes

GIC CDDI is a 64-bit System instruction.

Field descriptions

The GIC CDDI bit assignments are:



Bits [63:33]

Reserved, RES0.

Bit [32]

Reserved, WI.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b001	PPI
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC CDDI

This instruction has no effect if the INTID is unreachable

This instruction is treated as a NOP if the operation applies to the Virtual Interrupt Domain, the interrupt type is not a PPI and there is no resident VPE.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC CDDI, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b000	0b1100	0b0010	0b000

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGITR_EL2.GICCDDI == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_DI);
    else
        GIC(X[t, 64], GICInstrDomain_CD, GICInstr_DI);
elsif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_DI);
elsif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_DI);

```

8.1.3 GIC CDDIS, Interrupt Disable in the Current Interrupt Domain

The GIC CDDIS characteristics are:

Purpose

Disables the specified INTID in the Current Interrupt Domain.

If the Current Interrupt Domain is the Virtual Interrupt Domain this operation applies to the virtual INTID in the VM identified by the resident VPE.

Configuration


This instruction is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to GIC CDDIS are UNDEFINED.

Attributes

GIC CDDIS is a 64-bit System instruction.

Field descriptions

The GIC CDDIS bit assignments are:



Bits [63:32]

Reserved, RES0.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Valued not defined above are reserved

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC CDDIS

This instruction has no effect if the INTID is unreachable

This instruction is treated as a NOP if the operation applies to the Virtual Interrupt Domain and there is no resident VPE.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC CDDIS, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b000	0b1100	0b0001	0b000

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGITR_EL2.GICCDDIS == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_DIS);
    else
        GIC(X[t, 64], GICInstrDomain_CD, GICInstr_DIS);
elsif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_DIS);
elsif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_DIS);

```

8.1.4 GIC CDEN, Interrupt Enable in the Current Interrupt Domain

The GIC CDEN characteristics are:

Purpose

Enables the specified INTID in the Current Interrupt Domain.

If the Current Interrupt Domain is the Virtual Interrupt Domain, this operation applies to the virtual INTID in the VM identified by the resident VPE.

Configuration


This instruction is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to GIC CDEN are UNDEFINED.

Attributes

GIC CDEN is a 64-bit System instruction.

Field descriptions

The GIC CDEN bit assignments are:



Bits [63:32]

Reserved, RES0.

TYPE, bits [31:29]

Type of the interrupt.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC CDEN

This instruction has no effect if the INTID is unreachable

This instruction is treated as a NOP if the operation applies to the Virtual Interrupt Domain, the interrupt type is not a PPI and there is no resident VPE.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC CDEN, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b000	0b1100	0b0001	0b001

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGITR_EL2.GICCDEN == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_EN);
    else
        GIC(X[t, 64], GICInstrDomain_CD, GICInstr_EN);
elseif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_EN);
elseif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_EN);

```

8.1.5 GIC CDEOI, Priority Drop in the Current Interrupt Domain

The GIC CDEOI characteristics are:

Purpose

Performs a Priority Drop of the running priority at the CPU interface in the Current Interrupt Domain.

If the Current Interrupt Domain is the Virtual Interrupt Domain, the Priority Drop applies to the virtual running priority.

Otherwise, it applies to the physical running priority in the Current Physical Interrupt Domain.

Configuration


This instruction is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to GIC CDEOI are UNDEFINED.

Attributes

GIC CDEOI is a 64-bit System instruction.

Field descriptions

The GIC CDEOI bit assignments are:



Bits [63:0]

Reserved, RES0.

Accessing GIC CDEOI

The Rt field should be set to 0b11111. If the Rt field is not set to 0b11111, it is CONSTRAINED UNPREDICTABLE whether:

- The instruction is UNDEFINED.
- The instruction behaves as if the Rt field is set to 0b11111.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC CDEOI, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b000	0b1100	0b0001	0b111

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGITR_EL2.GICCDEOI == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_EOI);
    else

```


Chapter 8. System instructions

8.1. System instructions for the Current Interrupt Domain

```
        GIC(X[t, 64], GICInstrDomain_CD, GICInstr_EOI);  
elseif PSTATE.EL == EL2 then  
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_EOI);  
elseif PSTATE.EL == EL3 then  
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_EOI);
```

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC CDHM

This instruction has no effect if the INTID is unreachable

This instruction is treated as a NOP if the operation applies to the Virtual Interrupt Domain and there is no resident VPE.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC CDHM, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b000	0b1100	0b0010	0b001

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGITR_EL2.GICCDHM == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_HM);
    else
        GIC(X[t, 64], GICInstrDomain_CD, GICInstr_HM);
elseif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_HM);
elseif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_HM);

```

8.1.7 GIC CDPEND, Interrupt Set/Clear Pending state in the Current Interrupt Domain

The GIC CDPEND characteristics are:

Purpose

Generates a SET or CLEAR event for the specified INTID in the Current Interrupt Domain.
If the Current Interrupt Domain is the Virtual Interrupt Domain this operation applies to the virtual INTID in the VM identified by the resident VPE.

Configuration


This instruction is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to GIC CDPEND are UNDEFINED.

Attributes

GIC CDPEND is a 64-bit System instruction.

Field descriptions

The GIC CDPEND bit assignments are:



Bits [63:33]

Reserved, RES0.

PENDING, bit [32]

Pending status

PENDING	Meaning
0b0	Generate CLEAR event to the IRS.
0b1	Generate SET event to the IRS.

TYPE, bits [31:29]

Type of the interrupt.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC CDPEND

This instruction has no effect if the INTID is unreachable

This instruction is treated as a NOP if the operation applies to the Virtual Interrupt Domain and there is no resident VPE.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC CDPEND, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b000	0b1100	0b0001	0b100

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGITR_EL2.GICCDPEND == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_PEND);
    else
        GIC(X[t, 64], GICInstrDomain_CD, GICInstr_PEND);
elsif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_PEND);
elsif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_PEND);

```

8.1.8 GIC CDPRI, Interrupt Set priority in the Current Interrupt Domain

The GIC CDPRI characteristics are:

Purpose

Sets the priority for the specified INTID in the Current Interrupt Domain.

If the Current Interrupt Domain is the Virtual Interrupt Domain this operation applies to the virtual INTID in the VM identified by the resident VPE.

Configuration


This instruction is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to GIC CDPRI are UNDEFINED.

Attributes

GIC CDPRI is a 64-bit System instruction.

Field descriptions

The GIC CDPRI bit assignments are:



Bits [63:40]

Reserved, RES0.

PRIORITY, bits [39:35]

The priority of the specified INTID

Only the upper N bits of the Priority field are implemented where $N = (ICC_IDR0_EL1.PRI_BITS + 1)$. Unimplemented bits are RES0.

Bits [34:32]

Reserved, RES0.

TYPE, bits [31:29]

Type of the interrupt.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC CDPRI

This instruction has no effect if the INTID is unreachable

This instruction is treated as a NOP if the operation applies to the Virtual Interrupt Domain and there is no resident VPE.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC CDPRI, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b000	0b1100	0b0001	0b010

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGITR_EL2.GICCDPRI == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_PRI);
    else
        GIC(X[t, 64], GICInstrDomain_CD, GICInstr_PRI);
elseif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_PRI);
elseif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_PRI);

```

8.1.9 GIC CDRCFG, Request Interrupt Configuration in the Current Interrupt Domain

The GIC CDRCFG characteristics are:

Purpose

Request to read configuration of the specified INTID in the Current Interrupt Domain into ICC_ICSR_EL1.

If the Current Interrupt Domain is the Virtual Interrupt Domain, this operation applies to the virtual INTID in the VM identified by the resident VPE.

Configuration


This instruction is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to GIC CDRCFG are UNDEFINED.

Attributes

GIC CDRCFG is a 64-bit System instruction.

Field descriptions

The GIC CDRCFG bit assignments are:



Bits [63:32]

Reserved, RES0.

TYPE, bits [31:29]

Type of the interrupt.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC CDRCFG

This instruction is treated as a NOP if the operation applies to the Virtual Interrupt Domain and there is no resident VPE.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC CDRCFG, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b000	0b1100	0b0001	0b101

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGITR_EL2.GICCDCRCFG == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_RCFG);
    else
        GIC(X[t, 64], GICInstrDomain_CD, GICInstr_RCFG);
elseif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_RCFG);
elseif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_CD, GICInstr_RCFG);

```

8.1.10 GICR CDIA, Interrupt Acknowledge in the Current Interrupt Domain

The GICR CDIA characteristics are:

Purpose

Acknowledges the HPPI in the Current Interrupt Domain.

Configuration


This instruction is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to GICR CDIA are UNDEFINED.

Attributes

GICR CDIA is a 64-bit System instruction.

Field descriptions

The GICR CDIA bit assignments are:



Bits [63:33]

Reserved, RES0.

VALID, bit [32]

Indicates whether the instruction successfully acknowledged an interrupt.

VALID	Meaning
0b0	No interrupt was acknowledged.
0b1	The HPPI was acknowledged.

When this field is 1, the instruction acknowledges an HPPI of Sufficient priority that is not an NMI.

TYPE, bits [31:29]

The type of the acknowledged interrupt.

TYPE	Meaning
0b001	PPI
0b010	LPI
0b011	SPI

If VALID is 0, this field is RES0.

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the acknowledged interrupt.

If VALID is 0, this field is RES0.

Accessing GICR CDIA

Accesses to this instruction use the following encodings in the System instruction encoding space:

GICR <Xt>, CDIA

op0	op1	CRn	CRm	op2
0b01	0b000	0b1100	0b0011	0b000

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGITR_EL2.GICRCDIA == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = GICR(GICInstrDomain_VD, GICInstr_IA);
    else
        X[t, 64] = GICR(GICInstrDomain_CD, GICInstr_IA);
elsif PSTATE.EL == EL2 then
    X[t, 64] = GICR(GICInstrDomain_CD, GICInstr_IA);
elsif PSTATE.EL == EL3 then
    X[t, 64] = GICR(GICInstrDomain_CD, GICInstr_IA);

```

8.1.11 GICR CDNMIA, Non-maskable Interrupt Acknowledge in the Current Interrupt Domain

The GICR CDNMIA characteristics are:

Purpose

Acknowledges the HPPI, if it is an NMI, in the Current Interrupt Domain.

Configuration


This instruction is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to GICR CDNMIA are UNDEFINED.

Attributes

GICR CDNMIA is a 64-bit System instruction.

Field descriptions

The GICR CDNMIA bit assignments are:



Bits [63:33]

Reserved, RES0.

VALID, bit [32]

Indicates whether the instruction successfully acknowledged an NMI.

VALID	Meaning
0b0	No interrupt was acknowledged.
0b1	The HPPI was acknowledged.

When this field is 1, the instruction acknowledges an HPPI of Sufficient priority that is an NMI.

TYPE, bits [31:29]

The type of the acknowledged interrupt.

TYPE	Meaning
0b001	PPI
0b010	LPI
0b011	SPI

If VALID is 0, this field is RES0.

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the acknowledged interrupt.

If VALID is 0, this field is RES0.

Accessing GICR CDNMA

Accesses to this instruction use the following encodings in the System instruction encoding space:

GICR <Xt>, CDNMA

op0	op1	CRn	CRm	op2
0b01	0b000	0b1100	0b0011	0b001

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGITR_EL2.GICRCDNMA == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = GICR(GICInstrDomain_VD, GICInstr_NMIA);
    else
        X[t, 64] = GICR(GICInstrDomain_CD, GICInstr_NMIA);
elsif PSTATE.EL == EL2 then
    X[t, 64] = GICR(GICInstrDomain_CD, GICInstr_NMIA);
elsif PSTATE.EL == EL3 then
    X[t, 64] = GICR(GICInstrDomain_CD, GICInstr_NMIA);

```

8.2 System instructions for the Virtual Interrupt Domain

8.2.1 GIC VDAFF, Interrupt Set Target in the Virtual Interrupt Domain

The GIC VDAFF characteristics are:

Purpose

Sets the routing mode and target VPE for the specified INTID in the Virtual Interrupt Domain.

This operation applies to the virtual INTID in the VM identified by the resident VPE.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to GIC VDAFF are UNDEFINED.

Attributes

GIC VDAFF is a 64-bit System instruction.

Field descriptions

The GIC VDAFF bit assignments are:



Bits [63:48]

Reserved, RES0.

IAFFID, bits [47:32]

The interrupt Affinity value.

The IRS may support fewer than 16 bits of IAFFID. Unimplemented upper bits are treated as RES0 by the IRS.

When IRM is 1, this field provides an IMPLEMENTATION DEFINED hint to 1ofN selection algorithm. A value of 0 means that no hint is provided.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

IRM, bit [28]

Controls how the interrupt is routed.

If the GIC IRS does not support 1ofN distribution of the specified interrupt Type, this configuration is treated as RES0.

IRM	Meaning
0b0	The interrupt Routing mode is Targeted
0b1	The interrupt Routing mode is 1ofN

Bits [27:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC VDAFF

This instruction has no effect if the INTID is unreachable.

This instruction is treated as a NOP if there is no resident VPE.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC VDAFF, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b100	0b1100	0b0001	0b011

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_VD, GICInstr_AFF);
elsif PSTATE.EL == EL3 then
    if !EL2Enabled() then
        UNDEFINED;
    else
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_AFF);

```


8.2.2 GIC VDDI, Interrupt Deactivate in the Virtual Interrupt Domain

The GIC VDDI characteristics are:

Purpose

Clears the Active state of the specified INTID in the Virtual Interrupt Domain.

This instruction applies to all interrupt types, including PPIs.

If the interrupt type is a PPI, this operation applies to the virtual PPI identified by the specified INTID.

If the interrupt type is not a PPI, this operation applies to the virtual INTID in the VM identified by the resident VPE.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to GIC VDDI are UNDEFINED.

Attributes

GIC VDDI is a 64-bit System instruction.

Field descriptions

The GIC VDDI bit assignments are:



Bits [63:33]

Reserved, RES0.

Bit [32]

Reserved, WI.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b001	PPI
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC VDDI

This instruction has no effect if the INTID is unreachable

This instruction is treated as a NOP if the operation applies to the Virtual Interrupt Domain, the interrupt type is not a PPI and there is no resident VPE.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC VDDI, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b100	0b1100	0b0010	0b000

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_VD, GICInstr_DI);
elsif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_VD, GICInstr_DI);

```

8.2.3 GIC VDDIS, Interrupt Disable in the Virtual Interrupt Domain

The GIC VDDIS characteristics are:

Purpose

Disables the specified INTID in the Virtual Interrupt Domain.
This operation applies to the virtual INTID in the VM identified by the resident VPE.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to GIC VDDIS are UNDEFINED.

Attributes

GIC VDDIS is a 64-bit System instruction.

Field descriptions

The GIC VDDIS bit assignments are:



Bits [63:32]

Reserved, RES0.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Valued not defined above are reserved

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.
ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC VDDIS

This instruction has no effect if the INTID is unreachable.
This instruction is treated as a NOP if there is no resident VPE.
Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC VDDIS, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b100	0b1100	0b0001	0b000

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_VD, GICInstr_DIS);
elsif PSTATE.EL == EL3 then
    if !EL2Enabled() then
        UNDEFINED;
    else
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_DIS);

```

8.2.4 GIC VDEN, Interrupt Enable in the Virtual Interrupt Domain

The GIC VDEN characteristics are:

Purpose

Enables the specified INTID in the Virtual Interrupt Domain.

This operation applies to the virtual INTID in the VM identified by the resident VPE.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to GIC VDEN are UNDEFINED.

Attributes

GIC VDEN is a 64-bit System instruction.

Field descriptions

The GIC VDEN bit assignments are:



Bits [63:32]

Reserved, RES0.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC VDEN

This instruction has no effect if the INTID is unreachable.

This instruction is treated as a NOP if there is no resident VPE.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC VDEN, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b100	0b1100	0b0001	0b001

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_VD, GICInstr_EN);
elsif PSTATE.EL == EL3 then
    if !EL2Enabled() then
        UNDEFINED;
    else
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_EN);

```

8.2.5 GIC VDHM, Interrupt Handling mode in the Virtual Interrupt Domain

The GIC VDHM characteristics are:

Purpose

Sets the Handling mode of the specified INTID in the Virtual Interrupt Domain.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to GIC VDHM are UNDEFINED.

Attributes

GIC VDHM is a 64-bit System instruction.

Field descriptions

The GIC VDHM bit assignments are:



Bits [63:33]

Reserved, RES0.

HM, bit [32]

Handling mode

HM	Meaning
0b0	Edge
0b1	Level

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC VDHM

This instruction has no effect if the INTID is unreachable.

This instruction is treated as a NOP if there is no resident VPE.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC VDHM, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b100	0b1100	0b0010	0b001

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elseif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_VD, GICInstr_HM);
elseif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_VD, GICInstr_HM);

```


8.2.6 GIC VDPEND, Interrupt Set/Clear Pending state in the Virtual Interrupt Domain

The GIC VDPEND characteristics are:

Purpose

Generates a SET or CLEAR event for the specified INTID in the Virtual Interrupt Domain.
This operation applies to the virtual INTID in the VM identified by the specified VM identifier.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to GIC VDPEND are UNDEFINED.

Attributes

GIC VDPEND is a 64-bit System instruction.

Field descriptions

The GIC VDPEND bit assignments are:



Bits [63:48]

Reserved, RES0.

VM, bits [47:33]

The Virtual Machine identifier.

PENDING, bit [32]

Pending status

PENDING	Meaning
0b0	Generate CLEAR event to the IRS.
0b1	Generate SET event to the IRS.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC VDPEND

This instruction has no effect if the INTID is unreachable.

This instruction has no effect if the VM is invalid.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC VDPEND, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b100	0b1100	0b0001	0b100

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_VD, GICInstr_PEND);
elsif PSTATE.EL == EL3 then
    if !EL2Enabled() then
        UNDEFINED;
    else
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_PEND);

```

8.2.7 GIC VDPRI, Interrupt Set priority in the Virtual Interrupt Domain

The GIC VDPRI characteristics are:

Purpose

Sets the priority for the specified INTID in the Virtual Interrupt Domain.
This operation applies to the virtual INTID in the VM identified by the resident VPE.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to GIC VDPRI are UNDEFINED.

Attributes

GIC VDPRI is a 64-bit System instruction.

Field descriptions

The GIC VDPRI bit assignments are:



Bits [63:40]

Reserved, RES0.

PRIORITY, bits [39:35]

The priority of the specified INTID
Only the upper N bits of the Priority field are implemented where $N = (ICC_IDR0_EL1.PRI_BITS + 1)$.
Unimplemented bits are RES0.

Bits [34:32]

Reserved, RES0.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.
ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented lower bits are RES0.

Accessing GIC VDPRI

This instruction has no effect if the INTID is unreachable.

This instruction is treated as a NOP if there is no resident VPE.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC VDPRI, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b100	0b1100	0b0001	0b010

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_VD, GICInstr_PRI);
elsif PSTATE.EL == EL3 then
    if !EL2Enabled() then
        UNDEFINED;
    else
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_PRI);

```

8.2.8 GIC VDRCFG, Request Interrupt Configuration in the Virtual Interrupt Domain

The GIC VDRCFG characteristics are:

Purpose

Request to read configuration of the specified INTID in the Virtual Interrupt Domain into ICC_ICSR_EL1.

This operation applies to the virtual INTID in the VM identified by the resident VPE.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to GIC VDRCFG are UNDEFINED.

Attributes

GIC VDRCFG is a 64-bit System instruction.

Field descriptions

The GIC VDRCFG bit assignments are:



Bits [63:32]

Reserved, RES0.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC VDRCFG

This instruction is treated as a NOP if there is no resident VPE.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC VDRCFG, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b100	0b1100	0b0001	0b101

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    GIC(X[t, 64], GICInstrDomain_VD, GICInstr_RCFG);
elsif PSTATE.EL == EL3 then
    if !EL2Enabled() then
        UNDEFINED;
    else
        GIC(X[t, 64], GICInstrDomain_VD, GICInstr_RCFG);

```

8.3 System instructions for the Logical Interrupt Domain

8.3.1 GIC LDAFF, Interrupt Set Target in the Logical Interrupt Domain

The GIC LDAFF characteristics are:

Purpose

Sets the routing mode and target PE for the specified INTID in the Logical Interrupt Domain associated with the Security state selected by the SCR_EL3.{NS, NSE} bits.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL3 is implemented. Otherwise, direct accesses to GIC LDAFF are UNDEFINED.

Attributes

GIC LDAFF is a 64-bit System instruction.

Field descriptions

The GIC LDAFF bit assignments are:



Bits [63:48]

Reserved, RES0.

IAFFID, bits [47:32]

The interrupt Affinity value.

The IRS may support fewer than 16 bits of IAFFID. Unimplemented upper bits are treated as RES0 by the IRS.

When IRM is 1, this field provides an IMPLEMENTATION DEFINED hint to 1ofN selection algorithm. A value of 0 means that no hint is provided.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

IRM, bit [28]

Controls how the interrupt is routed.

If the GIC IRS does not support 1ofN distribution of the specified interrupt Type, this configuration is treated as RES0.

IRM	Meaning
0b0	The interrupt Routing mode is Targeted
0b1	The interrupt Routing mode is 1ofN

Bits [27:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC LDAFF

When executed at EL3, if SCR_EL3.{NS, NSE} select a reserved value, the behavior is CONSTRAINED UNPREDICTABLE with the permitted options:

- SCR_EL3.{NS, NSE} is treated as having an UNKNOWN value.
- The instruction is treated as a NOP.

This instruction has no effect if the INTID is unreachable

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC LDAFF, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b110	0b1100	0b0001	0b011

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_LD, GICInstr_AFF);

```

8.3.2 GIC LDDI, Interrupt Deactivate in the Logical Interrupt Domain

The GIC LDDI characteristics are:

Purpose

Enables the specified INTID in the Logical Interrupt Domain associated with the Security state selected by the SCR_EL3.{NS, NSE} bits.

This instruction applies to all interrupt types, including PPIs.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL3 is implemented. Otherwise, direct accesses to GIC LDDI are UNDEFINED.

Attributes

GIC LDDI is a 64-bit System instruction.

Field descriptions

The GIC LDDI bit assignments are:



Bits [63:33]

Reserved, RES0.

Bit [32]

Reserved, WI.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b001	PPI
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC LDDI

When executed at EL3, if SCR_EL3.{NS, NSE} select a reserved value, the behavior is CONSTRAINED UNPREDICTABLE with the permitted options:

- SCR_EL3.{NS, NSE} is treated as having an UNKNOWN value.
- The instruction is treated as a NOP.

This instruction has no effect if the INTID is unreachable.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC LDDI, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b110	0b1100	0b0010	0b000

```
if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    UNDEFINED;
elseif PSTATE.EL == EL2 then
    UNDEFINED;
elseif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_LD, GICInstr_DI);
```

8.3.3 GIC LDDIS, Interrupt Disable in the Logical Interrupt Domain

The GIC LDDIS characteristics are:

Purpose

Disables the specified INTID in the Logical Interrupt Domain associated with the Security state selected by the SCR_EL3.{NS, NSE} bits.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL3 is implemented. Otherwise, direct accesses to GIC LDDIS are UNDEFINED.

Attributes

GIC LDDIS is a 64-bit System instruction.

Field descriptions

The GIC LDDIS bit assignments are:



Bits [63:32]

Reserved, RES0.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Valued not defined above are reserved

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC LDDIS

When executed at EL3, if SCR_EL3.{NS, NSE} select a reserved value, the behavior is CONSTRAINED UNPREDICTABLE with the permitted options:

- SCR_EL3.{NS, NSE} is treated as having an UNKNOWN value.
- The instruction is treated as a NOP.

This instruction has no effect if the INTID is unreachable.
Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC LDDIS, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b110	0b1100	0b0001	0b000

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_LD, GICInstr_DIS);

```

8.3.4 GIC LDEN, Interrupt Enable in the Logical Interrupt Domain

The GIC LDEN characteristics are:

Purpose

Enables the specified INTID in the Logical Interrupt Domain associated with the Security state selected by the SCR_EL3.{NS, NSE} bits.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL3 is implemented. Otherwise, direct accesses to GIC LDEN are UNDEFINED.

Attributes

GIC LDEN is a 64-bit System instruction.

Field descriptions

The GIC LDEN bit assignments are:



Bits [63:32]

Reserved, RES0.

TYPE, bits [31:29]

Type of the interrupt.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC LDEN

When executed at EL3, if SCR_EL3.{NS, NSE} select a reserved value, the behavior is CONSTRAINED UNPREDICTABLE with the permitted options:

- SCR_EL3.{NS, NSE} is treated as having an UNKNOWN value.
- The instruction is treated as a NOP.

This instruction has no effect if the INTID is unreachable.
Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC LDEN, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b110	0b1100	0b0001	0b001

```
if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_LD, GICInstr_EN);
```

8.3.5 GIC LDHM, Interrupt Handling mode in the Logical Interrupt Domain

The GIC LDHM characteristics are:

Purpose

Sets the Handling mode of the specified INTID in the Logical Interrupt Domain associated with the Security state selected by the SCR_EL3.{NS, NSE} bits.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL3 is implemented. Otherwise, direct accesses to GIC LDHM are UNDEFINED.

Attributes

GIC LDHM is a 64-bit System instruction.

Field descriptions

The GIC LDHM bit assignments are:



Bits [63:33]

Reserved, RES0.

HM, bit [32]

Handling mode

HM	Meaning
0b0	Edge
0b1	Level

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC LDHM

When executed at EL3, if SCR_EL3.{NS, NSE} select a reserved value, the behavior is CONSTRAINED UNPREDICTABLE with the permitted options:

- SCR_EL3.{NS, NSE} is treated as having an UNKNOWN value.
- The instruction is treated as a NOP.

This instruction has no effect if the INTID is unreachable.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC LDHM, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b110	0b1100	0b0010	0b001

```
if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_LD, GICInstr_HM);
```

8.3.6 GIC LDPEND, Interrupt Set/Clear Pending state in the Logical Interrupt Domain

The GIC LDPEND characteristics are:

Purpose

Generates a SET or CLEAR event for the specified INTID in the Logical Interrupt Domain associated with the Security state selected by the SCR_EL3.{NS, NSE} bits.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL3 is implemented. Otherwise, direct accesses to GIC LDPEND are UNDEFINED.

Attributes

GIC LDPEND is a 64-bit System instruction.

Field descriptions

The GIC LDPEND bit assignments are:



Bits [63:33]

Reserved, RES0.

PENDING, bit [32]

Pending status

PENDING	Meaning
0b0	Generate CLEAR event to the IRS.
0b1	Generate SET event to the IRS.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC LDPEND

When executed at EL3, if SCR_EL3.{NS, NSE} select a reserved value, the behavior is CONSTRAINED UNPREDICTABLE with the permitted options:

- SCR_EL3.{NS, NSE} is treated as having an UNKNOWN value.
- The instruction is treated as a NOP.

This instruction has no effect if the INTID is unreachable.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC LDPEND, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b110	0b1100	0b0001	0b100

```
if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_LD, GICInstr_PEND);
```

8.3.7 GIC LDPRI, Interrupt Set priority in the Logical Interrupt Domain

The GIC LDPRI characteristics are:

Purpose

Sets the priority for the specified INTID in the Logical Interrupt Domain associated with the Security state selected by the SCR_EL3.{NS, NSE} bits.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL3 is implemented. Otherwise, direct accesses to GIC LDPRI are UNDEFINED.

Attributes

GIC LDPRI is a 64-bit System instruction.

Field descriptions

The GIC LDPRI bit assignments are:



Bits [63:40]

Reserved, RES0.

PRIORITY, bits [39:35]

The priority of the specified INTID

Only the upper N bits of the Priority field are implemented where $N = (ICC_IDR0_EL1.PRI_BITS + 1)$. Unimplemented bits are RES0.

Bits [34:32]

Reserved, RES0.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC LDPRI

When executed at EL3, if SCR_EL3.{NS, NSE} select a reserved value, the behavior is CONSTRAINED UNPREDICTABLE with the permitted options:

- SCR_EL3.{NS, NSE} is treated as having an UNKNOWN value.
- The instruction is treated as a NOP.

This instruction has no effect if the INTID is unreachable.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC LDPRI, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b110	0b1100	0b0001	0b010

```
if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    UNDEFINED;
elseif PSTATE.EL == EL2 then
    UNDEFINED;
elseif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_LD, GICInstr_PRI);
```

8.3.8 GIC LDRCFG, Request Interrupt Configuration in the Logical Interrupt Domain

The GIC LDRCFG characteristics are:

Purpose

Request to read configuration of the specified INTID in the Logical Interrupt Domain associated with the Security state selected by the SCR_EL3.{NS, NSE} bits into ICC_ICSR_EL1.

Configuration


This instruction is present only when FEAT_GCIE is implemented and EL3 is implemented. Otherwise, direct accesses to GIC LDRCFG are UNDEFINED.

Attributes

GIC LDRCFG is a 64-bit System instruction.

Field descriptions

The GIC LDRCFG bit assignments are:



Bits [63:32]

Reserved, RES0.

TYPE, bits [31:29]

Type of the interrupt

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The ID of the interrupt.

ID and TYPE together form an INTID.

The number of ID bits implemented is reported in ICC_IDR0_EL1.ID_BITS. Unimplemented upper bits are RES0.

Accessing GIC LDRCFG

When executed at EL3, if SCR_EL3.{NS, NSE} select a reserved value, the behavior is CONSTRAINED UNPREDICTABLE with the permitted options:

- SCR_EL3.{NS, NSE} is treated as having an UNKNOWN value.
- The instruction is treated as a NOP.

ICC_ICSR_EL1.F is set to 1 and other fields become UNKNOWN after execution of this instruction if the INTID is unreachable.

Accesses to this instruction use the following encodings in the System instruction encoding space:

GIC LDRCFG, <Xt>

op0	op1	CRn	CRm	op2
0b01	0b110	0b1100	0b0001	0b101

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    GIC(X[t, 64], GICInstrDomain_LD, GICInstr_RCFG);

```

8.4 GIC synchronization barrier instructions

8.4.1 GSB SYS, GIC Synchronization Barrier System

The GSB SYS characteristics are:

Purpose

GIC Synchronization Barrier System ensures completion of the effects of GIC instructions and prevents loads, stores, and GIC instructions from executing part of their functionality before the GSB SYS.

See [2.12 Interrupt ordering model and synchronization requirements](#) for more information.

Configuration

This instruction is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to GSB SYS are UNDEFINED.

Accessing GSB SYS

Accesses to this instruction use the following encodings in the System instruction encoding space:

GSB SYS

op0	op1	CRn	CRm	op2
0b01	0b000	0b1100	0b0000	0b000

```


if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    GSB(GICInstr_SYS);
elsif PSTATE.EL == EL2 then
    GSB(GICInstr_SYS);
elsif PSTATE.EL == EL3 then
    GSB(GICInstr_SYS);

```

8.4.2 GSB ACK, GIC Synchronization Barrier Interrupt Acknowledge

The GSB ACK characteristics are:

Purpose

GIC Synchronization Barrier Interrupt Acknowledge ensures completion of the effects of the GICR CDIA and GICR CDNMIA instructions and prevents loads, stores, and GIC instructions from executing part of their functionality before the GSB ACK.

See [2.12 Interrupt ordering model and synchronization requirements](#) for more information.

Configuration

This instruction is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to GSB ACK are UNDEFINED.

Accessing GSB ACK

Accesses to this instruction use the following encodings in the System instruction encoding space:

GSB ACK

op0	op1	CRn	CRm	op2
0b01	0b000	0b1100	0b0000	0b001

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    GSB(GICInstr_ACK);
elsif PSTATE.EL == EL2 then
    GSB(GICInstr_ACK);
elsif PSTATE.EL == EL3 then
    GSB(GICInstr_ACK);

```

Chapter 9

System registers

This chapter describes the GICv5 system registers.

9.1 Synchronization requirements for GICv5 System registers

I_{CWLKD} The GICv5 System registers follow the requirements for System registers in the *Arm[®] Architecture Reference Manual, for A-profile architecture* [1].

R_{NWMPEV} Access to the following System registers are self-synchronizing:

- ICC_PCR_EL3
- ICC_PCR_EL1
- ICV_PCR_EL1

This means that the architectural side effects of direct writes to any of the above System registers are visible to instructions in program order after the direct write. The effects of a direct write include updates to the pending state of the IRQ, FIQ, vIRQ, and vFIQ asynchronous exceptions.

9.2 CPU interface registers

Configuration and state of the CPU interface.

9.2.1 ICC_APR_EL1, Interrupt Controller Physical Active Priorities Register

The ICC_APR_EL1 characteristics are:

Purpose

Records physical active priorities for the Non-secure, Realm, and Secure Interrupt Domains.

Configuration

There are separate banked copies of this register for Non-secure, Realm and Secure state.

This register is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to ICC_APR_EL1 are UNDEFINED.

Attributes


ICC_APR_EL1 is a 64-bit register.

This register has the following instances:

- ICC_APR_EL1, when EL3 is not implemented
- ICC_APR_EL1_NS, when EL3 is implemented
- ICC_APR_EL1_RL, when FEAT_RME is implemented
- ICC_APR_EL1_S, when (EL3 is implemented and FEAT_RME is not implemented) or (FEAT_RME is implemented and FEAT_SEL2 is implemented)

Field descriptions

The ICC_APR_EL1 bit assignments are:



Bits [63:32]

Reserved, RES0.

P<x>, bits [x], for x = 31 to 0

Provides access to the active priorities.

P<x>	Meaning
0b0	Priority not active
0b1	Priority active

Fields in this register are indexed using the 5-bit priority as an unsigned integer, P[Priority].

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- Access is **RES0** if all of the following are true:
 - $x \bmod 2 == 1$
 - `ICC_IDR0_EL1.PRI_BITS == 0b011`
- Otherwise, access to this field is **RW**

Accessing ICC_APR_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_APR_EL1

op0	op1	CRn	CRm	op2
0b11	0b001	0b1100	0b0000	0b000

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGRTR_EL2.ICC_APR_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_APR_EL1;
    elseif HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            X[t, 64] = ICC_APR_EL1_S;
        elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            X[t, 64] = ICC_APR_EL1_RL;
        elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            X[t, 64] = ICC_APR_EL1_NS;
        else
            UNDEFINED;
    else
        X[t, 64] = ICC_APR_EL1;
elseif PSTATE.EL == EL2 then
    if HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            X[t, 64] = ICC_APR_EL1_S;
        elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            X[t, 64] = ICC_APR_EL1_RL;
        elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            X[t, 64] = ICC_APR_EL1_NS;
        else
            UNDEFINED;
    else
        X[t, 64] = ICC_APR_EL1;
elseif PSTATE.EL == EL3 then
    if SCR_EL3.NS == '0' then
        X[t, 64] = ICC_APR_EL1_S;
    elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS == '1'
        ↪ then
        X[t, 64] = ICC_APR_EL1_RL;

```

```

elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
    X[t, 64] = ICC_APR_EL1_NS;
else
    UNDEFINED;

```

MSR ICC_APR_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b001	0b1100	0b0000	0b000

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGWTR_EL2.ICC_APR_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_APR_EL1 = X[t, 64];
    elseif HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            ICC_APR_EL1_S = X[t, 64];
        elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            ICC_APR_EL1_RL = X[t, 64];
        elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            ICC_APR_EL1_NS = X[t, 64];
        else
            UNDEFINED;
    else
        ICC_APR_EL1 = X[t, 64];
elseif PSTATE.EL == EL2 then
    if HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            ICC_APR_EL1_S = X[t, 64];
        elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            ICC_APR_EL1_RL = X[t, 64];
        elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            ICC_APR_EL1_NS = X[t, 64];
        else
            UNDEFINED;
    else
        ICC_APR_EL1 = X[t, 64];
elseif PSTATE.EL == EL3 then
    if SCR_EL3.NS == '0' then
        ICC_APR_EL1_S = X[t, 64];
    elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS == '1'
        ↪ then
        ICC_APR_EL1_RL = X[t, 64];
    elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
        ICC_APR_EL1_NS = X[t, 64];
    else
        UNDEFINED;

```


9.2.2 ICC_APR_EL3, Interrupt Controller Physical Active Priorities Register for EL3

The ICC_APR_EL3 characteristics are:

Purpose

Records active priorities for the EL3 Interrupt Domain.

Configuration

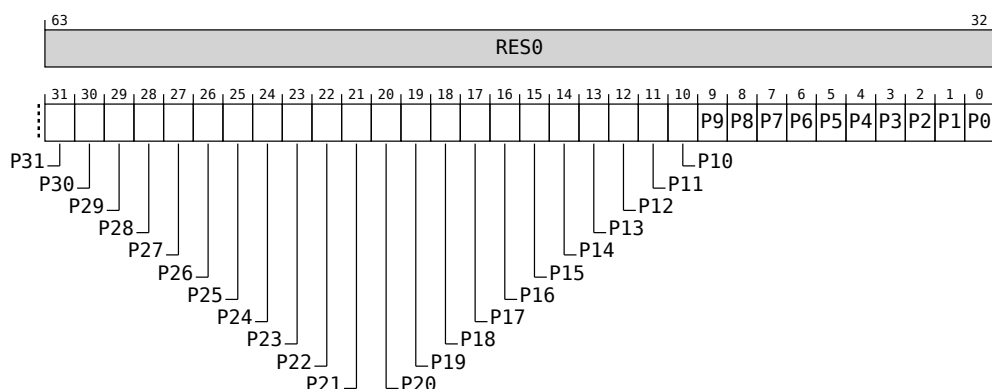
This register is present only when FEAT_GCIE is implemented and EL3 is implemented. Otherwise, direct accesses to ICC_APR_EL3 are UNDEFINED.

Attributes

ICC_APR_EL3 is a 64-bit register.

Field descriptions

The ICC_APR_EL3 bit assignments are:



Bits [63:32]

Reserved, RES0.

P<x>, bits [x], for x = 31 to 0

Provides access to the active priorities.

P<x>	Meaning
0b0	Priority not active
0b1	Priority active

Fields in this register are indexed using the 5-bit priority as an unsigned integer, P[Priority].

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- Access is **RES0** if all of the following are true:
 - x MOD 2 == 1
 - ICC_IDR0_EL1.PRI_BITS == 0b011
- Otherwise, access to this field is **RW**

Accessing ICC_APR_EL3

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_APR_EL3

op0	op1	CRn	CRm	op2
0b11	0b110	0b1100	0b1000	0b000

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_APR_EL3;

```

MSR ICC_APR_EL3, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b110	0b1100	0b1000	0b000

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    ICC_APR_EL3 = X[t, 64];

```

9.2.3 ICC_CR0_EL1, Interrupt Controller EL1 Physical Control Register

The ICC_CR0_EL1 characteristics are:

Purpose

Controls behavior of the physical CPU interface for the Non-secure, Realm, and Secure Interrupt Domains.

Configuration

There are separate banked copies of this register for Non-secure, Realm and Secure state.

This register is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to ICC_CR0_EL1 are UNDEFINED.

Attributes


ICC_CR0_EL1 is a 64-bit register.

This register has the following instances:

- ICC_CR0_EL1, when EL3 is not implemented
- ICC_CR0_EL1_NS, when EL3 is implemented
- ICC_CR0_EL1_RL, when FEAT_RME is implemented
- ICC_CR0_EL1_S, when (EL3 is implemented and FEAT_RME is not implemented) or (FEAT_RME is implemented and FEAT_SEL2 is implemented)

Field descriptions

The ICC_CR0_EL1 bit assignments are:



Bits [63:39]

Reserved, RES0.

PID, bit [38]

When EL3 is implemented:

Preemptive Interrupt Domain. Indicates whether the Interrupt Domain associated with the Security state selected by SCR_EL3.{NSE,NS} is the Preemptive Interrupt Domain.

Access to this field is **RO**.

Otherwise:

RAZ/WI

IPPT, bits [37:32]

When ICC_CR0_EL1.PID == 1:

Interrupt Preemptive Priority Threshold value for the Preemptive Interrupt Domain.

When fewer than 5 bits of priority is implemented, only bits [5:N] are implemented where N = (4 - ICC_IDR0_EL1.PRI_BITS). Unimplemented bits are RES0.

Note: ICC_CR0_EL1.IPPT is a 6 bits field to ensure it can be strictly higher than the interrupt priority, which is a 5-bit unsigned value.

See [2.9.3 Preemptive interrupts](#) for more information.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

Bits [31:1]

Reserved, RES0.

EN, bit [0]

Enable interrupts for the Interrupt Domain.

When this field is 0, there is no HPPI of Sufficient priority for the Interrupt Domain.

EN	Meaning
0b0	Disabled.
0b1	Enabled.

The reset behavior of this field is:

- On a Warm reset, this field resets to 0b0.

Accessing ICC_CR0_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_CR0_EL1

op0	op1	CRn	CRm	op2
0b11	0b001	0b1100	0b0000	0b001

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGTR_EL2.ICC_CR0_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_CR0_EL1;
    elsif HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            X[t, 64] = ICC_CR0_EL1_S;
        elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            X[t, 64] = ICC_CR0_EL1_RL;
        elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            X[t, 64] = ICC_CR0_EL1_NS;
        else
            UNDEFINED;
    else
        UNDEFINED;
    end
end

```

```

        X[t, 64] = ICC_CR0_EL1;
    elsif PSTATE.EL == EL2 then
        if HaveEL(EL3) then
            if SCR_EL3.NS == '0' then
                X[t, 64] = ICC_CR0_EL1_S;
            elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
                ↪ '1' then
                X[t, 64] = ICC_CR0_EL1_RL;
            elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
                X[t, 64] = ICC_CR0_EL1_NS;
            else
                UNDEFINED;
        else
            X[t, 64] = ICC_CR0_EL1;
    elsif PSTATE.EL == EL3 then
        if SCR_EL3.NS == '0' then
            X[t, 64] = ICC_CR0_EL1_S;
        elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS == '1'
            ↪ then
            X[t, 64] = ICC_CR0_EL1_RL;
        elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            X[t, 64] = ICC_CR0_EL1_NS;
        else
            UNDEFINED;

```

MSR ICC_CR0_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b001	0b1100	0b0000	0b001

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGWTR_EL2.ICC_CR0_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_CR0_EL1 = X[t, 64];
    elsif HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            ICC_CR0_EL1_S = X[t, 64];
        elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            ICC_CR0_EL1_RL = X[t, 64];
        elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            ICC_CR0_EL1_NS = X[t, 64];
        else
            UNDEFINED;
    else
        ICC_CR0_EL1 = X[t, 64];
elsif PSTATE.EL == EL2 then
    if HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            ICC_CR0_EL1_S = X[t, 64];
        elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then

```

```

        ICC_CR0_EL1_RL = X[t, 64];
    elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
        ICC_CR0_EL1_NS = X[t, 64];
    else
        UNDEFINED;
else
    ICC_CR0_EL1 = X[t, 64];
elsif PSTATE.EL == EL3 then
    if SCR_EL3.NS == '0' then
        ICC_CR0_EL1_S = X[t, 64];
    elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS == '1'
        ↪ then
        ICC_CR0_EL1_RL = X[t, 64];
    elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
        ICC_CR0_EL1_NS = X[t, 64];
    else
        UNDEFINED;

```

9.2.4 ICC_CR0_EL3, Interrupt Controller EL3 Physical Control Register

The ICC_CR0_EL3 characteristics are:

Purpose

Controls behavior of the CPU interface in the EL3 Interrupt Domain.

Configuration


This register is present only when FEAT_GCIE is implemented and EL3 is implemented. Otherwise, direct accesses to ICC_CR0_EL3 are UNDEFINED.

Attributes

ICC_CR0_EL3 is a 64-bit register.

Field descriptions

The ICC_CR0_EL3 bit assignments are:



Bits [63:34]

Reserved, RES0.

PID, bits [33:32]

Preemptive Interrupt Domain.

The PE behaves as if there is no Preemptive Interrupt Domain when any of the following are true:

- This field is set to 0b10.
- This field is set to an Interrupt Domain that is not implemented.

PID	Meaning
0b00	Secure Interrupt Domain
0b01	Non-secure Interrupt Domain
0b10	None
0b11	Realm Interrupt Domain

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Bits [31:1]

Reserved, RES0.

EN, bit [0]

Enable interrupts for the EL3 Interrupt Domain.

When this field is set to 0, the PE behaves as if there is no HPPI of Sufficient priority in the EL3 Interrupt Domain.

EN	Meaning
0b0	Disabled.
0b1	Enabled.

The reset behavior of this field is:

- On a Warm reset, this field resets to 0b0.

Accessing ICC_CR0_EL3

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_CR0_EL3

op0	op1	CRn	CRm	op2
0b11	0b110	0b1100	0b1001	0b000

```
if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_CR0_EL3;
```

MSR ICC_CR0_EL3, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b110	0b1100	0b1001	0b000

```
if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    ICC_CR0_EL3 = X[t, 64];
```


9.2.5 ICC_DOMHPPIR_EL3, Interrupt Controller Domain Highest Priority Pending Interrupt Register

The ICC_DOMHPPIR_EL3 characteristics are:

Purpose

Reports the HPPI for a Non-EL3 Interrupt Domain.

Configuration


This register is present only when FEAT_GCIE is implemented and EL3 is implemented. Otherwise, direct accesses to ICC_DOMHPPIR_EL3 are UNDEFINED.

Attributes

ICC_DOMHPPIR_EL3 is a 64-bit register.

Field descriptions

The ICC_DOMHPPIR_EL3 bit assignments are:



Bits [63:4]

Reserved, RES0.

P_HPPI, bit [3]

Preemptive HPPI for the Preemptive Interrupt Domain valid bit.

P_HPPI	Meaning
0b0	There is no Preemptive HPPI of Sufficient priority for the Preemptive Interrupt Domain.
0b1	There is an Preemptive HPPI of Sufficient priority for the Preemptive Interrupt Domain.

Accessing this field has the following behavior:

- **RO** if HavePreemptiveDomain()
- Otherwise, access to this field is **RES0**

RL_HPPI, bit [2]

HPPI for Realm Interrupt Domain valid bit.

RL_HPPI	Meaning
0b0	There is no HPPI of Sufficient priority for Realm Interrupt Domain.
0b1	There is an HPPI of Sufficient priority for Realm Interrupt Domain.

Accessing this field has the following behavior:

- **RO** if HaveDomain(GICDomain_RL)
- Otherwise, access to this field is **RES0**

S_HPPI, bit [1]

HPPI for Secure Interrupt Domain valid bit.

S_HPPI	Meaning
0b0	There is no HPPI of Sufficient priority for Secure Interrupt Domain.
0b1	There is an HPPI of Sufficient priority for Secure Interrupt Domain.

Accessing this field has the following behavior:

- **RO** if HaveDomain(GICDomain_S)
- Otherwise, access to this field is **RES0**

NS_HPPI, bit [0]

HPPI for Non-secure Interrupt Domain valid bit.

NS_HPPI	Meaning
0b0	There is no HPPI of Sufficient priority for Non-secure Interrupt Domain.
0b1	There is an HPPI of Sufficient priority for Non-secure Interrupt Domain.

Accessing this field has the following behavior:

- **RO** if HaveDomain(GICDomain_NS)
- Otherwise, access to this field is **RES0**

Accessing ICC_DOMHPPIR_EL3

Read-only

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_DOMHPPIR_EL3

op0	op1	CRn	CRm	op2
0b11	0b110	0b1100	0b1000	0b010

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_DOMHPPIR_EL3;

```

9.2.6 ICC_HAPR_EL1, Interrupt Controller Physical Highest Active Priority Register

The ICC_HAPR_EL1 characteristics are:

Purpose

Reports the running priority of the Current Physical Interrupt Domain.

Configuration


This register is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to ICC_HAPR_EL1 are UNDEFINED.

Attributes

ICC_HAPR_EL1 is a 64-bit register.

Field descriptions

The ICC_HAPR_EL1 bit assignments are:



Bits [63:8]

Reserved, RES0.

PRIORITY, bits [7:0]

The running priority for the Current Physical Interrupt Domain.

If there are no active priorities on the CPU interface in the applicable Interrupt Domain, or all active priorities have undergone a priority drop, the value returned is the Idle priority.

Accessing ICC_HAPR_EL1

Read-only

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_HAPR_EL1

op0	op1	CRn	CRm	op2
0b11	0b001	0b1100	0b0000	0b011

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGRTR_EL2.ICC_HAPR_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_HAPR_EL1;
    else
        X[t, 64] = ICC_HAPR_EL1;
elseif PSTATE.EL == EL2 then
    X[t, 64] = ICC_HAPR_EL1;
  
```

```
elseif PSTATE.EL == EL3 then
    X[t, 64] = ICC_HAPR_EL1;
```

9.2.7 ICC_HPPIR_EL1, Interrupt Controller Physical Highest Priority Pending Interrupt Register

The ICC_HPPIR_EL1 characteristics are:

Purpose

Reports the HPPI for the Physical Interrupt Domain associated with the Security state selected by SCR_EL3.{NS, NSE}.

Configuration


This register is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to ICC_HPPIR_EL1 are UNDEFINED.

Attributes

ICC_HPPIR_EL1 is a 64-bit register.

Field descriptions

The ICC_HPPIR_EL1 bit assignments are:



Bits [63:33]

Reserved, RES0.

HPPIV, bit [32]

HPPI valid.

There is an HPPI with Sufficient priority for the Interrupt Domain.

HPPIV	Meaning
0b0	Invalid: There is no HPPI with Sufficient priority for the Interrupt Domain.
0b1	VALID: There is an HPPI with Sufficient priority for the Interrupt Domain.

If ICC_HPPIR_EL1.HPPIV is 1, ID and TYPE together form the INTID of the HPPI for the Interrupt Domain.

TYPE, bits [31:29]

The encoding of this field depends on the value of HPPIV as described below:

- If ICC_HPPIR_EL1.HPPIV is 0, TYPE is RES0.
- If ICC_HPPIR_EL1.HPPIV is 1, TYPE specifies the Type of the interrupt.

TYPE	Meaning
0b001	PPI
0b010	LPI

TYPE	Meaning
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The encoding of this field depends on the value of HPPIV as described below:

- If ICC_HPPIR_EL1.HPPIV is 0, ID is RES0.
- If ICC_HPPIR_EL1.HPPIV is 1, ID specifies the interrupt ID.

Accessing ICC_HPPIR_EL1

Read-only

When accessed at EL3, if SCR_EL3.{NS, NSE} select a reserved value, the behavior is CONSTRAINED UNPREDICTABLE with the permitted options:

- SCR_EL3.{NS, NSE} is treated as having an UNKNOWN value.
- The access is treated as a NOP.

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_HPPIR_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1010	0b011

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGTR_EL2.ICC_HPPIR_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_HPPIR_EL1;
    else
        X[t, 64] = ICC_HPPIR_EL1;
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICC_HPPIR_EL1;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_HPPIR_EL1;

```

9.2.8 ICC_HPPIR_EL3, Interrupt Controller Physical Highest Priority Pending Interrupt Register

The ICC_HPPIR_EL3 characteristics are:

Purpose

Reports the HPPI for the EL3 Interrupt Domain.

Configuration


This register is present only when FEAT_GCIE is implemented and EL3 is implemented. Otherwise, direct accesses to ICC_HPPIR_EL3 are UNDEFINED.

Attributes

ICC_HPPIR_EL3 is a 64-bit register.

Field descriptions

The ICC_HPPIR_EL3 bit assignments are:



Bits [63:33]

Reserved, RES0.

HPPIV, bit [32]

HPPI valid.

There is an HPPI with Sufficient priority for the Interrupt Domain.

HPPIV	Meaning
0b0	Invalid: There is no HPPI with Sufficient priority for the Interrupt Domain.
0b1	VALID: There is an HPPI with Sufficient priority for the Interrupt Domain.

If ICC_HPPIR_EL3.HPPIV is 1, ID and TYPE together form the INTID of the HPPI for the Interrupt Domain.

TYPE, bits [31:29]

The encoding of this field depends on the value of HPPIV as described below:

- If ICC_HPPIR_EL3.HPPIV is 0, TYPE is RES0.
- If ICC_HPPIR_EL3.HPPIV is 1, TYPE specifies the Type of the interrupt.

TYPE	Meaning
0b001	PPI
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The encoding of this field depends on the value of HPPIV as described below:

- If ICC_HPPIR_EL3.HPPIV is 0, ID is RES0.
- If ICC_HPPIR_EL3.HPPIV is 1, ID specifies the interrupt ID.

Accessing ICC_HPPIR_EL3

Read-only

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_HPPIR_EL3

op0	op1	CRn	CRm	op2
0b11	0b110	0b1100	0b1001	0b001

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_HPPIR_EL3;

```


9.2.9 ICC_IAFFIDR_EL1, Interrupt Controller PE Interrupt Affinity ID Register

The ICC_IAFFIDR_EL1 characteristics are:

Purpose

Reports the PE interrupt Affinity ID for the PE.

Configuration


This register is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to ICC_IAFFIDR_EL1 are UNDEFINED.

Attributes

ICC_IAFFIDR_EL1 is a 64-bit register.

Field descriptions

The ICC_IAFFIDR_EL1 bit assignments are:



Bits [63:16]

Reserved, RES0.

IAFFID, bits [15:0]

PE interrupt Affinity ID.

Accessing ICC_IAFFIDR_EL1

Read-only

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_IAFFIDR_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1010	0b101

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGRTR_EL2.ICC_IAFFIDR_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        X[t, 64] = ICC_IAFFIDR_EL1;
elseif PSTATE.EL == EL2 then
    X[t, 64] = ICC_IAFFIDR_EL1;
elseif PSTATE.EL == EL3 then
    X[t, 64] = ICC_IAFFIDR_EL1;
  
```

9.2.10 ICC_ICSR_EL1, Interrupt Controller Interrupt Configuration and State Register

The ICC_ICSR_EL1 characteristics are:

Purpose

Reports the configuration and state of the INTID specified to a previous GIC request interrupt configuration instruction.

Configuration


This register is present only when FEAT_GDIE is implemented. Otherwise, direct accesses to ICC_ICSR_EL1 are UNDEFINED.

Attributes

ICC_ICSR_EL1 is a 64-bit register.

Field descriptions

The ICC_ICSR_EL1 bit assignments are:



Bits [63:48]

Reserved, RES0.

IAFFID, bits [47:32]

The interrupt Affinity value.

The IRS may support fewer than 16 bits of IAFFID. Upper bits not implemented by the IRS are returned as zero.

When IRM is 1, this field is IMPLEMENTATION SPECIFIC.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Bits [31:16]

Reserved, RES0.

Priority, bits [15:11]

The interrupt's Priority.

Bits [4:N] of the priority value are implemented where N = (4 - ICC_IDR0_EL1.PRI_BITS). Unimplemented bits are RES0.

This means that when fewer than 5 bits of priority is implemented, Priority[14 - ICC_IDR0_EL1.PRI_BITS:11] are RES0.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Bits [10:6]

Reserved, RES0.

HM, bit [5]

The interrupt's Handling mode.

HM	Meaning
0b0	Edge-triggered
0b1	Level-sensitive

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Active, bit [4]

The interrupt's Active state.

Active	Meaning
0b0	Inactive
0b1	Active

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

IRM, bit [3]

Interrupt Routing mode.

IRM	Meaning
0b0	The interrupt Routing mode is Targeted.
0b1	The interrupt Routing mode is 1ofN.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Pending, bit [2]

The interrupt's Pending state.

Pending	Meaning
0b0	Not pending
0b1	Pending

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Enabled, bit [1]

The interrupt's individual enable.

Enabled	Meaning
0b0	Disabled
0b1	Enabled

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

F, bit [0]

Indicates whether the IRS returned valid data.

F	Meaning
0b0	Request completed successfully.
0b1	Request did not complete successfully.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing ICC_ICSR_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_ICSR_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1010	0b100

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGRTR_EL2.ICC_ICSR_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        X[t, 64] = ICC_ICSR_EL1;
elseif PSTATE.EL == EL2 then
    X[t, 64] = ICC_ICSR_EL1;
elseif PSTATE.EL == EL3 then
    X[t, 64] = ICC_ICSR_EL1;

```

MSR ICC_ICSR_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1010	0b100

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGWTR_EL2.ICC_ICSR_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        ICC_ICSR_EL1 = X[t, 64];
    elsif PSTATE.EL == EL2 then
        ICC_ICSR_EL1 = X[t, 64];
    elsif PSTATE.EL == EL3 then
        ICC_ICSR_EL1 = X[t, 64];

```

9.2.11 ICC_IDR0_EL1, Interrupt Controller ID Register 0

The ICC_IDR0_EL1 characteristics are:

Purpose

Contains read-only fields with information about the CPU interface.

Configuration


This register is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to ICC_IDR0_EL1 are UNDEFINED.

Attributes

ICC_IDR0_EL1 is a 64-bit register.

Field descriptions

The ICC_IDR0_EL1 bit assignments are:



Bits [63:12]

Reserved, RES0.

GCIE_LEGACY, bits [11:8]

Indicates support for legacy GICv3.3 virtual CPU interface.

GCIE_LEGACY	Meaning
0b0000	Legacy GICv3.3 virtual CPU interface is not implemented
0b0001	Legacy GICv3.3 virtual CPU interface is implemented

FEAT_GCIE_LEGACY extension implements the functionality identified by the value 1.

PRI_BITS, bits [7:4]

The number of priority bits implemented, minus one.

PRI_BITS	Meaning
0b0011	4 bits of priority
0b0100	5 bits of priority

Values not defined above are reserved.

When FEAT_GCIE_LEGACY is implemented, the only permitted value of this field is 0b100.

ID_BITS, bits [3:0]

Identifier bits.

Read-only and writes are ignored.

The number of interrupt identifier bits supported.

ID_BITS	Meaning
0b0000	16 bits
0b0001	24 bits

Accessing ICC_IDR0_EL1

Read-only

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_IDR0_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1010	0b010

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGTR_EL2.ICC_IDRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        X[t, 64] = ICC_IDR0_EL1;
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICC_IDR0_EL1;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_IDR0_EL1;

```

9.2.12 ICC_PCR_EL1, Interrupt Controller Physical Interrupt Priority Control Register

The ICC_PCR_EL1 characteristics are:

Purpose

Reports the Physical priority mask for the Non-secure, Realm, and Secure Interrupt Domains.

Configuration

There are separate banked copies of this register for Non-secure, Realm and Secure state.

This register is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to ICC_PCR_EL1 are UNDEFINED.

Attributes


ICC_PCR_EL1 is a 64-bit register.

This register has the following instances:

- ICC_PCR_EL1, when EL3 is not implemented
- ICC_PCR_EL1_NS, when EL3 is implemented
- ICC_PCR_EL1_RL, when FEAT_RME is implemented
- ICC_PCR_EL1_S, when (EL3 is implemented and FEAT_RME is not implemented) or (FEAT_RME is implemented and FEAT_SEL2 is implemented)

Field descriptions

The ICC_PCR_EL1 bit assignments are:



Bits [63:5]

Reserved, RES0.

PRIORITY, bits [4:0]

The priority mask for the Interrupt Domain.

When fewer than 5 bits of priority is implemented, only bits [4:N] are implemented where N = (4 - ICC_IDR0_EL1.PRI_BITS). Unimplemented bits are RES0.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing ICC_PCR_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PCR_EL1

op0	op1	CRn	CRm	op2
0b11	0b001	0b1100	0b0000	0b010


```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGRTR_EL2.ICC_PCR_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PCR_EL1;
    elseif HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            X[t, 64] = ICC_PCR_EL1_S;
        elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            X[t, 64] = ICC_PCR_EL1_RL;
        elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            X[t, 64] = ICC_PCR_EL1_NS;
        else
            UNDEFINED;
    else
        X[t, 64] = ICC_PCR_EL1;
elseif PSTATE.EL == EL2 then
    if HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            X[t, 64] = ICC_PCR_EL1_S;
        elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            X[t, 64] = ICC_PCR_EL1_RL;
        elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            X[t, 64] = ICC_PCR_EL1_NS;
        else
            UNDEFINED;
    else
        X[t, 64] = ICC_PCR_EL1;
elseif PSTATE.EL == EL3 then
    if SCR_EL3.NS == '0' then
        X[t, 64] = ICC_PCR_EL1_S;
    elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS == '1'
        ↪ then
        X[t, 64] = ICC_PCR_EL1_RL;
    elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
        X[t, 64] = ICC_PCR_EL1_NS;
    else
        UNDEFINED;

```

MSR ICC_PCR_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b001	0b1100	0b0000	0b010

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;

```

```

elseif EL2Enabled() && ICH_HFGWTR_EL2.ICC_PCR_EL1 == '0' then
    AArch64.SystemAccessTrap(EL2, 0x18);
elseif EL2Enabled() && HCR_EL2.IMO == '1' then
    ICC_PCR_EL1 = X[t, 64];
elseif HaveEL(EL3) then
    if SCR_EL3.NS == '0' then
        ICC_PCR_EL1_S = X[t, 64];
    elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
        ↪ '1' then
        ICC_PCR_EL1_RL = X[t, 64];
    elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
        ICC_PCR_EL1_NS = X[t, 64];
    else
        UNDEFINED;
else
    ICC_PCR_EL1 = X[t, 64];
elseif PSTATE.EL == EL2 then
    if HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            ICC_PCR_EL1_S = X[t, 64];
        elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            ICC_PCR_EL1_RL = X[t, 64];
        elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            ICC_PCR_EL1_NS = X[t, 64];
        else
            UNDEFINED;
    else
        ICC_PCR_EL1 = X[t, 64];
elseif PSTATE.EL == EL3 then
    if SCR_EL3.NS == '0' then
        ICC_PCR_EL1_S = X[t, 64];
    elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS == '1'
        ↪ then
        ICC_PCR_EL1_RL = X[t, 64];
    elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
        ICC_PCR_EL1_NS = X[t, 64];
    else
        UNDEFINED;

```

9.2.13 ICC_PCR_EL3, Interrupt Controller Interrupt Priority Control Register for EL3

The ICC_PCR_EL3 characteristics are:

Purpose

Reports the Physical priority mask for the physical CPU interface for the EL3 Interrupt Domain.

Configuration


This register is present only when FEAT_GCIE is implemented and EL3 is implemented. Otherwise, direct accesses to ICC_PCR_EL3 are UNDEFINED.

Attributes

ICC_PCR_EL3 is a 64-bit register.

Field descriptions

The ICC_PCR_EL3 bit assignments are:



Bits [63:5]

Reserved, RES0.

PRIORITY, bits [4:0]

The priority mask for the EL3 Interrupt Domain.

When fewer than 5 bits of priority is implemented, only bits [4:N] are implemented where N = (4 - ICC_IDR0_EL1.PRI_BITS). Unimplemented bits are RES0.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing ICC_PCR_EL3

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PCR_EL3

op0	op1	CRn	CRm	op2
0b11	0b110	0b1100	0b1000	0b001

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_PCR_EL3;
  
```

MSR ICC_PCR_EL3, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b110	0b1100	0b1000	0b001

```
if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    ICC_PCR_EL3 = X[t, 64];
```

9.2.14 ID_AA64PFR2_EL1, AArch64 Processor Feature Register 2

The ID_AA64PFR2_EL1 characteristics are:

Purpose

Provides additional information about implemented PE features in AArch64 state.

Configuration


Prior to the introduction of the features described by this register, this register was unnamed and reserved, RES0 from EL1, EL2, and EL3.

Attributes

ID_AA64PFR2_EL1 is a 64-bit register.

Field descriptions

The ID_AA64PFR2_EL1 bit assignments are:



Bits [63:36]

Reserved, RES0.

FPMR, bits [35:32]

Bits [31:20]

Reserved, RES0.

UINJ, bits [19:16]

GCIE, bits [15:12]

Support for the GICv5 CPU interface extension.

GCIE	Meaning
0b0000	GICv5 CPU interface registers not supported.
0b0001	GICv5 CPU interface registers supported.

All other values are reserved.

FEAT_GCIE implements the functionality identified by the value 0b0001.

MTEFAR, bits [11:8]

MTESTOREONLY, bits [7:4]

MTEPERM, bits [3:0]

Accessing ID_AA64PFR2_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ID_AA64PFR2_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b0000	0b0100	0b010

```

if PSTATE.EL == EL0 then
    if IsFeatureImplemented(FEAT_IDST) then
        if EL2Enabled() && HCR_EL2.TGE == '1' then
            AArch64.SystemAccessTrap(EL2, 0x18);
        else
            AArch64.SystemAccessTrap(EL1, 0x18);
        else
            UNDEFINED;
    elseif PSTATE.EL == EL1 then
        if EL2Enabled() && !FALSE && (IsFeatureImplemented(FEAT_FGT) || !IsZero(
            ↪ID_AA64PFR2_EL1) || boolean IMPLEMENTATION_DEFINED "ID_AA64PFR2_EL1
            ↪trapped by HCR_EL2.TID3") && HCR_EL2.TID3 == '1' then
            AArch64.SystemAccessTrap(EL2, 0x18);
        else
            X[t, 64] = ID_AA64PFR2_EL1;
    elseif PSTATE.EL == EL2 then
        X[t, 64] = ID_AA64PFR2_EL1;
    elseif PSTATE.EL == EL3 then
        X[t, 64] = ID_AA64PFR2_EL1;

```

9.3 Virtual CPU interface registers

$\mathbb{I}_{\text{TRMKL}}$	Each ICC_ system register that is accessible at EL1 and higher and whose state is specific to the Virtual Interrupt Domain, has a corresponding virtual ICV_ register. The ICV_ registers are accessed using the same system register encodings as their ICC_ counterparts.
$\mathbb{I}_{\text{TBPCJ}}$	ICC_IDR0 does not have an ICV_ equivalent register.
$\mathbb{I}_{\text{LBWL V}}$	When Legacy operation is disabled, the field layouts and interpretations of the ICV_ registers are identical to the ICC_ registers. This enables binary compatibility for software between the physical and virtual CPU interfaces.
$\mathbb{I}_{\text{MCZFN}}$	ICC_ICSR_EL1 does not have an ICV_ equivalent register.
$\mathbb{I}_{\text{NMXVC}}$	ICC_IAFFIDR_EL1 does not have an ICV_ equivalent register. When virtualization is being used, Arm expects EL2 software to trap EL1 accesses to ICC_IAFFIDR_EL1.

9.3.1 ICV_APR_EL1, Interrupt Controller Virtual Active Priorities Register

The ICV_APR_EL1 characteristics are:

Purpose

Records active priorities for the Virtual CPU interface.

Configuration

AArch64 system register ICV_APR_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICH_APR_EL2[63:0].

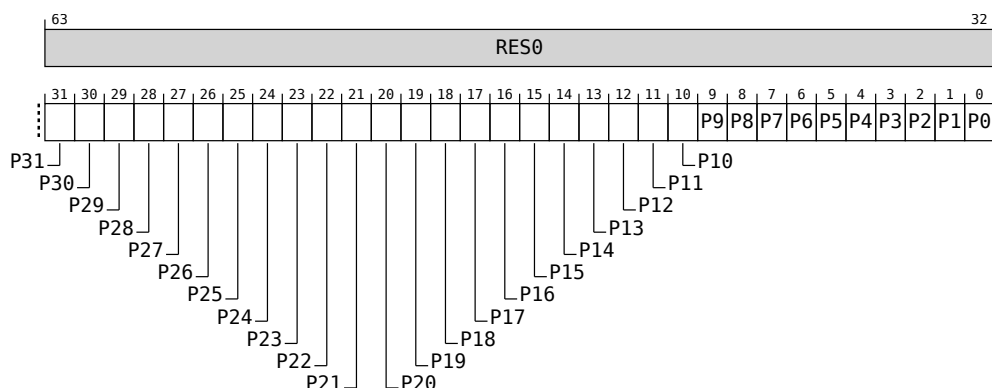
This register is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to ICV_APR_EL1 are UNDEFINED.

Attributes

ICV_APR_EL1 is a 64-bit register.

Field descriptions

The ICV_APR_EL1 bit assignments are:



Bits [63:32]

Reserved, RES0.

P<x>, bits [x], for x = 31 to 0

Provides access to the active priorities.

P<x>	Meaning
0b0	Priority not active
0b1	Priority active

Fields in this register are indexed using the 5-bit priority as an unsigned integer, P[Priority].

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- Access is **RES0** if all of the following are true:
 - $x \bmod 2 == 1$
 - $ICC_IDR0_EL1.PRI_BITS == 0b011$

- Otherwise, access to this field is **RW**

Accessing ICV_APR_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_APR_EL1

op0	op1	CRn	CRm	op2
0b11	0b001	0b1100	0b0000	0b000

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGRTR_EL2.ICC_APR_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_APR_EL1;
    elsif HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            X[t, 64] = ICC_APR_EL1_S;
        elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            X[t, 64] = ICC_APR_EL1_RL;
        elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            X[t, 64] = ICC_APR_EL1_NS;
        else
            UNDEFINED;
        else
            X[t, 64] = ICC_APR_EL1;
    elsif PSTATE.EL == EL2 then
        if HaveEL(EL3) then
            if SCR_EL3.NS == '0' then
                X[t, 64] = ICC_APR_EL1_S;
            elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
                ↪ '1' then
                X[t, 64] = ICC_APR_EL1_RL;
            elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
                X[t, 64] = ICC_APR_EL1_NS;
            else
                UNDEFINED;
            else
                X[t, 64] = ICC_APR_EL1;
    elsif PSTATE.EL == EL3 then
        if SCR_EL3.NS == '0' then
            X[t, 64] = ICC_APR_EL1_S;
        elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS == '1'
            ↪ then
            X[t, 64] = ICC_APR_EL1_RL;
        elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            X[t, 64] = ICC_APR_EL1_NS;
        else
            UNDEFINED;

```

MSR ICC_APR_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b001	0b1100	0b0000	0b000

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGWTR_EL2.ICC_APR_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_APR_EL1 = X[t, 64];
    elsif HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            ICC_APR_EL1_S = X[t, 64];
        elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            ICC_APR_EL1_RL = X[t, 64];
        elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            ICC_APR_EL1_NS = X[t, 64];
        else
            UNDEFINED;
    else
        ICC_APR_EL1 = X[t, 64];
    elsif PSTATE.EL == EL2 then
        if HaveEL(EL3) then
            if SCR_EL3.NS == '0' then
                ICC_APR_EL1_S = X[t, 64];
            elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
                ↪ '1' then
                ICC_APR_EL1_RL = X[t, 64];
            elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
                ICC_APR_EL1_NS = X[t, 64];
            else
                UNDEFINED;
        else
            ICC_APR_EL1 = X[t, 64];
    elsif PSTATE.EL == EL3 then
        if SCR_EL3.NS == '0' then
            ICC_APR_EL1_S = X[t, 64];
        elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS == '1'
            ↪ then
            ICC_APR_EL1_RL = X[t, 64];
        elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            ICC_APR_EL1_NS = X[t, 64];
        else
            UNDEFINED;

```

9.3.2 ICV_CR0_EL1, Interrupt Controller EL1 Virtual Control Register

The ICV_CR0_EL1 characteristics are:

Purpose

Controls behavior of the CPU interface in the Virtual Interrupt Domain.

Configuration


This register is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to ICV_CR0_EL1 are UNDEFINED.

Attributes

ICV_CR0_EL1 is a 64-bit register.

Field descriptions

The ICV_CR0_EL1 bit assignments are:



Bits [63:39]

Reserved, RES0.

Bit [38]

Reserved, RES0.

Bits [37:32]

Reserved, RES0.

Bits [31:1]

Reserved, RES0.

EN, bit [0]

Enable interrupts for the Interrupt Domain.

When this field is 0, there is no HPPI of Sufficient priority for the Interrupt Domain.

EN	Meaning
0b0	Disabled.
0b1	Enabled.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing ICV_CR0_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_CR0_EL1

op0	op1	CRn	CRm	op2
0b11	0b001	0b1100	0b0000	0b001

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGTR_EL2.ICC_CR0_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_CR0_EL1;
    elseif HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            X[t, 64] = ICC_CR0_EL1_S;
        elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            X[t, 64] = ICC_CR0_EL1_RL;
        elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            X[t, 64] = ICC_CR0_EL1_NS;
        else
            UNDEFINED;
    else
        X[t, 64] = ICC_CR0_EL1;
elseif PSTATE.EL == EL2 then
    if HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            X[t, 64] = ICC_CR0_EL1_S;
        elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            X[t, 64] = ICC_CR0_EL1_RL;
        elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            X[t, 64] = ICC_CR0_EL1_NS;
        else
            UNDEFINED;
    else
        X[t, 64] = ICC_CR0_EL1;
elseif PSTATE.EL == EL3 then
    if SCR_EL3.NS == '0' then
        X[t, 64] = ICC_CR0_EL1_S;
    elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS == '1'
        ↪ then
        X[t, 64] = ICC_CR0_EL1_RL;
    elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
        X[t, 64] = ICC_CR0_EL1_NS;
    else
        UNDEFINED;

```

MSR ICC_CR0_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b001	0b1100	0b0000	0b001

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGWTR_EL2.ICC_CR0_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICC_CR0_EL1 = X[t, 64];
    elsif HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            ICC_CR0_EL1_S = X[t, 64];
        elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            ICC_CR0_EL1_RL = X[t, 64];
        elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            ICC_CR0_EL1_NS = X[t, 64];
        else
            UNDEFINED;
    else
        ICC_CR0_EL1 = X[t, 64];
    elsif PSTATE.EL == EL2 then
        if HaveEL(EL3) then
            if SCR_EL3.NS == '0' then
                ICC_CR0_EL1_S = X[t, 64];
            elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
                ↪ '1' then
                ICC_CR0_EL1_RL = X[t, 64];
            elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
                ICC_CR0_EL1_NS = X[t, 64];
            else
                UNDEFINED;
        else
            ICC_CR0_EL1 = X[t, 64];
    elsif PSTATE.EL == EL3 then
        if SCR_EL3.NS == '0' then
            ICC_CR0_EL1_S = X[t, 64];
        elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS == '1'
            ↪ then
            ICC_CR0_EL1_RL = X[t, 64];
        elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            ICC_CR0_EL1_NS = X[t, 64];
        else
            UNDEFINED;

```

9.3.3 ICV_HAPR_EL1, Interrupt Controller Virtual Highest Active Priority Register

The ICV_HAPR_EL1 characteristics are:

Purpose

Reports the running priority of the Virtual Interrupt Domain.

Configuration


This register is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to ICV_HAPR_EL1 are UNDEFINED.

Attributes

ICV_HAPR_EL1 is a 64-bit register.

Field descriptions

The ICV_HAPR_EL1 bit assignments are:



Bits [63:8]

Reserved, RES0.

PRIORITY, bits [7:0]

The running priority for the Virtual Interrupt Domain.

If there are no active priorities on the CPU interface in the applicable Interrupt Domain, or all active priorities have undergone a priority drop, the value returned is the Idle priority.

Accessing ICV_HAPR_EL1

Read-only

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_HAPR_EL1

op0	op1	CRn	CRm	op2
0b11	0b001	0b1100	0b0000	0b011

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGRTR_EL2.ICC_HAPR_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_HAPR_EL1;
    else
        X[t, 64] = ICC_HAPR_EL1;
    elsif PSTATE.EL == EL2 then
        X[t, 64] = ICC_HAPR_EL1;

```

```
elseif PSTATE.EL == EL3 then
    X[t, 64] = ICC_HAPR_EL1;
```

9.3.4 ICV_HPPIR_EL1, Interrupt Controller Virtual Highest Priority Pending Interrupt Register

The ICV_HPPIR_EL1 characteristics are:

Purpose

Reports the HPPI for the Virtual Interrupt Domain.

Configuration

AArch64 system register ICV_HPPIR_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICH_HPPIR_EL2[63:0].


This register is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to ICV_HPPIR_EL1 are UNDEFINED.

Attributes

ICV_HPPIR_EL1 is a 64-bit register.

Field descriptions

The ICV_HPPIR_EL1 bit assignments are:



Bits [63:33]

Reserved, RES0.

HPPIV, bit [32]

HPPI valid.

There is an HPPI with Sufficient priority for the Interrupt Domain.

HPPIV	Meaning
0b0	Invalid: There is no HPPI with Sufficient priority for the Interrupt Domain.
0b1	VALID: There is an HPPI with Sufficient priority for the Interrupt Domain.

If ICV_HPPIR_EL1.HPPIV is 1, ID and TYPE together form the INTID of the HPPI for the Interrupt Domain.

TYPE, bits [31:29]

The encoding of this field depends on the value of HPPIV as described below:

- If ICV_HPPIR_EL1.HPPIV is 0, TYPE is RES0.
- If ICV_HPPIR_EL1.HPPIV is 1, TYPE specifies the Type of the interrupt.

TYPE	Meaning
0b001	PPI

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The encoding of this field depends on the value of HPPIV as described below:

- If ICV_HPPIR_EL1.HPPIV is 0, ID is RES0.
- If ICV_HPPIR_EL1.HPPIV is 1, ID specifies the interrupt ID.

Accessing ICV_HPPIR_EL1

Read-only

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_HPPIR_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1010	0b011

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGTR_EL2.ICC_HPPIR_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_HPPIR_EL1;
    else
        X[t, 64] = ICC_HPPIR_EL1;
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICC_HPPIR_EL1;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_HPPIR_EL1;

```

9.3.5 ICV_PCR_EL1, Interrupt Controller Virtual Interrupt Priority Control Register

The ICV_PCR_EL1 characteristics are:

Purpose

Reports the Virtual priority mask for the Virtual Interrupt Domain.

Configuration


This register is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to ICV_PCR_EL1 are UNDEFINED.

Attributes

ICV_PCR_EL1 is a 64-bit register.

Field descriptions

The ICV_PCR_EL1 bit assignments are:



Bits [63:5]

Reserved, RES0.

PRIORITY, bits [4:0]

The priority mask for the Interrupt Domain.

When fewer than 5 bits of priority is implemented, only bits [4:N] are implemented where N = (4 - ICC_IDR0_EL1.PRI_BITS). Unimplemented bits are RES0.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing ICV_PCR_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PCR_EL1

op0	op1	CRn	CRm	op2
0b11	0b001	0b1100	0b0000	0b010

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGRTR_EL2.ICC_PCR_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PCR_EL1;
    elseif HaveEL(EL3) then
        if SCR_EL3.NS == '0' then

```

```

        X[t, 64] = ICC_PCR_EL1_S;
    elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
        ↪ '1' then
        X[t, 64] = ICC_PCR_EL1_RL;
    elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
        X[t, 64] = ICC_PCR_EL1_NS;
    else
        UNDEFINED;
else
    X[t, 64] = ICC_PCR_EL1;
elseif PSTATE.EL == EL2 then
    if HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            X[t, 64] = ICC_PCR_EL1_S;
        elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            X[t, 64] = ICC_PCR_EL1_RL;
        elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            X[t, 64] = ICC_PCR_EL1_NS;
        else
            UNDEFINED;
    else
        X[t, 64] = ICC_PCR_EL1;
elseif PSTATE.EL == EL3 then
    if SCR_EL3.NS == '0' then
        X[t, 64] = ICC_PCR_EL1_S;
    elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS == '1'
        ↪ then
        X[t, 64] = ICC_PCR_EL1_RL;
    elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
        X[t, 64] = ICC_PCR_EL1_NS;
    else
        UNDEFINED;

```

MSR ICC_PCR_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b001	0b1100	0b0000	0b010

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGWTR_EL2.ICC_PCR_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_PCR_EL1 = X[t, 64];
    elseif HaveEL(EL3) then
        if SCR_EL3.NS == '0' then
            ICC_PCR_EL1_S = X[t, 64];
        elseif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
            ↪ '1' then
            ICC_PCR_EL1_RL = X[t, 64];
        elseif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            ICC_PCR_EL1_NS = X[t, 64];

```

```

        else
            UNDEFINED;
    else
        ICC_PCR_EL1 = X[t, 64];
    elsif PSTATE.EL == EL2 then
        if HaveEL(EL3) then
            if SCR_EL3.NS == '0' then
                ICC_PCR_EL1_S = X[t, 64];
            elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS ==
                ↪ '1' then
                ICC_PCR_EL1_RL = X[t, 64];
            elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
                ICC_PCR_EL1_NS = X[t, 64];
            else
                UNDEFINED;
        else
            ICC_PCR_EL1 = X[t, 64];
    elsif PSTATE.EL == EL3 then
        if SCR_EL3.NS == '0' then
            ICC_PCR_EL1_S = X[t, 64];
        elsif IsFeatureImplemented(FEAT_RME) && SCR_EL3.NSE == '1' && SCR_EL3.NS == '1'
            ↪ then
            ICC_PCR_EL1_RL = X[t, 64];
        elsif SCR_EL3.NSE == '0' && SCR_EL3.NS == '1' then
            ICC_PCR_EL1_NS = X[t, 64];
        else
            UNDEFINED;

```

9.4 PPI registers

Configuration and state of PPIs.

9.4.1 ICC_PPI_CACTIVER<n>_EL1, Interrupt Controller Physical PPI Clear Active Registers, n = 0 - 1

The ICC_PPI_CACTIVER<n>_EL1 characteristics are:

Purpose

Clear Active state for physical PPIs.

Configuration

AArch64 system register ICC_PPI_CACTIVER<n>_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICC_PPI_SACTIVER<n>_EL1[63:0].


This register is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to ICC_PPI_CACTIVER<n>_EL1 are UNDEFINED.

Attributes

ICC_PPI_CACTIVER<n>_EL1 is a 64-bit register.

Field descriptions

The ICC_PPI_CACTIVER<n>_EL1 bit assignments are:



ACTIVE<x>, bits [x], for x = 63 to 0

Configures whether PPIs are Active.

Reads return the Active state of the INTID.

Writing 1 clears the Active state of the INTID. Writing 0 has no effect.

ACTIVE<x>	Meaning
0b0	Inactive
0b1	Active

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 64) + x)
- Access is **RAZ/WI** if all of the following are true:
 - !IsPPIAssignedToCurrentDomain((n * 64) + x)
 - PSTATE.EL IN {EL2, EL1}
- Otherwise, access to this field is **W1C**

Accessing ICC_PPI_CACTIVER<n>_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_CACTIVER<n>_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b00:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGRTR_EL2.ICC_PPI_ACTIVERN_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PPI_CACTIVER_EL1[n];
    else
        X[t, 64] = ICC_PPI_CACTIVER_EL1[n];
    elsif PSTATE.EL == EL2 then
        X[t, 64] = ICC_PPI_CACTIVER_EL1[n];
    elsif PSTATE.EL == EL3 then
        X[t, 64] = ICC_PPI_CACTIVER_EL1[n];

```

MSR ICC_PPI_CACTIVER<n>_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b00:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;

```

```
    elsif EL2Enabled() && ICH_HFGWTR_EL2.ICC_PPI_ACTIVERn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_PPI_CACTIVER_EL1[n] = X[t, 64];
    else
        ICC_PPI_CACTIVER_EL1[n] = X[t, 64];
    elsif PSTATE.EL == EL2 then
        ICC_PPI_CACTIVER_EL1[n] = X[t, 64];
    elsif PSTATE.EL == EL3 then
        ICC_PPI_CACTIVER_EL1[n] = X[t, 64];
```

9.4.2 ICC_PPI_CPENDR<n>_EL1, Interrupt Controller Physical PPI Clear Pending State Registers, n = 0 - 1

The ICC_PPI_CPENDR<n>_EL1 characteristics are:

Purpose

Clear pending state for physical PPIs.

Configuration

AArch64 system register ICC_PPI_CPENDR<n>_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICC_PPI_SPENDR<n>_EL1[63:0].


This register is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to ICC_PPI_CPENDR<n>_EL1 are UNDEFINED.

Attributes

ICC_PPI_CPENDR<n>_EL1 is a 64-bit register.

Field descriptions

The ICC_PPI_CPENDR<n>_EL1 bit assignments are:



PEND<x>, bits [x], for x = 63 to 0

Controls the Pending state of PPIs.

Reads return the current state of the INTIDs.

Writing 1 to a field clears the Pending state of the corresponding INTID. Writing 0 has no effect.

PEND<x>	Meaning
0b0	Not pending
0b1	Pending

The reset behavior of this field is:

- On a Warm reset, this field resets to 0b0.

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 64) + x)
- Access is **RAZ/WI** if all of the following are true:
 - !IsPPIAssignedToCurrentDomain((n * 64) + x)
 - PSTATE.EL IN {EL2, EL1}
- When ICC_PPI_HMR<n>_EL1.HM<x> == 1, access to this field is **RO**
- Otherwise, access to this field is **W1C**

Accessing ICC_PPI_CPENDR<n>_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_CPENDR<n>_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b10:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGRTR_EL2.ICC_PPI_PENDRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PPI_CPENDR_EL1[n];
    else
        X[t, 64] = ICC_PPI_CPENDR_EL1[n];
    elsif PSTATE.EL == EL2 then
        X[t, 64] = ICC_PPI_CPENDR_EL1[n];
    elsif PSTATE.EL == EL3 then
        X[t, 64] = ICC_PPI_CPENDR_EL1[n];

```

MSR ICC_PPI_CPENDR<n>_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b10:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then

```

```
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGWTR_EL2.ICC_PPI_PENDRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_PPI_CPENDR_EL1[n] = X[t, 64];
    else
        ICC_PPI_CPENDR_EL1[n] = X[t, 64];
    elsif PSTATE.EL == EL2 then
        ICC_PPI_CPENDR_EL1[n] = X[t, 64];
    elsif PSTATE.EL == EL3 then
        ICC_PPI_CPENDR_EL1[n] = X[t, 64];
```

9.4.3 ICC_PPI_DOMAINR<n>_EL3, Interrupt Controller PPI Domain Registers, n = 0 - 3

The ICC_PPI_DOMAINR<n>_EL3 characteristics are:

Purpose

Controls which Interrupt Domain PPI sources are assigned to.

Configuration

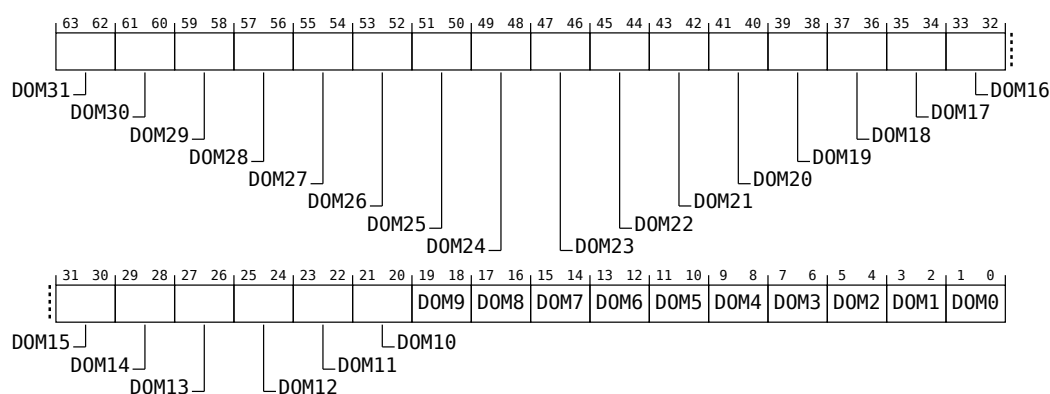
This register is present only when FEAT_GCIE is implemented and EL3 is implemented. Otherwise, direct accesses to ICC_PPI_DOMAINR<n>_EL3 are UNDEFINED.

Attributes

ICC_PPI_DOMAINR<n>_EL3 is a 64-bit register.

Field descriptions

The ICC_PPI_DOMAINR<n>_EL3 bit assignments are:



DOM<x>, bits [2x+1:2x], for x = 31 to 0

Controls the Physical Interrupt Domain that a PPI is assigned to.

Encodings for unimplemented Interrupt Domains are reserved.

DOM<x>	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- RES0** if !IsPPIImplemented((n * 32) + x)
- Otherwise, access to this field is **RW**

Accessing ICC_PPI_DOMAINR<n>_EL3

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_DOMAINR<n>_EL3

op0	op1	CRn	CRm	op2
0b11	0b110	0b1100	0b1000	0b1:n[1:0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_PPI_DOMAINR_EL3[n];

```

MSR ICC_PPI_DOMAINR<n>_EL3, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b110	0b1100	0b1000	0b1:n[1:0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    UNDEFINED;
elsif PSTATE.EL == EL2 then
    UNDEFINED;
elsif PSTATE.EL == EL3 then
    ICC_PPI_DOMAINR_EL3[n] = X[t, 64];

```

9.4.4 ICC_PPI_ENABLER<n>_EL1, Interrupt Controller Physical PPI Enable Registers, n = 0 - 1

The ICC_PPI_ENABLER<n>_EL1 characteristics are:

Purpose

Access to Enable state for physical PPIs.

Configuration


This register is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to ICC_PPI_ENABLER<n>_EL1 are UNDEFINED.

Attributes

ICC_PPI_ENABLER<n>_EL1 is a 64-bit register.

Field descriptions

The ICC_PPI_ENABLER<n>_EL1 bit assignments are:



EN<x>, bits [x], for x = 63 to 0

Configures whether PPIs are enabled.

Reads return the current state of the INTID.

EN<x>	Meaning
0b0	Disabled
0b1	Enabled

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 64) + x)
- Access is **RAZ/WI** if all of the following are true:
 - !IsPPIAssignedToCurrentDomain((n * 64) + x)
 - PSTATE.EL IN {EL2, EL1}
- Otherwise, access to this field is **RW**

Accessing ICC_PPI_ENABLER<n>_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_ENABLER<n>_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1010	0b11:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGRTR_EL2.ICC_PPI_ENABLERn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PPI_ENABLER_EL1[n];
    else
        X[t, 64] = ICC_PPI_ENABLER_EL1[n];
elseif PSTATE.EL == EL2 then
    X[t, 64] = ICC_PPI_ENABLER_EL1[n];
elseif PSTATE.EL == EL3 then
    X[t, 64] = ICC_PPI_ENABLER_EL1[n];

```

MSR ICC_PPI_ENABLER<n>_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1010	0b11:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGWTR_EL2.ICC_PPI_ENABLERn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_PPI_ENABLER_EL1[n] = X[t, 64];
    else
        ICC_PPI_ENABLER_EL1[n] = X[t, 64];
elseif PSTATE.EL == EL2 then

```

Chapter 9. System registers

9.4. PPI registers

```
    ICC_PPI_ENABLER_EL1[n] = X[t, 64];  
elseif PSTATE.EL == EL3 then  
    ICC_PPI_ENABLER_EL1[n] = X[t, 64];
```

9.4.5 ICC_PPI_HMR<n>_EL1, Interrupt Controller Physical PPI Handling mode Registers, n = 0 - 1

The ICC_PPI_HMR<n>_EL1 characteristics are:

Purpose

Report whether physical PPIs are Edge or Level.

Configuration


This register is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to ICC_PPI_HMR<n>_EL1 are UNDEFINED.

Attributes

ICC_PPI_HMR<n>_EL1 is a 64-bit register.

Field descriptions

The ICC_PPI_HMR<n>_EL1 bit assignments are:



HM<x>, bits [x], for x = 63 to 0

The PPI Handling mode.

HM<x>	Meaning
0b0	Edge
0b1	Level

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 64) + x)
- Otherwise, access to this field is **RO**

Accessing ICC_PPI_HMR<n>_EL1

Read-only

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_HMR<n>_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1010	0b00:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGRTR_EL2.ICC_PPI_HMRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PPI_HMR_EL1[n];
    else
        X[t, 64] = ICC_PPI_HMR_EL1[n];
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICC_PPI_HMR_EL1[n];
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_PPI_HMR_EL1[n];

```

9.4.6 ICC_PPI_PRIORITYR<n>_EL1, Interrupt Controller Physical PPI Priority Registers, n = 0 - 15

The ICC_PPI_PRIORITYR<n>_EL1 characteristics are:

Purpose

Configures the priority of physical PPIs.

Configuration

This register is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to ICC_PPI_PRIORITYR<n>_EL1 are UNDEFINED.

Attributes

ICC_PPI_PRIORITYR<n>_EL1 is a 64-bit register.

Field descriptions

The ICC_PPI_PRIORITYR<n>_EL1 bit assignments are:

63	61	60	56	55	53	52	48	47	45	44	40	39	37	36	32
RES0	PRIORITY7	RES0	PRIORITY6	RES0	PRIORITY5	RES0	PRIORITY4	RES0	PRIORITY3	RES0	PRIORITY2	RES0	PRIORITY1	RES0	PRIORITY0

Bits [63:61, 55:53, 47:45, 39:37, 31:29, 23:21, 15:13, 7:5]

Reserved, RES0.

PRIORITY<x>, bits [60:56, 52:48, 44:40, 36:32, 28:24, 20:16, 12:8, 4:0], for x = 7 to 0, where each field is 5 bits wide

Configures the priority of the corresponding PPI.

Only the upper N bits of each 5-bit Priority field are implemented where N = (ICC_IDR0_EL1.PRI_BITS + 1). Unimplemented bits are RES0.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- RES0 if !IsPPIImplemented((n * 8) + x)
- Access is **RAZ/WI** if all of the following are true:
 - !IsPPIAssignedToCurrentDomain((n * 8) + x)
 - PSTATE.EL IN {EL2, EL1}
- Otherwise, access to this field is **RW**

Accessing ICC_PPI_PRIORITYR<n>_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_PRIORITYR<n>_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b111:n[3]	n[2:0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGRTR_EL2.ICC_PPI_PRIORITYRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PPI_PRIORITYR_EL1[n];
    else
        X[t, 64] = ICC_PPI_PRIORITYR_EL1[n];
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICC_PPI_PRIORITYR_EL1[n];
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_PPI_PRIORITYR_EL1[n];

```

MSR ICC_PPI_PRIORITYR<n>_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b111:n[3]	n[2:0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGWTR_EL2.ICC_PPI_PRIORITYRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_PPI_PRIORITYR_EL1[n] = X[t, 64];
    else
        ICC_PPI_PRIORITYR_EL1[n] = X[t, 64];
elsif PSTATE.EL == EL2 then
    ICC_PPI_PRIORITYR_EL1[n] = X[t, 64];
elsif PSTATE.EL == EL3 then
    ICC_PPI_PRIORITYR_EL1[n] = X[t, 64];

```

9.4.7 ICC_PPI_SACTIVER<n>_EL1, Interrupt Controller Physical PPI Set Active Registers, n = 0 - 1

The ICC_PPI_SACTIVER<n>_EL1 characteristics are:

Purpose

Set Active state for physical PPIs.

Configuration

AArch64 system register ICC_PPI_SACTIVER<n>_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICC_PPI_CACTIVER<n>_EL1[63:0].


This register is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to ICC_PPI_SACTIVER<n>_EL1 are UNDEFINED.

Attributes

ICC_PPI_SACTIVER<n>_EL1 is a 64-bit register.

Field descriptions

The ICC_PPI_SACTIVER<n>_EL1 bit assignments are:



ACTIVE<x>, bits [x], for x = 63 to 0

Configures whether PPIs are Active.

Reads return the Active state of the INTID.

Writing 1 sets the Active state of the INTID. Writing 0 has no effect.

ACTIVE<x>	Meaning
0b0	Inactive
0b1	Active

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 64) + x)
- Access is **RAZ/WI** if all of the following are true:
 - !IsPPIAssignedToCurrentDomain((n * 64) + x)
 - PSTATE.EL IN {EL2, EL1}
- Otherwise, access to this field is **W1S**

Accessing ICC_PPI_SACTIVER<n>_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_SACTIVER<n>_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b01:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGTR_EL2.ICC_PPI_SACTIVERn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PPI_SACTIVER_EL1[n];
    else
        X[t, 64] = ICC_PPI_SACTIVER_EL1[n];
    elsif PSTATE.EL == EL2 then
        X[t, 64] = ICC_PPI_SACTIVER_EL1[n];
    elsif PSTATE.EL == EL3 then
        X[t, 64] = ICC_PPI_SACTIVER_EL1[n];

```

MSR ICC_PPI_SACTIVER<n>_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b01:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;

```

```
    elsif EL2Enabled() && ICH_HFGWTR_EL2.ICC_PPI_ACTIVERn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_PPI_SACTIVER_EL1[n] = X[t, 64];
    else
        ICC_PPI_SACTIVER_EL1[n] = X[t, 64];
    elsif PSTATE.EL == EL2 then
        ICC_PPI_SACTIVER_EL1[n] = X[t, 64];
    elsif PSTATE.EL == EL3 then
        ICC_PPI_SACTIVER_EL1[n] = X[t, 64];
```

9.4.8 ICC_PPI_SPENDR<n>_EL1, Interrupt Controller Physical PPI Set Pending State Registers, n = 0 - 1

The ICC_PPI_SPENDR<n>_EL1 characteristics are:

Purpose

Set pending state for Physical PPIs.

Configuration

AArch64 system register ICC_PPI_SPENDR<n>_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICC_PPI_CPENDR<n>_EL1[63:0].


This register is present only when FEAT_GCIE is implemented. Otherwise, direct accesses to ICC_PPI_SPENDR<n>_EL1 are UNDEFINED.

Attributes

ICC_PPI_SPENDR<n>_EL1 is a 64-bit register.

Field descriptions

The ICC_PPI_SPENDR<n>_EL1 bit assignments are:



PEND<x>, bits [x], for x = 63 to 0

Controls the Pending state of PPIs.

Reads return the current state of the INTIDs.

Writing 1 to a field sets the Pending state of the corresponding INTID. Writing 0 has no effect.

PEND<x>	Meaning
0b0	Not pending
0b1	Pending

The reset behavior of this field is:

- On a Warm reset, this field resets to 0b0.

Accessing this field has the following behavior:

- RES0** if !IsPPIImplemented((n * 64) + x)
- Access is **RAZ/WI** if all of the following are true:
 - !IsPPIAssignedToCurrentDomain((n * 64) + x)
 - PSTATE.EL IN {EL2, EL1}
- When ICC_PPI_HMR<n>_EL1.HM<x> == 1, access to this field is **RO**
- Otherwise, access to this field is **WIS**

Accessing ICC_PPI_SPENDR<n>_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_SPENDR<n>_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b11:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGRTR_EL2.ICC_PPI_PENDRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PPI_SPENDR_EL1[n];
    else
        X[t, 64] = ICC_PPI_SPENDR_EL1[n];
    elsif PSTATE.EL == EL2 then
        X[t, 64] = ICC_PPI_SPENDR_EL1[n];
    elsif PSTATE.EL == EL3 then
        X[t, 64] = ICC_PPI_SPENDR_EL1[n];

```

MSR ICC_PPI_SPENDR<n>_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b11:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then

```

```
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGWTR_EL2.ICC_PPI_PENDRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_PPI_SPENDR_EL1[n] = X[t, 64];
    else
        ICC_PPI_SPENDR_EL1[n] = X[t, 64];
    elsif PSTATE.EL == EL2 then
        ICC_PPI_SPENDR_EL1[n] = X[t, 64];
    elsif PSTATE.EL == EL3 then
        ICC_PPI_SPENDR_EL1[n] = X[t, 64];
```

9.5 Virtual PPI registers

ICV_PPI Each ICC_PPI_ system register which is accessible at EL1 has a corresponding virtual ICV_PPI_ register. The ICV_PPI_ registers are accessed using the same system register encodings as their ICC_PPI_ counterparts.

Configuration and state of virtual PPIs.

9.5.1 ICV_PPI_CACTIVER<n>_EL1, Interrupt Controller Virtual PPI Clear Active Registers, n = 0 - 1

The ICV_PPI_CACTIVER<n>_EL1 characteristics are:

Purpose

Clear Active state for virtual PPIs.

Configuration

AArch64 system register ICV_PPI_CACTIVER<n>_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICH_PPI_ACTIVETER<n>_EL2[63:0].

AArch64 system register ICV_PPI_CACTIVER<n>_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICV_PPI_SACTIVER<n>_EL1[63:0].


This register is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to ICV_PPI_CACTIVER<n>_EL1 are UNDEFINED.

Attributes

ICV_PPI_CACTIVER<n>_EL1 is a 64-bit register.

Field descriptions

The ICV_PPI_CACTIVER<n>_EL1 bit assignments are:



ACTIVE<x>, bits [x], for x = 63 to 0

Configures whether PPIs are Active.

Reads return the Active state of the INTID.

Writing 1 clears the Active state of the INTID. Writing 0 has no effect.

ACTIVE<x>	Meaning
0b0	Inactive
0b1	Active

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 64) + x)
- Otherwise, access to this field is **W1C**

Accessing ICV_PPI_CACTIVER<n>_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_CACTIVER<n>_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b00:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGTR_EL2.ICC_PPI_ACTIVERN_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PPI_CACTIVER_EL1[n];
    else
        X[t, 64] = ICC_PPI_CACTIVER_EL1[n];
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICC_PPI_CACTIVER_EL1[n];
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_PPI_CACTIVER_EL1[n];

```

MSR ICC_PPI_CACTIVER<n>_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b00:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGWTR_EL2.ICC_PPI_ACTIVERN_EL1 == '0' then

```

```
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_PPI_CACTIVER_EL1[n] = X[t, 64];
    else
        ICC_PPI_CACTIVER_EL1[n] = X[t, 64];
elsif PSTATE.EL == EL2 then
    ICC_PPI_CACTIVER_EL1[n] = X[t, 64];
elsif PSTATE.EL == EL3 then
    ICC_PPI_CACTIVER_EL1[n] = X[t, 64];
```

9.5.2 ICV_PPI_CPENDR<n>_EL1, Interrupt Controller Virtual PPI Clear Pending State Registers, n = 0 - 1

The ICV_PPI_CPENDR<n>_EL1 characteristics are:

Purpose

Clear pending state for virtual PPIs.

Configuration

AArch64 system register ICV_PPI_CPENDR<n>_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICH_PPI_PENDR<n>_EL2[63:0].

AArch64 system register ICV_PPI_CPENDR<n>_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICV_PPI_SPENDR<n>_EL1[63:0].


This register is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to ICV_PPI_CPENDR<n>_EL1 are UNDEFINED.

Attributes

ICV_PPI_CPENDR<n>_EL1 is a 64-bit register.

Field descriptions

The ICV_PPI_CPENDR<n>_EL1 bit assignments are:



PEND<x>, bits [x], for x = 63 to 0

Controls the Pending state of PPIs.

Reads return the current state of the INTIDs.

Writing 1 to a field clears the Pending state of the corresponding INTID. Writing 0 has no effect.

PEND<x>	Meaning
0b0	Not pending
0b1	Pending

When the Pending state of a physical PPI is directly injected to the Pending state of virtual PPI <x>, all of the following are true:

- Reads of Pend<x> return the value of the field corresponding to the physical PPI in ICC_PPI_CPENDR<n>_EL1.
- Writes to Pend<x> have the same effect as writes to the field corresponding to the physical PPI in ICC_PPI_CPENDR<n>_EL1.

Otherwise, all of the following are true:

- Reads of Pend<x> return the value of ICH_PPI_PENDR<n>_EL2.Pend<x>.
- Writes to Pend<x> update the value of ICH_PPI_PENDR<n>_EL2.Pend<x>.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 64) + x)
- When ICV_PPI_HMR<n>_EL1.HM<x> == 1, access to this field is **RO**
- Otherwise, access to this field is **W1C**

Accessing ICV_PPI_CPENDR<n>_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_CPENDR<n>_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b10:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGTR_EL2.ICC_PPI_PENDRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PPI_CPENDR_EL1[n];
    else
        X[t, 64] = ICC_PPI_CPENDR_EL1[n];
elseif PSTATE.EL == EL2 then
    X[t, 64] = ICC_PPI_CPENDR_EL1[n];
elseif PSTATE.EL == EL3 then
    X[t, 64] = ICC_PPI_CPENDR_EL1[n];

```


MSR ICC_PPI_CPENDR<n>_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b10:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGWTR_EL2.ICC_PPI_PENDRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_PPI_CPENDR_EL1[n] = X[t, 64];
    else
        ICC_PPI_CPENDR_EL1[n] = X[t, 64];
elsif PSTATE.EL == EL2 then
    ICC_PPI_CPENDR_EL1[n] = X[t, 64];
elsif PSTATE.EL == EL3 then
    ICC_PPI_CPENDR_EL1[n] = X[t, 64];

```

9.5.3 ICV_PPI_ENABLER<n>_EL1, Interrupt Controller Virtual PPI Clear Enable Registers, n = 0 - 1

The ICV_PPI_ENABLER<n>_EL1 characteristics are:

Purpose

Access to Enabled state for virtual PPIs.

Configuration

AArch64 system register ICV_PPI_ENABLER<n>_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICH_PPI_ENABLER<n>_EL2[63:0].


This register is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to ICV_PPI_ENABLER<n>_EL1 are UNDEFINED.

Attributes

ICV_PPI_ENABLER<n>_EL1 is a 64-bit register.

Field descriptions

The ICV_PPI_ENABLER<n>_EL1 bit assignments are:



EN<x>, bits [x], for x = 63 to 0

Configures whether PPIs are enabled.

Reads return the current state of the INTID.

EN<x>	Meaning
0b0	Disabled
0b1	Enabled

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 64) + x)
- Otherwise, access to this field is **RW**

Accessing ICV_PPI_ENABLER<n>_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_ENABLER<n>_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1010	0b11:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGTR_EL2.ICC_PPI_ENABLERn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PPI_ENABLER_EL1[n];
    else
        X[t, 64] = ICC_PPI_ENABLER_EL1[n];
elseif PSTATE.EL == EL2 then
    X[t, 64] = ICC_PPI_ENABLER_EL1[n];
elseif PSTATE.EL == EL3 then
    X[t, 64] = ICC_PPI_ENABLER_EL1[n];

```

MSR ICC_PPI_ENABLER<n>_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1010	0b11:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGWTR_EL2.ICC_PPI_ENABLERn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then

```

```
        ICC_PPI_ENABLER_EL1[n] = X[t, 64];  
    else  
        ICC_PPI_ENABLER_EL1[n] = X[t, 64];  
elseif PSTATE.EL == EL2 then  
    ICC_PPI_ENABLER_EL1[n] = X[t, 64];  
elseif PSTATE.EL == EL3 then  
    ICC_PPI_ENABLER_EL1[n] = X[t, 64];
```

9.5.4 ICV_PPI_HMR<n>_EL1, Interrupt Controller Virtual PPI Handling mode Registers, n = 0 - 1

The ICV_PPI_HMR<n>_EL1 characteristics are:

Purpose

Report whether virtual PPIs are Edge or Level.

Configuration


This register is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to ICV_PPI_HMR<n>_EL1 are UNDEFINED.

Attributes

ICV_PPI_HMR<n>_EL1 is a 64-bit register.

Field descriptions

The ICV_PPI_HMR<n>_EL1 bit assignments are:



HM<x>, bits [x], for x = 63 to 0

The PPI Handling mode.

HM<x>	Meaning
0b0	Edge
0b1	Level

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 64) + x)
- Otherwise, access to this field is **RO**

Accessing ICV_PPI_HMR<n>_EL1

Read-only

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_HMR<n>_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1010	0b00:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGRTR_EL2.ICC_PPI_HMRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PPI_HMR_EL1[n];
    else
        X[t, 64] = ICC_PPI_HMR_EL1[n];
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICC_PPI_HMR_EL1[n];
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_PPI_HMR_EL1[n];

```

9.5.5 ICV_PPI_PRIORITYR<n>_EL1, Interrupt Controller Virtual PPI Priority Registers, n = 0 - 15

The ICV_PPI_PRIORITYR<n>_EL1 characteristics are:

Purpose

Configures the priority of virtual PPIs.

Configuration

AArch64 system register ICV_PPI_PRIORITYR<n>_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICH_PPI_PRIORITYR<n>_EL2[63:0].

This register is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to ICV_PPI_PRIORITYR<n>_EL1 are UNDEFINED.

Attributes

ICV_PPI_PRIORITYR<n>_EL1 is a 64-bit register.

Field descriptions

The ICV_PPI_PRIORITYR<n>_EL1 bit assignments are:

63	61	60	56	55	53	52	48	47	45	44	40	39	37	36	32
RES0		PRIORITY7		RES0		PRIORITY6		RES0		PRIORITY5		RES0		PRIORITY4	

31	29	28	24	23	21	20	16	15	13	12	8	7	5	4	0
RES0		PRIORITY3		RES0		PRIORITY2		RES0		PRIORITY1		RES0		PRIORITY0	

Bits [63:61, 55:53, 47:45, 39:37, 31:29, 23:21, 15:13, 7:5]

Reserved, RES0.

PRIORITY<x>, bits [60:56, 52:48, 44:40, 36:32, 28:24, 20:16, 12:8, 4:0], for x = 7 to 0, where each field is 5 bits wide

Configures the priority of the corresponding PPI.

Only the upper N bits of each 5-bit Priority field are implemented where $N = (ICC_IDR0_EL1.PRI_BITS + 1)$. Unimplemented bits are RES0.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 8) + x)
- Otherwise, access to this field is **RW**

Accessing ICV_PPI_PRIORITYR<n>>_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_PRIORITYR<n>_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b111:n[3]	n[2:0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGRTR_EL2.ICC_PPI_PRIORITYRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PPI_PRIORITYR_EL1[n];
    else
        X[t, 64] = ICC_PPI_PRIORITYR_EL1[n];
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICC_PPI_PRIORITYR_EL1[n];
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_PPI_PRIORITYR_EL1[n];

```

MSR ICC_PPI_PRIORITYR<n>_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b111:n[3]	n[2:0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGWTR_EL2.ICC_PPI_PRIORITYRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_PPI_PRIORITYR_EL1[n] = X[t, 64];
    else
        ICC_PPI_PRIORITYR_EL1[n] = X[t, 64];
elsif PSTATE.EL == EL2 then
    ICC_PPI_PRIORITYR_EL1[n] = X[t, 64];
elsif PSTATE.EL == EL3 then
    ICC_PPI_PRIORITYR_EL1[n] = X[t, 64];

```

9.5.6 ICV_PPI_SACTIVER<n>_EL1, Interrupt Controller Virtual PPI Set Active Registers, n = 0 - 1

The ICV_PPI_SACTIVER<n>_EL1 characteristics are:

Purpose

Set Active state for virtual PPIs.

Configuration

AArch64 system register ICV_PPI_SACTIVER<n>_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICH_PPI_ACTIVER<n>_EL2[63:0].

AArch64 system register ICV_PPI_SACTIVER<n>_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICV_PPI_CACTIVER<n>_EL1[63:0].


This register is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to ICV_PPI_SACTIVER<n>_EL1 are UNDEFINED.

Attributes

ICV_PPI_SACTIVER<n>_EL1 is a 64-bit register.

Field descriptions

The ICV_PPI_SACTIVER<n>_EL1 bit assignments are:



ACTIVE<x>, bits [x], for x = 63 to 0

Configures whether PPIs are Active.

Reads return the Active state of the INTID.

Writing 1 sets the Active state of the INTID. Writing 0 has no effect.

ACTIVE<x>	Meaning
0b0	Inactive
0b1	Active

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 64) + x)
- Otherwise, access to this field is **W1S**

Accessing ICV_PPI_SACTIVER<n>_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_SACTIVER<n>_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b01:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGRTR_EL2.ICC_PPI_ACTIVERn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PPI_SACTIVER_EL1[n];
    else
        X[t, 64] = ICC_PPI_SACTIVER_EL1[n];
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICC_PPI_SACTIVER_EL1[n];
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICC_PPI_SACTIVER_EL1[n];

```

MSR ICC_PPI_SACTIVER<n>_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b01:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elsif EL2Enabled() && ICH_HFGWTR_EL2.ICC_PPI_ACTIVERn_EL1 == '0' then

```

```
        AArch64.SystemAccessTrap(EL2, 0x18);
    elsif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_PPI_SACTIVER_EL1[n] = X[t, 64];
    else
        ICC_PPI_SACTIVER_EL1[n] = X[t, 64];
elsif PSTATE.EL == EL2 then
    ICC_PPI_SACTIVER_EL1[n] = X[t, 64];
elsif PSTATE.EL == EL3 then
    ICC_PPI_SACTIVER_EL1[n] = X[t, 64];
```

9.5.7 ICV_PPI_SPENDR<n>_EL1, Interrupt Controller Virtual PPI Set Pending State Registers, n = 0 - 1

The ICV_PPI_SPENDR<n>_EL1 characteristics are:

Purpose

Set pending state for virtual PPIs.

Configuration

AArch64 system register ICV_PPI_SPENDR<n>_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICH_PPI_PENDR<n>_EL2[63:0].

AArch64 system register ICV_PPI_SPENDR<n>_EL1 bits [63:0] are architecturally mapped to AArch64 system register ICV_PPI_CPENDR<n>_EL1[63:0].


This register is present only when FEAT_GCIE is implemented and EL2 is implemented. Otherwise, direct accesses to ICV_PPI_SPENDR<n>_EL1 are UNDEFINED.

Attributes

ICV_PPI_SPENDR<n>_EL1 is a 64-bit register.

Field descriptions

The ICV_PPI_SPENDR<n>_EL1 bit assignments are:



PEND<x>, bits [x], for x = 63 to 0

Controls the Pending state of PPIs.

Reads return the current state of the INTIDs.

Writing 1 to a field sets the Pending state of the corresponding INTID. Writing 0 has no effect.

PEND<x>	Meaning
0b0	Not pending
0b1	Pending

When the Pending state of a physical PPI is directly injected to the Pending state of virtual PPI <x>, all of the following are true:

- Reads of Pend<x> return the value of the field corresponding to the physical PPI in ICC_PPI_SPENDR<n>_EL1.
- Writes to Pend<x> have the same effect as writes to the field corresponding to the physical PPI in ICC_PPI_SPENDR<n>_EL1.

Otherwise, all of the following are true:

- Reads of Pend<x> return the value of ICH_PPI_PENDR<n>_EL2.Pend<x>.
- Writes to Pend<x> update the value of ICH_PPI_PENDR<n>_EL2.Pend<x>.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 64) + x)
- When ICV_PPI_HMR<n>_EL1.HM<x> == 1, access to this field is **RO**
- Otherwise, access to this field is **W1S**

Accessing ICV_PPI_SPENDR<n>_EL1

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICC_PPI_SPENDR<n>_EL1

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b11:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGTR_EL2.ICC_PPI_PENDRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        X[t, 64] = ICV_PPI_SPENDR_EL1[n];
    else
        X[t, 64] = ICC_PPI_SPENDR_EL1[n];
elseif PSTATE.EL == EL2 then
    X[t, 64] = ICC_PPI_SPENDR_EL1[n];
elseif PSTATE.EL == EL3 then
    X[t, 64] = ICC_PPI_SPENDR_EL1[n];

```

MSR ICC_PPI_SPENDR<n>_EL1, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b000	0b1100	0b1101	0b11:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.IMO == '1' && ICH_VCTLR_EL2.V3 == '1' then
        UNDEFINED;
    elseif EL2Enabled() && ICH_HFGWTR_EL2.ICC_PPI_PENDRn_EL1 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    elseif EL2Enabled() && HCR_EL2.IMO == '1' then
        ICV_PPI_SPENDR_EL1[n] = X[t, 64];
    else
        ICC_PPI_SPENDR_EL1[n] = X[t, 64];
elseif PSTATE.EL == EL2 then
    ICC_PPI_SPENDR_EL1[n] = X[t, 64];
elseif PSTATE.EL == EL3 then
    ICC_PPI_SPENDR_EL1[n] = X[t, 64];

```

9.6 Hypervisor control registers

Registers to manage operation of the GICv5 and Legacy virtual CPU interface.

9.6.1 ICH_APR_EL2, Interrupt Controller Active Virtual Priorities Register

The ICH_APR_EL2 characteristics are:

Purpose

Records active priorities for the virtual interrupt domain.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

AArch64 system register ICH_APR_EL2 bits [63:0] are architecturally mapped to AArch64 system register ICV_APR_EL1[63:0].

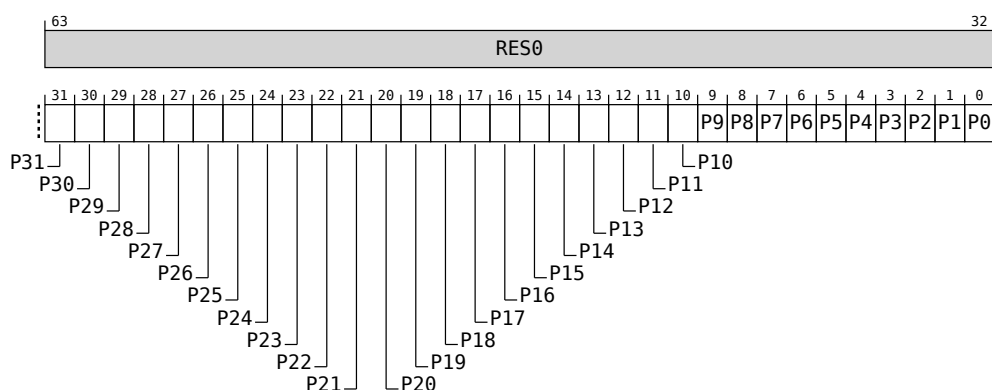
This register is present only when FEAT_GCIE is implemented and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_APR_EL2 are UNDEFINED.

Attributes

ICH_APR_EL2 is a 64-bit register.

Field descriptions

The ICH_APR_EL2 bit assignments are:



Bits [63:32]

Reserved, RES0.

P<x>, bits [x], for x = 31 to 0

Provides access to the active virtual priorities.

P<x>	Meaning
0b0	Priority not active
0b1	Priority active

This field is an alias of the equivalent field in ICV_APR_EL1.

Accessing this field has the following behavior:

- Access is **RES0** if all of the following are true:
 - x MOD 2 == 1
 - ICC_IDR0_EL1.PRI_BITS == 0b011
- Otherwise, access to this field is **RW**

Accessing ICH_APR_EL2

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_APR_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1000	0b100

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        X[t, 64] = NVMem[0xB00];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICH_APR_EL2;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICH_APR_EL2;

```

MSR ICH_APR_EL2, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1000	0b100

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0xB00] = X[t, 64];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    ICH_APR_EL2 = X[t, 64];
elsif PSTATE.EL == EL3 then
    ICH_APR_EL2 = X[t, 64];

```

9.6.2 ICH_CONTEXTR_EL2, Interrupt Controller Virtual Context Register

The ICH_CONTEXTR_EL2 characteristics are:

Purpose

Selects the current resident VPE.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.


This register is present only when FEAT_GCIE is implemented and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_CONTEXTR_EL2 are UNDEFINED.

Attributes

ICH_CONTEXTR_EL2 is a 64-bit register.

Field descriptions

The ICH_CONTEXTR_EL2 bit assignments are:



V, bit [63]

Indicates whether a VPE is currently resident.

If EL2 is not enabled in the Security state identified by SCR_EL3.{NSE,NS}, the Effective value of this field is 0.

V	Meaning
0b0	No VPE is resident.
0b1	A VPE is resident.

The reset behavior of this field is:

- On a Warm reset, this field resets to 0b0.

F, bit [62]

Indicates whether the last write that set V to 1 succeeded in making a VPE resident.

F	Meaning
0b0	Success
0b1	Fault

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

IRICHPPIDIS, bit [61]

Specifies whether the candidate HPPI presented by the IRI for the resident VPE is considered when selecting the HPPI for the Virtual Interrupt Domain.

When V is 0, this field is IGNORED.

IRICHPPIDIS	Meaning
0b0	The CPU interface considers both the virtual PPIs and the candidate HPPI presented by the IRS when determining the HPPI for the Virtual Interrupt Domain.
0b1	The CPU interface only considers the virtual PPIs when determining the HPPI for the Virtual Interrupt Domain.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

DB, bit [60]

Specifies whether a doorbell interrupt is requested when making a VPE non-resident.

DB	Meaning
0b0	No doorbell requested.
0b1	Doorbell requested.

On a write that changes V from 1 to 0, this field specifies whether a doorbell interrupt is requested for the previously resident VPE.

For all other writes, and for reads, this field is RES0.

If the current value of SCR_EL3.{NSE,NS} is different from the value at the time the VPE was made resident, this field behaves as if set to 0.

This field has no effect if EL2 is not enabled in the Security state identified by SCR_EL3.{NSE,NS}.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

DBPM, bits [59:55]

Doorbell priority mask.

On writes, when DB is 1, this field specifies the minimum priority for a virtual interrupt to trigger the VPE's doorbell.

On writes, when DB is 0, this field is IGNORED.

On reads, this field is UNKNOWN.

This field has no effect if EL2 is not enabled in the Security state identified by SCR_EL3.{NSE,NS}.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Bits [54:48]

Reserved, RES0.

VPE, bits [47:32]

When V is 1, identifies the resident VPE.

When V is 0, this field is IGNORED.

This field has no effect if EL2 is not enabled in the Security state identified by SCR_EL3.{NSE,NS}.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Bits [31:16]

Reserved, RES0.

VM, bits [15:0]

When V is 1, identifies the resident VM.

When V is 0, this field is IGNORED.

This field has no effect if EL2 is not enabled in the Security state identified by SCR_EL3.{NSE,NS}.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing ICH_CONTEXTR_EL2

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_CONTEXTR_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1011	0b110

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        X[t, 64] = NVMem[0xB08];
    elseif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elseif PSTATE.EL == EL2 then
    X[t, 64] = ICH_CONTEXTR_EL2;
elseif PSTATE.EL == EL3 then
    X[t, 64] = ICH_CONTEXTR_EL2;

```

MSR ICH_CONTEXTR_EL2, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1011	0b110

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0xB08] = X[t, 64];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    ICH_CONTEXTR_EL2 = X[t, 64];
elsif PSTATE.EL == EL3 then
    ICH_CONTEXTR_EL2 = X[t, 64];

```

9.6.3 ICH_HFGITR_EL2, Hypervisor GIC Fine-Grained Instruction Trap Register

The ICH_HFGITR_EL2 characteristics are:

Purpose

Provides instruction trap controls for GIC System instructions.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

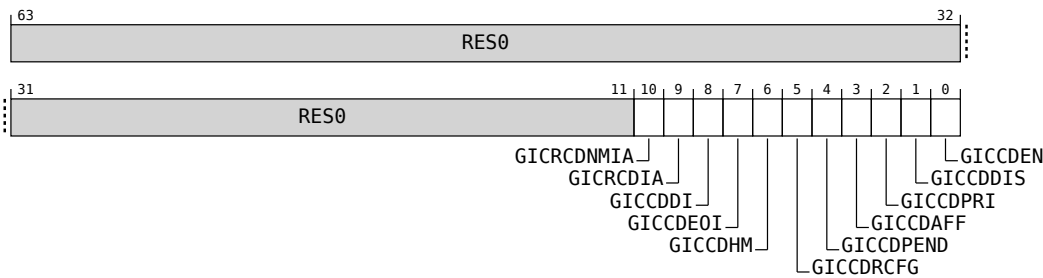
This register is present only when FEAT_GCIE is implemented and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_HFGITR_EL2 are UNDEFINED.

Attributes

ICH_HFGITR_EL2 is a 64-bit register.

Field descriptions

The ICH_HFGITR_EL2 bit assignments are:



Bits [63:11]

Reserved, RES0.

GICRCDNMIA, bit [10]

Trap execution of GICR CDN Mia at EL1 to EL2.

GICRCDNMIA	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then execution of GICR CDN Mia at EL1 is trapped to EL2 and reported with EC syndrome value 0x18, unless the instruction generates a higher priority exception.
0b1	Execution of GICR CDN Mia is not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

GICRCDIA, bit [9]

Trap execution of GICR CDIA at EL1 to EL2.

GICRCDIA	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then execution of GICR CDIA at EL1 is trapped to EL2 and reported with EC syndrome value 0x18, unless the instruction generates a higher priority exception.
0b1	Execution of GICR CDIA is not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

GICCDDI, bit [8]

Trap execution of GIC CDDI at EL1 to EL2.

GICCDDI	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then execution of GIC CDDI at EL1 is trapped to EL2 and reported with EC syndrome value 0x18, unless the instruction generates a higher priority exception.
0b1	Execution of GIC CDDI is not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

GICCDEOI, bit [7]

Trap execution of GIC CDEOI at EL1 to EL2.

GICCDEOI	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then execution of GIC CDEOI at EL1 is trapped to EL2 and reported with EC syndrome value 0x18, unless the instruction generates a higher priority exception.
0b1	Execution of GIC CDEOI is not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

GICCDHM, bit [6]

Trap execution of GIC CDHM at EL1 to EL2.

GICCDHM	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then execution of GIC CDHM at EL1 is trapped to EL2 and reported with EC syndrome value 0x18, unless the instruction generates a higher priority exception.
0b1	Execution of GIC CDHM is not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

GICCDRCFG, bit [5]

Trap execution of GIC CDRCFG at EL1 to EL2.

GICCDRCFG	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then execution of GIC CDRCFG at EL1 is trapped to EL2 and reported with EC syndrome value 0x18, unless the instruction generates a higher priority exception.
0b1	Execution of GIC CDRCFG is not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

GICCDPEND, bit [4]

Trap execution of GIC CDPEND at EL1 to EL2.

GICCDPEND	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then execution of GIC CDPEND at EL1 is trapped to EL2 and reported with EC syndrome value 0x18, unless the instruction generates a higher priority exception.
0b1	Execution of GIC CDPEND is not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

GICCDAFF, bit [3]

Trap execution of GIC CDAFF at EL1 to EL2.

GICDAFF	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then execution of GIC CDAFF at EL1 is trapped to EL2 and reported with EC syndrome value 0x18, unless the instruction generates a higher priority exception.
0b1	Execution of GIC CDAFF is not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

GICDPRI, bit [2]

Trap execution of GIC CDPRI at EL1 to EL2.

GICDPRI	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then execution of GIC CDPRI at EL1 is trapped to EL2 and reported with EC syndrome value 0x18, unless the instruction generates a higher priority exception.
0b1	Execution of GIC CDPRI is not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

GICCDDIS, bit [1]

Trap execution of GIC CDDIS at EL1 to EL2.

GICCDDIS	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then execution of GIC CDDIS at EL1 is trapped to EL2 and reported with EC syndrome value 0x18, unless the instruction generates a higher priority exception.
0b1	Execution of GIC CDDIS is not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

GICCDEN, bit [0]

Trap execution of GIC CDEN at EL1 to EL2.

GICCEN	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then execution of GIC CDEN at EL1 is trapped to EL2 and reported with EC syndrome value 0x18, unless the instruction generates a higher priority exception.
0b1	Execution of GIC CDEN is not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing ICH_HFGITR_EL2

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_HFGITR_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1001	0b111

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        X[t, 64] = NVMem[0xB10];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICH_HFGITR_EL2;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICH_HFGITR_EL2;

```

MSR ICH_HFGITR_EL2, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1001	0b111

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0xB10] = X[t, 64];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    ICH_HFGITR_EL2 = X[t, 64];
elsif PSTATE.EL == EL3 then
    ICH_HFGITR_EL2 = X[t, 64];

```

9.6.4 ICH_HFGRTR_EL2, Hypervisor GIC Fine-Grained Read Trap Register

The ICH_HFGRTR_EL2 characteristics are:

Purpose

Provides controls for traps of MRS reads of GIC System registers.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.


This register is present only when FEAT_GCIE is implemented and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_HFGRTR_EL2 are UNDEFINED.

Attributes

ICH_HFGRTR_EL2 is a 64-bit register.

Field descriptions

The ICH_HFGRTR_EL2 bit assignments are:



Bits [63:21]

Reserved, RES0.

ICC_PPI_ACTIVEn_EL1, bit [20]

Trap MRS reads of ICC_PPI_CACTIVER<n>_EL1, ICC_PPI_SACTIVER<n>_EL1, ICV_PPI_CACTIVER<n>_EL1 and ICV_PPI_SACTIVER<n>_EL1 at EL1 using AArch64 to EL2.

ICC_PPI_ACTIVEn_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MRS reads of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MRS reads of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_PPI_PRIORITYRn_EL1, bit [19]

Trap MRS reads of ICC_PPI_PRIORITYR<n>_EL1 and ICV_PPI_PRIORITYR<n>_EL1 at EL1 using AArch64 to EL2.

ICC_PPI_PRIORITYRn_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MRS reads of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MRS reads of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_PPI_PENDRn_EL1, bit [18]

Trap MRS reads of ICC_PPI_CPENDR<n>_EL1, ICC_PPI_SPENDR<n>_EL1 ICV_PPI_CPENDR<n>_EL1 and ICV_PPI_SPENDR<n>_EL1 at EL1 using AArch64 to EL2.

ICC_PPI_PENDRn_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MRS reads of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MRS reads of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_PPI_ENABLERn_EL1, bit [17]

Trap MRS reads of ICC_PPI_ENABLER<n>_EL1 and ICV_PPI_ENABLER<n>_EL1 at EL1 using AArch64 to EL2.

ICC_PPI_ENABLERn_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MRS reads of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MRS reads of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_PPI_HMRn_EL1, bit [16]

Trap MRS reads of ICC_PPI_HMR<n>_EL1 and ICV_PPI_HMR<n>_EL1 at EL1 using AArch64 to EL2.

ICC_PPI_HMRn_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MRS reads of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MRS reads of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Bits [15:8]

Reserved, RES0.

ICC_IAFFIDR_EL1, bit [7]

Trap MRS reads of ICC_IAFFIDR_EL1 at EL1 using AArch64 to EL2.

ICC_IAFFIDR_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MRS reads of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MRS reads of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_ICSR_EL1, bit [6]

Trap MRS reads of ICC_ICSR_EL1 at EL1 using AArch64 to EL2.

ICC_ICSR_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MRS reads of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MRS reads of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_PCR_EL1, bit [5]

Trap MRS reads of ICC_PCR_EL1 and ICV_PCR_EL1 at EL1 using AArch64 to EL2.

ICC_PCR_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MRS reads of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MRS reads of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_HPPIR_EL1, bit [4]

Trap MRS reads of ICC_HPPIR_EL1 and ICV_HPPIR_EL1 at EL1 using AArch64 to EL2.

ICC_HPPIR_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MRS reads of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MRS reads of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_HAPR_EL1, bit [3]

Trap MRS reads of ICC_HAPR_EL1 and ICV_HAPR_EL1 at EL1 using AArch64 to EL2.

ICC_HAPR_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MRS reads of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MRS reads of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_CR0_EL1, bit [2]

Trap MRS reads of ICC_CR0_EL1 and ICV_CR0_EL1 at EL1 using AArch64 to EL2.

ICC_CR0_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MRS reads of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MRS reads of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_IDRn_EL1, bit [1]

Trap MRS reads of ICC_IDR0_EL1 at EL1 using AArch64 to EL2.

ICC_IDRn_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MRS reads of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MRS reads of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_APR_EL1, bit [0]

Trap MRS reads of ICC_APR_EL1 and ICV_APR_EL1 at EL1 using AArch64 to EL2.

ICC_APR_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MRS reads of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MRS reads of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing ICH_HFGRTR_EL2

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_HFGRTR_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1001	0b100

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        X[t, 64] = NVMem[0xB18];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICH_HFGRTR_EL2;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICH_HFGRTR_EL2;

```

MSR ICH_HFGRTR_EL2, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1001	0b100

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0xB18] = X[t, 64];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    ICH_HFGRTR_EL2 = X[t, 64];
elsif PSTATE.EL == EL3 then
    ICH_HFGRTR_EL2 = X[t, 64];

```


9.6.5 ICH_HFGWTR_EL2, Hypervisor GIC Fine-Grained Write Trap Register

The ICH_HFGWTR_EL2 characteristics are:

Purpose

Provides controls for traps of MSR writes of GIC System registers.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.


This register is present only when FEAT_GCIE is implemented and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_HFGWTR_EL2 are UNDEFINED.

Attributes

ICH_HFGWTR_EL2 is a 64-bit register.

Field descriptions

The ICH_HFGWTR_EL2 bit assignments are:



Bits [63:21]

Reserved, RES0.

ICC_PPI_ACTIVERn_EL1, bit [20]

Trap MSR writes of ICC_PPI_CACTIVER<n>_EL1, ICC_PPI_SACTIVER<n>_EL1, ICV_PPI_CACTIVER<n>_EL1 and ICV_PPI_SACTIVER<n>_EL1 at EL1 using AArch64 to EL2.

ICC_PPI_ACTIVERn_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MSR writes of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MSR writes of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_PPI_PRIORITYRn_EL1, bit [19]

Trap MSR writes of ICC_PPI_PRIORITYR<n>_EL1 and ICV_PPI_PRIORITYR<n>_EL1 at EL1 using AArch64 to EL2.

ICC_PPI_PRIORITYRn_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MSR writes of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MSR writes of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_PPI_PENDRn_EL1, bit [18]

Trap MSR writes of ICC_PPI_CPENDR<n>_EL1, ICC_PPI_SPENDR<n>_EL1, ICC_PPI_CPENDR<n>_EL1 and ICC_PPI_SPENDR<n>_EL1 at EL1 using AArch64 to EL2.

ICC_PPI_PENDRn_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MSR writes of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MSR writes of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_PPI_ENABLERn_EL1, bit [17]

Trap MSR writes of ICC_PPI_ENABLER<n>_EL1 and ICC_PPI_ENABLER<n>_EL1 at EL1 using AArch64 to EL2.

ICC_PPI_ENABLERn_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MSR writes of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MSR writes of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Bits [16:7]

Reserved, RES0.

ICC_ICSR_EL1, bit [6]

Trap MSR writes of ICC_ICSR_EL1 at EL1 using AArch64 to EL2.

ICC_ICSR_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MSR writes of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MSR writes of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

ICC_PCR_EL1, bit [5]

Trap MSR writes of ICC_PCR_EL1 and ICV_PCR_EL1 at EL1 using AArch64 to EL2.

ICC_PCR_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MSR writes of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MSR writes of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Bits [4:3]

Reserved, RES0.

ICC_CR0_EL1, bit [2]

Trap MSR writes of ICC_CR0_EL1 and ICV_CR0_EL1 at EL1 using AArch64 to EL2.

ICC_CR0_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MSR writes of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.

ICC_CR0_EL1	Meaning
0b1	MSR writes of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Bit [1]

Reserved, RES0.

ICC_APR_EL1, bit [0]

Trap MSR writes of ICC_APR_EL1 and ICV_APR_EL1 at EL1 using AArch64 to EL2.

ICC_APR_EL1	Meaning
0b0	If EL2 is implemented and enabled in the current Security state, then MSR writes of the specified registers at EL1 using AArch64 are trapped to EL2 and reported with EC syndrome value 0x18, unless the read generates a higher priority exception.
0b1	MSR writes of the specified registers are not trapped by this mechanism.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing ICH_HFGWTR_EL2

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_HFGWTR_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1001	0b110

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        X[t, 64] = NVMem[0xB20];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICH_HFGWTR_EL2;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICH_HFGWTR_EL2;

```

MSR ICH_HFGWTR_EL2, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1001	0b110

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0xB20] = X[t, 64];
    elseif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elseif PSTATE.EL == EL2 then
    ICH_HFGWTR_EL2 = X[t, 64];
elseif PSTATE.EL == EL3 then
    ICH_HFGWTR_EL2 = X[t, 64];

```

9.6.6 ICH_HPPIR_EL2, Interrupt Controller Hypervisor Highest Priority Pending Interrupt Register

The ICH_HPPIR_EL2 characteristics are:

Purpose

Reports the HPPI for the Virtual Interrupt Domain.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

AArch64 system register ICH_HPPIR_EL2 bits [63:0] are architecturally mapped to AArch64 system register ICV_HPPIR_EL1[63:0].


This register is present only when FEAT_GCIE is implemented and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_HPPIR_EL2 are UNDEFINED.

Attributes

ICH_HPPIR_EL2 is a 64-bit register.

Field descriptions

The ICH_HPPIR_EL2 bit assignments are:



Bits [63:33]

Reserved, RES0.

HPPIV, bit [32]

HPPI valid.

There is an HPPI with Sufficient priority for the Interrupt Domain.

HPPIV	Meaning
0b0	Invalid: There is no HPPI with Sufficient priority for the Interrupt Domain.
0b1	VALID: There is an HPPI with Sufficient priority for the Interrupt Domain.

If ICH_HPPIR_EL2.HPPIV is 1, ID and TYPE together form the INTID of the HPPI for the Interrupt Domain.

TYPE, bits [31:29]

The encoding of this field depends on the value of HPPIV as described below:

- If ICH_HPPIR_EL2.HPPIV is 0, TYPE is RES0.
- If ICH_HPPIR_EL2.HPPIV is 1, TYPE specifies the Type of the interrupt.

TYPE	Meaning
0b001	PPI
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The encoding of this field depends on the value of HPPIV as described below:

- If ICH_HPPIR_EL2.HPPIV is 0, ID is RES0.
- If ICH_HPPIR_EL2.HPPIV is 1, ID specifies the interrupt ID.

Accessing ICH_HPPIR_EL2

Read-only

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_HPPIR_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1000	0b101

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICH_HPPIR_EL2;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICH_HPPIR_EL2;

```

9.6.7 ICH_PPI_ACTIVER<n>_EL2, Interrupt Controller Virtual Interrupt Active Registers, n = 0 - 1

The ICH_PPI_ACTIVER<n>_EL2 characteristics are:

Purpose

Access to Active state for virtual PPIs.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

AArch64 system register ICH_PPI_ACTIVER<n>_EL2 bits [63:0] are architecturally mapped to AArch64 system register ICV_PPI_CACTIVER<n>_EL1[63:0].

AArch64 system register ICH_PPI_ACTIVER<n>_EL2 bits [63:0] are architecturally mapped to AArch64 system register ICV_PPI_SACTIVER<n>_EL1[63:0].


This register is present only when FEAT_GCIE is implemented and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_PPI_ACTIVER<n>_EL2 are UNDEFINED.

Attributes

ICH_PPI_ACTIVER<n>_EL2 is a 64-bit register.

Field descriptions

The ICH_PPI_ACTIVER<n>_EL2 bit assignments are:



ACTIVE<x>, bits [x], for x = 63 to 0

Accesses the Active state of the virtual PPI.

ACTIVE<x>	Meaning
0b0	Inactive
0b1	Active

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 64) + x)
- Otherwise, access to this field is **RW**

Accessing ICH_PPI_ACTIVER<n>_EL2

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_PPI_ACTIVER<n>_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1010	0b11:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        X[t, 64] = NVMem[0xB30 + (8 * n)];
    elseif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elseif PSTATE.EL == EL2 then
    X[t, 64] = ICH_PPI_ACTIVER_EL2[n];
elseif PSTATE.EL == EL3 then
    X[t, 64] = ICH_PPI_ACTIVER_EL2[n];

```

MSR ICH_PPI_ACTIVER<n>_EL2, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1010	0b11:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0xB30 + (8 * n)] = X[t, 64];
    elseif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elseif PSTATE.EL == EL2 then

```

```
    ICH_PPI_ACTIVER_EL2[n] = X[t, 64];  
elseif PSTATE.EL == EL3 then  
    ICH_PPI_ACTIVER_EL2[n] = X[t, 64];
```

9.6.8 ICH_PPI_DVIR<n>_EL2, Interrupt Controller PPI Direct-inject Virtual Interrupt Registers, n = 0 - 1

The ICH_PPI_DVIR<n>_EL2 characteristics are:

Purpose

Controls whether Pending physical PPIs are directly injected as virtual PPIs to the Virtual Interrupt Domain.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.


This register is present only when FEAT_GCIE is implemented and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_PPI_DVIR<n>_EL2 are UNDEFINED.

Attributes

ICH_PPI_DVIR<n>_EL2 is a 64-bit register.

Field descriptions

The ICH_PPI_DVIR<n>_EL2 bit assignments are:



DVI<x>, bits [x], for x = 63 to 0

Controls whether the physical PPI Pending state is directly injected to a virtual PPI in the Virtual Interrupt Domain.

DVI<x>	Meaning
0b0	Physical PPI <x> is not directly injected as a virtual PPI.
0b1	Physical PPI <x> is directly injected as a virtual PPI.

The Effective value of DVI<x> is 0 if any of the following are true:

- Physical PPI <x> is not assigned to Current Physical Interrupt Domain.
- Physical PPI <x> is assigned to the EL3 Interrupt Domain.
- Physical PPI <x> is assigned to the Secure Interrupt Domain and SCR_EL3.EEL2 is 0.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 64) + x)
- Access is **RAZ/WI** if all of the following are true:
 - !IsPPIAssignedToCurrentDomain((n * 64) + x)
 - PSTATE.EL IN {EL2, EL1}
- Otherwise, access to this field is **RW**

Accessing ICH_PPI_DVIR<n>_EL2

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_PPI_DVIR<n>_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1010	0b00:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        X[t, 64] = NVMem[0xB40 + (8 * n)];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICH_PPI_DVIR_EL2[n];
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICH_PPI_DVIR_EL2[n];

```

MSR ICH_PPI_DVIR<n>_EL2, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1010	0b00:n[0]

```
if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0xB40 + (8 * n)] = X[t, 64];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    ICH_PPI_DVIR_EL2[n] = X[t, 64];
elsif PSTATE.EL == EL3 then
    ICH_PPI_DVIR_EL2[n] = X[t, 64];
```

9.6.9 ICH_PPI_ENABLER<n>_EL2, Interrupt Controller Virtual Interrupt Enable Registers, n = 0 - 1

The ICH_PPI_ENABLER<n>_EL2 characteristics are:

Purpose

Access to Enable state for virtual PPIs.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

AArch64 system register ICH_PPI_ENABLER<n>_EL2 bits [63:0] are architecturally mapped to AArch64 system register ICV_PPI_ENABLER<n>_EL1[63:0].


This register is present only when FEAT_GCIE is implemented and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_PPI_ENABLER<n>_EL2 are UNDEFINED.

Attributes

ICH_PPI_ENABLER<n>_EL2 is a 64-bit register.

Field descriptions

The ICH_PPI_ENABLER<n>_EL2 bit assignments are:



EN<x>, bits [x], for x = 63 to 0

Alias of equivalent ICV_PPI_ENABLER<n>_EL1 field.

Accessing this field has the following behavior:

- RES0 if !IsPPIImplemented((n * 64) + x)

- Otherwise, access to this field is **RW**

Accessing ICH_PPI_ENABLER<n>_EL2

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_PPI_ENABLER<n>_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1010	0b01:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        X[t, 64] = NVMem[0xB50 + (8 * n)];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICH_PPI_ENABLER_EL2[n];
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICH_PPI_ENABLER_EL2[n];

```

MSR ICH_PPI_ENABLER<n>_EL2, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1010	0b01:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0xB50 + (8 * n)] = X[t, 64];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    ICH_PPI_ENABLER_EL2[n] = X[t, 64];
elsif PSTATE.EL == EL3 then
    ICH_PPI_ENABLER_EL2[n] = X[t, 64];

```

9.6.10 ICH_PPI_PENDR<n>_EL2, Interrupt Controller Virtual Interrupt Pending State Registers, n = 0 - 1

The ICH_PPI_PENDR<n>_EL2 characteristics are:

Purpose

Pending state for virtual PPIs.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

AArch64 system register ICH_PPI_PENDR<n>_EL2 bits [63:0] are architecturally mapped to AArch64 system register ICV_PPI_CPENDR<n>_EL1[63:0].

AArch64 system register ICH_PPI_PENDR<n>_EL2 bits [63:0] are architecturally mapped to AArch64 system register ICV_PPI_SPENDR<n>_EL1[63:0].


This register is present only when FEAT_GCIE is implemented and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_PPI_PENDR<n>_EL2 are UNDEFINED.

Attributes

ICH_PPI_PENDR<n>_EL2 is a 64-bit register.

Field descriptions

The ICH_PPI_PENDR<n>_EL2 bit assignments are:



PEND<x>, bits [x], for x = 63 to 0

Pend<x> accesses the Pending state of virtual PPI <x>.

PEND<x>	Meaning
0b0	Not pending
0b1	Pending

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 64) + x)
- Otherwise, access to this field is **RW**

Accessing ICH_PPI_PENDR<n>_EL2

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_PPI_PENDR<n>_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1010	0b10:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICH_PPI_PENDR_EL2[n];
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICH_PPI_PENDR_EL2[n];

```

MSR ICH_PPI_PENDR<n>_EL2, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1010	0b10:n[0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    ICH_PPI_PENDR_EL2[n] = X[t, 64];
elsif PSTATE.EL == EL3 then
    ICH_PPI_PENDR_EL2[n] = X[t, 64];

```

9.6.11 ICH_PPI_PRIORITYR<n>_EL2, Interrupt Controller Virtual Interrupt Priority Registers, n = 0 - 15

The ICH_PPI_PRIORITYR<n>_EL2 characteristics are:

Purpose

Priority of virtual PPIs.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

AArch64 system register ICH_PPI_PRIORITYR<n>_EL2 bits [63:0] are architecturally mapped to AArch64 system register ICV_PPI_PRIORITYR<n>_EL1[63:0].

This register is present only when FEAT_GCIE is implemented and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_PPI_PRIORITYR<n>_EL2 are UNDEFINED.

Attributes

ICH_PPI_PRIORITYR<n>_EL2 is a 64-bit register.

Field descriptions

The ICH_PPI_PRIORITYR<n>_EL2 bit assignments are:

63	61	60	56	55	53	52	48	47	45	44	40	39	37	36	32
RES0	PRIORITY7	RES0	PRIORITY6	RES0	PRIORITY5	RES0	PRIORITY4	RES0	PRIORITY3	RES0	PRIORITY2	RES0	PRIORITY1	RES0	PRIORITY0

Bits [63:61, 55:53, 47:45, 39:37, 31:29, 23:21, 15:13, 7:5]

Reserved, RES0.

PRIORITY<x>, bits [60:56, 52:48, 44:40, 36:32, 28:24, 20:16, 12:8, 4:0], for x = 7 to 0, where each field is 5 bits wide

Alias of equivalent ICV_PPI_PRIORITYR<n>_EL1 field.

Fields corresponding to unimplemented INTIDs are RES0.

Accessing this field has the following behavior:

- **RES0** if !IsPPIImplemented((n * 8) + x)
- Otherwise, access to this field is **RW**

Accessing ICH_PPI_PRIORITYR<n>_EL2

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_PPI_PRIORITYR<n>_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b111:n[3]	n[2:0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then

```

```

        X[t, 64] = NVMem[0xB80 + (8 * n)];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
    elsif PSTATE.EL == EL2 then
        X[t, 64] = ICH_PPI_PRIORITYR_EL2[n];
    elsif PSTATE.EL == EL3 then
        X[t, 64] = ICH_PPI_PRIORITYR_EL2[n];

```

MSR ICH_PPI_PRIORITYR<n>_EL2, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b111:n[3]	n[2:0]

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0xB80 + (8 * n)] = X[t, 64];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
    elsif PSTATE.EL == EL2 then
        ICH_PPI_PRIORITYR_EL2[n] = X[t, 64];
    elsif PSTATE.EL == EL3 then
        ICH_PPI_PRIORITYR_EL2[n] = X[t, 64];

```

9.6.12 ICH_VCTLR_EL2, Interrupt Controller Virtual CPU interface Control Register

The ICH_VCTLR_EL2 characteristics are:

Purpose

Controls behavior of the Virtual CPU interface.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.


This register is present only when FEAT_GCIE is implemented and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_VCTLR_EL2 are UNDEFINED.

Attributes

ICH_VCTLR_EL2 is a 64-bit register.

Field descriptions

The ICH_VCTLR_EL2 bit assignments are:



Bits [63:2]

Reserved, RES0.

V3, bit [1]

When FEAT_GCIE_LEGACY is implemented:

Enable bit for Legacy operation.

If EL2 is not enabled in the Security state identified by SCR_EL3.{NSE,NS}, the Effective value of this field is 0.

When the Effective value of this field is 0, the Effective value of ICH_HCR_EL2.En is 0.

See ‘Legacy virtual CPU interface’ for more information.

V3	Meaning
0b0	Legacy operation disabled.
0b1	Legacy operation enabled.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally UNKNOWN value.

Otherwise:

RES0

EN, bit [0]

When *ICH_VCTLR_EL2.V3* == 0:

Enable interrupts for the Virtual Interrupt Domain.

If EL2 is not enabled in the Security state identified by *SCR_EL3.{NSE,NS}*, the Effective value of this field is 0.

EN	Meaning
0b0	Disabled.
0b1	Enabled.

When this field is 0, there is no HPPI of Sufficient priority for the Virtual Interrupt Domain.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally UNKNOWN value.

Otherwise:

RES0

Accessing ICH_VCTLR_EL2

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_VCTLR_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1011	0b100

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        X[t, 64] = NVMem[0xB28];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    X[t, 64] = ICH_VCTLR_EL2;
elsif PSTATE.EL == EL3 then
    X[t, 64] = ICH_VCTLR_EL2;

```

MSR ICH_VCTLR_EL2, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1011	0b100

```
if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0xB28] = X[t, 64];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    ICH_VCTLR_EL2 = X[t, 64];
elsif PSTATE.EL == EL3 then
    ICH_VCTLR_EL2 = X[t, 64];
```

9.6.13 ICH_VMCR_EL2, Interrupt Controller Virtual Machine Control Register

The ICH_VMCR_EL2 characteristics are:

Purpose

Enables the hypervisor to save and restore the state of the Virtual CPU interface.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

This register is present only when (FEAT_GICv3 is implemented or FEAT_GCIE is implemented) and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_VMCR_EL2 are UNDEFINED.


Attributes

ICH_VMCR_EL2 is a 64-bit register.

Field descriptions

The ICH_VMCR_EL2 bit assignments are:

When ICH_VCTLR_EL2.V3 == 0:



Bits [63:32]

Reserved, RES0.

VPMR, bits [31:27]

Alias of ICV_PCR_EL1.Priority.


Bits [26:1]

Reserved, RES0.

EN, bit [0]

Alias of ICV_CR0_EL1.EN.

When ICH_VCTLR_EL2.V3 == 1:



Fields listed below are as described in the ICH_VMCR_EL2 register in [3].

Bits [63:32]

Reserved, RES0.

VPMR, bits [31:24]

Alias of ICV_PMR_EL1.Priority.

VBPR0, bits [23:21]

Alias of ICV_BPR0_EL1.BinaryPoint.

VBPR1, bits [20:18]

Alias of ICV_BPR1_EL1.BinaryPoint.

Bits [17:10]

Reserved, RES0.

VEOIM, bit [9]

Alias of ICV_CTLR_EL1.EOI mode.

Bits [8:5]

Reserved, RES0.

VCBPR, bit [4]

Alias of ICV_CTLR_EL1.CBPR.

VFIQEn, bit [3]

Virtual FIQ enable.

This field is RES1 as GICv2 compatibility is not supported.

VAckCtl, bit [2]

Virtual AckCtl.

VENG1, bit [1]

Alias ICV_IGRPEN1_EL1.Enable.

VENG0, bit [0]

Alias of ICV_IGRPEN0_EL1.Enable.

Accessing ICH_VMCR_EL2

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_VMCR_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1011	0b111

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        X[t, 64] = NVMem[0x4C8];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL2.SRE == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        X[t, 64] = ICH_VMCR_EL2;
elsif PSTATE.EL == EL3 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL3.SRE == '0' then
        AArch64.SystemAccessTrap(EL3, 0x18);
    else
        X[t, 64] = ICH_VMCR_EL2;

```

MSR ICH_VMCR_EL2, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1011	0b111

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0x4C8] = X[t, 64];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL2.SRE == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        ICH_VMCR_EL2 = X[t, 64];
elsif PSTATE.EL == EL3 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL3.SRE == '0' then
        AArch64.SystemAccessTrap(EL3, 0x18);
    else
        ICH_VMCR_EL2 = X[t, 64];

```

9.6.14 Nested virtualization

I_{XDPKJ}

The Arm architecture’s enhanced support for nested virtualization (FEAT_NV2) provides a mechanism for hardware to transform reads and writes from System registers into reads and writes from memory.

See “Enhanced support for nested virtualization” in the Architecture Reference Manual for A-profile [1].

R_{KZTGX}

When system register accesses are transformed to memory accesses Table D8-63 (Memory address offset associated with transformed register access) in the Architecture Reference Manual for A-profile [1] describes the offsets used for each affected register. This table is extended as follows to cover the GICv5 system registers:

Register access if HCR_EL2.NV1 is 0	Register access if HCR_EL2.NV1 is 1	Memory offset
ICH_APR_EL2	ICH_APR_EL2	0xB00
ICH_CONTEXTR_EL2	ICH_CONTEXTR_EL2	0xB08
ICH_HFGITR_EL2	ICH_HFGITR_EL2	0xB10
ICH_HFGRTR_EL2	ICH_HFGRTR_EL2	0xB18
ICH_HFGWTR_EL2	ICH_HFGWTR_EL2	0xB20
ICH_PPI_ACTIVER<n>_EL2	ICH_PPI_ACTIVER<n>_EL2	0xB30 + 8*n
ICH_PPI_DVIR<n>_EL2	ICH_PPI_DVIR<n>_EL2	0xB40 + 8*n
ICH_PPI_ENABLER<n>_EL2	ICH_PPI_ENABLER<n>_EL2	0xB50 + 8*n
ICH_PPI_PRIORITYR<n>_EL2	ICH_PPI_PRIORITYR<n>_EL2	0xB80 + 8*n
ICH_VCTLR_EL2	ICH_VCTLR_EL2	0xB28
ICH_VMCR_EL2	ICH_VMCR_EL2	0x4C8

Offsets for registers introduced by GICv3 or GICv4 are unchanged.

9.7 Legacy hypervisor control registers

Registers to manage operation of the Legacy virtual CPU interface.

9.7.1 ICH_AP0R<n>_EL2, Interrupt Controller Active Virtual Priorities Registers 0, n = 0 - 3

The ICH_AP0R<n>_EL2 characteristics are:

Purpose

Records active Group 0 virtual priorities in the Legacy virtual CPU interface.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

This register has no effect if EL2 is not enabled in the current Security state.

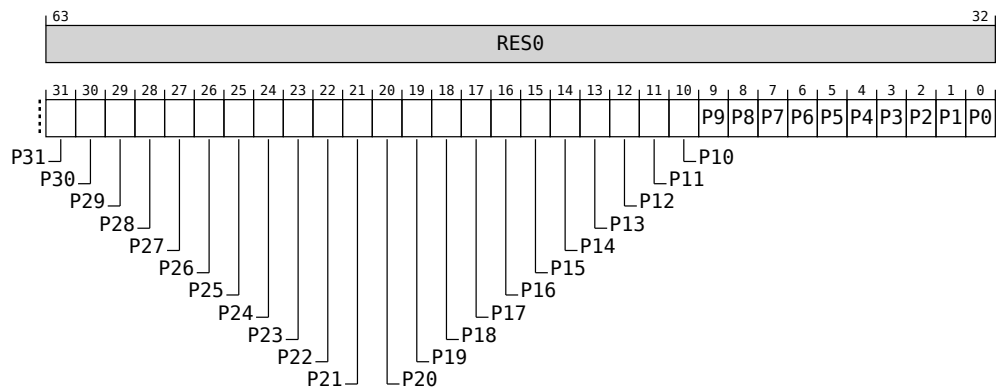
This register is present only when (FEAT_GICv3 is implemented or FEAT_GCIE_LEGACY is implemented) and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_AP0R<n>_EL2 are UNDEFINED.

Attributes

ICH_AP0R<n>_EL2 is a 64-bit register.

Field descriptions

The ICH_AP0R<n>_EL2 bit assignments are:



All fields listed below are as described in the ICH_AP0R<n>_EL2 register description in [3].

Bits [63:32]

Reserved, RES0.

P<x>, bits [x], for x = 31 to 0

Provides the access to the active virtual Group 0 priorities.

Accessing ICH_AP0R<n>_EL2

When FEAT_GCIE_LEGACY is implemented, only 32 priority levels is supported meaning that ICH_AP0R1_EL2, ICH_AP0R2_EL2, and ICH_AP0R3_EL2 are not implemented.

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_AP0R<m>_EL2 ; Where m = 0-3

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1000	0b0:m[1:0]

```

integer m = UInt(op2<1:0>);

if m == 1 && NUM_GIC_PREEMPTION_BITS < 6 then
    UNDEFINED;
elsif (m == 2 || m == 3) && NUM_GIC_PREEMPTION_BITS < 7 then
    UNDEFINED;
elsif PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        X[t, 64] = NVMem[0x480 + (8 * m)];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL2.SRE == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        X[t, 64] = ICH_AP0R_EL2[m];
elsif PSTATE.EL == EL3 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL3.SRE == '0' then
        AArch64.SystemAccessTrap(EL3, 0x18);
    else
        X[t, 64] = ICH_AP0R_EL2[m];

```

MSR ICH_AP0R<m>_EL2, <Xt> ; Where m = 0-3

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1000	0b0:m[1:0]

```

integer m = UInt(op2<1:0>);

if m == 1 && NUM_GIC_PREEMPTION_BITS < 6 then
    UNDEFINED;
elsif (m == 2 || m == 3) && NUM_GIC_PREEMPTION_BITS < 7 then
    UNDEFINED;
elsif PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0x480 + (8 * m)] = X[t, 64];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL2.SRE == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        ICH_AP0R_EL2[m] = X[t, 64];
elsif PSTATE.EL == EL3 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL3.SRE == '0' then
        AArch64.SystemAccessTrap(EL3, 0x18);
    else
        ICH_AP0R_EL2[m] = X[t, 64];

```

9.7.2 ICH_AP1R<n>_EL2, Interrupt Controller Active Virtual Priorities Registers 1, n = 0 - 3

The ICH_AP1R<n>_EL2 characteristics are:

Purpose

Records active Group 1 virtual priorities in the Legacy virtual CPU interface.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

This register has no effect if EL2 is not enabled in the current Security state.

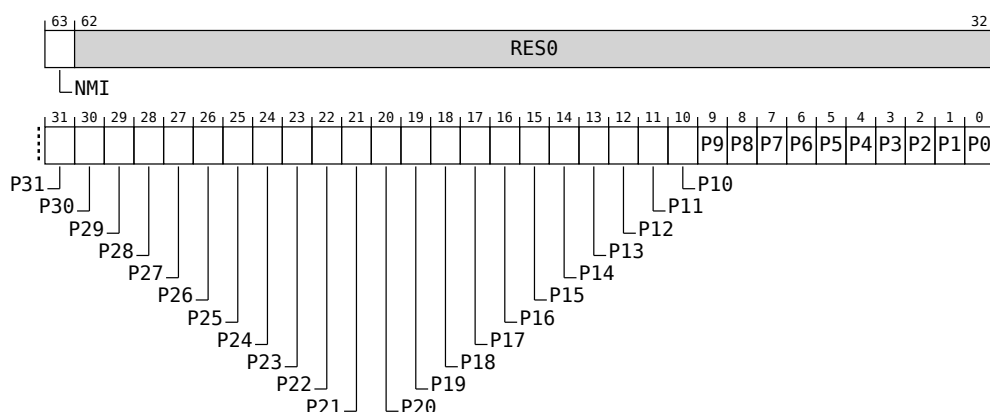
This register is present only when (FEAT_GICv3 is implemented or FEAT_GCIE_LEGACY is implemented) and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_AP1R<n>_EL2 are UNDEFINED.

Attributes

ICH_AP1R<n>_EL2 is a 64-bit register.

Field descriptions

The ICH_AP1R<n>_EL2 bit assignments are:



All fields listed below are as described in the ICH_AP1R<n>_EL2 register description in [3].

NMI, bit [63]

When (FEAT_GICv3_NMI is implemented or FEAT_GCIE_LEGACY is implemented) and n == 0:

Indicates whether there is an active virtual NMI priority.

NMI	Meaning
0b0	There is no active Group 1 NMI, or all active Group 1 NMIs have undergone priority drop.
0b1	There is an active Group 1 NMI.

The reset behavior of this field is:

- On a Warm reset, this field resets to 0b0.

Otherwise:

RES0

Bits [62:32]

Reserved, RES0.

P<x>, bits [x], for x = 31 to 0

Provides the access to the active virtual Group 1 priorities.

If FEAT_GCIE_LEGACY is implemented, this field accesses the same state as ICH_APR_EL2

The number of implemented priority bits is reported by ICC_IDR0_EL1.PRI_BITS. Fields corresponding to unimplemented priority levels are RES0.

Accessing ICH_AP1R<n>_EL2

When FEAT_GCIE_LEGACY is implemented, only 32 priority levels is supported meaning that ICH_AP1R1_EL2, ICH_AP1R2_EL2, and ICH_AP1R3_EL2 are not implemented.

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_AP1R<m>_EL2 ; Where m = 0-3

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1001	0b0:m[1:0]

```
integer m = UInt(op2<1:0>);
```

```
if m == 1 && NUM_GIC_PREEMPTION_BITS < 6 then
    UNDEFINED;
elseif (m == 2 || m == 3) && NUM_GIC_PREEMPTION_BITS < 7 then
    UNDEFINED;
elseif PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        X[t, 64] = NVMem[0x480 + (8 * m)];
    elseif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elseif PSTATE.EL == EL2 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL2.SRE == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        X[t, 64] = ICH_AP1R_EL2[m];
elseif PSTATE.EL == EL3 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL3.SRE == '0' then
        AArch64.SystemAccessTrap(EL3, 0x18);
    else
        X[t, 64] = ICH_AP1R_EL2[m];
```

MSR ICH_AP1R<m>_EL2, <Xt> ; Where m = 0-3

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1001	0b0:m[1:0]

```

integer m = UInt(op2<1:0>);

if m == 1 && NUM_GIC_PREEMPTION_BITS < 6 then
    UNDEFINED;
elseif (m == 2 || m == 3) && NUM_GIC_PREEMPTION_BITS < 7 then
    UNDEFINED;
elseif PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0x480 + (8 * m)] = X[t, 64];
    elseif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elseif PSTATE.EL == EL2 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL2.SRE == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        ICH_AP1R_EL2[m] = X[t, 64];
elseif PSTATE.EL == EL3 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL3.SRE == '0' then
        AArch64.SystemAccessTrap(EL3, 0x18);
    else
        ICH_AP1R_EL2[m] = X[t, 64];

```

9.7.3 ICH_EISR_EL2, Interrupt Controller End of Interrupt Status Register

The ICH_EISR_EL2 characteristics are:

Purpose

Enables a hypervisor to determine which List registers have outstanding EOI maintenance interrupts.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

This register has no effect if EL2 is not enabled in the current Security state.

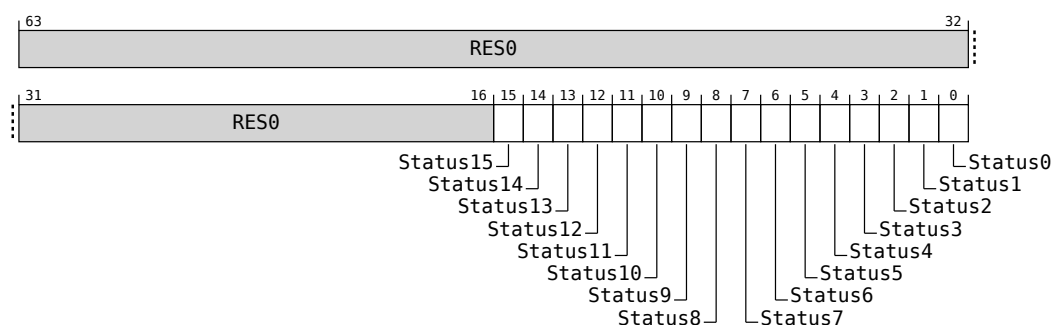
This register is present only when (FEAT_GICv3 is implemented or FEAT_GCIE_LEGACY is implemented) and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_EISR_EL2 are UNDEFINED.

Attributes

ICH_EISR_EL2 is a 64-bit register.

Field descriptions

The ICH_EISR_EL2 bit assignments are:



All fields listed below are as described in the ICH_EISR_EL2 register description in [3].

Bits [63:16]

Reserved, RES0.

Status<x>, bits [x], for x = 15 to 0

EOI maintenance interrupt status bit for List register.

Accessing ICH_EISR_EL2

Read-only

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_EISR_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1011	0b011

```
if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elseif PSTATE.EL == EL2 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL2.SRE == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        X[t, 64] = ICH_EISR_EL2;
elseif PSTATE.EL == EL3 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL3.SRE == '0' then
        AArch64.SystemAccessTrap(EL3, 0x18);
    else
        X[t, 64] = ICH_EISR_EL2;
```

9.7.4 ICH_ELRSR_EL2, Interrupt Controller Empty List Register Status Register

The ICH_ELRSR_EL2 characteristics are:

Purpose

Enables a hypervisor to locate a usable List register to deliver an interrupt to a VM.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

This register has no effect if EL2 is not enabled in the current Security state.

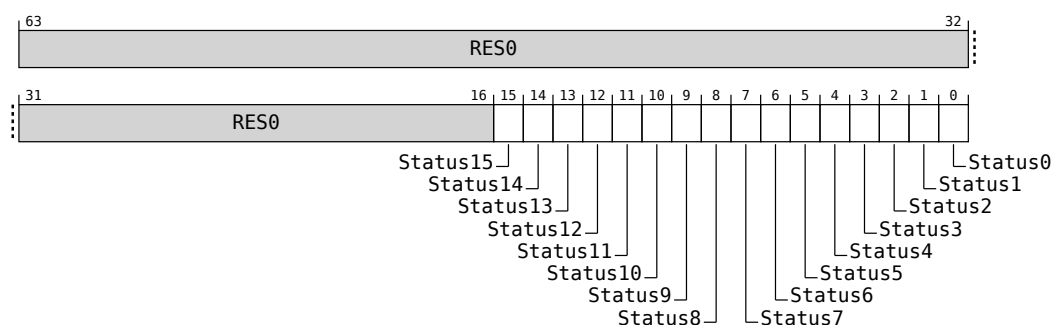
This register is present only when (FEAT_GICv3 is implemented or FEAT_GCIE_LEGACY is implemented) and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_ELRSR_EL2 are UNDEFINED.

Attributes

ICH_ELRSR_EL2 is a 64-bit register.

Field descriptions

The ICH_ELRSR_EL2 bit assignments are:



All fields listed below are as described in the ICH_ELRSR_EL2 register description in [3].

Bits [63:16]

Reserved, RES0.

Status<x>, bits [x], for x = 15 to 0

Status bit for List register.

Accessing ICH_ELRSR_EL2

Read-only

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_ELRSR_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1011	0b101

```
if PSTATE.EL == EL0 then
    UNDEFINED;
elseif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elseif PSTATE.EL == EL2 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL2.SRE == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        X[t, 64] = ICH_ELRSR_EL2;
elseif PSTATE.EL == EL3 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL3.SRE == '0' then
        AArch64.SystemAccessTrap(EL3, 0x18);
    else
        X[t, 64] = ICH_ELRSR_EL2;
```

9.7.5 ICH_HCR_EL2, Interrupt Controller Hyp Control Register

The ICH_HCR_EL2 characteristics are:

Purpose

Enables the hypervisor to control the behavior of the Virtual CPU interface in Legacy operation

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

This register has no effect if EL2 is not enabled in the current Security state.

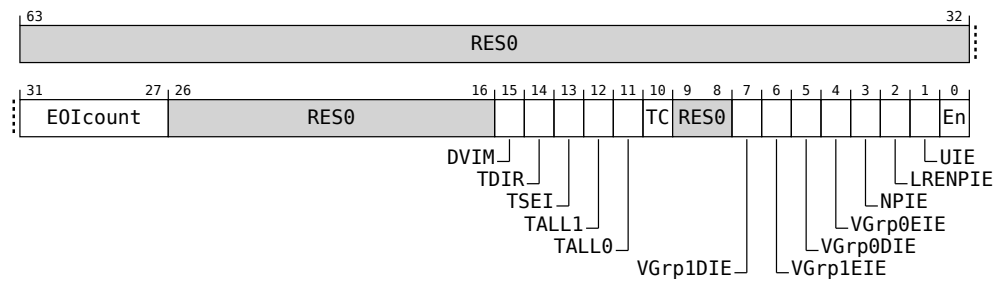
This register is present only when (FEAT_GICv3 is implemented or FEAT_GCIE_LEGACY is implemented) and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_HCR_EL2 are UNDEFINED.

Attributes

ICH_HCR_EL2 is a 64-bit register.

Field descriptions

The ICH_HCR_EL2 bit assignments are:



All fields listed below are as described in the ICH_HCR_EL2 register description in [3].

Bits [63:32]

Reserved, RES0.

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.

Bits [26:16]

Reserved, RES0.

DVIM, bit [15]

When FEAT_GICv3 is implemented:

Directly-injected Virtual Interrupt Mask.

Otherwise:

RES0

TDIR, bit [14]

Trap EL1 writes to ICx_DIR_EL1 to EL2.

When ICH_VCTLR_EL2.V3 == '0', the Effective value of this field is 0.

TSEI, bit [13]

When FEAT_GICv3 is implemented:

Trap all locally generated SEIs.

Otherwise:

RES0

TALL1, bit [12]

Trap all EL1 accesses to ICC_* and ICV_* System registers for Group 1 interrupts to EL2.

When ICH_VCTLR_EL2.V3 == '0', the Effective value of this field is 0.

TALL0, bit [11]

Trap all EL1 accesses to ICC_* and ICV_* System registers for Group 0 interrupts to EL2.

When ICH_VCTLR_EL2.V3 == '0', the Effective value of this field is 0.

TC, bit [10]

Trap all EL1 accesses to System registers that are common to Group 0 and Group 1 to EL2.

When ICH_VCTLR_EL2.V3 == '0', the Effective value of this field is 0.

Bits [9:8]

Reserved, RES0.

VGrp1DIE, bit [7]

VM Group 1 Disabled Interrupt Enable.

When ICH_VCTLR_EL2.V3 == '0', the Effective value of this field is 0.

VGrp1EIE, bit [6]

VM Group 1 Enabled Interrupt Enable.

When ICH_VCTLR_EL2.V3 == '0', the Effective value of this field is 0.

VGrp0DIE, bit [5]

VM Group 0 Disabled Interrupt Enable.

When ICH_VCTLR_EL2.V3 == '0', the Effective value of this field is 0.

VGrp0EIE, bit [4]

VM Group 0 Enabled Interrupt Enable.

When ICH_VCTLR_EL2.V3 == '0', the Effective value of this field is 0.

NPiE, bit [3]

No Pending Interrupt Enable.

When ICH_VCTLR_EL2.V3 == '0', the Effective value of this field is 0.

LRNPiE, bit [2]

List Register Entry Not Present Interrupt Enable.

When ICH_VCTLR_EL2.V3 == '0', the Effective value of this field is 0.

UIE, bit [1]

Underflow Interrupt Enable.

When ICH_VCTLR_EL2.V3 == '0', the Effective value of this field is 0.

En, bit [0]

Global enable bit for the virtual CPU interface.

When ICH_VCTLR_EL2.V3 == '0', the Effective value of this field is 0.

Accessing ICH_HCR_EL2

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_HCR_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1011	0b000

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        X[t, 64] = NVMem[0x4C0];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL2.SRE == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        X[t, 64] = ICH_HCR_EL2;
elsif PSTATE.EL == EL3 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL3.SRE == '0' then
        AArch64.SystemAccessTrap(EL3, 0x18);
    else
        X[t, 64] = ICH_HCR_EL2;

```

MSR ICH_HCR_EL2, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1011	0b000

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0x4C0] = X[t, 64];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL2.SRE == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else

```

```
        ICH_HCR_EL2 = X[t, 64];
    elsif PSTATE.EL == EL3 then
        if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL3.SRE == '0' then
            AArch64.SystemAccessTrap(EL3, 0x18);
        else
            ICH_HCR_EL2 = X[t, 64];
```

9.7.6 ICH_LR<n>_EL2, Interrupt Controller List Registers, n = 0 - 15

The ICH_LR<n>_EL2 characteristics are:

Purpose

Enables a hypervisor to provide interrupt context information to the Legacy virtual CPU interface.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

If list register n is not implemented, then accesses to this register are UNDEFINED.

This register has no effect if EL2 is not enabled in the current Security state.


This register is present only when (FEAT_GICv3 is implemented or FEAT_GCIE_LEGACY is implemented) and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_LR<n>_EL2 are UNDEFINED.

Attributes

ICH_LR<n>_EL2 is a 64-bit register.

Field descriptions

The ICH_LR<n>_EL2 bit assignments are:



All fields listed below are as described in the ICH_LR<n>_EL2 register description in [3].

State, bits [63:62]

See the register description in [3].

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.

HW	Meaning
0b0	The interrupt is triggered entirely by software.
0b1	The interrupt maps directly to a hardware interrupt. Deactivating the virtual interrupt also deactivates the physical interrupt specified in pINTID.

The reset behavior of this field is:

- On a Warm reset, this field resets to an UNKNOWN value.

Group, bit [60]

See the register description in [3].

NMI, bit [59]

See the register description in [3].

Bits [58:56]

Reserved, RES0.

Priority, bits [55:48]

See the register description in [3].

When FEAT_GCIE_LEGACY is implemented, 5 bits of priority are implemented meaning that bits[50:48] are RES0.

Bits [47:45]

Reserved, RES0.

pINTID, bits [44:32]

When FEAT_GICv3 is implemented

pINTID, bits [12:0] of bits [44:32]

See the register description in [3].

When FEAT_GCIE_LEGACY is implemented

pINTID, bits [12:0] of bits [44:32]

Physical PPI INTID, for hardware interrupts.

When ICH_LR<n>_EL2.HW is 0 (there is no corresponding physical interrupt), this field has the following meaning:

- Bits[44:42]: RES0.
- 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]: RES0.

When ICH_LR<n>_EL2.HW is 1 (there is a corresponding physical interrupt):

- This field specifies the ID of a physical PPI in the Current Physical Interrupt Domain.
- This field is only required to implement enough bits to hold a valid value for the implemented physical PPIs. Any unused higher order bits are RES0.
- If pINTID specifies an unreachable PPI, no physical interrupt is deactivated.

Otherwise:

RES0

vINTID, bits [31:0]

See the register description in [3].

Accessing ICH_LR<n>_EL2

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_LR<n>_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b110:n[3]	n[2:0]

```

if n >= NUM_GIC_LIST_REGS then
    UNDEFINED;
elsif PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        X[t, 64] = NVMem[0x400 + (8 * n)];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL2.SRE == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        X[t, 64] = ICH_LR_EL2[n];
elsif PSTATE.EL == EL3 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL3.SRE == '0' then
        AArch64.SystemAccessTrap(EL3, 0x18);
    else
        X[t, 64] = ICH_LR_EL2[n];

```

MSR ICH_LR<n>_EL2, <Xt>

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b110:n[3]	n[2:0]

```

if n >= NUM_GIC_LIST_REGS then
    UNDEFINED;
elsif PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.<NV2,NV> == '11' && ICH_VCTLR_EL2.V3 == '0' then
        NVMem[0x400 + (8 * n)] = X[t, 64];
    elsif EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL2.SRE == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        ICH_LR_EL2[n] = X[t, 64];
elsif PSTATE.EL == EL3 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL3.SRE == '0' then
        AArch64.SystemAccessTrap(EL3, 0x18);
    else
        ICH_LR_EL2[n] = X[t, 64];

```

9.7.7 ICH_MISR_EL2, Interrupt Controller Maintenance Interrupt State Register

The ICH_MISR_EL2 characteristics are:

Purpose

Indicates which maintenance interrupts are asserted.

Configuration

If EL2 is not implemented, this register is RES0 from EL3.

This register has no effect if EL2 is not enabled in the current Security state.


This register is present only when (FEAT_GICv3 is implemented or FEAT_GCIE_LEGACY is implemented) and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_MISR_EL2 are UNDEFINED.

Attributes

ICH_MISR_EL2 is a 64-bit register.

Field descriptions

The ICH_MISR_EL2 bit assignments are:



All fields listed below are as described in the ICH_MISR_EL2 register description in [3].

Bits [63:8]

Reserved, RES0.

VGrp1D, bit [7]

vPE Group 1 Disabled.

VGrp1E, bit [6]

vPE Group 1 Enabled.

VGrp0D, bit [5]

vPE Group 0 Disabled.

VGrp0E, bit [4]

vPE Group 0 Enabled.

NP, bit [3]

No Pending.

LRNP, bit [2]

List Register Entry Not Present.

U, bit [1]

Underflow.

EOI, bit [0]

End of Interrupt.

Accessing ICH_MISR_EL2

Read-only

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_MISR_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1011	0b010

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL2.SRE == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        X[t, 64] = ICH_MISR_EL2;
elsif PSTATE.EL == EL3 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL3.SRE == '0' then
        AArch64.SystemAccessTrap(EL3, 0x18);
    else

```

$X[t, 64] = \text{ICH_MISR_EL2};$

9.7.8 ICH_VTR_EL2, Interrupt Controller VGIC Type Register

The ICH_VTR_EL2 characteristics are:

Purpose

Reports supported GIC virtualization features.

Configuration

If EL2 is not implemented, all bits in this register are RES0 from EL3, except for nV4, which is RES1 from EL3 when EL2 is not implemented.

This register has no effect if EL2 is not enabled in the current Security state.


This register is present only when (FEAT_GICv3 is implemented or FEAT_GCIE_LEGACY is implemented) and (EL2 is implemented or EL3 is implemented). Otherwise, direct accesses to ICH_VTR_EL2 are UNDEFINED.

Attributes

ICH_VTR_EL2 is a 64-bit register.

Field descriptions

The ICH_VTR_EL2 bit assignments are:



All fields listed below are as described in the ICH_VTR_EL2 register description in [3].

Bits [63:32]

Reserved, RES0.

PRIbits, bits [31:29]

The number of virtual priority bits implemented, minus one.

If FEAT_GCIE_LEGACY is implemented, this field returns 4 (5 bits of priority).

PREbits, bits [28:26]

The number of virtual preemption bits implemented, minus one.

If FEAT_GCIE_LEGACY is implemented, this field returns 4 (5 bits of preemption).

IDbits, bits [25:23]

The number of virtual interrupt identifier bits supported.

IDbits	Meaning
0b000	16 bits
0b001	24 bits

SEIS, bit [22]

When FEAT_GICv3 is implemented:

SEI Support.

Otherwise:

RES0

A3V, bit [21]

Affinity 3 VALID.

A3V	Meaning
0b0	The virtual CPU interface logic only supports zero values of Affinity 3 in SGI generation System registers.
0b1	The virtual CPU interface logic supports non-zero values of Affinity 3 in SGI generation System registers.

nV4, bit [20]

Direct injection of virtual interrupts not supported.

This bit is RES1.

TDS, bit [19]

When FEAT_GICv3 is implemented:

Support for ICH_HCR_EL2.TDIR.

Otherwise:

RES1

DVIM, bit [18]

When FEAT_GICv3 is implemented:

Masking of directly-injected virtual interrupts.

Otherwise:

RES0

Bits [17:5]

Reserved, RES0.

ListRegs, bits [4:0]

The number of List Registers minus 1.

Accessing ICH_VTR_EL2

Read-only

Accesses to this register use the following encodings in the System register encoding space:

MRS <Xt>, ICH_VTR_EL2

op0	op1	CRn	CRm	op2
0b11	0b100	0b1100	0b1011	0b001

```

if PSTATE.EL == EL0 then
    UNDEFINED;
elsif PSTATE.EL == EL1 then
    if EL2Enabled() && HCR_EL2.NV == '1' && ICH_VCTLR_EL2.V3 == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        UNDEFINED;
elsif PSTATE.EL == EL2 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL2.SRE == '0' then
        AArch64.SystemAccessTrap(EL2, 0x18);
    else
        X[t, 64] = ICH_VTR_EL2;
elsif PSTATE.EL == EL3 then
    if IsFeatureImplemented(FEAT_GICv3) && ICC_SRE_EL3.SRE == '0' then
        AArch64.SystemAccessTrap(EL3, 0x18);
    else
        X[t, 64] = ICH_VTR_EL2;

```

9.7.9 Nested virtualization

‡_{HPNKC}

The Arm architecture’s enhanced support for nested virtualization (FEAT_NV2) provides a mechanism for hardware to transform reads and writes from System registers into reads and writes from memory.

When accesses to system registers introduced by GICv3 are transformed into memory accesses, the offsets are unchanged by GICv5. See “Enhanced support for nested virtualization” in the Architecture Reference Manual for A-profile [1].

9.8 Legacy virtual CPU interface registers

Configuration and state of Legacy virtual CPU interface accessible from a GICv3.3 VM.

9.8.1 AArch64 Legacy virtual CPU interface registers

R_{WSWHT}

When FEAT_GICv3 or FEAT_GCIE_LEGACY is implemented, the following registers are present:

- ICV_AP0R0_EL1, Interrupt Controller Virtual Active Priorities Group 0 Registers
- ICV_AP1R0_EL1, Interrupt Controller Virtual Active Priorities Group 1 Registers
- ICV_BPR0_EL1, Interrupt Controller Virtual Binary Point Register 0
- ICV_BPR1_EL1, Interrupt Controller Virtual Binary Point Register 1
- ICV_CTLR_EL1, Interrupt Controller Virtual Control Register
- ICV_DIR_EL1, Interrupt Controller Deactivate Virtual Interrupt Register
- ICV_EOIR0_EL1, Interrupt Controller Virtual End Of Interrupt Register 0
- ICV_EOIR1_EL1, Interrupt Controller Virtual End Of Interrupt Register 1
- ICV_HPPIR0_EL1, Interrupt Controller Virtual Highest Priority Pending Interrupt Register 0
- ICV_HPPIR1_EL1, Interrupt Controller Virtual Highest Priority Pending Interrupt Register 1
- ICV_IAR0_EL1, Interrupt Controller Virtual Interrupt Acknowledge Register 0
- ICV_IAR1_EL1, Interrupt Controller Virtual Interrupt Acknowledge Register 1
- ICV_IGRPEN0_EL1, Interrupt Controller Virtual Interrupt Group 0 Enable Register
- ICV_IGRPEN1_EL1, Interrupt Controller Virtual Interrupt Group 1 Enable Register
- ICV_NMIAR1_EL1, Interrupt Controller Virtual Non-maskable Interrupt Acknowledge Register 1
- ICV_PMR_EL1, Interrupt Controller Virtual Interrupt Priority Mask Register
- ICV_RPR_EL1, Interrupt Controller Virtual Running Priority Register

The above registers are as described in [3].

R_{XKKDD}

When either FEAT_GICv3 or FEAT_GCIE_LEGACY is implemented, the registers listed in Table 9.184 are present.

When FEAT_GCIE_LEGACY is implemented, all of the following are true:

- When HCR_EL2.IMO is 1 and Legacy operation is enabled, all of the following are true:
 - Accesses at EL1 to the registers listed in Table 9.184 generate a Trap exception taken to EL2 and reported using EC syndrome value 0x18.
 - Accesses at any Exception level other than EL1 to the registers listed in Table 9.184 are UNDEFINED.
- When HCR_EL2.IMO is 0 or Legacy operation is disabled, the registers listed in Table 9.184 are UNDEFINED at any Exception level.

Table 9.184: GICv3 CPU Interface registers

Register
ICC_SGI0R_EL1
ICC_SGI1R_EL1
ICC_ASGI1R_EL1

The registers listed in Table 9.184 are described in [3].

R_{MCTMX}

When either FEAT_GICv3 or FEAT_GCIE_LEGACY is implemented, ICC_SRE_EL1 is present.

When FEAT_GCIE_LEGACY is implemented, all of the following are true:

- When HCR_EL2.IMO is 1 and Legacy operation is enabled, all of the following are true:
 - For an access at EL1 to ICC_SRE_EL1, accesses to the following fields are RAO/WI:

- * ICC_SRE_EL1.SRE
- * ICC_SRE_EL1.DFB
- * ICC_SRE_EL1.DIB

– Access at any Exception level other than EL1 to ICC_SRE_EL1 is UNDEFINED.

- When HCR_EL2.IMO is 0 or Legacy operation is disabled, ICC_SRE_EL1 is UNDEFINED at any Exception level.

ICC_SRE_EL1 [Table 9.184](#) is described in [\[3\]](#).

R_{PWDWZ}

When either FEAT_GICv3 or FEAT_GCIE_LEGACY is implemented, the following fields are present:

- HFGTR_EL2.ICC_IGRPENn_EL1
- HFGWTR_EL2.ICC_IGRPENn_EL1

The above fields are described in [\[1\]](#).

Chapter 10

Registers and memory maps

This chapter describes the registers and memory mapped interfaces of GICv5 components.

10.1 Memory-mapped programmer's model

R _{JYKRN}	<p>Accesses to registers that are not implemented is CONSTRAINED UNPREDICTABLE and results in either:</p> <ul style="list-style-type: none"> • The access has RAZ/WI behavior. • The location has the same behavior as RES0.
R _{JGRJR}	<p>The access type definitions for the memory-mapped register interface are:</p> <ul style="list-style-type: none"> • RW: Read and write. • RO: Read only. Writes are ignored. • WO: Write only. Reads return an UNKNOWN value. • WI: Write ignore. Reads return an UNKNOWN value. • UNKNOWN/WI: The same as WI. • RAZ/WI: Read-As-Zero, Writes Ignored. • RES0: RES0 as defined in [1]. • RES1: RES1 as defined in [1].
R _{DQMPD}	<p>In the register definitions, references to other registers are always to a register in the same register frame as where the access is made.</p>
R _{HQTXF}	<p>All registers for the IWB, the ITS, and IRS are little-endian.</p>
R _{WRLMJ}	<p>All registers are either 32-bit or 64-bit.</p> <p>An implementation supports aligned 32-bit accesses to all registers and aligned 64-bit accesses to 64-bit registers. When a 32-bit access occurs to a 64-bit register, bits[63:32] of the register are accessed at offset +4 and bits[31:0] at offset 0.</p> <p>All aligned single-copy atomic accesses to a register of the same size as the access are treated as single-copy atomic accesses by the GIC.</p> <p>Aligned 16-bit halfword accesses are supported only to the lower half of the following registers:</p> <ul style="list-style-type: none"> • ITS_TRANSLATER • ITS_RL_TRANSLATER <p>All other accesses are <i>illegal accesses</i>.</p>
R _{VXDYH}	<p>The behavior of an illegal access is CONSTRAINED UNPREDICTABLE to one of the following:</p> <ul style="list-style-type: none"> • The access is RAZ/WI. • The access completes on the GIC and one of the following is true: <ul style="list-style-type: none"> – The access is a read and the read returns an UNKNOWN value. – The access is a write and the write sets any field of the accessed register, including fields outside the access, to an UNKNOWN value. It is CONSTRAINED UNPREDICTABLE whether side effects of a write occur or not. • The access generates an abort. <p>A 64-bit access to two adjacent 32-bit registers is CONSTRAINED UNPREDICTABLE and has one of the following behaviors:</p> <ul style="list-style-type: none"> • The access is RAZ/WI. • A read returns the value of both registers and a write updates both registers, as though two 32-bit accesses were performed in an UNPREDICTABLE order. • One of the pair of registers is read or written and the other register is RAZ/WI, as though a single 32-bit access was performed to an UNPREDICTABLE one of the pair of registers. • The access generates an abort.
D _{RFLQG}	<p>Some register fields are <i>Guarded</i> by another field in the same or another register. In the GICv5 specification, a Guarded register field is RO or WI unless the field by which it is Guarded is in a certain state.</p>

For example, IRS_IST_BASER.ADDR is Guarded by IRS_IST_BASER.VALID and IRS_PE_SEL.R.IAFFID is Guarded by IRS_PE_STATUSR.IDLE.

R_{BRRWQ} Writing to Guarded register fields that are **RO** or **WI** have no effects beyond optionally reporting an error in the software error reporting mechanism.

S_{MXQGP} Software must ensure proper ordering when updating Guarded fields.

For a write that updates the field by which other fields are Guarded, all of the following are true:

- In order to update the Guarded fields, software must ensure that the GIC has observed the write before it observes a write to the Guarded fields .
- This ordering can be enforced through the use of an appropriate memory type for accesses and by issuing appropriate barriers, or by reading back the write before making another write to the Guarded fields.

For a write that updates Guarded fields, all of the following are true:

- To make sure the GIC observes the write to the Guarded fields, software must ensure that the GIC has observed the write to the Guarded fields before it observes a write to the field by which they are Guarded.
- This ordering can be enforced through the use of an appropriate memory type for accesses and by issuing appropriate barriers, or by reading back the value written to the Guarded fields before making another write to the field by which they are Guarded.

D_{DYZRK} Some registers that have write side-effects define an *action bit*. The side effects take place when a write occurs that sets the action bit to 1.

Registers with an action bit are always tracked with an IDLE bit in a status register to indicate when the write side-effects are complete.

I_{TWSRH} Following a write that sets the action bit to 1, when the corresponding IDLE bit is 1, another write that sets the action bit to 1, without any intervening write that sets the action bit to 0, will cause the side-effects defined by the register.

I_{XXQBX} If a write occurs to a register that has write side-effects and define an action bit, and the write does not write 1 to the action bit, the write side-effects will not take place and other registers controlled by the register with the action bit are not affected.

For example, if a write occurs to IRS_VPE_SEL.R that does not set IRS_VPE_SEL.R.S to 1, updates to any of the IRS_VPE_x registers apply to the VPE selected the last time a write occurred that set IRS_VPE_SEL.R.S to 1.

I_{FFYYH} 64-bit registers that have write side-effects, and support 32-bit accesses, define an action bit in either of the 32-bit halves.

The write side-effects can use the full 64-bit register value as input to their operation when a write follows this sequence:

1. Write to the 32-bit half of the register which does not contain the action bit.
2. Ensure the write is observed before a following write by using a barrier instruction or reading back the value written to the 32-bit half of the register.
3. Write to the other 32-bit half of the register, which contains the action bit, setting the action bit to 1.

R_{SCDGF} Registers are not required to support being the target of exclusive or atomic read-modify-write update operations.

I_{VBHBM} Register frames associated with an Interrupt Domain are accessible via the PAS associated with that Interrupt Domain.

For register frames not associated with the MPPAS, except for the ITS_TRANSLATE_FRAME and IRS_SETLPI_FRAME, it is IMPLEMENTATION DEFINED whether the register frames are also accessible in the MPPAS at the same addresses.

I_{WJTJR} Any memory-mapped access to a GICv5 register region is defined to be beyond the PE.

S_{DGVJR}

Armv9 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.

Software can 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 a GICv5 implementation.

Reads and writes of the memory-mapped registers complete in the order in which they arrive at the GIC.

For accesses to different register locations, software can determine the order in which they arrive at the GIC by doing all of the following:

- Accessing the GIC using the Device-nGnRnE or Device-nGnRE memory types.
- Using the appropriate memory barriers.

Software can determine the completion of a write by doing one of the following:

- Accessing the GIC using the Device-nGnRnE memory type and executing a DSB barrier.
- Reading back the value written.

For more information on memory types and barriers ensuring completion, see [1].

10.2 IRS register frames

I_{DLZJM}

For each Interrupt Domain, the IRS exposes the following register frames:

1. The IRS configuration frame.
2. Optionally, the SETLPI register frame.

Each register frame size is 64K.

I_{JWWVY}

In the IRS register definitions, references to other registers are always to a register in the register frame for the same Interrupt Domain on the same IRS as where the access is made.

I_{GGXDW}

If IRS_IDR0.SETLPI is 0, the SETLPI register frame is not present.

10.2.1 IRS_CONFIG_FRAME, IRS configuration register frame

The IRS_CONFIG_FRAME characteristics are:

Purpose

Contains control registers for an IRS for an Interrupt Domain.

An IRS configuration register frame is present for each supported Interrupt Domain on each IRS.

This register frame is accessible in the PAS associated with the Interrupt Domain.

It is IMPLEMENTATION DEFINED whether this register frame is also accessible in the MPPAS at the same address.

The base address is distinct from the base address of any other GIC register frame, including the configuration register frames for other Interrupt Domains on all IRSs.

The base address is aligned to 64KB.

Attributes

The IRS_CONFIG_FRAME block is of size 64KB

Default access

Accesses to the address space of this block that do not resolve to a register or a further block are treated as RAZ/WI.

Contents

Offset	Name	Notes
0x0000	IRS_IDR0	Most permissive access: RO
0x0004	IRS_IDR1	Most permissive access: RO
0x0008	IRS_IDR2	Most permissive access: RO
0x000C	IRS_IDR3	Most permissive access: RO
0x0010	IRS_IDR4	Most permissive access: RO
0x0014	IRS_IDR5	Most permissive access: RO
0x0018	IRS_IDR6	Most permissive access: RO
0x001C	IRS_IDR7	Most permissive access: RO
0x0040	IRS_IIDR	Most permissive access: RO
0x0044	IRS_AIDR	Most permissive access: RO
0x0080	IRS_CR0	Most permissive access: RW
0x0084	IRS_CR1	Most permissive access: RW
0x00C0	IRS_SYNCR	Most permissive access: WO
0x00C4	IRS_SYNC_STATUSR	Most permissive access: RO
0x0100	IRS_SPI_VMR	Most permissive access: RW
0x0108	IRS_SPI_SEL	Most permissive access: WO
0x010C	IRS_SPI_DOMAINR	Most permissive access: RW
0x0110	IRS_SPI_RESAMPLER	Most permissive access: WO
0x0114	IRS_SPI_CFG	Most permissive access: RW

Offset	Name	Notes
0x0118	IRS_SPI_STATUSR	Most permissive access: RO
0x0140	IRS_PE_SEL	Most permissive access: WO
0x0144	IRS_PE_STATUSR	Most permissive access: RO
0x0148	IRS_PE_CR0	Most permissive access: RW
0x0180	IRS_IST_BASER	Most permissive access: RW
0x0190	IRS_IST_CFGR	Most permissive access: RW
0x0194	IRS_IST_STATUSR	Most permissive access: RO
0x01C0	IRS_MAP_L2_ISTR	Most permissive access: WO
0x0200	IRS_VMT_BASER	Most permissive access: RW
0x0210	IRS_VMT_CFGR	Most permissive access: RW
0x0214	IRS_VMT_STATUSR	Most permissive access: RO
0x0240	IRS_VPE_SEL	Most permissive access: RW
0x0248	IRS_VPE_DBR	Most permissive access: RW
0x0250	IRS_VPE_HPPIR	Most permissive access: RO
0x0258	IRS_VPE_CR0	Most permissive access: RW
0x025C	IRS_VPE_STATUSR	Most permissive access: RO
0x0280	IRS_VM_DBR	Most permissive access: RW
0x0288	IRS_VM_SEL	Most permissive access: WO
0x028C	IRS_VM_STATUSR	Most permissive access: RO
0x02C0	IRS_VMAP_L2_VMTR	Most permissive access: RW
0x02C8	IRS_VMAP_VMR	Most permissive access: RW
0x02D0	IRS_VMAP_VISTR	Most permissive access: RW
0x02D8	IRS_VMAP_L2_VISTR	Most permissive access: RW
0x02E0	IRS_VMAP_VPER	Most permissive access: RW
0x0300	IRS_SAVE_VMR	Most permissive access: RW
0x0308	IRS_SAVE_VM_STATUSR	Most permissive access: RO
0x0340	IRS_MEC_IDR	Most permissive access: RO
0x0344	IRS_MEC_MECID_R	Most permissive access: RW
0x0380	IRS_MPAM_IDR	Most permissive access: RO
0x0384	IRS_MPAM_PARTID_R	Most permissive access: RW
0x03C0	IRS_SWERR_STATUSR	Most permissive access: RW
0x03C8	IRS_SWERR_SYNDROMER0	Most permissive access: RO
0x03D0	IRS_SWERR_SYNDROMER1	Most permissive access: RO
0x0E00 + (4 * n) for n in ↪ 63:0	-	Most permissive access: ImplementationDefined

10.2.1.1 IRS_AIDR

The IRS_AIDR characteristics are:

Purpose

IRS Architecture Identification Register. Identifies the GIC architecture version to which the implementation conforms.


Attributes

IRS_AIDR is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_AIDR bit assignments are:



Bits [31:12]

Reserved, RES0.

Component, bits [11:8]

GIC component

Component	Meaning
0b0000	IRS
0b0001	ITS
0b0010	IWB

ArchMajorRev, bits [7:4]

Major Architecture revision.

ArchMajorRev	Meaning
0b0000	GICv5.x

ArchMinorRev, bits [3:0]

Minor Architecture revision.

ArchMinorRev	Meaning
0b0000	GICv5.0

Accessing IRS_AIDR

Accesses to this register use the following encodings:

Accessible at address 0x0044

Access on this interface is **RO**.

10.2.1.2 IRS_CR0

The IRS_CR0 characteristics are:

Purpose

IRS control register 0.


Attributes

IRS_CR0 is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_CR0 bit assignments are:



Bits [31:2]

Reserved, RES0.

IDLE, bit [1]

Whether the transition between enabled and disabled states of the IRS for the Interrupt Domain is complete.

IDLE	Meaning
0b0	The effects of updating IRSEN are not guaranteed to have completed.
0b1	The effects of updating IRSEN are have completed.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b1.

Access to this field is **RO**.

IRSEN, bit [0]

Controls whether the IRS Forwards and Recalls candidate HPPIs to PEs for the Interrupt Domain.

IRSEN	Meaning
0b0	The IRS is disabled for the Interrupt Domain.
0b1	The IRS is enabled for the Interrupt Domain.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b0.

Accessing this field has the following behavior:

- When IRS_CR0.IDLE == 0, access to this field is **RO**
- Otherwise, access to this field is **RW**

Accessing IRS_CR0

Accesses to this register use the following encodings:

Accessible at address 0x0080

Access on this interface is **RW**.

10.2.1.3 IRS_CR1

The IRS_CR1 characteristics are:

Purpose

IRS configuration register 1


Attributes

IRS_CR1 is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_CR1 bit assignments are:



Bits [31:16]

Reserved, RES0.

VPED_WA, bit [15]

When *IRS_IDR0.VIRT* == 1:

Write-Allocate hint for the VPE descriptors.

VPED_WA	Meaning
0b0	No Write-Allocate.
0b1	Write-Allocate.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

VPED_RA, bit [14]

When *IRS_IDR0.VIRT* == 1:

Read-Allocate hint for the VPE descriptors.

VPED_RA	Meaning
0b0	No Read-Allocate.
0b1	Read-Allocate.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

VMD_WA, bit [13]

When *IRS_IDR0.VIRT* == 1:

Write-Allocate hint for the VM descriptors.

VMD_WA	Meaning
0b0	No Write-Allocate.
0b1	Write-Allocate.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

VMD_RA, bit [12]

When *IRS_IDR0.VIRT* == 1:

Read-Allocate hint for the VM descriptors.

VMD_RA	Meaning
0b0	No Read-Allocate.
0b1	Read-Allocate.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

VPET_WA, bit [11]

When *IRS_IDR0.VIRT* == 1:

Write-Allocate hint for the VPE table.

VPET_WA	Meaning
0b0	No Write-Allocate.
0b1	Write-Allocate.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

VPET_RA, bit [10]

When IRS_IDR0.VIRT == 1:

Read-Allocate hint for the VPE table.

VPET_RA	Meaning
0b0	No Read-Allocate.
0b1	Read-Allocate.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

VMT_WA, bit [9]

When IRS_IDR0.VIRT == 1:

Write-Allocate hint for the VM table.

VMT_WA	Meaning
0b0	No Write-Allocate.
0b1	Write-Allocate.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

VMT_RA, bit [8]

When IRS_IDR0.VIRT == 1:

Read-Allocate hint for the VM table.

VMT_RA	Meaning
0b0	No Read-Allocate.
0b1	Read-Allocate.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

IST_WA, bit [7]

Write-Allocate hint for ISTs.

When IRS_IDR0.VIRT is 1, this control applies to the virtual ISTs as well as the physical IST.

IST_WA	Meaning
0b0	No Write-Allocate.
0b1	Write-Allocate.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

IST_RA, bit [6]

Read-Allocate hint for ISTs.

When IRS_IDR0.VIRT is 1, this control applies to the virtual ISTs as well as the physical IST.

IST_RA	Meaning
0b0	No Read-Allocate.
0b1	Read-Allocate.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

IC, bits [5:4]

Controls the Inner Cacheability attribute used when the IRS accesses tables as a requester.

IC	Meaning
0b00	Non-cacheable.
0b01	Write-Back Cacheable.
0b10	Write-Through Cacheable.
0b11	Reserved, treated as 0b00.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

OC, bits [3:2]

Controls the Outer Cacheability attribute used when the IRS accesses tables as a requester.

OC	Meaning
0b00	Non-cacheable.

OC	Meaning
0b01	Write-Back Cacheable.
0b10	Write-Through Cacheable.
0b11	Reserved, treated as 0b00.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

SH, bits [1:0]

Controls the Shareability attribute used when the IRS accesses tables as a requester.

SH	Meaning
0b00	Non-shareable.
0b01	Reserved, treated as 0b00.
0b10	Outer Shareable.
0b11	Inner Shareable.

When IRS_CR1.OC is 0b00 and IRS_CR1.IC is 0b00, this field is IGNORED and behaves as Outer Shareable.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing IRS_CR1

Accesses to this register use the following encodings:

Accessible at address 0x0084

- When IRS_IST_BASER.VALID == 1 or IRS_IST_STATUSR.IDLE == 0, access on this interface is **RO**.
- When IRS_VMT_BASER.VALID == 1 or IRS_VMT_STATUSR.IDLE == 0, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.2.1.4 IRS_IDR0

The IRS_IDR0 characteristics are:

Purpose

IRS identification register 0. Contains read-only fields with information about the IRS GIC component.


Attributes

IRS_IDR0 is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_IDR0 bit assignments are:



IRSID, bits [31:16]

Unique identifier for this IRS in the system.

This value is the same across all Interrupt Domains for this IRS.

Bits [15:13]

Reserved, RES0.

SWE, bit [12]

Software error reporting support in the Interrupt Domain.

SWE	Meaning
0b0	Software error reporting is not supported.
0b1	Software error reporting is supported.

Support for Software error reporting is optional.

MPAM, bit [11]

Memory Partitioning And Monitoring (MPAM) support.

MPAM	Meaning
0b0	MPAM is not supported.
0b1	MPAM is supported.

Support for MPAM is optional.

MEC, bit [10]

When *IRS_IDR0.INT_DOM* == 0b11:

Support for Memory Encryption Contexts (MEC) for the Realm Interrupt Domain.

MEC	Meaning
0b0	Memory Encryption Contexts are not supported.
0b1	Memory Encryption Contexts are supported.

Support for MEC is optional.

Otherwise:

RES0

SETLPI, bit [9]

When *IRS_IDR2.LPI* == 1:

Whether the IRS implements the IRS SETLPI register.

The IRS SETLPI register allows an MSI to set a physical LPI Pending without translation by an ITS.

SETLPI	Meaning
0b0	The IRS does not implement the IRS SETLPI register.
0b1	The IRS implements the IRS SETLPI register.

Otherwise:

RES0

VIRT_ONE_N, bit [8]

Whether virtual 1ofN is supported.

VIRT_ONE_N	Meaning
0b0	Virtual 1ofN is not supported.
0b1	Virtual 1ofN is supported.

This field is RES0, if any of the following are true:

- Virt is 0.
- ONE_N is 0.

ONE_N, bit [7]

ONE_N	Meaning
0b0	1ofN is not supported.
0b1	1ofN is supported.

VIRT, bit [6]

This field is RES0 for the EL3 Interrupt Domain

VIRT	Meaning
0b0	Virtualization is not supported.
0b1	Virtualization is supported.

PA_RANGE, bits [5:2]

Physical Address range supported.

The physical address range corresponds to the system physical address size.

The value of this field is the same for all Interrupt Domains across all IRSs in the system.

PA_RANGE	Meaning
0b0000	32 bits, 4GB
0b0001	36 bits, 64GB
0b0010	40 bits, 1TB
0b0011	42 bits, 4TB
0b0100	44 bits, 16TB
0b0101	48 bits, 256TB
0b0110	52 bits, 4PB
0b0111	56 bits, 64PB

Values not defined above are reserved.

INT_DOM, bits [1:0]

The Interrupt Domain that the register frame containing this register controls.

INT_DOM	Meaning
0b00	Secure
0b01	Non-secure

INT_DOM	Meaning
0b10	EL3
0b11	Realm

Accessing IRS_IDR0

Accesses to this register use the following encodings:

Accessible at address 0x0000

Access on this interface is **RO**.

10.2.1.5 IRS_IDR1

The IRS_IDR1 characteristics are:

Purpose

IRS identification register 1. Contains read-only fields with information about the IRS GIC component.


Attributes

IRS_IDR1 is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_IDR1 bit assignments are:



Bits [31:23]

Reserved, RES0.

PRI_BITS, bits [22:20]

The number of priority bits implemented, minus one.

PRI_BITS	Meaning
0b000	1 bit of priority.
0b001	2 bits of priority.
0b010	3 bits of priority.
0b011	4 bits of priority.
0b100	5 bits of priority.

Values not defined above are reserved.

When fewer than 5 bits are implemented, the lower order priority bits are RES0.

IAFFID_BITS, bits [19:16]

Number of bits of IAFFID supported - 1.

Unimplemented upper bits of IAFFID are RES0 when sending and receiving commands between a PE and the IRS.

PE_CNT, bits [15:0]

The number of PEs connected to this IRS.

Accessing IRS_IDR1

Accesses to this register use the following encodings:

Accessible at address 0x0004

Access on this interface is **RO**.

10.2.1.6 IRS_IDR2

The IRS_IDR2 characteristics are:

Purpose

IRS identification register 2. Contains read-only fields with information about the IRS GIC component.


Attributes

IRS_IDR2 is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_IDR2 bit assignments are:



Bits [31:20]

Reserved, RES0.

ISTMD_SZ, bits [19:15]

Describes the minimum number of INTID.ID bits which requires a level 2 ISTE size of 16 bytes to store metadata.

When the configured number of INTID.ID bits is smaller than ISTMD_SZ, the minimum required size of a level 2 ISTE is 8 bytes.

When the configured number of INTID.ID bits is larger than or equal to ISTMD_SZ, the minimum required size of a level 2 ISTE is 16 bytes.

When ISTMD is 0, this field is RES0.

ISTMD, bit [14]

Whether the IST entries contain metadata storage.

When the IST entries contain metadata storage, the size of a level 2 ISTE is 8 bytes or 16 bytes.

The size of an IST entry containing metadata depends on the configured ID range and the values reported in ISTMD_SZ.

ISTMD	Meaning
0b0	The IST entries do not require storage for metadata and the size of a level 2 ISTE is 4 bytes.
0b1	The IST entries require storage for metadata and the size of a level 2 ISTE is 8 bytes or 16 bytes.

IST_L2SZ, bits [13:11]

Supported level 2 IST sizes when a 2-level IST structure is used.

IST_L2SZ	Meaning
0b001	Level 2 IST sizes supported: 4KB
0b010	Level 2 IST sizes supported: 16KB
0b011	Level 2 IST sizes supported: 4KB and 16KB
0b100	Level 2 IST sizes supported: 64KB
0b101	Level 2 IST sizes supported: 4KB and 64KB
0b110	Level 2 IST sizes supported: 16KB and 64KB
0b111	Level 2 IST sizes supported: 4KB, 16KB, and 64KB

Values not defined above are reserved.

When IST_LEVELS is 0, this field is RES0.

When LPI is 0, this field is RES0.

IST_LEVELS, bit [10]

Levels supported for the IST.

When LPI is 0, this field is RES0.

IST_LEVELS	Meaning
0b0	Only a single level linear structure is supported.
0b1	2-level structure is supported in addition to the linear structure.

MIN_LPI_ID_BITS, bits [9:6]

The minimum number of LPI_ID_BITS supported.

The maximum value supported for this field is 14.

When LPI is 0, this field is RES0.

LPI, bit [5]

Whether physical LPIs are implemented.

When physical LPIs are not supported, the physical IST registers are not implemented.

ID_BITS, bits [4:0]

The maximum number of INTID.ID bits that the IRS supports.

The maximum supported value of this field is 24.

Accessing IRS_IDR2

Accesses to this register use the following encodings:

Accessible at address 0x0008

Access on this interface is **RO**.

10.2.1.7 IRS_IDR3

The IRS_IDR3 characteristics are:

Purpose

IRS identification register 3. Contains read-only fields with information about the IRS GIC component.


Attributes

IRS_IDR3 is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_IDR3 bit assignments are:



Bits [31:11]

Reserved, RES0.

VMT_LEVELS, bit [10]

When *IRS_IDR0.VIRT* == 1:

Levels supported for the VM table.

VMT_LEVELS	Meaning
0b0	Only a single level linear structure is supported.
0b1	2-level structure is supported in addition to the linear structure.

Otherwise:

RES0

VM_ID_BITS, bits [9:5]

When *IRS_IDR0.VIRT* == 1:

The maximum number of VM ID bits that the IRS supports.

The minimum supported value of this field is 8 and the maximum supported value is 16.

Otherwise:

RES0

VMD_SZ, bits [4:1]

When *IRS_IDR0.VIRT* == 1:

Specifies the size in bytes of a VM descriptor.

Each VM descriptor is 2^(VMD_SZ) bytes.

The minimum valid value for this field is 3 and the maximum valid value is 12.

This corresponds to a minimum descriptor size of 8 bytes and a maximum descriptor size of 4096 bytes.

Otherwise:

RES0

VMD, bit [0]

When *IRS_IDR0.VIRT* == 1:

Whether each VM requires an IMPLEMENTATION DEFINED memory area.

VMD	Meaning
0b0	The VMs do not require VM descriptor areas.
0b1	Each VM requires a VM descriptor area.

Otherwise:

RES0

Accessing *IRS_IDR3*

Accesses to this register use the following encodings:

Accessible at address 0x000C

Access on this interface is **RO**.

10.2.1.8 IRS_IDR4

The IRS_IDR4 characteristics are:

Purpose

IRS identification register 4. Contains read-only fields with information about the IRS GIC component.

Attributes

IRS_IDR4 is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_IDR4 bit assignments are:



Bits [31:10]

Reserved, RES0.

VPE_ID_BITS, bits [9:6]

The number of VPE ID bits supported - 1.

The minimum supported value of this field is 8 and the maximum supported value is 16.

Accessing this field has the following behavior:

- When IRS_IDR0.VIRT == 0, access to this field is **RES0**
- Otherwise, access to this field is **RO**

VPED_SZ, bits [5:0]

Specifies the size in bytes of a VPE Descriptor.

Each VPE Descriptor is $2^{(VPED_SZ)}$ bytes.

The minimum valid value for this field is 3 and the maximum valid value is 12.

This corresponds to a minimum descriptor size of 8 bytes and a maximum descriptor size of 4096 bytes.

Accessing this field has the following behavior:

- When IRS_IDR0.VIRT == 0, access to this field is **RES0**
- Otherwise, access to this field is **RO**

Accessing IRS_IDR4

Accesses to this register use the following encodings:

Accessible at address 0x0010

Access on this interface is **RO**.

10.2.1.9 IRS_IDR5

The IRS_IDR5 characteristics are:

Purpose

IRS identification register 5. Contains read-only fields with information about the IRS GIC component.


Attributes

IRS_IDR5 is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_IDR5 bit assignments are:



Bits [31:25]

Reserved, RES0.

SPI_RANGE, bits [24:0]

The total number of SPIs in the system.

The maximum value reported in this field is 2^{24} .

The number reported in this register is the same across all Interrupt Domains and across all IRSs in a system.

Accessing IRS_IDR5

Accesses to this register use the following encodings:

Accessible at address 0x0014

Access on this interface is **RO**.

10.2.1.10 IRS_IDR6

The IRS_IDR6 characteristics are:

Purpose

IRS identification register 6. Contains read-only fields with information about the range of SPI INTIDs managed by this IRS.

Attributes

IRS_IDR6 is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_IDR6 bit assignments are:



Bits [31:25]

Reserved, RES0.

SPI_IRS_RANGE, bits [24:0]

The number of SPI INTID.ID managed on this IRS.

The value reported in this field is less than or equal to the value reported in IRS_IDR5.SPI_RANGE.

If IRS_IDR5.SPI_RANGE is 0, this field is RES0.

Accessing IRS_IDR6

Accesses to this register use the following encodings:

Accessible at address 0x0018

Access on this interface is **RO**.

10.2.1.11 IRS_IDR7

The IRS_IDR7 characteristics are:

Purpose

IRS identification register 7. Contains read-only fields with information about the minimum SPI INTID.ID value implemented on this IRS.

Attributes

IRS_IDR7 is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_IDR7 bit assignments are:



Bits [31:24]

Reserved, RES0.

SPI_BASE, bits [23:0]

The minimum SPI INTID.ID implemented on this IRS.

If IRS_IDR5.SPI_RANGE is 0 or IRS_IDR6.SPI_IRS_RANGE is 0, this field is RES0.

Accessing IRS_IDR7

Accesses to this register use the following encodings:

Accessible at address 0x001C

Access on this interface is **RO**.

10.2.1.12 IRS_IIDR

The IRS_IIDR characteristics are:

Purpose

IRS Implementer Identification Register. Provides information about the implementation and implementer of the GIC, and architecture version supported.

Attributes

IRS_IIDR is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_IIDR bit assignments are:

31	20	19	16	15	12	11	0
ProductID				Variant	Revision	Implementer	

ProductID, bits [31:20]

IMPLEMENTATION DEFINED value identifying the GIC part

When the IRS_PIDR{0,1} registers are present, Arm expects that the IRS_PIDR{0,1}.PART_{0,1} fields match the value of IRS_IIDR.ProductID.

If required, however, an implementation is permitted to provide values for IRS_PIDR.{0,1}.PART_{0,1} that do not match the value of IRS_IIDR.ProductID

Variant, bits [19:16]

IMPLEMENTATION DEFINED value used to distinguish product variants, or major revisions of the product

Revision, bits [15:12]

IMPLEMENTATION DEFINED value used to distinguish minor revisions of the product

Implementer, bits [11:0]

Contains the JEP106 code of the company that implemented the GIC

For an Arm implementation, the JEP106 code is 0x43B

When the IRS_PIDR{1,2,4} registers are present, Arm expects that the IRS_PIDR{0,1}.PART_{0,1} fields match the value of IRS_IIDR.Implementer.

Accessing IRS_IIDR

Accesses to this register use the following encodings:

Accessible at address 0x0040

Access on this interface is **RO**.

10.2.1.13 IRS_IST_BASER

The IRS_IST_BASER characteristics are:

Purpose

IRS IST base address register

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_IST_STATUSR.IDLE is 1.


Attributes

IRS_IST_BASER is a 64-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_IST_BASER bit assignments are:



Bits [63:56]

Reserved, RES0.

ADDR, bits [55:6]

Bits[55:6] of the base physical address of the IST.

When IRS_IST_CFGR.STRUCTURE is 0, ADDR points to a linear IST and all of the following are true:

- Bits[N:0] of the resulting address are 0 where N depends on the ISTSZ and LPI_ID_BITS fields in IRS_IST_CFGR as follows:
$$N = \text{Max}(5, (\text{ISTSZ} + 1) + \text{LPI_ID_BITS}).$$
- This means that the level 2 ISTE array is aligned to the size of the array or to a 64-byte boundary when its size is smaller than 64 bytes.

When IRS_IST_CFGR.STRUCTURE is 1, ADDR points to the level 1 table in a 2-level IST and all of the following are true:

- Bits[N:0] of the resulting address are 0 where N depends on L2SZ, ISTSZ, and LPI_ID_BITS in IRS_IST_CFGR as follows:
$$N = \text{Max}(5, \text{LPI_ID_BITS} - ((10 - \text{ISTSZ}) + (2 * \text{L2SZ})) + 2)$$
- This means that the level 1 IST is aligned to the size of the level 1 table or to a 64-byte boundary when its size is smaller than 64 bytes.

See [4.7 The interrupt state table \(IST\)](#) for more information.

Access to any level of the IST and any additional memory accesses occurring as a result of the address in this field are performed using the PAS of the Interrupt Domain where this register is accessed.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in IRS_IDR0.PA_RANGE.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- When IRS_IST_BASER.VALID == 1, access to this field is **RO**
- Otherwise, access to this field is **RW**

Bits [5:1]

Reserved, RES0.

VALID, bit [0]

Whether the ADDR points to a valid IST.

VALID	Meaning
0b0	The IST address is not valid.
0b1	The IST address is valid.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b0.

Accessing IRS_IST_BASER

Accesses to this register use the following encodings:

Accessible at address 0x0180

- When IRS_IDR2.LPI == 0, access on this interface is **RAZ/WI**.
- When IRS_IST_STATUSR.IDLE == 0, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.2.1.14 IRS_IST_CFGR

The IRS_IST_CFGR characteristics are:

Purpose

IRS IST configuration register


Attributes

IRS_IST_CFGR is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_IST_CFGR bit assignments are:



Bits [31:17]

Reserved, RES0.

STRUCTURE, bit [16]

Whether the IST uses a linear or 2-level structure.

STRUCTURE	Meaning
0b0	A linear IST structure is used.
0b1	A 2-level IST structure is used.

When IRS_IDR2.IST_LEVELS is 0, this field is RES0.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Bits [15:9]

Reserved, RES0.

ISTSZ, bits [8:7]

The size of each level 2 ISTE.

Values not defined above are reserved.

If this field is programmed to specify a size smaller than the minimum required size or programmed to a reserved value, it is treated as having the value corresponding to the minimum required size for all other purposes than a direct read of the register.

See [4.7 The interrupt state table \(IST\)](#) for more information about the minimum required size.

ISTSZ	Meaning
0b00	The size of a level 2 ISTE is 4 bytes.
0b01	The size of a level 2 ISTE is 8 bytes.

ISTSZ	Meaning
0b10	The size of a level 2 ISTE is 16 bytes.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

L2SZ, bits [6:5]

Level 2 IST size when a 2-level IST structure is used.

L2SZ	Meaning
0b00	The level 2 IST size is 4KB.
0b01	The level 2 IST size is 16KB.
0b10	The level 2 IST size is 64KB.

Values not defined above are reserved.

IRS_IDR2.IST_L2SZ reports the supported values.

If $LPI_ID_BITS \leq ((10 - ISTSZ) + (2 * L2SZ))$ and STRUCTURE is 1, all of the following are true:

- The IST consists of a single L1_ISTE and a single level 2 IST.
- The level 2 IST contains $(2^{LPI_ID_BITS})$ entries.
- The IRS Domain is allowed to access the full level 2 IST size as specified by L2SZ.

Arm recommends that STRUCTURE is 0 when $LPI_ID_BITS \leq ((10 - ISTSZ) + (2 * L2SZ))$.

See [4.7 The interrupt state table \(IST\)](#) for more information.

If programming a reserved value or an unsupported value, the IRS Domain behavior is CONSTRAINED UNPREDICTABLE to any behavior which could be achieved by programming a valid and supported value.

When STRUCTURE is 0, this field is RES0.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

LPI_ID_BITS, bits [4:0]

The number of LPIs for the IRS Domain.

The IST contains $2^{LPI_ID_BITS}$ level 2 IST entries in total.

The minimum value for this field is IRS_IDR2.MIN_LPI_ID_BITS.

If programmed to a value smaller than the minimum, the field is treated as having the minimum value for all other purposes than reading back the field.

The maximum value for this field is IRS_IDR2.ID_BITS.

If this field is programmed to a value larger than the maximum, it is treated as having the maximum value for all other purposes than a direct read of the register. The maximum value is reported in IRS_IDR2.ID_BITS.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing IRS_IST_CFGR

Accesses to this register use the following encodings:

Accessible at address 0x0190

- When IRS_IDR2.LPI == 0, access on this interface is **RAZ/WI**.
- When IRS_IST_BASER.Valid == 1 or IRS_IST_STATUSR.IDLE == 0, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.2.1.15 IRS_IST_STATUSR

The IRS_IST_STATUSR characteristics are:

Purpose

IRS physical IST management status register.


Attributes

IRS_IST_STATUSR is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_IST_STATUSR bit assignments are:



Bits [31:1]

Reserved, RES0.

IDLE, bit [0]

Whether the effects of any of the following writes are complete:

- A write that updates IRS_IST_BASER.VALID.
- A write to IRS_MAP_L2_ISTR.ID.

IDLE	Meaning
0b0	The effects of the write are not complete.
0b1	The effects of the write are complete.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b1.

Access to this field is **RO**.

Accessing IRS_IST_STATUSR

Accesses to this register use the following encodings:

Accessible at address 0x0194

- When IRS_IDR2.LPI == 0, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RO**.

10.2.1.16 IRS_MAP_L2_ISTR

The IRS_MAP_L2_ISTR characteristics are:

Purpose

IRS map physical level 2 IST register

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_IST_STATUSR.IDLE is 1.


Attributes

IRS_MAP_L2_ISTR is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_MAP_L2_ISTR bit assignments are:



A write to this register makes the specified physical level 2 IST valid.

There are no effects of a write to this register, if any of the following are true:

- The physical IST is invalid.
- The physical IST uses a linear structure.
- The LPI INTID.ID is outside the configured physical LPI range.
- The level 2 IST is already valid.

Bits [31:24]

Reserved, RES0.

ID, bits [23:0]

An LPI INTID.ID covered by the level 1 ISTE corresponding to the level 2 IST that is made valid.

If unimplemented upper bits of the INTID.ID are not zero, it is IMPLEMENTATION DEFINED whether the upper bits are treated as 0 or the interrupt message is ignored by the IRS.

The reset behavior of this field is:

- This field resets to an UNKNOWN value.

Accessing IRS_MAP_L2_ISTR

Accesses to this register use the following encodings:

Accessible at address 0x01C0

- When IRS_IDR2.LPI == 0, access on this interface is **RAZ/WI**.
- When IRS_IST_STATUSR.IDLE == 0, access on this interface is **UNKNOWN/WI**.
- When IRS_IST_BASER.VALID == 0, access on this interface is **UNKNOWN/WI**.
- When IRS_IST_CFGR.STRUCTURE == 0, access on this interface is **UNKNOWN/WI**.
- Otherwise, access on this interface is **WO**.

10.2.1.17 IRS_MEC_IDR

The IRS_MEC_IDR characteristics are:

Purpose

IRS MEC identification register. Contains read-only fields with information about the IRS support for MEC.


Attributes

IRS_MEC_IDR is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_MEC_IDR bit assignments are:



Bits [31:4]

Reserved, RES0.

MECIDSIZE, bits [3:0]

When *IRS_IDR0.MEC* == 1:

The number of bits minus one of MECID supported by the IRS.

The maximum permitted value is 0xF which indicates a MECID width of 16 bits.

The value 0x0 is a valid encoding and indicates that one bit of MECID is supported.

Otherwise:

RES0

Accessing IRS_MEC_IDR

Accesses to this register use the following encodings:

Accessible at address 0x0340

- When *IRS_IDR0.INT_DOM* != 0b11, access on this interface is **RAZ/WI**.
- When *IRS_IDR0.MEC* != 1, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RO**.

10.2.1.18 IRS_MEC_MECID_R

The IRS_MEC_MECID_R characteristics are:

Purpose

IRS MEC MECID register for the Realm PAS.

Attributes

IRS_MEC_MECID_R is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_MEC_MECID_R bit assignments are:



Bits [31:16]

Reserved, RES0.

MECID, bits [15:0]

MECID for IRS access to Realm PA space for:

- Accesses to physical and virtual IST entries.
- Accesses to VM table entries and VPE table entrie.
- Accesses to VM descriptors and VPE descriptors.

Bits above the supported MECID size, indicated in IRS_MEC_IDR.MECIDSIZE are RES0.

If MECIDSIZE is less than 0xF, the IRS treats bits [15:MECIDSIZE+1] of this field as zero.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing IRS_MEC_MECID_R

Accesses to this register use the following encodings:

Accessible at address 0x0344

- When IRS_IDR0.INT_DOM != 0b11, access on this interface is **RAZ/WI**.
- When IRS_IDR0.MEC != 1, access on this interface is **RAZ/WI**.
- When IRS_IST_BASER.VALID == 1 or IRS_IST_STATUSR.IDLE == 0, access on this interface is **RO**.
- When IRS_VMT_BASER.VALID == 1 or IRS_VMT_STATUSR.IDLE == 0, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.2.1.19 IRS_MPAM_IDR

The IRS_MPAM_IDR characteristics are:

Purpose

IRS MPAM identification register. Contains read-only fields with information about the IRS support for MPAM.


Attributes

IRS_MPAM_IDR is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_MPAM_IDR bit assignments are:



Bits [31:25]

Reserved, RES0.

HAS_MPAM_SP, bit [24]

Whether the IRS has support for MPAM PARTID space selection for the Interrupt Domain.

If HAS_MPAM_SP is 1, the IRS uses the MPAM PARTID specified by IRS_MPAM_PARTID_R.MPAM_SP.

If HAS_MPAM_SP is 0, the following PARTID space is used for IRS accesses to memory:

- Accesses made for the Secure Interrupt Domain use the Secure PARTID space.
- Accesses made for the Non-secure Interrupt Domain use the Non-secure PARTID space.
- Accesses made for the EL3 Interrupt Domain use the Root or Secure PARTID space.
- Accesses made for the Realm Interrupt Domain use the Realm PARTID space.

The value of this field is the same across all Interrupt Domains for an IRS.

PMG_MAX, bits [23:16]

The maximum PMG value that is permitted to be used for the IRS for the Interrupt Domain.

The PMG bit width is defined as the bit position of the most significant 1 in PMG_MAX[7:0], plus one, or is defined as zero if PMG_MAX is zero.

For example, if PMG_MAX == 0x0f, the PMG bit width is 4.

This field is permitted to be zero-sized.

PARTID_MAX, bits [15:0]

The maximum PARTID value that is permitted to be used for the IRS for the Interrupt Domain.

The PARTID bit width is defined as the bit position of the most significant 1 in PARTID_MAX[15:0], plus one, or is defined as zero if PARTID_MAX is zero.

For example, if PARTID_MAX == 0x0034, the PARTID bit width is 6.

This field is permitted to be zero-sized, but Arm recommends that it is non-zero when MPAM is implemented.

Accessing IRS_MPAM_IDR

Accesses to this register use the following encodings:

Accessible at address 0x0380

- When IRS_IDR0.MPAM != 1, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RO**.

10.2.1.20 IRS_MPAM_PARTID_R

The IRS_MPAM_PARTID_R characteristics are:

Purpose

IRS MPAM PARTID and PMG register.


Attributes

IRS_MPAM_PARTID_R is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_MPAM_PARTID_R bit assignments are:



IDLE, bit [31]

Whether the effects of the previous write to this register are complete.

IDLE	Meaning
0b0	The effects of the previous write to this register are not complete.
0b1	The effects of the previous write to this register are complete.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b1.

Access to this field is **RO**.

Bits [30:26]

Reserved, RES0.

MPAM_SP, bits [25:24]

When *IRS_MPAM_IDR.HAS_MPAM_SP == 1* and *IRS_IDR0.INT_DOM == 0b00*

MPAM_SP, bits [1:0] of bits [25:24]

MPAM PARTID space for the IRS for the Secure Interrupt Domain.

MPAM_SP	Meaning
0b00	Secure PARTID space.
0b01	Non-secure PARTID space.

Values not defined above are reserved.

Programming a reserved value results in the IRS using an UNKNOWN PARTID space.

The reset behavior of this field is:

- On a GIC reset, this field resets to an IMPLEMENTATION DEFINED value.

Accessing this field has the following behavior:

- When `IRS_MPAM_PARTID_R.IDLE == 0`, access to this field is **RO**
- Otherwise, access to this field is **RW**

When `IRS_MPAM_IDR.HAS_MPAM_SP == 1` and `IRS_IDR0.INT_DOM == 0b01`

MPAM_SP, bits [1:0] of bits [25:24]

MPAM PARTID space for the IRS for the Non-secure Interrupt Domain.

MPAM_SP	Meaning
0b01	Non-secure PARTID space.

Values not defined above are reserved.

Programming a reserved value results in the IRS using an UNKNOWN PARTID space.

The reset behavior of this field is:

- On a GIC reset, this field resets to an IMPLEMENTATION DEFINED value.

Accessing this field has the following behavior:

- When `IRS_MPAM_PARTID_R.IDLE == 0`, access to this field is **RO**
- Otherwise, access to this field is **RW**

When `IRS_MPAM_IDR.HAS_MPAM_SP == 1` and `IRS_IDR0.INT_DOM == 0b10`

MPAM_SP, bits [1:0] of bits [25:24]

MPAM PARTID space for the IRS for the EL3 Interrupt Domain.

MPAM_SP	Meaning
0b00	Secure PARTID space.
0b01	Non-secure PARTID space.
0b10	Root PARTID space.
0b11	Realm PARTID space.

Values not defined above are reserved.

Programming a reserved value results in the IRS using an UNKNOWN PARTID space.

The reset behavior of this field is:

- On a GIC reset, this field resets to an IMPLEMENTATION DEFINED value.

Accessing this field has the following behavior:

- When `IRS_MPAM_PARTID_R.IDLE == 0`, access to this field is **RO**
- Otherwise, access to this field is **RW**

When `IRS_MPAM_IDR.HAS_MPAM_SP == 1` and `IRS_IDR0.INT_DOM == 0b11`

MPAM_SP, bits [1:0] of bits [25:24]

MPAM PARTID space for the IRS for the Realm Interrupt Domain.

MPAM_SP	Meaning
0b01	Non-secure PARTID space.
0b11	Realm PARTID space.

Values not defined above are reserved.

Programming a reserved value results in the IRS using an UNKNOWN PARTID space.

The reset behavior of this field is:

- On a GIC reset, this field resets to an IMPLEMENTATION DEFINED value.

Accessing this field has the following behavior:

- When IRS_MPAM_PARTID_R.IDLE == 0, access to this field is **RO**
- Otherwise, access to this field is **RW**

Otherwise:

RES0

PMG, bits [23:16]

PMG for accesses to memory by the IRS for the Interrupt Domain.

Bits above the supported PMG bit width, as indicated by IRS_MPAM_IDR.PMG_MAX, are RES0.

If a value greater than IRS_MPAM_IDR.PMG_MAX is programmed, an UNKNOWN PMG is used.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0x00.

Accessing this field has the following behavior:

- When IRS_MPAM_PARTID_R.IDLE == 0, access to this field is **RO**
- Otherwise, access to this field is **RW**

PARTID, bits [15:0]

PARTID for accesses to memory by the IRS for the Interrupt Domain.

Bits above the supported PARTID bit width, as indicated by IRS_MPAM_IDR.PARTID_MAX, are RES0.

If a value greater than IRS_MPAM_IDR.PARTID_MAX is programmed, an UNKNOWN PARTID is used.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0x0000.

Accessing this field has the following behavior:

- When IRS_MPAM_PARTID_R.IDLE == 0, access to this field is **RO**
- Otherwise, access to this field is **RW**

Accessing IRS_MPAM_PARTID_R

Accesses to this register use the following encodings:

Accessible at address 0x0384

- When IRS_IDR0.MPAM != 1, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RW**.

10.2.1.21 IRS_PE_CR0

The IRS_PE_CR0 characteristics are:

Purpose

IRS PE control register 0.

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_PE_STATUSR.IDLE is 1.

Following a write to this register, IRS_PE_STATUSR.V is updated.


Attributes

IRS_PE_CR0 is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_PE_CR0 bit assignments are:



The fields in this register provide access to the configuration of the PE specified in the last write to IRS_PE_SELRL. The PE configuration applies to the IRS Domain corresponding to the PAS that this register is accessed in.

Bits [31:1]

Reserved, RES0.

DPS, bit [0]

Disable 1ofN PE selection.

DPS	Meaning
0b0	1ofN PE selection is enabled. An interrupt configured to use the 1ofN Routing mode is permitted to select this PE.
0b1	1ofN PE selection is disabled. An interrupt configured to use the 1ofN Routing mode is not permitted to select this PE.

This field determines whether the selected PE is permitted to be selected for 1ofN interrupt delivery.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- When IRS_IDR0.ONE_N == 0, access to this field is **RES0**
- Otherwise, access to this field is **RW**

Accessing IRS_PE_CR0

Accesses to this register use the following encodings:

Accessible at address 0x0148

- When IRS_PE_STATUSR.[V,IDLE] != 0b11, access on this interface is **UNKNOWN/WI**.
- Otherwise, access on this interface is **RW**.

10.2.1.22 IRS_PE_SEL

The IRS_PE_SEL characteristics are:

Purpose

IRS PE selection register.

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_PE_STATUSR.IDLE is 1.


Attributes

IRS_PE_SEL is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_PE_SEL bit assignments are:



Bits [31:16]

Reserved, RES0.

IAFFID, bits [15:0]

PE interrupt Affinity ID.

Selects a PE whose configuration can be accessed via IRS_PE_CR0.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing IRS_PE_SEL

Accesses to this register use the following encodings:

Accessible at address 0x0140

- When IRS_PE_STATUSR.IDLE == 0, access on this interface is **WI**.
- Otherwise, access on this interface is **WO**.

10.2.1.23 IRS_PE_STATUSR

The IRS_PE_STATUSR characteristics are:

Purpose

IRS PE status register


Attributes

IRS_PE_STATUSR is a 32-bit register.

This register is part of the IRS_CONFIG_FRAME block.

Field descriptions

The IRS_PE_STATUSR bit assignments are:



The fields in this register return information about the PE specified in the last write to IRS_PE_SELRL and the status of the last write to IRS_PE_CR0.

Bits [31:3]

Reserved, RES0.

ONLINE, bit [2]

Whether the PE is online or offline.

When the IRS determines that there is a candidate HPPI for the PE and the PE is offline the IRS generates a Wake Request to the PE.

When the IRS determines that there is a candidate HPPI for the PE and the PE is online the IRS Forwards the candidate HPPI to the PE.

When {V, IDLE} is not '0b11', the value of this field is UNKNOWN.

ONLINE	Meaning
0b0	The PE is offline.
0b1	The PE is online.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

V, bit [1]

Whether the value last written to IRS_PE_SELRL selected a valid PE.

When IDLE is 0, the value of this field is UNKNOWN.

V	Meaning
0b0	The PE selected using IRS_PE_SELRL is not valid.
0b1	The PE selected using IRS_PE_SELRL is valid.

This field resets to 0 to indicate that there was no write to IRS_PE_SEL_R that selected a valid PE since the IRS was reset.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b0.

IDLE, bit [0]

Following a write to IRS_PE_SEL_R, when this field returns 1 and V is 1, a read from IRS_PE_CR0 returns the configuration value for the PE when the accesses are performed on the same IRS.

Following a write to IRS_PE_CR0, when this field returns 1, the effects of the write have completed.

IDLE	Meaning
0b0	The effects of writing to IRS_PE_SEL_R and IRS_PE_CR0 are not guaranteed to be complete.
0b1	The effects of writing to IRS_PE_SEL_R and IRS_PE_CR0 are complete.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b1.

Accessing IRS_PE_STATUSR

Accesses to this register use the following encodings:

Accessible at address 0x0144

Access on this interface is **RO**.

10.2.1.24 IRS_SAVE_VMR

The IRS_SAVE_VMR characteristics are:

Purpose

IRS save VM register

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_SAVE_VM_STATUSR.IDLE is 1.


Attributes

IRS_SAVE_VMR is a 64-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_SAVE_VMR bit assignments are:



S, bit [63]

Writing 1 to this field saves the state of virtual interrupts to the virtual ISTs for the VM specified by VM_ID.

S	Meaning
0b0	The write has no effect on the virtual ISTs.
0b1	The write saves the state of virtual interrupts to the virtual ISTs.

Access to this field is **WO/UNKNOWN**.

Q, bit [62]

Writing 1 to this field queries whether the VM specified by VM_ID is Quiesced on all IRSs since the last write that set S to 1 and returns the result in IRS_SAVE_VM_STATUSR.Q when IRS_SAVE_VM_STATUSR.IDLE is 1.

When a write sets S to 1, the Effective value written to this field is 1.

Following a write that updates VM_ID, if there has been no write that set S to 1, the value returned in IRS_SAVE_VM_STATUSR.Q is UNKNOWN.

Q	Meaning
0b0	The write has no effect on the value returned in IRS_SAVE_VM_STATUSR.Q.
0b1	The write queries whether the VM is Quiesced on all IRSs since the last write that set S to 1.

Access to this field is **WO/UNKNOWN**.

Bits [61:16]

Reserved, RES0.

VM_ID, bits [15:0]

The VM ID specifying the VM whose virtual interrupt state should be written to the ISTs.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing IRS_SAVE_VMR

Accesses to this register use the following encodings:

Accessible at address 0x0300

- When IRS_IDR0.VIRT == 0, access on this interface is **RAZ/WI**.
- When IRS_SAVE_VM_STATUSR.IDLE == 0, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

This field resets to 1 to allow initial writes to registers Guarded by this field. Because there has been no write at reset, this does not imply that any invalidate operation is complete.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b1.

Accessing IRS_SAVE_VM_STATUSR

Accesses to this register use the following encodings:

Accessible at address 0x0308

- When IRS_IDR0.VIRT == 0, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RO**.

10.2.1.27 IRS_SPI_DOMAINR

The IRS_SPI_DOMAINR characteristics are:

Purpose

IRS SPI Interrupt Domain configuration Register.

Allows software to read and update the Interrupt Domain assignment of the SPI selected by IRS_SPI_SEL.

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_SPI_STATUSR.IDLE is 1.

Following a write to this register, IRS_SPI_STATUSR.V is updated.


Attributes

IRS_SPI_DOMAINR is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_SPI_DOMAINR bit assignments are:



This register is updated after a write to IRS_SPI_SEL when IRS_SPI_STATUSR.{V, IDLE} is {1, 1}.

Bits [31:2]

Reserved, RES0.

DOMAIN, bits [1:0]

Configures the Interrupt Domain associated with the SPI.

Some SPIs may be statically assigned to a Domain, in which case this field always returns the statically assigned Domain for the SPI.

To check if a SPI can be dynamically assigned to a Domain, software must read back the value in this field after attempting an update to establish if the update was successful once IRS_SPI_STATUSR.IDLE is 1.

DOMAIN	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

Programming an Interrupt Domain not supported by the IRS results in CONSTRAINED UNPREDICTABLE behavior with the following options:

- The SPI behaves as if it is not assigned to any Interrupt Domain.
- The SPI is treated as being assigned to another supported Interrupt Domain for all other purposes than reading back this field

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing IRS_SPI_DOMAINR

Accesses to this register use the following encodings:

Accessible at address 0x010C

- When IRS_IDR0.INT_DOM != 0b10, access on this interface is **RAZ/WI**.
- When IRS_SPI_STATUSR.[V,IDLE] != 0b11, access on this interface is **UNKNOWN/WI**.
- Otherwise, access on this interface is **RW**.

10.2.1.28 IRS_SPI_RESAMPLER

The IRS_SPI_RESAMPLER characteristics are:

Purpose

IRS SPI resample register.
Resample an SPI signal.

Attributes

IRS_SPI_RESAMPLER is a 32-bit register.
This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_SPI_RESAMPLER bit assignments are:



This register allows resampling an SPI.

Bits [31:24]

Reserved, RES0.

SPI_ID, bits [23:0]

SPI INTID.ID of the SPI to resample.

Following a write to this register, if all of the following are true, the SPI is resampled:

- The SPI is managed on this IRS.
- The SPI can be accessed in this IRS Domain.

Otherwise, the write to this register has no effect.

See ‘Physical SPIs’ for more information.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing IRS_SPI_RESAMPLER

Accesses to this register use the following encodings:

Accessible at address 0x0110

Access on this interface is **WO**.

10.2.1.29 IRS_SPI_SEL

The IRS_SPI_SEL characteristics are:

Purpose

IRS SPI selection register.

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_SPI_STATUSR.IDLE is 1.


Attributes

IRS_SPI_SEL is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_SPI_SEL bit assignments are:



Bits [31:24]

Reserved, RES0.

ID, bits [23:0]

Selects the SPI that the following registers access:

- IRS_SPI_CFGR
- IRS_SPI_DOMAINR
- IRS_SPI_VMR

Only implemented SPIs with INTID.ID in the range from IRS_IDR7.SPI_BASE to (IRS_IDR7.SPI_BASE + IRS_IDR6.SPI_IRS_RANGE - 1) may be selected.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing IRS_SPI_SEL

Accesses to this register use the following encodings:

Accessible at address 0x0108

- When IRS_SPI_STATUSR.IDLE == 0, access on this interface is **WI**.
- Otherwise, access on this interface is **WO**.

10.2.1.30 IRS_SPI_STATUSR

The IRS_SPI_STATUSR characteristics are:

Purpose

IRS SPI status register.


Attributes

IRS_SPI_STATUSR is a 32-bit register.

This register is part of the IRS_CONFIG_FRAME block.

Field descriptions

The IRS_SPI_STATUSR bit assignments are:



Bits [31:2]

Reserved, RES0.

V, bit [1]

Whether the value last written to IRS_SPI_SEL_R selects a valid SPI that is managed on this IRS and can be accessed in this IRS Domain.

This field is updated if and only if a write occurs to any of the following registers in this IRS Domain:

- IRS_SPI_CFG_R
- IRS_SPI_SEL_R
- IRS_SPI_VMR

Following a write to any of the above registers, if this field returns 0, the write to the register has no effects other than updating the values in this register.

SPIs with INTID.ID in the range described by IRS_IDR7.SPI_BASE and IRS_IDR6.SPI_IRS_RANGE are managed on this IRS.

If the in IRS_SPI_SEL_R specifies an SPI and all of the following are true, this field returns 1:

- The selected SPI is implemented.
- The selected SPI is managed by this IRS.
- The selected SPI is assigned to this IRS Domain or this IRS Domain is the EL3 IRS Domain.

Otherwise, this field returns 0.

When IDLE is 0, the value of this field is UNKNOWN.

V	Meaning
0b0	The SPI selected using IRS_SPI_SEL_R is not a valid SPI that is managed on this IRS and can be accessed in this IRS Domain.
0b1	The SPI selected using IRS_SPI_SEL_R is a valid SPI that is managed on this IRS and can be accessed in this IRS Domain.

This field resets to 0 to indicate that there was no write to IRS_SPI_SEL_R that selected a valid SPI that is managed on this IRS since the IRS was reset.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b0.

IDLE, bit [0]

Whether writes to any of the following registers have completed:

- IRS_SPI_CFGR
- IRS_SPI_DOMAINR
- IRS_SPI_SEL
- IRS_SPI_VMR

Following a write to IRS_SPI_SEL, when {IDLE, V} is {1, 1}, a read from any of the other registers returns the configuration value for the SPI.

Following a write to any of the registers listed above, on this IRS, when {IDLE, V} is {1, 1}, the effects of the write are complete.

IDLE	Meaning
0b0	The effects of writing to the SPI selection and configuration registers are not guaranteed to have completed.
0b1	The effects of writing to the SPI selection and configuration registers have completed.

This field resets to 1 to allow initial writes to registers Guarded by this field. Because there has been no write at reset, this does not imply that any invalidate operation is complete.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b1.

Accessing IRS_SPI_STATUSR

Accesses to this register use the following encodings:

Accessible at address 0x0118

Access on this interface is **RO**.

Accesses to this register use the following encodings:

Accessible at address 0x0100

- When `IRS_IDR0.VIRT == 0`, access on this interface is **RAZ/WI**.
- When `IRS_SPI_STATUSR.[V,IDLE] != 0b11`, access on this interface is **UNKNOWN/WI**.
- When `!IsSPIDVISupported(IRS_SPI_SEL.ID)`, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RW**.

10.2.1.32 IRS_SWERR_STATUSR

The IRS_SWERR_STATUSR characteristics are:

Purpose

IRS software error status register. Specifies whether a software error has been reported. If an error is reported, it contains syndrome information for the error.


Attributes

IRS_SWERR_STATUSR is a 64-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_SWERR_STATUSR bit assignments are:



Bits [63:32]

Reserved, RES0.

IMP_EC, bits [31:24]

IMPLEMENTATION DEFINED error code when IRS_SWERR_STATUSR.EC == 0.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- Access is **RES0** if any of the following are true:
 - IRS_SWERR_STATUSR.V == 0
 - IRS_SWERR_STATUSR.EC != 0
- Otherwise, access to this field is **RO**

EC, bits [23:16]

Specifies the error code that software can use to triage and handle the error.

EC	Meaning
0x00	An error was reported because of an IMPLEMENTATION DEFINED reason.
0x01	Failed lookup of L1_ISTE due to an external abort.
0x02	Failed lookup of L2_ISTE due to an external abort.
0x03	Failed lookup of L1_VMTE due to an external abort.
0x04	Failed lookup of L2_VMTE due to an external abort.
0x05	Failed lookup of VPE Table due to an external abort.
0x06	Failed lookup of VPE descriptor due to an external abort.
0x07	Failed lookup of VM descriptor due to an external abort.

EC	Meaning
0x08	A physical LPI could not be processed because IRS_IST_BASER.VALID is 0.
0x09	A physical LPI could not be processed because L1_ISTE.VALID is 0.
0x0A	A physical LPI was not processed because the LPI INTID.ID > $(2 \wedge \text{IRS_IST_CFGR.LPI_ID_BITS}) - 1$.
0x0B	A virtual interrupt could not be processed because L1_ISTE.VALID is 0
0x0C	A virtual interrupt could not be processed because L1_VMTE.VALID is 0.
0x0D	A virtual interrupt could not be processed because L2_VMTE.VALID is 0.
0x0E	A virtual LPI was not processed because L2_VMTE.LPI_IST_VALID is 0.
0x0F	A virtual SPI was not processed because L2_VMTE.SPI_IST_VALID is 0.
0x10	A virtual LPI was not processed because the LPI INTID.ID > $(2 \wedge \text{L2_VMTE.LPI_ID_BITS}) - 1$.
0x11	A virtual SPI was not processed because the SPI INTID.ID > $(2 \wedge \text{L2_VMTE.SPI_ID_BITS}) - 1$.
0x12	A virtual LPI signaled by an ITS was not processed because the VM ID > $(2 \wedge \text{IRS_VMT_CFGR.VM_ID_BITS}) - 1$.
0x13	A virtual interrupt was not processed because the VPE ID > $(2 \wedge \text{L2_VMTE.VPE_ID_BITS}) - 1$.
0x14	A virtual interrupt was signaled in a non-EL3 Interrupt Domain by an ITS, or via the GIC VDPEND system instruction, and the interrupt was unreachable because IRS_IDR0.VIRT is 0 or the VM was invalid.
0x15	A physical interrupt whose Routing mode is Targeted specified an invalid IAFFID.
0x16	A virtual interrupt whose Routing mode is Targeted specified a VPE ID that is < $2 \wedge \text{L2_VMTE.VPE_ID_BITS}$ but invalid otherwise.
0x17	The IRS was trying to signal a VPE doorbell but the VPE doorbell INTID is unreachable.
0x19	The IRS has detected corrupt metadata in the L2_ISTE.
0x1A	The IRS has detected corrupt metadata in a VPE descriptor.
0x1B	The IRS has detected corrupt metadata in a VM descriptor.
0x1C	IRS_MAP_L2_ISTR was written when IRS_IST_STATUSR.IDLE == 0.
0x1D	IRS_VMAP_L2_VMTR was written when IRS_VMT_STATUSR.IDLE == 0.

EC	Meaning
0x1E	IRS_VMAP_VMR was written when IRS_VMT_STATUSR.IDLE == 0.
0x1F	IRS_VMAP_VISTR was written when IRS_VMT_STATUSR.IDLE == 0.
0x20	IRS_VMAP_L2_VISTR was written when IRS_VMT_STATUSR.IDLE == 0.
0x21	IRS_VMAP_VPER was written when IRS_VMT_STATUSR.IDLE == 0.
0x22	IRS_IST_BASER was written when IRS_IST_STATUSR.IDLE == 0.
0x23	IRS_PE_SEL_R was written when IRS_PE_STATUSR.IDLE == 0.
0x24	IRS_SAVE_VMR was written when IRS_SAVE_VM_STATUSR.IDLE == 0.
0x25	IRS_SPI_SEL_R was written when IRS_SPI_STATUSR.IDLE == 0.
0x26	IRS_SYNC_R was written when IRS_SYNC_STATUSR.IDLE == 0.
0x27	IRS_VM_SEL_R was written when IRS_VM_STATUSR.IDLE == 0.
0x28	IRS_VMT_BASER was written when IRS_VMT_STATUSR.IDLE == 0.
0x29	IRS_VPE_SEL_R was written when IRS_VPE_STATUSR.IDLE == 0.
0x2A	The IRS has detected that L2_ISTE.IRM is 1 and 1ofN is not supported.
0x2B	A resample request for an SPI from a write to IRS_SPI_RESAMPLER failed.

All other values are reserved.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- When IRS_SWERR_STATUSR.V == 0, access to this field is **UNKNOWN/WI**
- Otherwise, access to this field is **RO**

Bits [15:4]

Reserved, RES0.

OF, bit [3]

Specifies whether multiple software errors have been detected.

When this field is 1, the syndrome information reports information about the error that last caused IRS_SWERR_STATUSR.V to transition from 0 to 1.

OF	Meaning
0b0	No errors have been detected, since the error that was reported when IRS_SWERR_STATUSR.V last transitioned from 0 to 1.
0b1	At least one error has been detected, since the error that was reported when IRS_SWERR_STATUSR.V last transitioned from 0 to 1.

When clearing IRS_SWERR_STATUSR.V to 0, if this field is nonzero, software writes 1 to clear this field to zero.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- When IRS_SWERR_STATUSR.V == 0, access to this field is **UNKNOWN/WI**
- Otherwise, access to this field is **W1C**

S1V, bit [2]

Specifies whether IRS_SWERR_SYNDROMER1 is valid.

S1V	Meaning
0b0	IRS_SWERR_SYNDROMER1 is not valid.
0b1	IRS_SWERR_SYNDROMER1 is valid.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- When IRS_SWERR_STATUSR.V == 0, access to this field is **UNKNOWN/WI**
- Otherwise, access to this field is **RO**

S0V, bit [1]

Specifies whether IRS_SWERR_SYNDROMER0 is valid.

S0V	Meaning
0b0	IRS_SWERR_SYNDROMER0 is not valid.
0b1	IRS_SWERR_SYNDROMER0 is valid.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- When IRS_SWERR_STATUSR.V == 0, access to this field is **UNKNOWN/WI**
- Otherwise, access to this field is **RO**

V, bit [0]

Specifies whether other fields in this register are valid and at least one software error has been reported.

V	Meaning
0b0	IRS_SWERR_STATUSR is not valid.
0b1	IRS_SWERR_STATUSR is valid.

Software writes 1 to this field to clear it to zero.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Access to this field is **W1C**.

Accessing IRS_SWERR_STATUSR

After reading IRS_SWERR_STATUSR, software clears the valid fields in the register to allow new errors to be reported.

However, between reading the register and clearing the valid fields, a new error might have overwritten the register.

To prevent this error being lost by software, the register prevents updates to fields that might have been updated by a new error.

This is done by ensuring a write to the register is ignored if all of the following are true:

- Any of IRS_SWERR_STATUSR.{V, OF} are nonzero before the write.
- The write does not clear the nonzero IRS_SWERR_STATUSR.{V, OF} fields to zero by writing ones to the applicable field or fields.

Accesses to this register use the following encodings:

Accessible at address 0x03C0

- When IRS_IDR0.SWE != 1, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RW**.

10.2.1.33 IRS_SWERR_SYNDROMER0

The IRS_SWERR_SYNDROMER0 characteristics are:

Purpose

IRS software error syndrome register 0. Records IRS specific software error syndrome information.


Attributes

IRS_SWERR_SYNDROMER0 is a 64-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_SWERR_SYNDROMER0 bit assignments are:



VIRTUAL, bit [63]

Specifies whether the error syndrome information is for a physical or virtual interrupt.

VIRTUAL	Meaning
0b0	The error syndrome information is for a physical interrupt
0b1	The error syndrome information is for a virtual interrupt

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

TYPE, bits [62:60]

The type of the interrupt for which an error was detected.

TYPE	Meaning
0b010	LPI
0b011	SPI

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Bits [59:56]

Reserved, RES0.

ID, bits [55:32]

ID of the interrupt for which an error was detected.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Bits [31:16]

Reserved, RES0.

VM_ID, bits [15:0]

Bits[15:0] of the VM ID of the virtual interrupt that could not be routed by the IRS as described below.

- For a virtual LPI, this is the VM ID specified in the incoming interrupt event.
- For a virtual SPI, this is the VM ID to which the virtual SPI was assigned.

When VIRTUAL is 0, this field is RES0.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing IRS_SWERR_SYNDROMER0

Accesses to this register use the following encodings:

Accessible at address 0x03C8

- When IRS_IDR0.SWE != 1, access on this interface is **RAZ/WI**.
- When IRS_SWERR_STATUSR.V == 1 and IRS_SWERR_STATUSR.S0V == 1, access on this interface is **RO**.
- Otherwise, access on this interface is **UNKNOWN/WI**.

10.2.1.34 IRS_SWERR_SYNDROMER1

The IRS_SWERR_SYNDROMER1 characteristics are:

Purpose

IRS software error syndrome register 1. Records IRS specific software error syndrome information.


Attributes

IRS_SWERR_SYNDROMER1 is a 64-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_SWERR_SYNDROMER1 bit assignments are:



Bits [63:56]

Reserved, RES0.

ADDR, bits [55:3]

Bits[55:3] of the physical address of a translation structure associated with the detected error.

The address in this field is associated with the PAS of the Interrupt Domain where the error is detected.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in IRS_IDR0.PA_RANGE.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Bits [2:0]

Reserved, RES0.

Accessing IRS_SWERR_SYNDROMER1

Accesses to this register use the following encodings:

Accessible at address 0x03D0

- When IRS_IDR0.SWE != 1, access on this interface is **RAZ/WI**.
- When IRS_SWERR_STATUSR.V == 1 and IRS_SWERR_STATUSR.S1V == 1, access on this interface is **RO**.
- Otherwise, access on this interface is **UNKNOWN/WI**.

10.2.1.35 IRS_SYNCR

The IRS_SYNCR characteristics are:

Purpose

IRS synchronize interrupt events register.

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_SYNC_STATUSR.IDLE is 1.


Attributes

IRS_SYNCR is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_SYNCR bit assignments are:



SYNC, bit [31]

Writing 1 to this field requests synchronization of interrupt events for the IRS Domain.

Writing 0 to this field has no effect.

See 'IRS synchronization requests' for more information.

SYNC	Meaning
0b0	Ignored.
0b1	Issue synchronization request.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Bits [30:0]

Reserved, RES0.

Accessing IRS_SYNCR

Accesses to this register use the following encodings:

Accessible at address 0x00C0

- When IRS_SYNC_STATUSR.IDLE == 0, access on this interface is **WI**.
- Otherwise, access on this interface is **WO**.

10.2.1.36 IRS_SYNC_STATUSR

The IRS_SYNC_STATUSR characteristics are:

Purpose

IRS synchronize interrupt events status register.


Attributes

IRS_SYNC_STATUSR is a 32-bit register.

This register is part of the IRS_CONFIG_FRAME block.

Field descriptions

The IRS_SYNC_STATUSR bit assignments are:



Bits [31:1]

Reserved, RES0.

IDLE, bit [0]

Whether the effects of the last write to IRS_SYNCR have completed.

IDLE	Meaning
0b0	The effects of writing to IRS_SYNCR not guaranteed to have completed.
0b1	The effects of writing to IRS_SYNCR have completed.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b1.

Access to this field is **RO**.

Accessing IRS_SYNC_STATUSR

This register is read-only.

Accesses to this register use the following encodings:

Accessible at address 0x00C4

Access on this interface is **RO**.

10.2.1.37 IRS_VM_DBR

The IRS_VM_DBR characteristics are:

Purpose

IRS VM 1ofN doorbell configuration register.

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_VM_STATUSR.IDLE is 1.

Following a write to this register, IRS_VM_STATUSR.V is updated


Attributes

IRS_VM_DBR is a 64-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_VM_DBR bit assignments are:



The fields in this register return information about the VM selected in IRS_VM_SEL R.

This register is updated after a write to IRS_VM_SEL R when IRS_VM_STATUSR.{V, IDLE} is {1, 1}.

EN, bit [63]

Whether the doorbell settings are valid for the VM.

EN	Meaning
0b0	1ofN doorbells are disabled for the VM.
0b1	1ofN doorbells are enabled for the VM.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Bits [62:16]

Reserved, RES0.

VPE_ID, bits [15:0]

The VPE ID of the 1ofN VPE doorbell target.

See ‘VPE doorbells for 1ofN interrupts’ for more information.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing IRS_VM_DBR

Accesses to this register use the following encodings:

Accessible at address 0x0280

- When IRS_IDR0.VIRT == 0 or IRS_IDR0.VIRT_ONE_N == 0, access on this interface is **RAZ/WI**.
- When IRS_VM_STATUSR.[V,IDLE] != 0b11, access on this interface is **UNKNOWN/WI**.
- Otherwise, access on this interface is **RW**.

10.2.1.38 IRS_VM_SEL

The IRS_VM_SEL characteristics are:

Purpose

IRS VM selection register.

Selects a VM whose 1ofN doorbell configuration can be accessed via IRS_VM_DBR.

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_VM_STATUSR.IDLE is 1.


Attributes

IRS_VM_SEL is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_VM_SEL bit assignments are:



Bits [31:16]

Reserved, RES0.

VM_ID, bits [15:0]

Identifies the VM.

Accessing IRS_VM_SEL

Accesses to this register use the following encodings:

Accessible at address 0x0288

- When IRS_IDR0.VIRT == 0 or IRS_IDR0.VIRT_ONE_N == 0, access on this interface is **RAZ/WI**.
- When IRS_VM_STATUSR.IDLE == 0, access on this interface is **WI**.
- Otherwise, access on this interface is **WO**.

This field resets to 1 to allow initial writes to registers Guarded by this field. Because there has been no write at reset, this does not imply that any invalidate operation is complete.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b1.

Accessing IRS_VM_STATUSR

Accesses to this register use the following encodings:

Accessible at address 0x028C

- When IRS_IDR0.VIRT == 0 or IRS_IDR0.VIRT_ONE_N == 0, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RO**.

10.2.1.40 IRS_VMAP_L2_VISTR

The IRS_VMAP_L2_VISTR characteristics are:

Purpose

IRS map level 2 virtual IST register

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_VMT_STATUSR.IDLE is 1.


Attributes

IRS_VMAP_L2_VISTR is a 64-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_VMAP_L2_VISTR bit assignments are:



M, bit [63]

Writing 1 to this field makes the virtual level 2 IST specified by TYPE and ID for the VM specified by VM_ID valid.

M	Meaning
0b0	The write to this field has no effect.
0b1	The write to this field makes the specified virtual level 2 IST valid.

There are no effects to the virtual IST on a write to this register, if any of the following are true:

- The VM specified by VM_ID is invalid.
- The virtual IST is invalid.
- The virtual IST uses a linear format.
- The INTID.ID is outside the configured virtual interrupt range for the specified VM.
- The level 2 IST is valid.
- The write to this register does not set this field to 1.

If this field is not set to 1 on a write to this register, updates to any other field in this register has no effect beyond updating the value of that field.

Access to this field is **WO/UNKNOWN**.

Bits [62:48]

Reserved, RES0.

VM_ID, bits [47:32]

The VM ID specifying the VM for which a virtual level 2 IST is being made valid.

The reset behavior of this field is:

- This field resets to an UNKNOWN value.

TYPE, bits [31:29]

Whether the operation applies to the virtual LPI IST or virtual SPI IST.

Values not defined are reserved.

When programming a reserved value, it is CONSTRAINED UNPREDICTABLE whether the invalidate is IGNORED or an UNKNOWN value is used for this field.

TYPE	Meaning
0b010	LPI
0b011	SPI

The reset behavior of this field is:

- This field resets to an UNKNOWN value.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

An INTID.ID covered by the level 1 ISTE corresponding to the level 2 IST that is made valid.

If unimplemented upper bits of the INTID.ID are not zero, it is CONSTRAINED UNPREDICTABLE whether the upper bits are treated as 0 or a request to invalidate is IGNORED.

The reset behavior of this field is:

- This field resets to an UNKNOWN value.

Accessing IRS_VMAP_L2_VISTR

Accesses to this register use the following encodings:

Accessible at address 0x02D8

- When IRS_IDR0.VIRT == 0, access on this interface is **RAZ/WI**.
- When IRS_VMT_STATUSR.IDLE == 0, access on this interface is **RO**.
- When IRS_VMT_BASER.VALID == 0, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RW**.

10.2.1.41 IRS_VMAP_L2_VMTR

The IRS_VMAP_L2_VMTR characteristics are:

Purpose

IRS map level 2 VM table register

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_VMT_STATUSR.IDLE is 1.


Attributes

IRS_VMAP_L2_VMTR is a 64-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_VMAP_L2_VMTR bit assignments are:



M, bit [63]

Writing 1 to this field makes the level 2 VM table specified by VM_ID valid.

M	Meaning
0b0	The write to this field has no effect.
0b1	The write to this field makes the specified level 2 VM table valid.

There are no effects to the VM table on a write to this register, if any of the following are true:

- The VM table is invalid.
- The VM table uses a linear structure.
- The VM ID specified by VM_ID is outside the configured VM ID range.
- The level 2 VM table is valid.
- The write to this register does not set this field to 1.

If this field is not set to 1 on a write to this register, updates to any other field in this register has no effect beyond updating the value of that field.

Access to this field is **WO/UNKNOWN**.

Bits [62:16]

Reserved, RES0.

VM_ID, bits [15:0]

The VM ID specifying the level 2 VM table being made valid.

Accessing IRS_VMAP_L2_VMTR

Accesses to this register use the following encodings:

Accessible at address 0x02C0

- When IRS_IDR0.VIRT == 0, access on this interface is **RAZ/WI**.
- When IRS_VMT_STATUSR.IDLE == 0, access on this interface is **RAZ/WI**.
- When IRS_VMT_CFGR.STRUCTURE == 0, access on this interface is **RAZ/WI**.
- When IRS_VMT_BASER.VALID == 0, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RW**.

10.2.1.42 IRS_VMAP_VISTR

The IRS_VMAP_VISTR characteristics are:

Purpose

IRS map virtual IST register

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_VMT_STATUSR.IDLE is 1.


Attributes

IRS_VMAP_VISTR is a 64-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_VMAP_VISTR bit assignments are:



M, bit [63]

Writing 1 to this field makes the virtual IST specified by TYPE for the VM specified by VM_ID valid or invalid as specified by U.

M	Meaning
0b0	The write to this field has no effect.
0b1	The write to this field makes the specified virtual IST valid.

There are no effects to the validity of the virtual IST on a write to this register, if any of the following are true:

- The VM specified by VM_ID is invalid.
- The VM ID is outside the configured VM ID range.
- The virtual IST is valid and U is 0.
- The virtual IST is invalid and U is 1.
- The write to this register does not set this field to 1.

If this field is not set to 1 on a write to this register, updates to any other field in this register has no effect beyond updating the value of that field.

Access to this field is **WO/UNKNOWN**.

U, bit [62]

Whether a write that sets M to 1 makes the specified virtual IST valid or invalid.

U	Meaning
0b0	A write that sets M to 1 makes the specified virtual IST valid.
0b1	A write that sets M to 1 makes the specified virtual IST invalid.

Bits [61:48]

Reserved, RES0.

VM_ID, bits [47:32]

The VM ID specifying the VM for which a virtual IST is being made valid.

The reset behavior of this field is:

- This field resets to an UNKNOWN value.

TYPE, bits [31:29]

Whether the operation applies to the virtual LPI IST or virtual SPI IST.

Values not defined are reserved.

When programming a reserved value, it is CONSTRAINED UNPREDICTABLE whether the invalidate is IGNORED or an UNKNOWN value is used for this field.

TYPE	Meaning
0b010	LPI
0b011	SPI

The reset behavior of this field is:

- This field resets to an UNKNOWN value.

Bits [28:0]

Reserved, RES0.

Accessing IRS_VMAP_VISTR

Accesses to this register use the following encodings:

Accessible at address 0x02D0

- When IRS_IDR0.VIRT == 0, access on this interface is **RAZ/WI**.
- When IRS_VMT_STATUSR.IDLE == 0, access on this interface is **RO**.
- When IRS_VMT_BASER.VALID == 0, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RW**.

10.2.1.43 IRS_VMAP_VMR

The IRS_VMAP_VMR characteristics are:

Purpose

IRS map VM register

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_VMT_STATUSR.IDLE is 1.


Attributes

IRS_VMAP_VMR is a 64-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_VMAP_VMR bit assignments are:



M, bit [63]

Writing 1 to this field makes the VM specified by VM_ID valid or invalid as specified by U.

M	Meaning
0b0	The write to this field has no effect.
0b1	The write to this field makes the specified VM valid or invalid.

There are no effects to the validity of the VM on a write to this register, if any of the following are true:

- The level 2 VM table containing the entry corresponding to the VM specified by VM_ID is invalid.
- The VM ID is outside the configured VM ID range.
- The VM valid and U is 0.
- The VM invalid and U is 1.
- The write to this register does not set this field to 1.

If this field is not set to 1 on a write to this register, updates to any other field in this register has no effect beyond updating the value of that field.

Access to this field is **WO/UNKNOWN**.

U, bit [62]

Whether a write that sets M to 1 makes the specified VM valid or invalid.

U	Meaning
0b0	A write that sets M to 1 makes the specified VM valid.
0b1	A write that sets M to 1 makes the specified VM invalid.

Bits [61:16]

Reserved, RES0.

VM_ID, bits [15:0]

The VM ID specifying the VM being made valid or invalid.

Accessing IRS_VMAP_VMR

Accesses to this register use the following encodings:

Accessible at address 0x02C8

- When IRS_IDR0.VIRT == 0, access on this interface is **RAZ/WI**.
- When IRS_VMT_STATUSR.IDLE == 0, access on this interface is **RAZ/WI**.
- When IRS_VMT_BASER.VALID == 0, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RW**.

10.2.1.44 IRS_VMAP_VPER

The IRS_VMAP_VPER characteristics are:

Purpose

IRS map VPE register

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_VMT_STATUSR.IDLE is 1.


Attributes

IRS_VMAP_VPER is a 64-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_VMAP_VPER bit assignments are:



M, bit [63]

Writing 1 to this field makes the VPE specified by VM_ID and VPE_ID valid.

M	Meaning
0b0	The write to this field has no effect.
0b1	The write to this field makes the specified VPE valid.

There are no effects to the validity of the VPE on a write to this register, if any of the following are true:

- The VM specified by VM_ID is invalid.
- The VM ID is outside the configured VM ID range.
- The VPE ID is outside the configured VPE range for the VM.
- The VPE is valid.
- The write to this register does not set this field to 1.

If this field is not set to 1 on a write to this register, updates to any other field in this register has no effect beyond updating the value of that field.

Access to this field is **WO/UNKNOWN**.

Bits [62:48]

Reserved, RES0.

VM_ID, bits [47:32]

Identifies the VM.

The reset behavior of this field is:

- This field resets to an UNKNOWN value.

Bits [31:16]

Reserved, RES0.

VPE_ID, bits [15:0]

Identifies the VPE.

The reset behavior of this field is:

- This field resets to an UNKNOWN value.

Accessing IRS_VMAP_VPER

Accesses to this register use the following encodings:

Accessible at address 0x02E0

- When IRS_IDR0.VIRT == 0, access on this interface is **RAZ/WI**.
- When IRS_VMT_STATUSR.IDLE == 0, access on this interface is **RO**.
- When IRS_VMT_BASER.VALID == 0, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RW**.

10.2.1.45 IRS_VMT_BASER

The IRS_VMT_BASER characteristics are:

Purpose

IRS VM table base address register

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_VMT_STATUSR.IDLE is 1.


Attributes

IRS_VMT_BASER is a 64-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_VMT_BASER bit assignments are:



Bits [63:56]

Reserved, RES0.

ADDR, bits [55:3]

Bits[55:3] of the VM table base physical address.

When IRS_VMT_CFGR.STRUCTURE is 0, ADDR points to a linear VM table and all of the following are true:

- Bits[N:0] of the resulting address are 0 where $N = 4 + \text{IRS_VMT_CFGR.VM_ID_BITS}$.
- This means that the level 2 VMTE array is aligned to the size of the array.

When IRS_VMT_CFGR.STRUCTURE is 1, ADDR points to the level 1 table in a 2-level VM table and all of the following are true:

- Bits[N:0] of the resulting address are 0 where N depends on VM_ID_BITS in IRS_VMT_CFGR as follows:
 $N = \text{Max}(2, \text{VM_ID_BITS} - 7 + 2)$
- This means that the level 1 VMT is aligned to the size of the level 1 VMTE array.

Access to any level of the VM table and any additional memory accesses occurring as a result of the address in this field are performed using the PAS of the IRS Domain where this register is accessed.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in IRS_IDR0.PA_RANGE.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- When IRS_VMT_BASER.VALID == 1, access to this field is **RO**
- Otherwise, access to this field is **RW**

Bits [2:1]

Reserved, RES0.

VALID, bit [0]

Whether the ADDR points to a valid VM table.

VALID	Meaning
0b0	The VM table address is not valid.
0b1	The VM table address is valid.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b0.

Accessing IRS_VMT_BASER

Accesses to this register use the following encodings:

Accessible at address 0x0200

- When IRS_IDR0.VIRT == 0, access on this interface is **RAZ/WI**.
- When IRS_VMT_STATUSR.IDLE == 0, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.2.1.46 IRS_VMT_CFGR

The IRS_VMT_CFGR characteristics are:

Purpose

IRS VM table configuration register


Attributes

IRS_VMT_CFGR is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_VMT_CFGR bit assignments are:



Bits [31:17]

Reserved, RES0.

STRUCTURE, bit [16]

Whether the VM table uses a linear or 2-level structure.

STRUCTURE	Meaning
0b0	A linear VM table structure is used.
0b1	A 2-level VM table structure is used.

If VM_ID_BITS < 8 and STRUCTURE is 1, all of the following are true:

- The IST consists of a single level 1 VM table entry and a single level 2 VM table.
- The level 2 VM table contains (2 ^ VM_ID_BITS) entries.
- The IRS is allowed to access the full 4KB level 2 VM table.

Arm recommends that STRUCTURE is 0 when VM_ID_BITS < 8.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Bits [15:5]

Reserved, RES0.

VM_ID_BITS, bits [4:0]

The number of VM_ID bits for the IRS Domain.

If this field is programmed to a value larger than the maximum, it is treated as having the maximum value for all other purposes than a direct read of the register. The maximum value is reported in IRS_IDR3.VM_ID_BITS.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing IRS_VMT_CFGR

Accesses to this register use the following encodings:

Accessible at address 0x0210

- When IRS_IDR0.VIRT == 0, access on this interface is **RAZ/WI**.
- When IRS_VMT_BASER.Valid == 1 or IRS_VMT_STATUSR.IDLE == 0, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.2.1.47 IRS_VMT_STATUSR

The IRS_VMT_STATUSR characteristics are:

Purpose

IRS virtualization data structures management status register.


Attributes

IRS_VMT_STATUSR is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_VMT_STATUSR bit assignments are:



Bits [31:1]

Reserved, RES0.

IDLE, bit [0]

Whether the effects of any of the following writes are complete:

- A write that updates IRS_VMT_BASER.Valid.
- A write that sets IRS_VMAP_L2_VMTR.M to 1.
- A write that sets IRS_VMAP_VMR.M to 1.
- A write that sets IRS_VMAP_VPER.M to 1.
- A write that sets IRS_VMAP_VISTR.M to 1.
- A write that sets IRS_VMAP_L2_VISTR.M to 1.

IDLE	Meaning
0b0	The effects a write to any of the virtualization data structure management registers are not complete
0b1	The effects a write to any of the virtualization data structure management registers are complete

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b1.

Access to this field is **RO**.

Accessing IRS_VMT_STATUSR

Accesses to this register use the following encodings:

Accessible at address 0x0214

- When IRS_IDR0.VIRT == 0, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RO**.

10.2.1.48 IRS_VPE_CR0

The IRS_VPE_CR0 characteristics are:

Purpose

IRS VPE Control Register 0

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_VPE_STATUSR.IDLE is 1.

Following a write to this register, IRS_VPE_STATUSR.V is updated.


Attributes

IRS_VPE_CR0 is a 32-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_VPE_CR0 bit assignments are:



This register is updated after a write that sets IRS_VPE_SEL.R.S to 1 when IRS_VPE_STATUSR.{V, IDLE} is {1, 1}.

Bits [31:1]

Reserved, RES0.

DPS, bit [0]

Disable 1ofN PE selection.

DPS	Meaning
0b0	1ofN PE selection is enabled. A virtual interrupt configured to use the 1ofN Routing mode is permitted to select this VPE.
0b1	1ofN PE selection is disabled. A virtual interrupt configured to use the 1ofN Routing mode is not permitted to select this VPE.

This field determines whether the selected VPE is permitted to be selected for 1ofN interrupt delivery.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing IRS_VPE_CR0

Accesses to this register use the following encodings:

Accessible at address 0x0258

- When IRS_IDR0.VIRT == 0, access on this interface is **RAZ/WI**.
- When IRS_VPE_STATUSR.[V,IDLE] != 0b11, access on this interface is **UNKNOWN/WI**.
- Otherwise, access on this interface is **RW**.

10.2.1.49 IRS_VPE_DBR

The IRS_VPE_DBR characteristics are:

Purpose

IRS VPE doorbell settings register.

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_VPE_STATUSR.IDLE is 1.

Following a write to this register, IRS_VPE_STATUSR.V is updated.


Attributes

IRS_VPE_DBR is a 64-bit register.

This register is part of the [IRS_CONFIG_FRAME](#) block.

Field descriptions

The IRS_VPE_DBR bit assignments are:



This register is updated after a write that sets IRS_VPE_SEL.R.S to 1 when IRS_VPE_STATUSR.{V, IDLE} is {1, 1}.

DBV, bit [63]

Whether the doorbell settings for the VPE are valid.

DBV	Meaning
0b0	The doorbell settings for the VPE are not valid.
0b1	The doorbell settings for the VPE are valid.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

REQ_DB, bit [62]

Whether a doorbell is requested for the VPE.

When a VPE doorbell is requested for the VPE, the doorbell event is generated when the VPE doorbell conditions are met.

See ‘VPE doorbells’ for more information.

REQ_DB	Meaning
0b0	A doorbell is not requested for the VPE
0b1	A doorbell is requested for the VPE.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- When `IRS_VPE_DBR.DBV == 1`, access to this field is **RW**
- Otherwise, access to this field is **WO/UNKNOWN**

Bits [61:37]

Reserved, RES0.

DBPM, bits [36:32]

Doorbell priority mask.

This field specifies the minimum priority for a virtual interrupt to trigger the VPE's doorbell.

Accessing this field has the following behavior:

- When `IRS_VPE_DBR.REQ_DB == 1`, access to this field is **RW**
- Otherwise, access to this field is **WO/UNKNOWN**

Bits [31:24]

Reserved, RES0.

INTID, bits [23:0]

If `ext-IRS_IDR2.LPI` is 1, this field specifies the LPI `INTID.ID` of the VPE doorbell.

The number of ID bits implemented is reported in `IRS_IDR2.ID_BITS`. Unimplemented upper bits are RES0.

If `ext-IRS_IDR2.LPI` is 0, this field is IMPLEMENTATION DEFINED.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- When `IRS_VPE_DBR.DBV == 1`, access to this field is **RW**
- Otherwise, access to this field is **WO/UNKNOWN**

Accessing IRS_VPE_DBR

Accesses to this register use the following encodings:

Accessible at address 0x0248

- When `IRS_IDR0.VIRT == 0`, access on this interface is **RAZ/WI**.
- When `IRS_VPE_STATUSR.[V,IDLE] != 0b11`, access on this interface is **UNKNOWN/WI**.
- Otherwise, access on this interface is **RW**.

10.2.1.50 IRS_VPE_HPPIR

The IRS_VPE_HPPIR characteristics are:

Purpose

IRS VPE HPPI register.


Attributes

IRS_VPE_HPPIR is a 64-bit register.

This register is part of the IRS_CONFIG_FRAME block.

Field descriptions

The IRS_VPE_HPPIR bit assignments are:



This register reports information about the candidate HPPI for the VPE selected by IRS_VPE_SEL R.

This register is updated after a write that sets IRS_VPE_SEL R.S to 1 when IRS_VPE_STATUSR.{V, IDLE} is {1, 1}.

If there is a change in the HPPI for the VPE following the write that set IRS_VPE_SEL R.S to 1, it is IMPLEMENTATION DEFINED whether the old or the new HPPI is reported in this register.

Bits [63:33]

Reserved, RES0.

HPPIV, bit [32]

Whether there is a candidate HPPI for the VPE.

HPPIV	Meaning
0b0	Invalid: There is no candidate HPPI for the VPE.
0b1	VALID: There is a candidate HPPI for the VPE.

When the value of this field is 1, ID and TYPE together form the INTID of the candidate HPPI for the VPE.

TYPE, bits [31:29]

The encoding of this field depends upon the value in HPPIV as described below:

- If HPPIV is 0, this field is RES0.
- If HPPIV is 1, this field contains valid information.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

The encoding of this field depends upon the value in HPPIV as described below:

- If HPPIV is 0, this field is RES0.
- If HPPIV is 1, this field contains valid information.

Accessing IRS_VPE_HPPIR

Accesses to this register use the following encodings:

Accessible at address 0x0250

- When IRS_IDR0.VIRT == 0, access on this interface is **RAZ/WI**.
- When IRS_VPE_STATUSR.[V,IDLE] != 0b11, access on this interface is **UNKNOWN/WI**.
- Otherwise, access on this interface is **RO**.

10.2.1.51 IRS_VPE_SEL

The IRS_VPE_SEL characteristics are:

Purpose

IRS VPE selection register.

Configuration

The effects of a write to this register are not guaranteed to have completed before IRS_VPE_STATUSR.IDLE is 1.


Attributes

IRS_VPE_SEL is a 64-bit register.

This register is part of the IRS_CONFIG_FRAME block.

Field descriptions

The IRS_VPE_SEL bit assignments are:



S, bit [63]

Writing 1 to this field selects a VPE whose configuration can be accessed via the following registers:

- IRS_VPE_DBR
- IRS_VPE_HPPIR
- IRS_VPE_CR0

S	Meaning
0b0	The write to this register has no effect.
0b1	The write to this register selects a VPE.

If this field is not set to 1 on a write to this register, updates to any other field in this register has no effect beyond updating the value of that field.

This means that if IRS_VPE_STATUSR.V is updated as a result of a write to another register, the value returned reflects the VPE selected when this field was set to 1.

Access to this field is **WO/UNKNOWN**.

Bits [62:48]

Reserved, RES0.

VPE_ID, bits [47:32]

Identifies the VPE.

The reset behavior of this field is:

- This field resets to an UNKNOWN value.

Bits [31:16]

Reserved, RES0.

VM_ID, bits [15:0]

Identifies the VM.

The reset behavior of this field is:

- This field resets to an UNKNOWN value.

Accessing IRS_VPE_SEL

Accesses to this register use the following encodings:

Accessible at address 0x0240

- When IRS_IDR0.VIRT == 0, access on this interface is **RAZ/WI**.
- When IRS_VPE_STATUSR.IDLE == 0, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.2.1.52 IRS_VPE_STATUSR

The IRS_VPE_STATUSR characteristics are:

Purpose


IRS VPE status register

Attributes

IRS_VPE_STATUSR is a 32-bit register.
This register is part of the IRS_CONFIG_FRAME block.

Field descriptions

The IRS_VPE_STATUSR bit assignments are:



Bits [31:2]

Reserved, RES0.

V, bit [1]

Whether the value last written to IRS_VPE_SEL R selects a valid VPE.

This field is updated if and only if any of following occur:

- A write to IRS_VPE_CR0.
- A write to IRS_VPE_DBR.
- A write that sets IRS_VPE_SEL R.S to 1.

When IDLE is 0, the value of this field is UNKNOWN.

V	Meaning
0b0	The VPE selected using IRS_VPE_SEL R is not valid.
0b1	The VPE selected using IRS_VPE_SEL R is valid.

This field resets to 0 to indicate that there was no write to IRS_VPE_SEL R that selected a valid VPE that is managed on this IRS since the IRS was reset.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b0.

IDLE, bit [0]

A write that sets IRS_VPE_SEL R.S to 1 is complete when this field is 1.

When this field is 1 and V is 1, a read from any of the following registers returns values for the selected VPE:

- IRS_VPE_DBR.
- IRS_VPE_HPPIR.
- IRS_VPE_CR0.

Following a write to one of the following registers, when this field returns 1, the effects of the write have completed:

- IRS_VPE_DBR.
- IRS_VPE_CR0.

IDLE	Meaning
0b0	The effects of the last write that set IRS_VPE_SEL.R.S to 1 and any write to VPE configuration registers are not guaranteed to be complete.
0b1	The effects of the last write that set IRS_VPE_SEL.R.S to 1 and any write to VPE configuration registers are complete.

This field resets to 1 to allow initial writes to registers Guarded by this field. Because there has been no write at reset, this does not imply that any invalidate operation is complete.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b1.

Accessing IRS_VPE_STATUSR

Accesses to this register use the following encodings:

Accessible at address 0x025C

- When IRS_IDR0.VIRT == 0, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RO**.

10.2.2 IRS_SETLPI_FRAME, IRS SETLPI register frame

The IRS_SETLPI_FRAME characteristics are:

Purpose

Contains the IRS_SETLPIR register used to generate a SET_EDGE message for an LPI without going through an ITS.

An IRS SETLPI register frame is present for each supported Interrupt Domain on each IRS where IRS_IDR0.SETLPI is 1.

Arm strongly recommends that this register frame is not accessible by PEs. This is because a write to this register can require the IRS access to memory, leading to in-out dependencies than can potentially lead to deadlocks in the system. If this register frame is not accessible by PEs, the behavior on an attempted access from a PE is IMPLEMENTATION DEFINED and is likely to result in an External abort.

This register frame is only accessible in the PAS associated with the Interrupt Domain.

The base address is distinct from addresses of registers accessible in any other PAS.

The base address is aligned to 64KB.

Configuration

This Register Block is present only when IRS_IDR0.SETLPI == 1.

Attributes

The IRS_SETLPI_FRAME block is of size 64KB

Default access

Accesses to the address space of this block that do not resolve to a register or a further block are treated as RAZ/WI.

Contents

Offset	Name	Notes
0x0000	IRS_SETLPIR	Most permissive access: WO

10.2.2.1 IRS_SETLPIR

The IRS_SETLPIR characteristics are:

Purpose

IRS SETLPI register.

A write to this register generates a SET_EDGE message for the LPI ID specified as part of the write.

Configuration

This register is present only when IRS_IDR0.SETLPI == 1. Otherwise, direct accesses to IRS_SETLPIR are RAZ/WI.


Attributes

IRS_SETLPIR is a 32-bit register.

This register is part of the [IRS_SETLPI_FRAME](#) block.

Field descriptions

The IRS_SETLPIR bit assignments are:



Bits [31:24]

Reserved, RES0.

ID, bits [23:0]

Bits[23:0] of the LPI INTID.ID to set generate a SET_EDGE message for.

If unimplemented upper bits of the INTID.ID are not zero, it is IMPLEMENTATION DEFINED whether the upper bits are treated as 0 or the interrupt message is ignored by the IRS.

Accessing IRS_SETLPIR

Writes to this register are ignored if the IRS is not enabled for the Interrupt Domain.

Accesses to this register use the following encodings:

Accessible at address 0x0000

Access on this interface is **WO**.

10.3 ITS register frames

$\mathbb{I}_{\text{DQKQD}}$

Each ITS Domain exposes the following two register frames:

- The ITS configuration register frame.
- The ITS translation register frame.

Different from GICv3, the GICv5 architecture does not require the ITS register frames to be contiguous with respect to each other.

$\mathbb{I}_{\text{BHGQG}}$

In the ITS register definitions, references to other registers are always to a register in the register frame for the same ITS Domain as where the access is made.

10.3.1 ITS_CONFIG_FRAME, ITS configuration register frame

The ITS_CONFIG_FRAME characteristics are:

Purpose

Contains control registers for an ITS Domain.

An ITS configuration register frame is present for each supported ITS Domain.

This register frame is accessible in the PAS associated with the ITS Domain.

It is IMPLEMENTATION DEFINED whether this register frame is also accessible in the MPPAS at the same address.

The base address is distinct from the base address of any other GIC register frame, including the ITS configuration register frames for other ITS Domains.

The base address is aligned to 64KB.

Attributes

The ITS_CONFIG_FRAME block is of size 64KB

Default access

Accesses to the address space of this block that do not resolve to a register or a further block are treated as RAZ/WI.

Contents

Offset	Name	Notes
0x0000	ITS_IDR0	Most permissive access: RO
0x0004	ITS_IDR1	Most permissive access: RO
0x0008	ITS_IDR2	Most permissive access: RO
0x0040	ITS_IIDR	Most permissive access: RO
0x0044	ITS_AIDR	Most permissive access: RO
0x0080	ITS_CR0	Most permissive access: RW
0x0084	ITS_CR1	Most permissive access: RW
0x00C0	ITS_DT_BASER	Most permissive access: RW
0x00D0	ITS_DT_CFGR	Most permissive access: RW
0x0100	ITS_DIDR	Most permissive access: RW
0x0108	ITS_EIDR	Most permissive access: RW
0x010C	ITS_INV_EVENTR	Most permissive access: WO
0x0110	ITS_INV_DEVICER	Most permissive access: WO
0x0114	ITS_READ_EVENTR	Most permissive access: WO
0x0118	ITS_READ_EVENT_DATAR	Most permissive access: RO
0x0120	ITS_STATUSR	Most permissive access: RO
0x0140	ITS_SYNCR	Most permissive access: WO
0x0148	ITS_SYNC_STATUSR	Most permissive access: RO
0x0180	ITS_GEN_EVENT_DIDR	Most permissive access: RW

Offset	Name	Notes
0x0188	ITS_GEN_EVENT_EIDR	Most permissive access: RW
0x018C	ITS_GEN_EVENTR	Most permissive access: WO
0x0190	ITS_GEN_EVENT_STATUSR	Most permissive access: RO
0x01C0	ITS_MEC_IDR	Most permissive access: RO
0x01C4	ITS_MEC_MECID_R	Most permissive access: RW
0x0200	ITS_MPAM_IDR	Most permissive access: RO
0x0204	ITS_MPAM_PARTID_R	Most permissive access: RW
0x0240	ITS_SWERR_STATUSR	Most permissive access: RW
0x0248	ITS_SWERR_SYNDROMER0	Most permissive access: RO
0x0250	ITS_SWERR_SYNDROMER1	Most permissive access: RO
0x0E00 + (4 * n) for n in ↪ 63:0	-	Most permissive access: ImplementationDefined

10.3.1.1 ITS_AIDR

The ITS_AIDR characteristics are:

Purpose

ITS Architecture Identification Register. Identifies the GIC architecture version to which the implementation conforms.


Attributes

ITS_AIDR is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_AIDR bit assignments are:



Bits [31:12]

Reserved, RES0.

Component, bits [11:8]

GIC component

Component	Meaning
0b0000	IRS
0b0001	ITS
0b0010	IWB

ArchMajorRev, bits [7:4]

Major Architecture revision.

ArchMajorRev	Meaning
0b0000	GICv5.x

ArchMinorRev, bits [3:0]

Minor Architecture revision.

ArchMinorRev	Meaning
0b0000	GICv5.0

Accessing ITS_AIDR

Accesses to this register use the following encodings:

Accessible at address 0x0044

Access on this interface is **RO**.

10.3.1.2 ITS_CR0

The ITS_CR0 characteristics are:

Purpose

ITS configuration register 0


Attributes

ITS_CR0 is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_CR0 bit assignments are:

**Bits [31:2]**

Reserved, RES0.

IDLE, bit [1]

Whether the transition between enabled and disabled states of the ITS Domain is complete.

IDLE	Meaning
0b0	The effects of updating ITSEN are not guaranteed to have completed.
0b1	The effects of updating ITSEN have completed.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b1.

Access to this field is **RO**.

ITSEN, bit [0]

Controls if the ITS Domain is enabled and whether it can generate interrupt messages to an IRS for the Interrupt Domain.

ITSEN	Meaning
0b0	Disabled. The ITS does not generate any interrupt messages for the Interrupt Domain
0b1	Enabled. The ITS may generate interrupt messages for the Interrupt Domain

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b0.

Accessing this field has the following behavior:

- When ITS_CR0.IDLE == 0, access to this field is **RO**
- Otherwise, access to this field is **RW**

Accessing ITS_CR0

Accesses to this register use the following encodings:

Accessible at address 0x0080

Access on this interface is **RW**.

10.3.1.3 ITS_CR1

The ITS_CR1 characteristics are:

Purpose

ITS configuration register 1


Attributes

ITS_CR1 is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_CR1 bit assignments are:



Bits [31:8]

Reserved, RES0.

ITT_RA, bit [7]

Read-Allocate hint for the interrupt translation tables.

ITT_RA	Meaning
0b0	No Read-Allocate.
0b1	Read-Allocate.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

DT_RA, bit [6]

Read-Allocate hint for the device table.

DT_RA	Meaning
0b0	No Read-Allocate.
0b1	Read-Allocate.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

IC, bits [5:4]

Controls the Inner Cacheability attribute used when the ITS accesses tables as a requester.

IC	Meaning
0b00	Non-cacheable.
0b01	Write-Back Cacheable.
0b10	Write-Through Cacheable.
0b11	Reserved, treated as 0b00.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

OC, bits [3:2]

Controls the Outer Cacheability attribute used when the ITS accesses tables as a requester.

OC	Meaning
0b00	Non-cacheable
0b01	Write-Back Cacheable.
0b10	Write-Through Cacheable.
0b11	Reserved, treated as 0b00.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

SH, bits [1:0]

Controls the Shareability attribute used when the ITS accesses tables as a requester.

SH	Meaning
0b00	Non-shareable.
0b01	Reserved, treated as 0b00.
0b10	Outer Shareable.
0b11	Inner Shareable.

When ITS_CR1.OC is 0b00 and ITS_CR1.IC is 0b00, this field is IGNORED and behaves as Outer Shareable.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing ITS_CR1

Accesses to this register use the following encodings:

Accessible at address 0x0084

- When ITS_CR0.ITSEN == 1 or ITS_CR0.IDLE == 0, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.3.1.4 ITS_DIDR

The ITS_DIDR characteristics are:

Purpose

ITS DeviceID Register.

This register is used to specify the DeviceID for the following requests:

- Invalidation of cached information for DeviceIDs issued by a write to ITS_INV_DEVICER.
- Invalidation of cached information for EventIDs issued by a write to ITS_INV_EVENTR.
- Requesting the read of the translation for an event by a write to ITS_READ_EVENTR.


Attributes

ITS_DIDR is a 64-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_DIDR bit assignments are:



Bits [63:32]

Reserved, RES0.

DEVICE_ID, bits [31:0]

The DeviceID as it should apply to a cache invalidation or reading an event translation request.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing ITS_DIDR

Accesses to this register use the following encodings:

Accessible at address 0x0100

- When ITS_CR0.[IDLE,ITSEN] != 0b11 or ITS_STATUSR.IDLE != 1, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.3.1.5 ITS_DT_BASER

The ITS_DT_BASER characteristics are:

Purpose

ITS device table base address register


Attributes

ITS_DT_BASER is a 64-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_DT_BASER bit assignments are:



Bits [63:56]

Reserved, RES0.

ADDR, bits [55:3]

Bits[55:3] of the DT base physical address.

When ITS_DT_CFGR.STRUCTURE is 0, ADDR points to a linear DT and all of the following are true:

- Bits[N:0] of the resulting address are 0 where $N = 2 + \text{ITS_DT_CFGR.DEVICEID_BITS}$.
- This means that the level 2 DTE array is aligned to the size of the array.

When ITS_DT_CFGR.STRUCTURE is 1, ADDR points to the level 1 table in a 2-level DT and all of the following are true:

- Bits[N:0] of the resulting address are 0 where N depends on L2SZ and DEVICEID_BITS in ITS_DT_CFGR as follows:
$$\text{Max}(2, \text{DEVICEID_BITS} - (9 + (2 * \text{L2SZ})) + 2)$$
- This means that the level 1 DT is aligned to the size of the table.

Access to any level of the DT and any additional memory accesses occurring as a result of the address in this field are performed using the PAS of the ITS Domain where this register is accessed.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in ITS_IDR0.PA_RANGE.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Bits [2:0]

Reserved, RES0.

Accessing ITS_DT_BASER

Accesses to this register use the following encodings:

Accessible at address 0x00C0

- When ITS_CR0.ITSEN == 1 or ITS_CR0.IDLE == 0, access on this interface is **RO**.

- Otherwise, access on this interface is **RW**.

10.3.1.6 ITS_DT_CFGR

The ITS_DT_CFGR characteristics are:

Purpose

ITS device table base address configuration register


Attributes

ITS_DT_CFGR is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_DT_CFGR bit assignments are:



Bits [31:17]

Reserved, RES0.

STRUCTURE, bit [16]

Whether the device table uses a linear or 2-level structure.

If ITS_IDR1.DT_LEVELS is 0, this field is RES0.

STRUCTURE	Meaning
0b0	A linear device table structure is used.
0b1	A 2-level device table structure is used.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Bits [15:8]

Reserved, RES0.

L2SZ, bits [7:6]

Level 2 device table size when a 2-level device table structure is used.

L2SZ	Meaning
0b00	A level 2 device table is maximum 4KB and resolves 9 bits of DeviceID.
0b01	A level 2 device table is maximum 16KB and resolves 11 bits of DeviceID.
0b10	A level 2 device table is maximum 64KB and resolves 13 bits of DeviceID.

Values not defined above are reserved.

If `DEVICEID_BITS` $\leq (9 + (2 * L2SZ))$ and `STRUCTURE` is 1, the DT consists of a single `L1_DTE` and a single `L2_DT`.

The `L2_DT` contains $(2 ^ \text{DEVICEID_BITS})$ entries.

When the value of `L1_DTE.SPAN` is programmed to a value larger than `DEVICEID_BITS`, the ITS is allowed to access memory in the range specified by `L1_DTE.SPAN`.

Arm recommends that `STRUCTURE` is 0 when `DEVICEID_BITS` $\leq (9 + (2 * L2SZ))$.

When programming a reserved value or an unsupported value, the ITS behavior is **CONSTRAINED UNPREDICTABLE** to any behavior which could be achieved by programming a valid and supported value.

When `STRUCTURE` is 0, this field is `RES0`.

The reset behavior of this field is:

- On a GIC reset, this field resets to an **UNKNOWN** value.

DEVICEID_BITS, bits [5:0]

The number of DeviceID bits which can be translated for the accessing Security state by the ITS.

If this field is programmed to a value larger than the maximum, it is treated as having the maximum value for all other purposes than a direct read of the register. The maximum value is reported in `ITS_IDR1.DEVICEID_BITS`.

The reset behavior of this field is:

- On a GIC reset, this field resets to an **UNKNOWN** value.

Accessing ITS_DT_CFGR

Accesses to this register use the following encodings:

Accessible at address 0x00D0

- When `ITS_CR0.ITSEN` == 1 or `ITS_CR0.IDLE` == 0, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.3.1.7 ITS_EIDR

The ITS_EIDR characteristics are:

Purpose

ITS EventID Register.

This register is used to specify the EventID for the following requests:

- Invalidation of cached information for EventIDs issued by a write to ITS_INV_EVENTR.
- Requesting the read of the configuration for an event by a write to ITS_READ_EVENTR.


Attributes

ITS_EIDR is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_EIDR bit assignments are:



Bits [31:16]

Reserved, RES0.

EVENT_ID, bits [15:0]

The EventID as it should apply to a cache invalidation or reading an event translation request.

The reset behavior of this field is:

- This field resets to an UNKNOWN value.

Accessing ITS_EIDR

Accesses to this register use the following encodings:

Accessible at address 0x0108

- When ITS_CR0.[IDLE,ITSEN] != 0b11 or ITS_STATUSR.IDLE != 1, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.3.1.8 ITS_GEN_EVENTR

The ITS_GEN_EVENTR characteristics are:

Purpose

ITS generate incoming event register.

This register is used to generate a SET_EDGE event for the EventID specified in ITS_GEN_EVENT_EIDR and the DeviceID specified in ITS_GEN_EVENT_DIDR.

Configuration

The effects of a write to this register are not guaranteed to have completed before ITS_GEN_EVENT_STATUSR.IDLE is 1.


Attributes

ITS_GEN_EVENTR is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_GEN_EVENTR bit assignments are:



R, bit [31]

Request the ITS to generate an incoming event.

R	Meaning
0b0	The write has no effect on the ITS.
0b1	Generate an incoming event with the following information: <ul style="list-style-type: none">• Event is associated with the Interrupt Domain associated with the PAS of the write.• The event is translated in the ITS Domain specified in TARGET_DOMAIN.• DeviceID is specified in the ITS_GEN_EVENT_DIDR register.• EventID is specified in the ITS_GEN_EVENT_EIDR register.

Bits [30:2]

Reserved, RES0.

TARGET_DOMAIN, bits [1:0]

Specifies the ITS Domain in which the incoming event is translated.

If the specified ITS Domain is not enabled, the write to this register does not generate an incoming event.

TARGET_DOMAIN	Meaning	Applies
0b00	The incoming event is translated in the ITS Domain specified by ITS_IDR0.INT_DOM.	
0b01	The incoming event is translated in the Realm ITS Domain.	When ITS_IDR0.INT_DOM == 0b01

Values not defined above are reserved.

Values corresponding to unimplemented Domains are reserved.

If a reserved value is programmed, the ITS behavior is CONSTRAINED UNPREDICTABLE to any behavior which could be achieved by programming a valid value.

Accessing ITS_GEN_EVENTR

Accesses to this register use the following encodings:

Accessible at address 0x018C

- When ITS_CR0.[IDLE,ITSEN] != 0b11 or ITS_GEN_EVENT_STATUSR.IDLE != 1, access on this interface is **WI**.
- Otherwise, access on this interface is **WO**.

10.3.1.9 ITS_GEN_EVENT_STATUSR

The ITS_GEN_EVENT_STATUSR characteristics are:

Purpose


- ITS generate incoming event status register.
- Reports the status of a request to generate an incoming event via a write to the ITS_GEN_EVENTNR register.

Attributes

- ITS_GEN_EVENT_STATUSR is a 32-bit register.
- This register is part of the ITS_CONFIG_FRAME block.

Field descriptions

The ITS_GEN_EVENT_STATUSR bit assignments are:



Bits [31:1]

Reserved, RES0.

IDLE, bit [0]

Reports whether the effects of the last write to ITS_GEN_EVENTNR are complete.

IDLE	Meaning
0b0	The effects of the last write to ITS_GEN_EVENTNR are not guaranteed to be complete.
0b1	The effects of the last write to ITS_GEN_EVENTNR are complete.

The reset behavior of this field is:

- This field resets to an UNKNOWN value.

Access to this field is **RO**.

Accessing ITS_GEN_EVENT_STATUSR

Accesses to this register use the following encodings:

Accessible at address 0x0190

Access on this interface is **RO**.

10.3.1.10 ITS_GEN_EVENT_EIDR

The ITS_GEN_EVENT_EIDR characteristics are:

Purpose

ITS generate incoming event EventID register.

This register is used to specify the EventID in a request to the ITS to generate an incoming event. The DeviceID of the event is specified in the ITS_GEN_EVENT_DIDR register.


Attributes

ITS_GEN_EVENT_EIDR is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_GEN_EVENT_EIDR bit assignments are:



Bits [31:16]

Reserved, RES0.

EVENT_ID, bits [15:0]

EventID in the ITS Domain containing the register.

The reset behavior of this field is:

- This field resets to an UNKNOWN value.

Accessing ITS_GEN_EVENT_EIDR

Accesses to this register use the following encodings:

Accessible at address 0x0188

- When ITS_CR0.[IDLE,ITSEN] != 0b11 or ITS_GEN_EVENT_STATUSR.IDLE != 1, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.3.1.11 ITS_GEN_EVENT_DIDR

The ITS_GEN_EVENT_DIDR characteristics are:

Purpose

ITS generate incoming event DeviceID register.

This register is used to specify the DeviceID in a request to the ITS to generate an incoming event. The EventID of the event is specified in the ITS_GEN_EVENT_EIDR register.


Attributes

ITS_GEN_EVENT_DIDR is a 64-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_GEN_EVENT_DIDR bit assignments are:



Bits [63:32]

Reserved, RES0.

DEVICE_ID, bits [31:0]

The DeviceID for the event generated as the result of a write to ITS_GEN_EVENTR.

The reset behavior of this field is:

- This field resets to an UNKNOWN value.

Accessing ITS_GEN_EVENT_DIDR

Accesses to this register use the following encodings:

Accessible at address 0x0180

- When ITS_CR0.[IDLE,ITSEN] != 0b11 or ITS_GEN_EVENT_STATUSR.IDLE != 1, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.3.1.12 ITS_IDR0

The ITS_IDR0 characteristics are:

Purpose

ITS identification register 0. Contains read-only fields with information about the ITS GIC component.


Attributes

ITS_IDR0 is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_IDR0 bit assignments are:



ITSID, bits [31:16]

Unique identifier for this ITS in the system.
This value is the same across all Interrupt Domains for this ITS.

Bits [15:9]

Reserved, RES0.

SWE, bit [8]

Software error reporting support in the Interrupt Domain.

SWE	Meaning
0b0	Software error reporting is not supported.
0b1	Software error reporting is supported.

Support for Software error reporting is optional.

MPAM, bit [7]

Memory Partitioning And Monitoring (MPAM) support.

MPAM	Meaning
0b0	MPAM is not supported.
0b1	MPAM is supported.

Support for MPAM is optional.

MEC, bit [6]

When ITS_IDR0.INT_DOM == 0b11:

Support for Memory Encryption Contexts (MEC) for the Realm ITS Domain.

MEC	Meaning
0b0	Memory Encryption Contexts are not supported.
0b1	Memory Encryption Contexts are supported.

Support for MEC is optional.

Otherwise:

RES0

PA_RANGE, bits [5:2]

Physical Address range supported.

The physical address range corresponds to the system physical address size.

The value of this field is the same for all ITS Domains across all ITSs in the system

PA_RANGE	Meaning
0b0000	32 bits, 4GB
0b0001	36 bits, 64GB
0b0010	40 bits, 1TB
0b0011	42 bits, 4TB
0b0100	44 bits, 16TB
0b0101	48 bits, 256TB
0b0110	52 bits, 4PB
0b0111	56 bits, 64PB

Values not defined above are reserved.

INT_DOM, bits [1:0]

The ITS Domain that the register frame containing this register controls.

INT_DOM	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

Accessing ITS_IDR0

Accesses to this register use the following encodings:

Accessible at address 0x0000

Access on this interface is **RO**.

10.3.1.13 ITS_IDR1

The ITS_IDR1 characteristics are:

Purpose

ITS identification register 1. Contains read-only fields with information about the ITS GIC component.


Attributes

ITS_IDR1 is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_IDR1 bit assignments are:



Bits [31:11]

Reserved, RES0.

L2SZ, bits [10:8]

Supported level 2 ITS table sizes when a 2-level table structure is used.

L2SZ	Meaning
0b1xx	Level 2 DT and ITT sizes supported: 64KiB
0bx1x	Level 2 DT and ITT sizes supported: 16KiB
0bxx1	Level 2 DT and ITT sizes supported: 4KiB

Values not defined above are reserved.

When both ITT_LEVELS and DT_LEVELS are 0, this field is RES0.

Arm strongly recommends that the ITS supports L2SZ value 0b111

ITT_LEVELS, bit [7]

Levels supported for the ITT.

ITT_LEVELS	Meaning
0b0	Only a linear ITT structure is supported
0b1	Both a 2-level and a linear ITT structure is supported.

DT_LEVELS, bit [6]

Levels supported for the device table.

DT_LEVELS	Meaning
0b0	Only a linear device table structure is supported
0b1	Both a 2-level and a linear device table structure is supported.

DEVICEID_BITS, bits [5:0]

The ITS can support translation of up to $(2^{\text{DEVICEID_BITS}})$ DeviceIDs.

The maximum permitted value of this field is 32.

Accessing ITS_IDR1

Accesses to this register use the following encodings:

Accessible at address 0x0004

Access on this interface is **RO**.

10.3.1.14 ITS_IDR2

The ITS_IDR2 characteristics are:

Purpose

ITS identification register 2. Contains read-only fields with information about how the ITS GIC component is implemented.

Attributes

ITS_IDR2 is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_IDR2 bit assignments are:



Bits [31:7]

Reserved, RES0.

XDMN_EVENTS, bits [6:5]

Indicates whether the ITS Domain supports translation of events associated with other Interrupt Domains.

XDMN_EVENTS	Meaning	Applies
0b00	The ITS Domain does not support translation of incoming events associated with other Interrupt Domains.	
0b01	The ITS Domain supports translation of incoming events associated with the Non-secure Interrupt Domain.	When ITS_IDR0.INT_DOM == 0b11

Values not defined above are reserved.

EVENTID_BITS, bits [4:0]

The ITS can support translation of up to $2^{EVENTID_BITS}$ EventIDs.

The maximum permitted value of this field is 16.

Accessing ITS_IDR2

Accesses to this register use the following encodings:

Accessible at address 0x0008

Access on this interface is **RO**.

10.3.1.15 ITS_IIDR

The ITS_IIDR characteristics are:

Purpose

ITS Implementer Identification Register. Provides information about the implementation and implementer of the GIC, and architecture version supported.

Attributes

ITS_IIDR is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_IIDR bit assignments are:

31	20	19	16	15	12	11	0
ProductID				Variant	Revision	Implementer	

ProductID, bits [31:20]

IMPLEMENTATION DEFINED value identifying the GIC part

When the ITS_PIDR{0,1} registers are present, Arm expects that the ITS_PIDR{0,1}.PART_{0,1} fields match the value of ITS_IIDR.ProductID.

If required, however, an implementation is permitted to provide values for ITS_PIDR.{0,1}.PART_{0,1} that do not match the value of ITS_IIDR.ProductID

Variant, bits [19:16]

IMPLEMENTATION DEFINED value used to distinguish product variants, or major revisions of the product

Revision, bits [15:12]

IMPLEMENTATION DEFINED value used to distinguish minor revisions of the product

Implementer, bits [11:0]

Contains the JEP106 code of the company that implemented the GIC

For an Arm implementation, the JEP106 code is 0x43B

When the ITS_PIDR{1,2,4} registers are present, Arm expects that the ITS_PIDR{0,1}.PART_{0,1} fields match the value of ITS_IIDR.Implementer.

Accessing ITS_IIDR

Accesses to this register use the following encodings:

Accessible at address 0x0040

Access on this interface is **RO**.

10.3.1.16 ITS_INV_DEVICER

The ITS_INV_DEVICER characteristics are:

Purpose

ITS cache invalidation register for a single device or a range of devices.

Configuration

The effects of a write to this register are not guaranteed to have completed before ITS_STATUSR.IDLE is 1.


Attributes

ITS_INV_DEVICER is a 32-bit register.

This register is part of the ITS_CONFIG_FRAME block.

Field descriptions

The ITS_INV_DEVICER bit assignments are:



I, bit [31]

Writing 1 to this field invalidates cached information from DTEs for the DeviceID specified in ITS_DIDR.DEVICE_ID.

I	Meaning
0b0	The write has no effect.
0b1	Invalidate cached information for the specified DeviceID.

Bits [30:6]

Reserved, RES0.

EVENTID_BITS, bits [5:1]

The range of EventIDs that the invalidation operation applies to for the specified DeviceID.

If this field is programmed to a value larger than the maximum, it is treated as having the maximum value for all other purposes than reading back the field.

The maximum value is reported in ITS_IDR2.EVENTID_BITS.

If ITS_INV_DEVICER.L1 is 1, the effective value of this field is the maximum value reported in ITS_IDR2.EVENTID_BITS.

L1, bit [0]

Controls which cached information the invalidation operation applies to.

When this field is 0, the operation invalidates cached information from the single level 2 DTE specified by ITS_DIDR.DEVICE_ID.

When this field is 1, the operation invalidates cached information from the level 1 DTE specified by ITS_DIDR.DEVICE_ID as well as any level 2 DTE covered by the range of DeviceIDs given by ITS_DT_CFGR.L2SZ for the DeviceID specified in ITS_DIDR.DEVICE_ID.

If ITS_DT_CFGR.STRUCTURE is 0, the effective value of this field is 0.

L1	Meaning
0b0	The cache invalidation operation invalidates information from the specified level 2 DTE.DEVICE_ID.
0b1	The cache invalidation operation invalidates information from the specified level 1 DTE and level 2 DTEs.

Accessing ITS_INV_DEVICER

Accesses to this register use the following encodings:

Accessible at address 0x0110

- When ITS_CR0.[IDLE,ITSEN] != 0b11 or ITS_STATUSR.IDLE != 1, access on this interface is **WI**.
- Otherwise, access on this interface is **WO**.

10.3.1.17 ITS_INV_EVENTR

The ITS_INV_EVENTR characteristics are:

Purpose

ITS cache invalidation register for a single event or a range of events.

Configuration

The effects of a write to this register are not guaranteed to have completed before ITS_STATUSR.IDLE is 1.


Attributes

ITS_INV_EVENTR is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_INV_EVENTR bit assignments are:



I, bit [31]

Writing 1 to this field invalidates cached information from ITTEs for the DeviceID and EventID specified in ITS_DIDR.DEVICE_ID and ITS_EIDR.EVENT_ID, respectively.

I	Meaning
0b0	The write has no effect.
0b1	Invalidate cached information for the specified DeviceID and EventID.

Bits [30:3]

Reserved, RES0.

ITT_L2SZ, bits [2:1]

The range of EventIDs that the invalidation operation applies to.

ITT_L2SZ	Meaning
0b00	The invalidate operation applies to cached information from any level 2 ITTE whose EventID bits [15:9] matches ITS_EIDR.EVENT_ID bits [15:9].
0b01	The invalidate operation applies to cached information from any level 2 ITTE whose EventID bits [15:11] matches ITS_EIDR.EVENT_ID bits [15:11].
0b10	The invalidate operation applies to cached information from any level 2 ITTE whose EventID bits [15:13] matches ITS_EIDR.EVENT_ID bits [15:13].

If ITS_INV_EVENTR.L1 is 0, this field is IGNORED.

Values not defined are reserved.

When programming a reserved value or an unsupported value, the ITS behavior is CONSTRAINED UNPREDICTABLE to any behavior which could be achieved by programming a valid and supported value.

L1, bit [0]

Controls which cached information the invalidation operation applies to.

When this field is 0, the operation invalidates cached information from the single level 2 ITTE specified by ITS_DIDR.DEVICE_ID and ITS_EIDR.EVENT_ID.

When this field is 1, the operation invalidates cached information from the level 1 ITTE specified by ITS_DIDR.DEVICE_ID and ITS_EIDR.EVENT_ID as well as any level 2 ITTE covered by the range of EventIDs given by ITS_EIDR.EVENT_ID and ITT_L2SZ for the DeviceID specified in ITS_DIDR.DEVICE_ID.

L1	Meaning
0b0	The cache invalidation operation invalidates information from the specified level 2 ITTE.
0b1	The cache invalidation operation invalidates information from the specified level 1 ITTE and level 2 ITTEs.

Accessing ITS_INV_EVENTR

Accesses to this register use the following encodings:

Accessible at address 0x010C

- When ITS_CR0.[IDLE,ITSEN] != 0b11 or ITS_STATUSR.IDLE != 1, access on this interface is **WI**.
- Otherwise, access on this interface is **WO**.

10.3.1.18 ITS_MEC_IDR

The ITS_MEC_IDR characteristics are:

Purpose

ITS MEC identification register. Contains read-only fields with information about the ITS support for MEC.


Attributes

ITS_MEC_IDR is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_MEC_IDR bit assignments are:

**Bits [31:4]**

Reserved, RES0.

MECIDSIZE, bits [3:0]

When ITS IDR0.MEC == 1:

The number of bits minus one of MECID supported by the ITS.

The maximum permitted value is 0xF which indicates a MECID width of 16 bits.

The value 0x0 is a valid encoding and indicates that one bit of MECID is supported.

Otherwise:

RES0

Accessing ITS_MEC_IDR

Accesses to this register use the following encodings:

Accessible at address 0x01C0

- When ITS_IDR0.INT_DOM != 0b11, access on this interface is **RAZ/WI**.
- When ITS_IDR0.MEC != 1, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RO**.

10.3.1.19 ITS_MEC_MECID_R

The ITS_MEC_MECID_R characteristics are:

Purpose

ITS MEC MECID register for the Realm PAS.


Attributes

ITS_MEC_MECID_R is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_MEC_MECID_R bit assignments are:



Bits [31:16]

Reserved, RES0.

MECID, bits [15:0]

MECID for ITS access to Realm PA space for:

- Fetches of device table entries.
- Fetches of interrupt translation table entries.

Bits above the supported MECID size, indicated in ITS_MEC_IDR.MECIDSIZE are RES0.

If MECIDSIZE is less than 0xF, the ITS treats bits [15:MECIDSIZE+1] of this field as zero.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing ITS_MEC_MECID_R

Accesses to this register use the following encodings:

Accessible at address 0x01C4

- When ITS_IDR0.INT_DOM != 0b11, access on this interface is **RAZ/WI**.
- When ITS_IDR0.MEC != 1, access on this interface is **RAZ/WI**.
- When ITS_CR0.ITSSEN == 1 or ITS_CR0.IDLE == 0, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.3.1.20 ITS_MPAM_IDR

The ITS_MPAM_IDR characteristics are:

Purpose

ITS MPAM identification register. Contains read-only fields with information about the ITS support for MPAM.


Attributes

ITS_MPAM_IDR is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_MPAM_IDR bit assignments are:



Bits [31:25]

Reserved, RES0.

HAS_MPAM_SP, bit [24]

Whether the ITS Domain has support for MPAM PARTID space selection.

If HAS_MPAM_SP is 1, the ITS uses the MPAM PARTID specified by ITS_MPAM_PARTID_R.MPAM_SP.

If HAS_MPAM_SP is 0, the following PARTID space is used for ITS accesses to memory:

- Accesses made by the Secure ITS Domain use the Secure PARTID space.
- Accesses made by the Non-secure ITS Domain use the Non-secure PARTID space.
- Accesses made by the EL3 ITS Domain use the Root or Secure PARTID space.
- Accesses made by the Realm ITS Domain use the Realm PARTID space.

The value of this field is the same across all ITS Domains for an ITS.

PMG_MAX, bits [23:16]

The maximum PMG value that is permitted to be used in the ITS Domain.

The PMG bit width is defined as the bit position of the most significant 1 in PMG_MAX[7:0], plus one, or is defined as zero if PMG_MAX is zero.

For example, if PMG_MAX == 0x0f, the PMG bit width is 4.

This field is permitted to be zero-sized.

PARTID_MAX, bits [15:0]

The maximum PARTID value that is permitted to be used in the ITS Domain.

The PARTID bit width is defined as the bit position of the most significant 1 in PARTID_MAX[15:0], plus one, or is defined as zero if PARTID_MAX is zero.

For example, if PARTID_MAX == 0x0034, the PARTID bit width is 6.

This field is permitted to be zero-sized, but Arm recommends that it is non-zero when MPAM is implemented.

Accessing ITS_MPAM_IDR

Accesses to this register use the following encodings:

Accessible at address 0x0200

- When ITS_IDR0.MPAM != 1, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RO**.

10.3.1.21 ITS_MPAM_PARTID_R

The ITS_MPAM_PARTID_R characteristics are:

Purpose

ITS MPAM PARTID and PMG register.


Attributes

ITS_MPAM_PARTID_R is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_MPAM_PARTID_R bit assignments are:



IDLE, bit [31]

Whether the effects of the previous write to this register are complete.

Following a write to this register, when this field is 1, the new values written to this register is guaranteed to be used for subsequent memory accesses by the ITS.

IDLE	Meaning
0b0	The effects of the previous write to this register are not complete.
0b1	The effects of the previous write to this register are complete.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b1.

Access to this field is **RO**.

Bits [30:26]

Reserved, RES0.

MPAM_SP, bits [25:24]

When *ITS_MPAM_IDR.HAS_MPAM_SP* == 1 and *ITS_IDR0.INT_DOM* == 0b00

MPAM_SP, bits [1:0] of bits [25:24]

MPAM PARTID space for the Secure ITS Domain.

MPAM_SP	Meaning
0b00	Secure PARTID space.
0b01	Non-secure PARTID space.

Values not defined above are reserved.

Programming a reserved value results in the ITS using an UNKNOWN PARTID space.

The reset behavior of this field is:

- On a GIC reset, this field resets to an IMPLEMENTATION DEFINED value.

Accessing this field has the following behavior:

- When ITS_MPAM_PARTID_R.IDLE == 0, access to this field is **RO**
- Otherwise, access to this field is **RW**

When ITS_MPAM_IDR.HAS_MPAM_SP == 1 and ITS_IDR0.INT_DOM == 0b01

MPAM_SP, bits [1:0] of bits [25:24]

MPAM PARTID space for the Non-secure ITS Domain.

MPAM_SP	Meaning
0b01	Non-secure PARTID space.

Values not defined above are reserved.

Programming a reserved value results in the ITS using an UNKNOWN PARTID space.

The reset behavior of this field is:

- On a GIC reset, this field resets to an IMPLEMENTATION DEFINED value.

Accessing this field has the following behavior:

- When ITS_MPAM_PARTID_R.IDLE == 0, access to this field is **RO**
- Otherwise, access to this field is **RW**

When ITS_MPAM_IDR.HAS_MPAM_SP == 1 and ITS_IDR0.INT_DOM == 0b10

MPAM_SP, bits [1:0] of bits [25:24]

MPAM PARTID space for the EL3 ITS Domain.

MPAM_SP	Meaning
0b00	Secure PARTID space.
0b01	Non-secure PARTID space.
0b10	Root PARTID space.
0b11	Realm PARTID space.

Values not defined above are reserved.

Programming a reserved value results in the ITS using an UNKNOWN PARTID space.

The reset behavior of this field is:

- On a GIC reset, this field resets to an IMPLEMENTATION DEFINED value.

Accessing this field has the following behavior:

- When ITS_MPAM_PARTID_R.IDLE == 0, access to this field is **RO**
- Otherwise, access to this field is **RW**

When *ITS_MPAM_IDR.HAS_MPAM_SP* == 1 and *ITS_IDR0.INT_DOM* == 0b11

MPAM_SP, bits [1:0] of bits [25:24]

MPAM PARTID space for the Realm ITS Domain.

MPAM_SP	Meaning
0b01	Non-secure PARTID space.
0b11	Realm PARTID space.

Values not defined above are reserved.

Programming a reserved value results in the ITS using an UNKNOWN PARTID space.

The reset behavior of this field is:

- On a GIC reset, this field resets to an IMPLEMENTATION DEFINED value.

Accessing this field has the following behavior:

- When *ITS_MPAM_PARTID_R.IDLE* == 0, access to this field is **RO**
- Otherwise, access to this field is **RW**

Otherwise:

RES0

PMG, bits [23:16]

PMG for accesses to memory by the ITS Domain.

Bits above the supported PMG bit width, as indicated by *ITS_MPAM_IDR.PMG_MAX*, are RES0.

If a value greater than *ITS_MPAM_IDR.PMG_MAX* is programmed, an UNKNOWN PMG is used.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0x00.

Accessing this field has the following behavior:

- When *ITS_MPAM_PARTID_R.IDLE* == 0, access to this field is **RO**
- Otherwise, access to this field is **RW**

PARTID, bits [15:0]

PARTID for accesses to memory by the ITS Domain.

Bits above the supported PARTID bit width, as indicated by *ITS_MPAM_IDR.PARTID_MAX*, are RES0.

If a value greater than *ITS_MPAM_IDR.PARTID_MAX* is programmed, an UNKNOWN PARTID is used.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0x0000.

Accessing this field has the following behavior:

- When *ITS_MPAM_PARTID_R.IDLE* == 0, access to this field is **RO**
- Otherwise, access to this field is **RW**

Accessing ITS_MPAM_PARTID_R

Accesses to this register use the following encodings:

Accessible at address 0x0204

- When ITS_IDR0.MPAM != 1, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RW**.

10.3.1.22 ITS_READ_EVENTR

The ITS_READ_EVENTR characteristics are:

Purpose

ITS read event request register.
This register is used to read the translation information for an EventID and DeviceID specified in ITS_EIDR and ITS_DIDR registers, respectively.

Configuration

The effects of a write to this register are not guaranteed to have completed before ITS_STATUSR.IDLE is 1.

Attributes

ITS_READ_EVENTR is a 32-bit register.
This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_READ_EVENTR bit assignments are:



R, bit [31]

Request ITS to read event translation information.

R	Meaning
0b0	The write has no effect on the ITS.
0b1	Read the translation information for the specified DeviceID and EventID into ITS_READ_EVENT_DATAR.

Bits [30:0]

Reserved, RES0.

Accessing ITS_READ_EVENTR

Accesses to this register use the following encodings:

Accessible at address 0x0114

- When ITS_CR0.[IDLE,ITSEN] != 0b11 or ITS_STATUSR.IDLE != 1, access on this interface is **WI**.
- Otherwise, access on this interface is **WO**.

10.3.1.23 ITS_READ_EVENT_DATAR

The ITS_READ_EVENT_DATAR characteristics are:

Purpose


ITS read event data register.
This register is used to return translation information for an EventID and DeviceID specified in ITS_EIDR and ITS_DIDR registers, respectively.

Attributes

ITS_READ_EVENT_DATAR is a 64-bit register.
This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_READ_EVENT_DATAR bit assignments are:



VIRT, bit [63]

When ITS_IDR0.INT_DOM != 0b10:

Specifies if the interrupt message generated by the ITS to the IRS in response to this event is for a physical or a virtual interrupt.

VIRT	Meaning
0b0	The interrupt message generated by the ITS is for a physical interrupt
0b1	The interrupt message generated by the ITS is for a virtual interrupt

If VALID is 0, the value of this field is UNKNOWN.
The reset behavior of this field is:
• On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

Bits [62:48]

Reserved, RES0.

VM_ID, bits [47:32]

Bits[15:0] of the VM ID passed to the IRS as part of the interrupt message targeting virtual interrupts.
If VIRT is 0, this field is RES0.
If VALID 0, the value of this field is UNKNOWN.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

VALID, bit [31]

Specifies whether the ITS has valid translation information for the specified EventID and DeviceID.

VALID	Meaning
0b0	The EventID does not have a valid translation.
0b1	The EventID has a valid translation.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Bits [30:24]

Reserved, RES0.

LPI_ID, bits [23:0]

Bits[23:0] of the LPI ID for the specified EventID.

If unimplemented upper bits of the LPI_ID are not zero, it is IMPLEMENTATION DEFINED whether the upper bits are treated as 0 or the interrupt message is ignored by the IRS.

If VALID 0, the value of this field is UNKNOWN.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing ITS_READ_EVENT_DATAR

Accesses to this register use the following encodings:

Accessible at address 0x0118

- When ITS_CR0.[IDLE,ITSEN] != 0b11 or ITS_STATUSR.IDLE != 1, access on this interface is UNKNOWN/WI.
- Otherwise, access on this interface is RO.

10.3.1.24 ITS_STATUSR

The ITS_STATUSR characteristics are:

Purpose

ITS Status Register.

Reports whether the effects of the last write to all of the following registers are complete:

- ITS_INV_DEVICER.
- ITS_INV_EVENTR.
- ITS_READ_EVENTR.


Attributes

ITS_STATUSR is a 32-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_STATUSR bit assignments are:



Bits [31:1]

Reserved, RES0.

IDLE, bit [0]

Reports the status of the last write to all of the following registers:

- ITS_INV_DEVICER.
- ITS_INV_EVENTR.
- ITS_READ_EVENTR.

IDLE	Meaning
0b0	The effects of the last write are not guaranteed to be complete.
0b1	The effects of the last write are guaranteed to be complete.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Access to this field is **RO**.

Accessing ITS_STATUSR

Accesses to this register use the following encodings:

Accessible at address 0x0120

Access on this interface is **RO**.

10.3.1.25 ITS_SWERR_STATUSR

The ITS_SWERR_STATUSR characteristics are:

Purpose

ITS software error status register. Specifies whether a software error has been reported. If an error is reported, it contains syndrome information for the error.


Attributes

ITS_SWERR_STATUSR is a 64-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_SWERR_STATUSR bit assignments are:



Bits [63:32]

Reserved, RES0.

IMP_EC, bits [31:24]

IMPLEMENTATION DEFINED error code when ITS_SWERR_STATUSR.EC == 0.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- Access is **RES0** if any of the following are true:
 - ITS_SWERR_STATUSR.V == 0
 - ITS_SWERR_STATUSR.EC != 0
- Otherwise, access to this field is **RO**

EC, bits [23:16]

Specifies the error code that software can use to triage and handle the error.

EC	Meaning
0x00	An error was reported because of an IMPLEMENTATION DEFINED reason.
0x01	Failed lookup of L1_DTE due to an external abort.
0x02	Failed lookup of L2_DTE due to an external abort.
0x03	Failed lookup of L1_ITTE due to an external abort.
0x04	Failed lookup of L2_ITTE due to an external abort.
0x05	An incoming event could not be translated because L1_DTE.VALID is 0.

EC	Meaning
0x06	An incoming event could not be translated because L2_DTE.VALID is 0.
0x07	An incoming event could not be translated because L1_ITTE.VALID is 0.
0x08	An incoming event could not be translated because L2_ITTE.VALID is 0.
0x09	An incoming event could not be translated because the DeviceID > (2 ^ ITS_DT_CFGR.DEVICEID_BITS) - 1.
0x0A	An incoming event could not be translated because the EventID > (2 ^ L2_DTE.EVENTID_BITS) - 1.
0x0B	An incoming event could not be translated because the DeviceID exceeds the L1_DTE.SPAN.
0x0C	An incoming event could not be translated because the EventID exceeds the L1_ITTE.SPAN.
0x0D	An incoming event could not be translated because the event is associated with the Non-secure Interrupt Domain and L2_ITTE.DAC = 0. The error is reported in the Realm ITS Domain.
0x0E	ITS_GEN_EVENTR was written when ITS_GEN_EVENT_STATUSR.IDLE == 0.
0x0F	ITS_READ_EVENTR was written when ITS_STATUSR.IDLE == 0.
0x10	ITS_INV_DEVICER was written when ITS_STATUSR.IDLE == 0.
0x11	ITS_INV_EVENTR was written when ITS_STATUSR.IDLE == 0.
0x12	ITS_SYNCR was written when ITS_SYNC_STATUSR.IDLE == 0.
0x13	An incoming event could not be translated because the EventID exceeds (2 ^ ITS_IDR2.EVENTID_BITS) - 1.
0x14	An incoming event specified an EventID with a write to a reserved location.

All other values are reserved.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- When ITS_SWERR_STATUSR.V == 0, access to this field is **UNKNOWN/WI**
- Otherwise, access to this field is **RO**

Bits [15:4]

Reserved, RES0.

OF, bit [3]

Specifies whether multiple software errors have been detected.

When this field is 1, the syndrome information reports information about the error that last caused ITS_SWERR_STATUSR.V to transition from 0 to 1.

OF	Meaning
0b0	No errors have been detected, since the error that was reported when ITS_SWERR_STATUSR.V last transitioned from 0 to 1.
0b1	At least one error has been detected, since the error that was reported when ITS_SWERR_STATUSR.V last transitioned from 0 to 1.

When clearing ITS_SWERR_STATUSR.V to 0, if this field is nonzero, software writes 1 to clear this field to zero.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- When ITS_SWERR_STATUSR.V == 0, access to this field is **UNKNOWN/WI**
- Otherwise, access to this field is **W1C**

S1V, bit [2]

Specifies whether ITS_SWERR_SYNDROMER1 is valid.

S1V	Meaning
0b0	ITS_SWERR_SYNDROMER1 is not valid.
0b1	ITS_SWERR_SYNDROMER1 is valid.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- When ITS_SWERR_STATUSR.V == 0, access to this field is **UNKNOWN/WI**
- Otherwise, access to this field is **RO**

S0V, bit [1]

Specifies whether ITS_SWERR_SYNDROMER0 is valid.

S0V	Meaning
0b0	ITS_SWERR_SYNDROMER0 is not valid.
0b1	ITS_SWERR_SYNDROMER0 is valid.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- When ITS_SWERR_STATUSR.V == 0, access to this field is **UNKNOWN/WI**
- Otherwise, access to this field is **RO**

V, bit [0]

Specifies whether ITS_SWERR_STATUSR is valid and at least one software error has been reported.

V	Meaning
0b0	ITS_SWERR_STATUSR is not valid.
0b1	ITS_SWERR_STATUSR is valid.

Software writes 1 to this field to clear it to zero.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Access to this field is **W1C**.

Accessing ITS_SWERR_STATUSR

After reading ITS_SWERR_STATUSR, software clears the valid fields in the register to allow new errors to be reported.

However, between reading the register and clearing the valid fields, a new error might have overwritten the register.

To prevent this error being lost by software, the register prevents updates to fields that might have been updated by a new error.

This is done by ensuring a write to the register is ignored if all of the following are true:

- Any of ITS_SWERR_STATUSR.{V, OF} are nonzero before the write.
- The write does not clear the nonzero ITS_SWERR_STATUSR.{V, OF} fields to zero by writing ones to the applicable field or fields.

Accesses to this register use the following encodings:

Accessible at address 0x0240

- When ITS_IDR0.SWE != 1, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RW**.

10.3.1.26 ITS_SWERR_SYNDROMER0

The ITS_SWERR_SYNDROMER0 characteristics are:

Purpose

ITS software error syndrome register 0. Records ITS specific software error syndrome information.


Attributes

ITS_SWERR_SYNDROMER0 is a 64-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_SWERR_SYNDROMER0 bit assignments are:



DEVICE_ID, bits [63:32]

The DeviceID for the incoming event that generated the software error.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Bits [31:16]

Reserved, RES0.

EVENT_ID, bits [15:0]

The EventID for the incoming event that generated the software error.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing ITS_SWERR_SYNDROMER0

Accesses to this register use the following encodings:

Accessible at address 0x0248

- When ITS_IDR0.SWE != 1, access on this interface is **RAZ/WI**.
- When ITS_SWERR_STATUSR.V == 1 and ITS_SWERR_STATUSR.S0V == 1, access on this interface is **RO**.
- Otherwise, access on this interface is **UNKNOWN/WI**.

10.3.1.27 ITS_SWERR_SYNDROMER1

The ITS_SWERR_SYNDROMER1 characteristics are:

Purpose

ITS software error syndrome register 1. Records ITS specific software error syndrome information.


Attributes

ITS_SWERR_SYNDROMER1 is a 64-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_SWERR_SYNDROMER1 bit assignments are:



Bits [63:56]

Reserved, RES0.

ADDR, bits [55:3]

Bits[55:3] of the physical address of a translation structure associated with the detected error.

The address in this field is associated with the PAS of the ITS Domain where the error is detected.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in ITS_IDR0.PA_RANGE.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Bits [2:0]

Reserved, RES0.

Accessing ITS_SWERR_SYNDROMER1

Accesses to this register use the following encodings:

Accessible at address 0x0250

- When ITS_IDR0.SWE != 1, access on this interface is **RAZ/WI**.
- When ITS_SWERR_STATUSR.V == 1 and ITS_SWERR_STATUSR.S1V == 1, access on this interface is **RO**.
- Otherwise, access on this interface is **UNKNOWN/WI**.

10.3.1.28 ITS_SYNCR

The ITS_SYNCR characteristics are:

Purpose

ITS synchronize translation events register.

Configuration

The effects of a write to this register are not guaranteed to have completed before ITS_SYNC_STATUSR.IDLE is 1.


Attributes

ITS_SYNCR is a 64-bit register.

This register is part of the [ITS_CONFIG_FRAME](#) block.

Field descriptions

The ITS_SYNCR bit assignments are:



SYNC, bit [63]

A synchronization request applies to the following Accepted events:

- Events that are generated from an IWB.
- Events that are generated from a system peripheral using an IMPLEMENTATION DEFINED mechanism.
- Events that are generated as a result of a write to ITS_GEN_EVENTR.
- Events that are generated as a result of a write to a register in an ITS translate register frame associated with the ITS Domain.

Writing 0 to this field has no effect.

SYNC	Meaning
0b0	The write is IGNORED.
0b1	The write issues a synchronization request to the ITS Domain.

All events covered by the synchronization request, that are Accepted at the time of the write to this register, are guaranteed to have been translated when ITS_SYNC_STATUSR.IDLE is 1.

For each such event, if the event has a valid translation, the ITS has generated an interrupt event that is Accepted by the associated IRS.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Bits [62:33]

Reserved, RES0.

SYNCALL, bit [32]

Whether a synchronization request issued by writing 1 to SYNC applies to all Accepted ITS events or is only required to apply to those with specified DeviceID.

SYNCALL	Meaning
0b0	Synchronize only events for the specified DeviceID.
0b1	Synchronize all events for the ITS Domain.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

DEVICE_ID, bits [31:0]

An ITS synchronization request synchronized all Accepted ITS events with the specified DeviceID.

When SYNCALL is 1, this field is IGNORED.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing ITS_SYNCR

Accesses to this register use the following encodings:

Accessible at address 0x0140

- When ITS_CR0.[IDLE,ITSEN] != 0b11 or ITS_SYNC_STATUSR.IDLE != 1, access on this interface is **WI**.
- Otherwise, access on this interface is **WO**.

10.3.1.29 ITS_SYNC_STATUSR

The ITS_SYNC_STATUSR characteristics are:

Purpose

ITS synchronize interrupt events status register.


Attributes

ITS_SYNC_STATUSR is a 32-bit register.

This register is part of the ITS_CONFIG_FRAME block.

Field descriptions

The ITS_SYNC_STATUSR bit assignments are:



Bits [31:1]

Reserved, RES0.

IDLE, bit [0]

Whether the effects of the last write to ITS_SYNCR have completed.

IDLE	Meaning
0b0	The effects of writing to ITS_SYNCR not guaranteed to have completed.
0b1	The effects of writing to ITS_SYNCR have completed.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b1.

Access to this field is **RO**.

Accessing ITS_SYNC_STATUSR

This register is read-only.

Accesses to this register use the following encodings:

Accessible at address 0x0148

Access on this interface is **RO**.

10.3.2 ITS_TRANSLATE_FRAME, ITS translate register frame

The ITS_TRANSLATE_FRAME characteristics are:

Purpose

Contains translate registers used to generate translated interrupts for an ITS Domain.

One or more ITS translate register frames are present for each supported ITS Domain.

Arm strongly recommends that translate register frame is not accessible by PEs. This is because a write to an translate register may require the ITS access to memory, leading to in-out dependencies that can potentially lead to deadlocks in the system. The translate register frames may be mapped in stage 2 translation tables shared between an SMMU and a PE, and software in a VM may perform an access to any of the mapped translate register frames. If the translate register frames are not accessible by PEs, the behavior on an attempted access from a PE is IMPLEMENTATION DEFINED and is likely to result in an External abort. If the translate register frames are accessible by PEs, writes from each PE must use a unique DeviceID that cannot spoof a write originating from a different PE or requesting device.

Each translate register frame is only accessible in the PAS associated with the ITS Domain.

The base address of each translate register frame is distinct from addresses of registers accessible in any other PAS.

The base address of each translate register frame is aligned to 64KB.

Attributes

The ITS_TRANSLATE_FRAME block is of size 64KB

Default access

Accesses to the address space of this block that do not resolve to a register or a further block are treated as RAZ/WI.

Contents

Offset	Name	Notes
0x0000	ITS_TRANSLATER	Most permissive access: WO
0x0008	ITS_RL_TRANSLATER	Most permissive access: WO

10.3.2.1 ITS_TRANSLATER

The ITS_TRANSLATER characteristics are:

Purpose

ITS translate event register.

A write to this register generates a SET_EDGE event for the DeviceID of the agent writing to the register and the EventID specified as part of the write.

Attributes

ITS_TRANSLATER is a 32-bit register.

This register is part of the [ITS_TRANSLATE_FRAME](#) block.

Field descriptions

The ITS_TRANSLATER bit assignments are:



Bits [31:16]

Reserved, RES0.

EVENT_ID, bits [15:0]

The EventID to translate.

Accessing ITS_TRANSLATER

This register is write-only.

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.

Writes to this register are ignored if the ITS Domain is not enabled.

See [Chapter 5 Interrupt translation service \(ITS\)](#) for more information about ordering and completion requirements for writes to this register.

Accesses to this register use the following encodings:

Accessible at address 0x0000

Access on this interface is **WO**.

10.3.2.2 ITS_RL_TRANSLATER

The ITS_RL_TRANSLATER characteristics are:

Purpose

ITS translate event in Realm ITS Domain register.

A write to this register that uses the Non-secure PAS, generates a SET_EDGE event for the DeviceID of the agent writing to the register and the EventID specified as part of the write.

The SET_EDGE event is associated with the Non-secure Interrupt Domain and processed by the ITS in the Realm ITS Domain.

If the translation of the SET_EDGE event is successful, a SET_EDGE interrupt event is generated for the IRS in the Realm Interrupt Domain


Attributes

ITS_RL_TRANSLATER is a 32-bit register.

This register is part of the [ITS_TRANSLATE_FRAME](#) block.

Field descriptions

The ITS_RL_TRANSLATER bit assignments are:



Bits [31:16]

Reserved, RES0.

EVENT_ID, bits [15:0]

The EventID to translate.

Accessing ITS_RL_TRANSLATER

This register is write-only.

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.

Writes to this register are IGNORED if any of the following are true:

- The Non-secure ITS Domain is not enabled.
- The Realm ITS Domain is not enabled.

See [Chapter 5 Interrupt translation service \(ITS\)](#) for more information about ordering and completion requirements for writes to this register.

Accesses to this register use the following encodings:

Accessible at address 0x0008

- When ITS_IDR0.INT_DOM == 0b01 and IsITSDomainImplemented(Realm), access on this interface is **WO**.
- Otherwise, access on this interface is **RAZ/WI**.

10.4 IWB register frames

10.4.1 IWB_CONFIG_FRAME, IWB configuration registers frame

The IWB_CONFIG_FRAME characteristics are:

Purpose

Contains control registers for an IWB.

The base address is aligned to 64KB.

Access to this register frame in a PAS associated with an unimplemented Interrupt Domain is CONSTRAINED UNPREDICTABLE with a choice of:

- The access is RAZ/WI
- The access is made to an implemented Interrupt Domain as follows:
 - An access in the Realm PAS is translated to an access in the Non-secure Interrupt Domain
 - An access in the Secure PAS is translated to an access in the Non-secure Interrupt Domain
 - An access in the Root PAS is translated to an access in the MPPAS of the IWB

Attributes

The IWB_CONFIG_FRAME block is of size 64KB

Default access

Accesses to the address space of this block that do not resolve to a register or a further block are treated as RAZ/WI.

Contents

Offset	Name	Notes
0x0000	IWB_IDR0	Most permissive access: RO
0x0040	IWB_IIDR	Most permissive access: RO
0x0044	IWB_AIDR	Most permissive access: RO
0x0080	IWB_CR0	Most permissive access: RW
0x00C0	IWB_WENABLE_STATUSR	Most permissive access: RO
0x00C4	IWB_WDOMAIN_STATUSR	Most permissive access: RO
0x00C8	IWB_WRESAMPLER	Most permissive access: WO
0x2000 + (4 * n)	IWB_WENABLER[n]	Most permissive access: RW
0x4000 + (4 * n)	IWB_WTMR[n]	Most permissive access: RW
0x6000 + (4 * n)	IWB_WDOMAINR[n]	Most permissive access: RW
0x0E00 + (4 * n) for n in ↪ 63:0	-	Most permissive access: ImplementationDefined

10.4.1.1 IWB_AIDR

The IWB_AIDR characteristics are:

Purpose

IWB Architecture Identification Register. Identifies the GIC architecture version to which the implementation conforms.


Attributes

IWB_AIDR is a 32-bit register.

This register is part of the IWB_CONFIG_FRAME block.

Field descriptions

The IWB_AIDR bit assignments are:



Bits [31:12]

Reserved, RES0.

Component, bits [11:8]

GIC component

Component	Meaning
0b0000	IRS
0b0001	ITS
0b0010	IWB

ArchMajorRev, bits [7:4]

Major Architecture revision.

ArchMajorRev	Meaning
0b0000	GICv5.x

ArchMinorRev, bits [3:0]

Minor Architecture revision.

ArchMinorRev	Meaning
0b0000	GICv5.0

Accessing IWB_AIDR

Accesses to this register use the following encodings:

Accessible at address 0x0044

Access on this interface is **RO**.

10.4.1.2 IWB_IDR0

The IWB_IDR0 characteristics are:

Purpose

IWB ID register 0. Contains read-only fields with information about the IWB GIC component.


Attributes

IWB_IDR0 is a 32-bit register.

This register is part of the IWB_CONFIG_FRAME block.

Field descriptions

The IWB_IDR0 bit assignments are:



Bits [31:15]

Reserved, RES0.

INT_DOMS, bits [14:11]

The Interrupt Domains supported by the IWB.

INT_DOMS	Meaning
0b0001	Only the Secure Interrupt Domain is supported
0b0010	Only the Non-secure Interrupt Domain is supported
0b0111	The EL3, Secure and Non-secure Interrupt Domains are supported
0b1110	The EL3, Realm, and Non-secure Interrupt Domains are supported
0b1111	The EL3, Realm, Secure, and Non-secure Interrupt Domains are supported

All values not listed are reserved.

IW_RANGE, bits [10:0]

Indicates the number of implemented wire control registers.

The number is reported as the highest numbered wire control field as multiple of 32 minus one.

For example:

- A value of 0 means that the IWB supports register accesses to control wires 0-31
- A value of 1 means that the IWB supports register accesses to control wires 0-63
- ...and so on until 0-65535.

Accessing IWB_IDR0

Accesses to this register use the following encodings:

Accessible at address 0x0000

Access on this interface is RO.

10.4.1.3 IWB_IIDR

The IWB_IIDR characteristics are:

Purpose

IWB Implementer Identification Register. Provides information about the implementation and implementer of the GIC, and architecture version supported.

Attributes

IWB_IIDR is a 32-bit register.

This register is part of the [IWB_CONFIG_FRAME](#) block.

Field descriptions

The IWB_IIDR bit assignments are:

31	20	19	16	15	12	11	0		
ProductID				Variant		Revision		Implementer	

ProductID, bits [31:20]

IMPLEMENTATION DEFINED value identifying the GIC part

When the IWB_PIDR{0,1} registers are present, Arm expects that the IWB_PIDR{0,1}.PART_{0,1} fields match the value of IWB_IIDR.ProductID.

If required, however, an implementation is permitted to provide values for IWB_PIDR.{0,1}.PART_{0,1} that do not match the value of IWB_IIDR.ProductID

Variant, bits [19:16]

IMPLEMENTATION DEFINED value used to distinguish product variants, or major revisions of the product

Revision, bits [15:12]

IMPLEMENTATION DEFINED value used to distinguish minor revisions of the product

Implementer, bits [11:0]

Contains the JEP106 code of the company that implemented the GIC

For an Arm implementation, the JEP106 code is 0x43B

When the IWB_PIDR{1,2,4} registers are present, Arm expects that the IWB_PIDR{0,1}.PART_{0,1} fields match the value of IWB_IIDR.Implementer.

Accessing IWB_IIDR

Accesses to this register use the following encodings:

Accessible at address 0x0040

Access on this interface is **RO**.

10.4.1.4 IWB_CR0

The IWB_CR0 characteristics are:

Purpose

IWB control register 0.


Attributes

IWB_CR0 is a 32-bit register.

This register is part of the [IWB_CONFIG_FRAME](#) block.

Field descriptions

The IWB_CR0 bit assignments are:



Bits [31:2]

Reserved, RES0.

IDLE, bit [1]

Whether the transition between enabled and disabled states of the IWB is complete.

IDLE	Meaning
0b0	The effects of updating IWBEN are not guaranteed to have completed.
0b1	The effects of updating IWBEN have completed.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b1.

Access to this field is **RO**.

IWBEN, bit [0]

Controls if the IWB is enabled.

IWBEN	Meaning
0b0	Disabled. The IWB does not generate any events its destination.
0b1	Enabled. The IWB may generate events to its destination.

The reset behavior of this field is:

- On a GIC reset, this field resets to 0b0.

Accessing this field has the following behavior:

- When IWB_CR0.IDLE == 0, access to this field is **RO**
- RO** if !IsAccessIWBMPAS()
- Otherwise, access to this field is **RW**

Accessing IWB_CR0

Accesses to this register use the following encodings:

Accessible at address 0x0080

Access on this interface is **RW**.

10.4.1.5 IWB_WDOMAIN_STATUSR

The IWB_WDOMAIN_STATUSR characteristics are:

Purpose

IWB wire assignment status register for an Interrupt Domain.


Attributes

IWB_WDOMAIN_STATUSR is a 32-bit register.

This register is part of the IWB_CONFIG_FRAME block.

Field descriptions

The IWB_WDOMAIN_STATUSR bit assignments are:



Bits [31:1]

Reserved, RES0.

IDLE, bit [0]

Tracks status of writes to IWB_WDOMAINR<n> on this IWB.

IDLE	Meaning
0b0	A write to IWB_WDOMAINR<n> is in progress.
0b1	No write to IWB_WDOMAINR<n> is in progress.

Accessing IWB_WDOMAIN_STATUSR

Accesses to this register use the following encodings:

Accessible at address 0x00C4

- When IsAccessIWBMPAS(), access on this interface is **RO**.
- Otherwise, access on this interface is **RAZ/WI**.

10.4.1.6 IWB_WDOMAINR<n>, n = 0 - 4095

The IWB_WDOMAINR<n> characteristics are:

Purpose

IWB wire Interrupt Domain selection register. Allows software to configure the Interrupt Domain that wires $n * 16$ through $((n * 16) + 15)$ are assigned to.

Configuration

The number of implemented IWB_WDOMAINR<n> registers is $(IWB_IDR0.IW_RANGE + 1) * 2$.

The effects of a write to this register are not guaranteed to have completed until IWB_WDOMAIN_STATUSR.IDLE is 1.


Attributes

IWB_WDOMAINR<n> is a 32-bit register.

This register is part of the [IWB_CONFIG_FRAME](#) block.

Field descriptions

The IWB_WDOMAINR<n> bit assignments are:



WDOM<x>, bits [2x+1:2x], for x = 15 to 0

Configures the Interrupt Domain that wire $((16 * n) + x)$ is assigned to.

Programming an Interrupt Domain not supported by the IWB results in CONSTRAINED UNPREDICTABLE behavior with the following options:

- No signals are generated by the IWB for a wire assigned to an unsupported Interrupt Domain.
- The wire is treated as being assigned to another supported Interrupt Domain and a read of this field returns the Interrupt Domain the wire is assigned to.

Access to control fields for wires which are not implemented are RAZ/WI.

WDOM<x>	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- **RO** if `IsWireDomainRO((n * 16) + x)`
- Otherwise, access to this field is **RW**

Accessing IWB_WDOMAINR<n>

When the Interrupt Domain assigned of a wire is fixed, access to the corresponding field is **RO**.

This register can only be accessed through the MPPAS of the IWB.

Accesses to this register use the following encodings:

Accessible at address $0x6000 + (4 * n)$

- When `IWB_CR0.IDLE == 0`, access on this interface is **RO**.
- When `!IsAccessIWBMPAS()`, access on this interface is **RAZ/WI**.
- Otherwise, access on this interface is **RW**.

10.4.1.7 IWB_WENABLE_STATUSR

The IWB_WENABLE_STATUSR characteristics are:

Purpose

IWB wire enable status register for an Interrupt Domain.


Attributes

IWB_WENABLE_STATUSR is a 32-bit register.

This register is part of the IWB_CONFIG_FRAME block.

Field descriptions

The IWB_WENABLE_STATUSR bit assignments are:



Bits [31:1]

Reserved, RES0.

IDLE, bit [0]

Tracks status of writes to IWB_WENABLER<n> on this IWB where the writes used the same PAS that was used to access this register.

IDLE	Meaning
0b0	A write to IWB_WENABLER<n> using the PAS that was used to access this register is in progress.
0b1	No write to IWB_WENABLER<n> using the PAS that was used to access this register is in progress.

Accessing IWB_WENABLE_STATUSR

Accesses to this register use the following encodings:

Accessible at address 0x00C0

Access on this interface is RO.

10.4.1.8 IWB_WENABLER<n>, n = 0 - 2047

The IWB_WENABLER<n> characteristics are:

Purpose

IWB wire enable register. Allows software to configure if individual wires are enabled or disabled.

Configuration

The number of implemented IWB_WENABLER<n> registers is IWB_IDR0.IW_RANGE + 1.

The effects of a write to this register are not guaranteed to have completed before IWB_WENABLE_STATUSR.IDLE is 1.


Attributes

IWB_WENABLER<n> is a 32-bit register.

This register is part of the [IWB_CONFIG_FRAME](#) block.

Field descriptions

The IWB_WENABLER<n> bit assignments are:



WEN<x>, bits [x], for x = 31 to 0

Configures if wire $((32 * n) + x)$ is enabled or disabled.

Access to control fields for wires which are not implemented are RAZ/WI.

WEN<x>	Meaning
0b0	Wire disabled
0b1	Wire enabled

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- RAZ/WI** if !IsWireAccessible(n, x)
- Otherwise, access to this field is **RW**

Accessing IWB_WENABLER<n>

When accessed using the MPPAS of the IWB, all fields are RW.

When accessed using any other PAS, a field is RW if all of the following are true:

- The wire corresponding to the field is assigned to an Interrupt Domain that is implemented by the IWB.
- The PAS corresponding to the access is same as the PAS associated with the Interrupt Domain to which the wire is assigned.

Otherwise, the field is RAZ/WI for that access.

Accesses to this register use the following encodings:

Accessible at address $0x2000 + (4 * n)$

- When IWB_CR0.IDLE == 0, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.4.1.9 IWB_WRESAMPLER

The IWB_WRESAMPLER characteristics are:

Purpose

IWB wire resample register.


Attributes

IWB_WRESAMPLER is a 32-bit register.

This register is part of the [IWB_CONFIG_FRAME](#) block.

Field descriptions

The IWB_WRESAMPLER bit assignments are:



Bits [31:16]

Reserved, RES0.

IWI, bits [15:0]

Input Wire Index. Specifies the wire to resample.

Following a write to this register, if all of the following are true, the wire is resampled:

- The access to this register is performed using the PAS associated with the Interrupt Domain that the wire is assigned to or the MPPAS of the IWB.
- The specified wire is enabled.

If the specified wire is disabled, the write to this register has no effect.

See ‘IWB wire control registers’ for more information.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing IWB_WRESAMPLER

Accesses to this register use the following encodings:

Accessible at address 0x00C8

- When IWB_CR0.IDLE == 0, access on this interface is **WI**.
- Otherwise, access on this interface is **WO**.

10.4.1.10 IWB_WTMR<n>, n = 0 - 2047

The IWB_WTMR<n> characteristics are:

Purpose

IWB Wire Trigger mode register. Allows software to configure if the wire signal is level-sensitive or edge-triggered.

Configuration

The number of implemented IWB_WTMR<n> registers is IWB_IDR0.IW_RANGE + 1.


Attributes

IWB_WTMR<n> is a 32-bit register.

This register is part of the [IWB_CONFIG_FRAME](#) block.

Field descriptions

The IWB_WTMR<n> bit assignments are:



TM<x>, bits [x], for x = 31 to 0

Configures if wire ((32 * n) + x) is level-sensitive or edge-triggered.

When the wire is configured as level-sensitive, a CLEAR event is sent to the ITS when the wire signal is de-asserted, and a SET_LEVEL event is sent to the ITS when the wire signal is asserted.

When the wire is configured as edge-triggered, only SET_EDGE event are generated to the ITS and only when the wire changes from de-asserted to asserted.

When IsWireConfigRO(n * 32 + x) is 0, this field resets to 0.

Access to control fields for wires which are not implemented are RAZ/WI.

TM<x>	Meaning
0b0	Edge-triggered
0b1	Level-sensitive

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing this field has the following behavior:

- Access is **RO** if any of the following are true:
 - !IsWireAccessible(n, x)
 - IsWireConfigRO((n * 32) + x)
- Otherwise, access to this field is **RW**

Accessing IWB_WTMR<n>

When the Trigger mode of a wire is software programmable, all of the following are true:

- When accessed using the MPPAS of the IWB, all fields are RW.
- When accessed using any other PAS, a field is RW if all of the following are true:
 - The wire corresponding to the field is assigned to an Interrupt Domain that is implemented by the IWB.
 - The PAS corresponding to the access is same as the PAS associated with the Interrupt Domain to which the wire is assigned.

Otherwise, the field is **RO** for that access.

Accesses to this register use the following encodings:

Accessible at address $0x4000 + (4 * n)$

- When IWB_CR0.IDLE == 0, access on this interface is **RO**.
- Otherwise, access on this interface is **RW**.

10.5 GIC PMU register frame

10.5.1 GIC_PMU_FRAME, GIC PMU register frame

The GIC_PMU_FRAME characteristics are:

Purpose

Contains GIC PMU registers.

The GIC PMU register frame is accessible in each PAS corresponding to an Interrupt Domain that the PMU can count events for.

The base address of the GIC PMU register frame is distinct from any other GIC register frame.

The base address of the GIC PMU register frame is aligned to 64KB.

Attributes

The GIC_PMU_FRAME block is of size 64KB

Default access

Accesses to the address space of this block that do not resolve to a register or a further block are treated as RAZ/WI.

Contents

Offset	Name	Notes
0x400 + (8 * n)	GIC_PMEVTYPEPER[n]	Most permissive access: RW
0x800 + (8 * n)	GIC_PMEVFILT2R[n]	Most permissive access: RW
0xA00 + (8 * n)	GIC_PMEVFILTR[n]	Most permissive access: RW
0xD80	GIC_PMIDR0	Most permissive access: RO

I_KCDRK

The GIC PMU frame describes the register frame of a GIC PMU compliant with the register formats and layouts defined by the *Arm® CoreSight™ Architecture Performance Monitoring Unit Architecture* [11].

Not all registers are shown. Only those registers where the ITS architecture specifies an architected behavior of an IMPLEMENTATION DEFINED field as specified in [11] are shown. These registers are renamed as their behavior is specified by the GICv5 architecture. For the remaining offsets and register definitions, see [11].

10.5.1.1 GIC_PMEVFILT2R<n>, n = 0 - 63

The GIC_PMEVFILT2R<n> characteristics are:

Purpose

GIC PMU Event Filter 2 Register

Attributes


GIC_PMEVFILT2R<n> is a 64-bit register.

This register is part of the [GIC_PMU_FRAME](#) block.

Field descriptions

The GIC_PMEVFILT2R<n> bit assignments are:

When GIC_PMEVTYPER<n>.PMEVTYPE IN {0b01x}:



Bits [63:61]

Reserved, RES0.

FILTER_EID_SPAN, bit [60]

When GIC_PMEVTYPER<n>.FS == 1, GIC_PMEVTYPER<n>.FSPAN == 1 and FILTER_EID == 1:

Controls if filtering using EventID uses exact matching or matches a range of values.

If the ITS being monitored only implements support for 1 bit of EventID, this field is treated as 0.

See ‘GIC Performance Monitoring Units’ for more information.

FILTER_EID_SPAN	Meaning
0b0	GIC_PMEVFILTR<n>.EVENT_ID filters EventID on an exact match.
0b1	GIC_PMEVFILTR<n>.EVENT_ID filters EventID on a range of values.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

FILTER_EID, bit [59]

When GIC_PMEVTYPER<n>.FS == 1 and FILTER_DID == 1:

Controls whether filtering using EventID is enabled.

Filtering on EventID is only supported when also filtering on DeviceID.

See ‘GIC Performance Monitoring Units’ for more information.

FILTER_EID	Meaning
0b0	Count events from all EventIDs for the matched DeviceIDs.
0b1	Count events that match the DeviceID and EventID filter programming.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

FILTER_DID_SPAN, bit [58]

When GIC_PMEVTYPER<n>.FS == 1, GIC_PMEVTYPER<n>.FSPAN == 1 and FILTER_DID == 1:

Controls if filtering using DeviceID uses exact matching or matches a range of values.

If the ITS being monitored only implements support for 1 bit of DeviceID, this field is treated as 0.

See ‘GIC Performance Monitoring Units’ for more information.

FILTER_DID_SPAN	Meaning
0b0	GIC_PMEVFILTR<n>.DEVICE_ID filters DeviceID on an exact match.
0b1	GIC_PMEVFILTR<n>.DEVICE_ID filters DeviceID on a range of values.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

FILTER_DID, bit [57]

When GIC_PMEVTYPER<n>.FS == 1:

Controls whether filtering using DeviceID is enabled.

FILTER_DID	Meaning
0b0	Count events from all DeviceIDs.
0b1	Count events that match the DeviceID filter programming.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

Bits [56:16]

Reserved, RES0.

ITSID, bits [15:0]


Events from the ITS that match this value are counted.

If the agent being monitored does not contain more than one ITS, this field is RES0.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

When $GIC_PMEVTYPER<n>.PMEVTYPE$ IN {0b00x}:



FILTER_VM_ID, bit [63]

When $GIC_PMEVTYPER<n>.FS == 1$ and $FILTER_VIRT == 0b01$:

Controls whether counting of virtual events is filtered based on their VM ID.

When an event is selected that does not support filtering on virtual events, this field is RES0 and has no impact on the counted events.

FILTER_VM_ID	Meaning
0b0	Events translated to virtual interrupt events are not filtered on VM ID.
0b1	Events translated to virtual interrupt events are only counted if the VM ID matches the value specified in $GIC_PMEVFILTR<n>.VM_ID$.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

FILTER_VIRT, bits [62:61]

When $GIC_PMEVTYPER<n>.FS == 1$:

Controls whether counting of events is filtered based on whether they are virtual or physical.

When an event is selected that does not support filtering on virtual events, this field is RES0 and has no impact on the counted events.

FILTER_VIRT	Meaning
0b00	Physical and virtual interrupt events are counted.
0b01	Only virtual interrupt events are counted.
0b10	Only physical interrupt events are counted.

Values not defined are reserved.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

FILTER_INTID_SPAN, bit [60]

When $GIC_PMEVTYPER<n>.FS == 1$, $GIC_PMEVTYPER<n>.FSPAN == 1$ and $FILTER_INTID == 1$:

Controls if filtering using INTID uses exact matching or matches a range of INTID.ID values.

When an event is selected that does not support filtering on INTID, this field is RES0 and has no impact on the counted events.

See ‘GIC Performance Monitoring Units’ for more information.

FILTER_INTID_SPAN	Meaning
0b0	GIC_PMEVFILTR<n>.ID filters INTID.ID on an exact match.
0b1	GIC_PMEVFILTR<n>.ID filters INTID.ID on a range of values.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

FILTER_INTID, bit [59]

When $GIC_PMEVTYPER<n>.FS == 1$:

Controls whether filtering using INTID is enabled.

When an event is selected that does not support filtering on INTID, this field is RES0 and has no impact on the counted events.

FILTER_INTID	Meaning
0b0	Count interrupt events for all INTIDs.
0b1	Count interrupt events that match the GIC_PMEVFILTR<n>.INTID filter programming.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

FILTER_RW, bits [58:57]

When $GIC_PMEVTYPER<n>.FS == 1$:

Controls whether counting of events is filtered based on whether they are read or write events.

When an event is selected that does not support filtering on reads or writes, this field is RES0 and has no impact on the counted events.

FILTER_RW	Meaning
0b00	Both read and write events are counted.
0b01	Only read events are counted.
0b10	Only write events are counted.

Values not defined are reserved.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

FILTER_SRC_PE, bit [56]

When $GIC_PMEVTYPER<n>.FS == 1$:

Controls whether filtering of PE commands using source PE IAFFID is enabled.

When an event is selected that does not support filtering on source PE IAFFID, this field is RES0 and has no impact on the counted events.

FILTER_SRC_PE	Meaning
0b0	Count PE commands from all PEs.
0b1	Count PE commands sent from PEs whose IAFFID match the $GIC_PMEVFILTR<n>.IAFFID$ filter value.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

Bits [55:16]

Reserved, RES0.

IAFFID, bits [15:0]

When $GIC_PMEVTYPER<n>.FS == 1$ and $FILTER_SRC_PE == 1$:

When filtering on source PE IAFFID, only PE commands from a PE whose IAFFID matches this value are counted.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

Accessing GIC_PMEVFILT2R<n>

Accesses to this register use the following encodings:

Accessible at address $0x800 + (8 * n)$

Access on this interface is **RW**.

10.5.1.2 GIC_PMEVFILTR<n>, n = 0 - 63

The GIC_PMEVFILTR<n> characteristics are:

Purpose

GIC PMU Event Filter Register

Attributes


GIC_PMEVFILTR<n> is a 64-bit register.

This register is part of the [GIC_PMU_FRAME](#) block.

Field descriptions

The GIC_PMEVFILTR<n> bit assignments are:

When GIC_PMEVTYPER<n>.PMEVTYPE IN {0b01x}:



Bits [63:48]

Reserved, RES0.

EVENT_ID, bits [47:32]

When filtering by EventID, the EventIDs that match this value are counted.

When GIC_PMEVFILT2R<n>.FILTER_EID is 0, this field is IGNORED.

When GIC_PMEVFILT2R<n>.FILTER_EID is 1, GIC_PMEVFILT2R<n>.FILTER_EID_SPAN controls how the EventID is matched against this value.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

DEVICE_ID, bits [31:0]

When filtering by DeviceID, the DeviceIDs that match this value are counted.


When GIC_PMEVFILT2R<n>.FILTER_DID is 0, this field is IGNORED.

When GIC_PMEVFILT2R<n>.FILTER_DID is 1, GIC_PMEVFILT2R<n>.FILTER_DID_SPAN controls how the DeviceID is matched against this value.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

When GIC_PMEVTYPER<n>.PMEVTYPE IN {0b00x}:



Bits [63:48]

Reserved, RES0.

VM_ID, bits [47:32]

When $GIC_PMEVFILT2R<n>.FILTER_VM_ID == 1$:

When filtering on VM_ID only events that are translated to virtual interrupt events matching this VM_ID are counted.

When an event is selected that does not support filtering on VM ID, this field is RES0 and has no impact on the counted events.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

TYPE, bits [31:29]

When $GIC_PMEVFILT2R<n>.FILTER_INTID == 1$:

Type of the interrupt.

When an event is selected that does not support filtering on INTID, this field is RES0 and has no impact on the counted events.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

Bits [28:24]

Reserved, RES0.

ID, bits [23:0]

When $GIC_PMEVFILT2R<n>.FILTER_INTID == 1$:

The ID of the interrupt.

ID and TYPE together form an INTID.

The monitored IRS may implement fewer than 24 bits of ID. Unimplemented upper bits are RES0.

When $GIC_PMEVFILT2R<n>.FILTER_INTID$ is 1, $GIC_PMEVFILT2R<n>.FILTER_INTID_SPAN$ controls how the INTID is matched against this value.

When an event is selected that does not support filtering on INTID, this field is RES0 and has no impact on the counted events.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

Accessing GIC_PMEVFLTR<n>

Accesses to this register use the following encodings:

Accessible at address $0xA00 + (8 * n)$

Access on this interface is **RW**.

10.5.1.3 GIC_PMEVTYPER<n>, n = 0 - 63

The GIC_PMEVTYPER<n> characteristics are:

Purpose

GIC PMU Event Type Select Register


Attributes

GIC_PMEVTYPER<n> is a 64-bit register.

This register is part of the [GIC_PMU_FRAME](#) block.

Field descriptions

The GIC_PMEVTYPER<n> bit assignments are:



Bits [63:32]

Reserved, RES0.

V, bit [31]

Whether PMEVTYPER and PMEVTID select a valid event.

V	Meaning
0b0	The event selected by PMEVTYPER and PMEVTID is invalid.
0b1	The event selected by PMEVTYPER and PMEVTID is valid.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Access to this field is **RO**.

FS, bit [30]

Whether filtering is supported for the event selected in GIC_PMEVTID.

When V is 0, the value of this field is UNKNOWN.

FS	Meaning
0b0	Filtering for the selected event is not supported.
0b1	Filtering for the selected event is supported.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Access to this field is **RO**.

FSPAN, bit [29]

Whether filtering using a range of values is supported for the event selected in GIC_PMEVTID. See ‘GIC Performance Monitoring Units’ for more information about filtering on a range of values.

When V is 0 or FS is 0, the value of this field is UNKNOWN.

FSPAN	Meaning
0b0	Filtering using a range of values is not supported.
0b1	Filtering using a range of values is supported.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Access to this field is **RO**.

RL, bit [28]

When GIC_PMIDR0.OACE == 1 and GIC_PMIDR0.DOM_RL == 1:

Configures matching the Realm Interrupt Domain.

RL	Meaning
0b0	Events or characteristics attributable to the Realm Interrupt Domain are unaffected by this bit.
0b1	Do not count events or monitor characteristics attributable to the Realm Interrupt Domain.

This bit is ignored by the PMU when the PMU is not allowed to observe events or characteristics attributable to Realm operation of the agent being monitored.

Accessing this bit has the following behavior:

- This bit reads-as-zero if the PMU is never allowed to observe events or characteristics attributable to Realm operation of the agent being monitored.
- Otherwise, this bit is read/write.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

EL3, bit [27]

When GIC_PMIDR0.OACE == 1 and GIC_PMIDR0.DOM_EL3 == 1:

Configures matching the EL3 Interrupt Domain.

EL3	Meaning
0b0	Events or characteristics attributable to the EL3 Interrupt Domain are unaffected by this bit.

EL3	Meaning
0b1	Do not count events or monitor characteristics attributable to the EL3 Interrupt Domain.

This bit is ignored by the PMU when the PMU is not allowed to observe events or characteristics attributable to EL3 operation of the agent being monitored.

Accessing this bit has the following behavior:

- This bit reads-as-zero if the PMU is never allowed to observe events or characteristics attributable to EL3 operation of the agent being monitored.
- Otherwise, this bit is read/write.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

NS, bit [26]

When $GIC_PMIDR0.OACE == 1$ and $GIC_PMIDR0.DOM_NS == 1$:

Configures matching the Non-secure Interrupt Domain.

NS	Meaning
0b0	Events or characteristics attributable to the Non-secure Interrupt Domain are unaffected by this bit.
0b1	Do not count events or monitor characteristics attributable to the Non-secure Interrupt Domain.

This bit is ignored by the PMU when the PMU is not allowed to observe events or characteristics attributable to Non-secure operation of the agent being monitored.

Accessing this bit has the following behavior:

- This bit reads-as-zero if the PMU is never allowed to observe events or characteristics attributable to Non-secure operation of the agent being monitored.
- Otherwise, this bit is read/write.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

S, bit [25]

When $GIC_PMIDR0.OACE == 1$ and $GIC_PMIDR0.DOM_S == 1$:

Configures matching the Secure Interrupt Domain.

S	Meaning
0b0	Events or characteristics attributable to the Secure Interrupt Domain are unaffected by this bit.
0b1	Do not count events or monitor characteristics attributable to the Secure Interrupt Domain.

This bit is ignored by the PMU when the PMU is not allowed to observe events or characteristics attributable to Secure operation of the agent being monitored.

Accessing this bit has the following behavior:

- This bit reads-as-zero if the PMU is never allowed to observe events or characteristics attributable to Secure operation of the agent being monitored.
- Otherwise, this bit is read/write.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Otherwise:

RES0

Bits [24:16]

Reserved, RES0.

PMEVTYPE, bits [15:12]

The performance monitor event type.

The value in this field selects the event for the type selected in PMEVTTYPE.

Values not defined are reserved.

PMEVTYPE	Meaning
0b0000	Architected IRS event.
0b0001	IMPLEMENTATION DEFINED IRS event.
0b0010	Architected ITS event.
0b0011	IMPLEMENTATION DEFINED ITS event.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

PMEVTID, bits [11:0]

The performance monitor event ID.

The value in this field selects the event for the type selected in PMEVTTYPE.

If the value in this field combined with the value in PMEVTTYPE selects an invalid event, all other fields in this register are IGNORED.

If less than 16 bits of event ID is implemented, unimplemented upper bits are RES0.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

Accessing GIC_PMEVTYPER<n>

Accesses to this register use the following encodings:

Accessible at address $0x400 + (8 * n)$

Access on this interface is **RW**.

10.5.1.4 GIC_P MIDR0

The GIC_P MIDR0 characteristics are:

Purpose

GIC PMU identification register 0. Contains read-only fields with information about the GIC PMU.


Attributes

GIC_P MIDR0 is a 32-bit register.

This register is part of the GIC_P MU_FRAME block.

Field descriptions

The GIC_P MIDR0 bit assignments are:



Bits [31:7]

Reserved, RES0.

DOM_RL, bit [6]

Whether a component in the agent being monitored supports the Realm Interrupt Domain.

DOM_RL	Meaning
0b0	The Realm Interrupt Domain is not supported by any component in the agent being monitored.
0b1	The Realm Interrupt Domain is supported by a component in the agent being monitored.

DOM_EL3, bit [5]

Whether a component in the agent being monitored supports the EL3 Interrupt Domain.

DOM_EL3	Meaning
0b0	The EL3 Interrupt Domain is not supported by any component in the agent being monitored.
0b1	The EL3 Interrupt Domain is supported by a component in the agent being monitored.

DOM_NS, bit [4]

Whether a component in the agent being monitored supports the Non-secure Interrupt Domain.

DOM_NS	Meaning
0b0	The Non-secure Interrupt Domain is not supported by any component in the agent being monitored.
0b1	The Non-secure Interrupt Domain is supported by a component in the agent being monitored.

DOM_S, bit [3]

Whether a component in the agent being monitored supports the Secure Interrupt Domain.

DOM_S	Meaning
0b0	The Secure Interrupt Domain is not supported by any component in the agent being monitored.
0b1	The Secure Interrupt Domain is supported by a component in the agent being monitored.

OACE, bit [2]

Whether the PMU implements the observability and access control extension.

OACE	Meaning
0b0	The PMU does not implement the observability and access control extension.
0b1	The PMU implements the observability and access control extension.

ITS_PMU, bit [1]

Whether this PMU counts events from an ITS.

This field can be interpreted in combination with IRS_PMU to determine the GIC PMU configuration.

See ‘GIC Performance Monitoring Unit (PMU)’ for more information.

ITS_PMU	Meaning
0b0	This PMU does not count events from an ITS.
0b1	This PMU counts events from at least one ITS.

IRS_PMU, bit [0]

Whether this PMU counts events from an ITS.

This field can be interpreted in combination with ITS_PMU to determine the GIC PMU configuration.

See ‘GIC Performance Monitoring Unit (PMU)’ for more information.

IRS_PMU	Meaning
0b0	This PMU does not count events from an IRS.
0b1	This PMU counts events from an IRS.

Accessing GIC_PMIDR0

Accesses to this register use the following encodings:

Accessible at address 0xD80

Access on this interface is **RO**.

10.6 Identification registers

R_{JTSPV} In the following register frames, offsets 0xFFD0–0xFFFC are defined as read-only identification register space:

- IRS_CONFIG_FRAME
- ITS_CONFIG_FRAME
- IWB_CONFIG_FRAME

I_{SMPQP} For Arm implementations, the following assignment of fields, consistent with CoreSight ID registers[12], is used:

Offset	Name	Field	Value	Meaning
0xFFFF0	x_CIDR0, Component ID0	[7:0]	0x0D	Preamble
0xFFFF4	x_CIDR1, Component ID1	[7:4]	0xF	CLASS
		[3:0]	0x0	Preamble
0xFFFF8	x_CIDR2, Component ID2	[7:0]	0x05	Preamble
0xFFFFC	x_CIDR3, Component ID3	[7:0]	0xB1	Preamble
0xFFE0	x_PIDR0, Peripheral ID0	[7:0]	IMPLEMENTATION DEFINED	Bits [7:0] of the Part number
0xFFE4	x_PIDR1, Peripheral ID1	[7:4]	IMPLEMENTATION DEFINED	Bits [3:0] of the JEP106 Designer code
		[3:0]	IMPLEMENTATION DEFINED	Bits [11:8] of the Part number
0xFFE8	x_PIDR2, Peripheral ID2	[7:4]	IMPLEMENTATION DEFINED	REVISION
		[3]	1	JEDEC-assigned value for DES always used
		[2:0]	IMPLEMENTATION DEFINED	Bits [6:4] bits of the JEP106 Designer code
0xFFEC	x_PIDR3, Peripheral ID3	[7:4]	IMPLEMENTATION DEFINED	REVRAND
		[3:0]	IMPLEMENTATION DEFINED	CMOD
0xFFD0	x_PIDR4, Peripheral ID4	[7:4]	0	SIZE
		[3:0]	IMPLEMENTATION DEFINED	JEP106 Designer continuation code
0xFFD4	x_PIDR5, Peripheral ID5	-	RES0	Reserved
0xFFD8	x_PIDR6, Peripheral ID6	-	RES0	Reserved
0xFFEC	x_PIDR7, Peripheral ID7	-	RES0	Reserved
0xFFBC	x_DEVARCH	[31:21]	0x23b	JEP106 continuation and identification codes
		[20]	1	DEVARCH is present.
		[19:16]	0	GICv5.0
		[15:0]	TBC	ARCHID, GICv5.

Where “x” is replaced with:

- IRS for the IRS_CONFIG_FRAME
- ITS for the ITS_CONFIG_FRAME
- IWB for the IWB_CONFIG_FRAME

Offsets and fields outside of those defined in this table are RES0.

Arm recommends that implementers use this scheme to provide a consistent software discovery model. If the CoreSight ID registers are not implemented, Arm recommends that all the locations in the table are RES0.

The Designer code fields for Arm-designed implementations use continuation code 0x4 and Designer code 0x3B. Non-Arm implementations that follow this CoreSight ID register layout set the Designer fields appropriate to the implementer.

I_{LFDIN}

For a full descriptions of the CoreSight ID registers see the Arm CoreSight Architecture Specification [12].

Chapter 11

Data structures

This chapter describes the GICv5 data structures.

11.1 ITS Data Structures

11.1.1 L1_DTE, Level 1 device table entry


The L1_DTE characteristics are:

Attributes

L1_DTE is a 8-byte structure.

Field descriptions

The L1_DTE bit assignments are:



VALID, bit [0]

Entry is valid

VALID	Meaning
0b0	This entry is invalid. All events identified by a DeviceID and EventID covered by this entry are ignored.
0b1	This entry is valid and L2_ADDR points to a level 2 device table.

Bits [2:1]

Reserved, RES0.

L2_ADDR, bits [55:3]

Bits[55:3] of the address of the start of the level 2 array of device table entries.

Bits[N:0] of the resulting address are 0 where $N = 2 + \text{SPAN}$.

This means that the level 2 DTE array is aligned to the size of the array.

The level 2 array is accessed using the same PAS as the level 1 DTE.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in ITS_IDR0.PA_RANGE.

Bits [59:56]

Reserved, RES0.

SPAN, bits [63:60]

Size of the structure pointed to by L2_ADDR

The level 2 device table pointed to by L2_ADDR contains 2^{SPAN} entries.

If this field is programmed to a value larger than the maximum, it is treated as having the maximum value for all other purposes than reading back the field. The maximum allowed value of this field is determined by ITS_DT_CFGR.L2SZ.

If ITS_DT_CFGR.L2SZ is 0b00, then the maximum value of this field is 9. If ITS_DT_CFGR.L2SZ is 0b01, then the maximum value of this field is 11. If ITS_DT_CFGR.L2SZ is 0b10, then the maximum value of this field is 13.

11.1.2 L2_DTE, Level 2 device table entry


The L2_DTE characteristics are:

Attributes

L2_DTE is a 8-byte structure.

Field descriptions

The L2_DTE bit assignments are:



VALID, bit [0]

Entry is valid

VALID	Meaning
0b0	This entry is invalid. All events identified by a DeviceID and EventID covered by this entry are ignored.
0b1	This entry is valid and ITT_ADDR points to a valid ITT.

ITT_L2SZ, bits [2:1]

Level 2 ITT size when a 2-level ITT structure is used.

ITT_L2SZ	Meaning
0b00	A level 2 ITT is maximum 4KB and resolves 9 bits of EventID.
0b01	A level 2 ITT is maximum 16KB and resolves 11 bits of EventID.
0b10	A level 2 ITT is maximum 64KB and resolves 13 bits of EventID.

Values not defined above are reserved.

If $\text{EVENTID_BITS} \leq (9 + (2 * \text{ITT_L2SZ}))$ and ITT_STRUCTURE is 1, the ITT consists of a single L1_ITTE and a single L2_ITT.

The L2_ITT contains $(2^{\text{EVENTID_BITS}})$ entries.

If ITT_STRUCTURE is 1, and $\text{EVENTID_BITS} < (9 + (2 * \text{ITT_L2SZ}))$, it is possible to program L1_ITTE.SPAN to a value larger than EVENTID_BITS. In this case, the number of EventIDs for the DeviceID is defined by $2^{\text{EVENTID_BITS}}$, but the ITS is permitted to access all the L2_ITTEs in the range described by L1_ITTE.SPAN.

Arm recommends that ITT_STRUCTURE is 0 when $\text{EVENTID_BITS} \leq (9 + (2 * \text{ITT_L2SZ}))$.

When programming a reserved value or an unsupported value, the ITS behavior is CONSTRAINED UNPREDICTABLE to any behavior which could be achieved by programming a valid and supported value.

ITT_ADDR, bits [55:3]

Bits[55:3] of the base physical address of the ITT for the device described by this entry.

When ITT_STRUCTURE is 0, ITT_ADDR points to a linear ITT and all of the following are true:

- Bits[N:0] of the resulting address are 0 where $N = 2 + \text{EVENTID_BITS}$.
- This means that the level 2 ITTE array is aligned to the size of the array.

When ITT_STRUCTURE is 1, ITT_ADDR points to a 2-level ITT and all of the following are true:

- Bits[N:0] of the resulting address are 0 where N is:
 $\text{Max}(2, (\text{EVENTID_BITS} - (9 + (2 * \text{ITT_L2SZ})) + 2))$
- This means that the level 1 ITT is aligned to the size of the table.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in ITS_IDR0.PA_RANGE.

Bit [56]

Reserved, RES0.

DSWE, bit [57]

For any ITS event generated by the device described by this entry, disable reporting of software errors for the following error codes in ITS_SWERR_STATUSR.EC:

- 0x03.
- 0x04.
- 0x06.
- 0x07.
- 0x08.
- 0x0A.
- 0x0C.
- 0x0D.
- 0x13.
- 0x14.

DSWE	Meaning
0b0	Error reporting for the specified error codes is disabled.
0b1	Error reporting for the specified error codes is enabled.

This field is RES0, if ITS_IDR0.SWE is 0.

ITT_STRUCTURE, bit [58]

Whether the the ITT pointed to by ITT_ADDR uses a linear or 2-level structure.

ITT_STRUCTURE	Meaning
0b0	A linear ITT structure is used.

ITT_STRUCTURE	Meaning
0b1	A 2-level ITT structure is used.

If ITS_IDR1.ITT_LEVELS is 0, this field is RES0.

EVENTID_BITS, bits [63:59]

The number of EventID bits which can be translated for this device.

The ITT pointed to by ITT_ADDR contains $2^{\text{EVENTID_BITS}}$ level 2 ITT entries in total.

If this field is programmed to a value larger than the maximum, it is treated as having the maximum value for all other purposes than reading back the field. The maximum value is reported in ITS_IDR2.EVENTID_BITS.

11.1.3 L1_ITTE, Level 1 interrupt translation table entry


The L1_ITTE characteristics are:

Attributes

L1_ITTE is a 8-byte structure.

Field descriptions

The L1_ITTE bit assignments are:



VALID, bit [0]

Entry is valid

VALID	Meaning
0b0	This entry is invalid. All events identified by a DeviceID and EventID covered by this entry are ignored.
0b1	This entry is valid and L2_ADDR points to a level 2 ITT.

Bits [2:1]

Reserved, RES0.

L2_ADDR, bits [55:3]

Bits[55:3] of the level 2 array of ITT entries.

Bits[N:0] of the resulting address are 0 where $N = 2 + \text{SPAN}$.

This means that the level 2 ITTE array is aligned to the size of the array.

The level 2 array is accessed using the same PAS as the level 1 ITTE.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in ITS_IDR0.PA_RANGE.

Bits [59:56]

Reserved, RES0.

SPAN, bits [63:60]

Size of structure pointed to by L2_ADDR

The level 2 ITT pointed to by L2_ADDR contains 2^{SPAN} entries.

If this field is programmed to a value larger than the maximum, it is treated as having the maximum value for all other purposes than reading back the field. The maximum allowed value of this field is constrained by L2DTE.ITT_L2SZ.

If L2_DTE.ITT_L2SZ is 0b00, then the maximum value of this field is 9. If L2_DTE.ITT_L2SZ is 0b01, then the maximum value of this field is 11. If L2_DTE.ITT_L2SZ is 0b10, then the maximum value of this field is 13.

11.1.4 L2_ITTE, Level 2 interrupt translation table entry


The L2_ITTE characteristics are:

Attributes

L2_ITTE is a 8-byte structure.

Field descriptions

The L2_ITTE bit assignments are:



LPI_ID, bits [23:0]

Bits[23:0] of the LPI ID for the event described by this entry.

If unimplemented upper bits of the LPI_ID are not zero, it is IMPLEMENTATION DEFINED whether the upper bits are treated as 0 or the interrupt message is ignored by the IRS.

Bits [27:24]

Reserved, RES0.

DAC, bits [29:28]

Controls whether the event described by this entry can be associated with an Interrupt Domain different from the one associated with this ITS Domain.

DAC	Meaning	Applies
0b00	The event described by this entry is not permitted to be translated when it is generated by a write to a register in the PAS associated with another ITS Domain.	
0b01	The event described by this entry is permitted to be translated when it is generated by a write to a register in the Non-secure PAS.	When ITS_IDR2.XDMN_EVENTS == 0b01

Values not defined above are reserved.

The event is ignored and not translated if any of the following are true:

- The event is not associated with the Interrupt Domain specified by this field.
- A reserved value is programmed into this field.

VIRTUAL, bit [30]

Controls if the interrupt message generated to the IRS is for a physical or a virtual interrupt.

VIRTUAL	Meaning
0b0	The interrupt message generated by the ITS is a physical interrupt.

VIRTUAL	Meaning
0b1	The interrupt message generated by the ITS is a virtual interrupt.

When the ITT is used in the EL3 ITS Domain, this field is RES0.

VALID, bit [31]

Interrupt translation entry is valid.

VALID	Meaning
0b0	This entry is invalid. All events identified by a DeviceID and EventID covered by this entry are ignored.
0b1	This entry is valid and events identified by a DeviceID and EventID covered by this entry are translated into events to the IRS using the VM ID and LPI ID above.

VM_ID, bits [47:32]

Bits[15:0] of the VM ID passed to the IRS as part of the interrupt message targeting virtual interrupts.

When Virtual is 0, this field is RES0.

Bits [63:48]

Reserved, RES0.

11.2 IRS Data Structures

11.2.1 L1_VMTE, Level 1 VM table entry


The L1_VMTE characteristics are:

Attributes

L1_VMTE is a 8-byte structure.

Field descriptions

The L1_VMTE bit assignments are:



VALID, bit [0]

Entry is valid

VALID	Meaning
0b0	This entry is invalid. All lookups of a VM ID covered by this entry are ignored.
0b1	This entry is valid and L2_ADDR points to a level 2 VMT.

Bits [11:1]

Reserved, RES0.

L2_ADDR, bits [55:12]

Bits[55:12] of the address of the start of the array of level 2 VMT entries.

This means that the level 2 VMTE array is aligned to the size of a level 2 VMT (4KB).

The level 2 array is accessed using the same PAS as the level 1 VMTE.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in IRS_IDR0.PA_RANGE.

Bits [63:56]

Reserved, RES0.

11.2.2 L2_VMTE, Level 2 VM table entry


The L2_VMTE characteristics are:

Attributes

L2_VMTE is a 32-byte structure.

Field descriptions

The L2_VMTE bit assignments are:



VALID, bit [0]

Entry is valid

VALID	Meaning
0b0	This entry describes an invalid VM.
0b1	This entry describes a valid VM.

Bits [2:1]

Reserved, RES0.

VMD_ADDR, bits [55:3]

When *IRS_IDR3.VMD* == 1:

Bits[55:3] of the base address of the VM Descriptor for the VM described by this entry.

The VM Descriptor address is aligned to its size as follows:

- Bits[N:0] of the resulting address are 0 where $N = \text{IRS_IDR3.VMD_SZ} - 1$.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in IRS_IDR0.PA_RANGE.

Otherwise:

RES0

Bits [63:56]

Reserved, RES0.

Bits [66:64]

Reserved, RES0.

VPET_ADDR, bits [119:67]

Bits[55:3] of the base physical address of the array of VPE table entries for the VM described by this entry.

The VPE table always uses a linear structure and is aligned as follows:

- Bits[N:0] of the resulting address are 0 where $N = 2 + \text{VPE_ID_BITS}$.
- This means that the level 2 VPETE array is aligned to the size of the array.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in IRS_IDR0.PA_RANGE.

Bits [122:120]

Reserved, RES0.

VPE_ID_BITS, bits [127:123]

The number of VPEs which are supported for this VM.

The VPE table must contain $(2^{\text{VPE_ID_BITS}})$ VPE table entries in total.

If this field is programmed to a value larger than the maximum, it is treated as having the maximum value for all other purposes than reading back the field.

The maximum value is reported in IRS_IDR4.VPE_ID_BITS.

LPI_IST_VALID, bit [128]

LPI_IST_ADDR points to a valid IST

LPI_IST_VALID	Meaning
0b0	LPI_IST_ADDR is not valid. All accesses to LPI configuration and state for the VM are dropped.
0b1	LPI_IST_ADDR is valid and points to a valid IST.

LPI_IST_L2SZ, bits [130:129]

Level 2 IST size when a 2-level IST structure is used.

LPI_IST_L2SZ	Meaning
0b00	A level 2 IST is 4KB.
0b01	A level 2 IST is 16KB.

LPI_IST_L2SZ	Meaning
0b10	A level 2 IST is 64KB.

Values not defined above are reserved.

IRS_IDR2.IST_L2SZ reports the supported values.

If $LPI_ID_BITS \leq (10 - LPI_ISTSZ) + (2 * LPI_IST_L2SZ)$ and $LPI_IST_STRUCTURE$ is 1, all of the following are true:

- The IST consists of a single level 1 IST entry and a single level 2 IST.
- The level 2 IST contains $(2^{LPI_ID_BITS})$ entries.
- The IRS is allowed to access the full level 2 IST size as specified by LPI_IST_L2SZ .

Arm recommends that $LPI_IST_STRUCTURE$ is 0 when $LPI_ID_BITS \leq (10 - LPI_ISTSZ) + (2 * LPI_IST_L2SZ)$.

If programming a reserved value or an unsupported value, the IRS behavior is CONSTRAINED UNPREDICTABLE to any behavior which could be achieved by programming a valid and supported value.

When $LPI_IST_STRUCTURE$ is 0, this field is RES0.

Bits [133:131]

Reserved, RES0.

LPI_IST_ADDR, bits [183:134]

Bits[55:6] of the base physical address of the LPI IST for the VM described by this entry.

When $LPI_IST_STRUCTURE$ is 0, LPI_IST_ADDR points to a linear IST and all of the following are true:

- Bits[N:0] of the resulting address are 0 where $N = \text{Max}(5, ((LPI_ISTSZ + 2) + LPI_ID_BITS) - 1)$.
- This means that the level 2 IST entry array is aligned to the size of the array or to a 64-bytes boundary when its size is smaller than 64 bytes.

When $LPI_IST_STRUCTURE$ is 1, LPI_IST_ADDR points to the level 1 table in a 2-level IST and all of the following are true:

- Bits[N:0] of the resulting address are 0 where N depends on LPI_IST_L2SZ and LPI_ID_BITS as follows:

$$N = \text{Max}(5, LPI_ID_BITS - ((10 - LPI_ISTSZ) + (2 * LPI_IST_L2SZ)) + 2)$$
- This means that the level 1 IST is aligned to the size of the level 1 IST entry array or to a 64-byte boundary when its size is smaller than 64 bytes.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in $IRS_IDR0.PA_RANGE$.

LPI_ISTSZ, bits [185:184]

The size of each level 2 IST entry in the virtual LPI IST.

Values not defined above are reserved.

If this field is programmed to specify a size smaller than the minimum required size or programmed to a reserved value, it is treated as having the value corresponding to the minimum required size for all other purposes than a read of the entry.

LPI_ISTSZ	Meaning
0b00	The size of a level 2 LPI IST entry is 4 bytes.
0b01	The size of a level 2 LPI IST entry is 8 bytes.
0b10	The size of a level 2 LPI IST entry is 16 bytes.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

LPI_IST_STRUCTURE, bit [186]

Structure of the LPI IST pointed to by LPI_IST_ADDR.

LPI_IST_STRUCTURE	Meaning
0b0	A linear IST structure is used.
0b1	A 2-level IST structure is used.

Values not defined above are reserved.

When programming a reserved value, the IRS behavior is CONSTRAINED UNPREDICTABLE to any behavior which could be achieved by programming a valid value.

When IRS_IDR2.IST_LEVELS is 0, this field is treated as 0 for all other purposes than reading back the field.

LPI_ID_BITS, bits [191:187]

The number of LPIs which are supported for this VM.

The IST must contain $2^{(LPI_ID_BITS)}$ level 2 IST entries in total.

The minimum value for this field is IRS_IDR2.MIN_LPI_ID_BITS.

If programmed to a value smaller than the minimum, the field is treated as having the minimum value for all other purposes than reading back the field.

The maximum value for this field is IRS_IDR2.ID_BITS.

If this field is programmed to a value larger than the maximum, it is treated as having the maximum value for all other purposes than reading back the field. The maximum value is reported in IRS_IDR2.ID_BITS.

SPI_IST_VALID, bit [192]

Whether SPI_IST_ADDR points to a valid IST.

SPI_IST_VALID	Meaning
0b0	SPI_IST_ADDR is not valid. All accesses to SPI configuration and state for the VM are dropped.
0b1	SPI_IST_ADDR is valid and points to a valid IST.

SPI_IST_L2SZ, bits [194:193]

Level 2 IST size when a 2-level SPI IST structure is used.

SPI_IST_L2SZ	Meaning
0b00	A level 2 IST is 4KB.
0b01	A level 2 IST is 16KB.
0b10	A level 2 IST is 64KB.

Values not defined above are reserved.

IRS_IDR2.IST_L2SZ reports the supported values.

If $SPI_ID_BITS \leq ((10 - SPI_ISTSZ) + (2 * SPI_IST_L2SZ))$ and $SPI_IST_STRUCTURE$ is 1, all of the following are true:

- The IST consists of a single level 1 IST entry and a single level 2 IST.
- The level 2 IST contains $(2^{SPI_ID_BITS})$ entries.
- The IRS is allowed to access the full level 2 IST size as specified by SPI_IST_L2SZ .

Arm recommends that $SPI_IST_STRUCTURE$ is 0 when $SPI_ID_BITS \leq (10 - SPI_ISTSZ) + (2 * SPI_IST_L2SZ)$.

If programming a reserved value or an unsupported value, the IRS behavior is CONSTRAINED UNPREDICTABLE to any behavior which could be achieved by programming a valid and supported value.

When $SPI_IST_STRUCTURE$ is 0, this field is RES0.

Bits [197:195]

Reserved, RES0.

SPI_IST_ADDR, bits [247:198]

Bits[55:6] of the base physical address of the SPI IST for the VM described by this entry.

When $SPI_IST_STRUCTURE$ is 0, SPI_IST_ADDR points to a linear IST and all of the following are true:

- Bits[N:0] of the resulting address are 0 where $N = \text{Max}(5, ((SPI_ISTSZ + 2) + SPI_ID_BITS) - 1)$.
- This means that the level 2 IST entry array is aligned to the size of the array or to a 64-byte boundary when its size is smaller than 64 bytes.

When $SPI_IST_STRUCTURE$ is 1, SPI_IST_ADDR points to the level 1 table in a 2-level IST and all of the following are true:

- Bits[N:0] of the resulting address are 0 where N depends on the SPI_IST_L2SZ and SPI_ID_BITS field as follows:

$$N = \text{Max}(5, SPI_ID_BITS - ((10 - SPI_ISTSZ) + (2 * SPI_IST_L2SZ)) + 2)$$
- This means that the level 1 IST is aligned to the size of the level 1 IST entry array or to a 64-byte boundary when its size is smaller than 64 bytes.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in $IRS_IDR0.PA_RANGE$.

SPI_ISTSZ, bits [249:248]

The size of each level 2 IST entry in the virtual SPI IST.

Values not defined above are reserved.

If this field is programmed to specify a size smaller than the minimum required size or programmed to a reserved value, it is treated as having the value corresponding to the minimum required size for all other purposes than a read of the entry.

SPI_ISTSZ	Meaning
0b00	The size of a level 2 SPI IST entry is 4 bytes.
0b01	The size of a level 2 SPI IST entry is 8 bytes.
0b10	The size of a level 2 SPI IST entry is 16 bytes.

The reset behavior of this field is:

- On a GIC reset, this field resets to an UNKNOWN value.

SPI_IST_STRUCTURE, bit [250]

Structure of the SPI IST pointed to by SPI_IST_ADDR.

SPI_IST_STRUCTURE	Meaning
0b0	A linear IST structure is used.
0b1	A 2-level IST structure is used.

Values not defined above are reserved.

When programming a reserved value, the IRS behavior is CONSTRAINED UNPREDICTABLE to any behavior which could be achieved by programming a valid value.

When IRS_IDR2.IST_LEVELS is 0, this field is treated as 0 for all other purposes than reading back the field.

SPI_ID_BITS, bits [255:251]

The number of SPIs which are supported for this VM.

The SPI IST must contain $2^{(SPI_ID_BITS)}$ level 2 IST entries in total.

The minimum value for this field is 0, which is equivalent to a minimum of 1 SPI in the VM.

If the VM has no SPIs, SPI_IST_VALID is 0.

If this field is programmed to a value larger than the maximum, it is treated as having the maximum value for all other purposes than reading back the field. The maximum value is reported in IRS_IDR2.ID_BITS.

11.2.3 L1_ISTE, Level 1 interrupt state table entry


The L1_ISTE characteristics are:

Attributes

L1_ISTE is a 8-byte structure.

Field descriptions

The L1_ISTE bit assignments are:



VALID, bit [0]

Whether the entry is valid.

VALID	Meaning
0b0	This entry is invalid. All INTIDs covered by this entry are unreachable.
0b1	This entry is valid and L2_ADDR points to a level 2 IST.

Bits [11:1]

Reserved, RES0.

L2_ADDR, bits [55:12]

Bits[55:12] of the address of the start of the array of level 2 IST entries.

Bits[N:0] of the resulting address are 0 where $N = 11 + (2 * L2SZ)$

This means that the level 2 ISTE array is aligned to the size of a level 2 IST.

The level 2 array is accessed using the same PAS as the level 1 ISTE.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in IRS_IDR0.PA_RANGE.

Bits [63:56]

Reserved, RES0.

11.2.4 L2_ISTE, Level 2 interrupt state table entry

The L2_ISTE characteristics are:

Purpose

This data structure shows the 4 lowest address bytes of a single level 2 ISTE.

The size of a complete level 2 ISTE can be 4 bytes, 8 bytes, or 16 bytes.

If the size of a level 2 ISTE is more than 4 bytes, the higher address bytes are RES0.


See [4.7 The interrupt state table \(IST\)](#) for more information.

Attributes

L2_ISTE is a 4-byte structure.

Field descriptions

The L2_ISTE bit assignments are:



Pending, bit [0]

Interrupt Pending state

Pending	Meaning
0b0	The interrupt is not Pending.
0b1	The interrupt is Pending.

Active, bit [1]

Interrupt Active state

Active	Meaning
0b0	The interrupt is not Active.
0b1	The interrupt is Active.

HM, bit [2]

Handling mode of the interrupt.

HM	Meaning
0b0	Edge: The interrupt becomes Pending when the interrupt is acknowledged.
0b1	Level: The interrupt Pending state is unmodified when the interrupts is acknowledged.

Enable, bit [3]

Interrupt Enabled setting

Enable	Meaning
0b0	The interrupt is Disabled.
0b1	The interrupt is Enabled.

IRM, bit [4]

Interrupt Routing mode.

IRM	Meaning
0b0	The interrupt Routing mode is Targeted.
0b1	The interrupt Routing mode is 1ofN.

This field is RES0, if any of the following are true:

- IRS_IDR0.ONE_N is 0.
- The entry is part of a virtual IST and IRS_IDR0.VIRT_ONE_N is 0.

Bits [8:5]

Reserved, RES0.

HWU, bits [10:9]

Reserved for hardware use.

This field should be zero when an IST becomes valid and is otherwise IMPLEMENTATION DEFINED.

Priority, bits [15:11]

The Priority value of the interrupt.

Bits [4:N] of the priority value are implemented, where $N = (4 - \text{IRS_IDR1.PRI_BITS})$. Unimplemented bits are RES0.

This means that when fewer than 5 bits of priority is implemented, Priority[14 - IRS_IDR1.PRI_BITS:11] are RES0.

IAFFID, bits [31:16]

Interrupt Affinity ID.

When this entry is part of a physical IST, this field specifies a PE.

When this entry is part of a virtual IST, this field specifies a VPE.

11.2.5 VPETE, VPE table entry


The VPETE characteristics are:

Attributes

VPETE is a 8-byte structure.

Field descriptions

The VPETE bit assignments are:



VALID, bit [0]

Entry is valid

VALID	Meaning
0b0	This entry describes an invalid VPE.
0b1	This entry describes a valid VPE.

Bits [2:1]

Reserved, RES0.

VPED_ADDR, bits [55:3]

Bits[55:3] of the base address of the VPE Descriptor for the VPE described by this entry if used.

The VPE Descriptor address is aligned to its size as follows:

- Bits[N:0] of the resulting address are 0 where $N = \text{IRS_IDR4.VPED_SZ} - 1$.

In implementations that support fewer than 56 bits of physical address, any unimplemented upper bits are RES0. The number of implemented address bits is reported in IRS_IDR0.PA_RANGE.

Bits [63:56]

Reserved, RES0.

11.2.6 VM_DESC, VM descriptor

The VM_DESC characteristics are:

Purpose

The size, format, and content of this structure is IMPLEMENTATION DEFINED.


This shows an example size of 8-byte.

Attributes

VM_DESC is a 8-byte structure.

Field descriptions

The VM_DESC bit assignments are:



IMPDEF, bits [63:0]

IMPLEMENTATION DEFINED

11.2.7 VPE_DESC, VPE descriptor

The VPE_DESC characteristics are:

Purpose

The size, format, and content of this structure is IMPLEMENTATION DEFINED.


This shows an example size of 32-byte.

Attributes

VPE_DESC is a 32-byte structure.

Field descriptions

The VPE_DESC bit assignments are:



IMPDEF, bits [255:0]

IMPLEMENTATION DEFINED

Part A

GLCv5 Stream Protocol interface

Chapter A1

GICv5 Stream Protocol overview

I _{TGPMY}	The GICv5 Stream Protocol is one approach to connecting an IRS and a CPU interface. The GICv5 Stream Protocol supports independent development of an IRS and a PE. Arm recommends that a GICv5 implementation uses the GICv5 Stream Protocol.
I _{VPFYV}	The GICv5 Stream Protocol is based on the GIC Stream Protocol defined by GICv3[3].
D _{KNBQZ}	An <i>end-point</i> means either a CPU interface or an interrupt source.
I _{WYMBT}	The GICv5 Stream Protocol supports the connection between an IRS and one or more end-points.
D _{NXTBL}	The connection between an IRS and a CPU interface is referred to as an <i>Interrupt Handling channel</i> .
D _{SBHHG}	The connection between an IRS and an interrupt source is referred to as an <i>Interrupt Signaling channel</i> .
I _{FSTMD}	<p>GICv5 Stream Protocol uses a transport layer formed of a pair of AMBA AXI5-Stream Protocol links:</p> <ul style="list-style-type: none">• From the IRS to one or more end-points.• From one or more end-points to the IRS. <p>See <i>AMBA® AXI Protocol Specification</i>[7] for information about the AMBA AXI5-Stream Protocol.</p>
D _{WVBHL}	<i>AXI Stream connection</i> refers to the transport layer connection between an IRS and one or more end-points.
D _{DCXRK}	A command sent by the IRS to an end-point is defined as a <i>Downstream</i> command.
D _{PTTTR}	A command sent by an end-point to the IRS is defined as an <i>Upstream</i> command.

Chapter A2

AMBA AXI5-Stream Transport Layer

1_VSYKR

The GICv5 Stream Protocol is based on the following unidirectional AXI5-Stream connections:

- A downstream AXI5-Stream Interface containing connections from an IRS to one or more end-points.
- An upstream AXI5-Stream Interface containing connections from one or more end-points to an IRS.

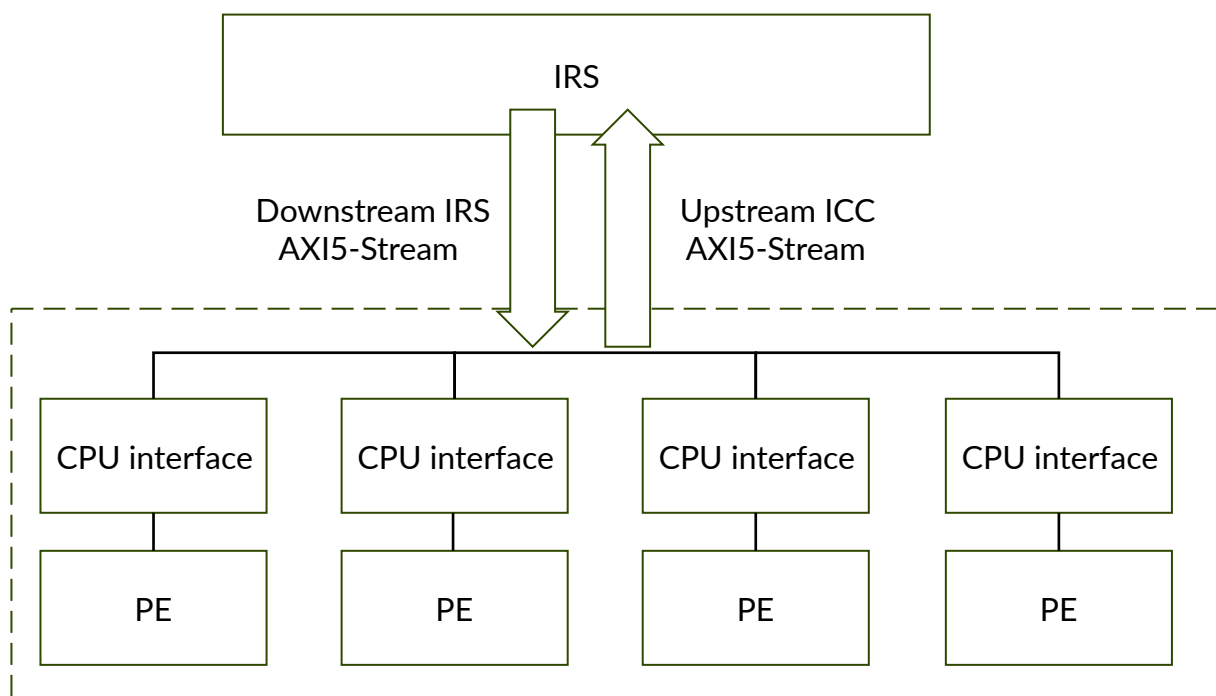


Figure A2.1: GICv5 Stream Protocol Interface

A GICv5 Stream Protocol connection connects a combination of end-point types:

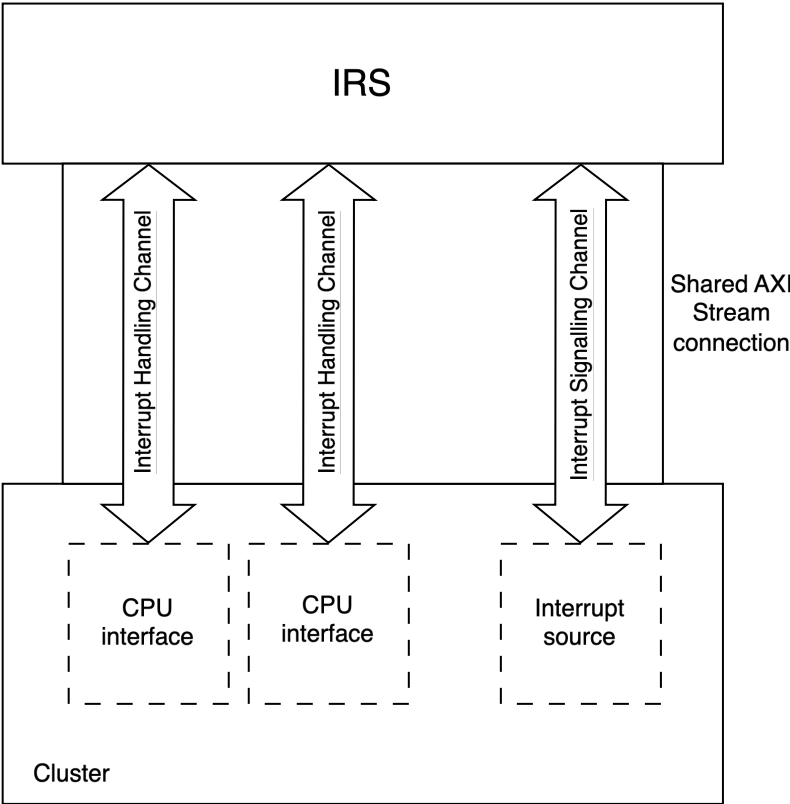


Figure A2.2: GICv5 Stream Protocol Interface

An IRS might have multiple GICv5 Stream Protocol connections to different groups of PEs, for example to connect to multiple clusters.




Figure A2.3: IRS with multiple GICv5 Stream Interfaces

R_{XR}VG D

An interconnect between an IRS 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.

A2.1 Signals

R _{JSSRX}	The interface requires a global clock, ACLK, and a reset signal, ARESETn.
I _{VPLGV}	For the GICv5 Stream Protocol, each stream interface is identified by a prefix to the AXI Stream signal names: <ul style="list-style-type: none"> Downstream signals from an IRS to the CPU interface are prefixed with the letters IRS. Upstream signals from the CPU interface to an IRS are prefixed with the letters ICC.
R _{GQLTR}	The following table illustrates the GICv5 Stream Protocol interface signals from the IRS to the downstream end-point:

Signal	Description
IRSTVALID	When set to 1, this signal indicates that the Requester is driving a valid transfer.
IRSTREADY	When set to 1, this signal indicates that the Completer can accept a transfer in the current cycle.
IRSTDATA[BN:0]	The interface data path.
IRSTLAST	When set to 1, this signal indicates the final transfer of a command.
IRSTID	This signal identifies the channel the transfer is associated with.
IRSTDEST[N:0]	This signal identifies the target CPU interface to provide routing information for the stream.
IRSTWAKEUP	Indicates if there is activity associated with the AXI5-Stream interface.

The following table illustrates the GICv5 Stream Protocol interface signals from the end-point to the upstream IRS:

Signal	Description
ICCTVALID	When set to 1, this signal indicates that the Requester is driving a valid transfer.
ICCTREADY	When set to 1, this signal indicates that the Completer 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 command.
ICCTID[N:0]	This signal identifies the originating CPU interface, to provide routing information for the stream.
ICCTDEST	This signal identifies the channel the transfer is associated with.
ICCTWAKEUP	Indicates if there is activity associated with the AXI5-Stream interface.

Where:

- 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 $\log_2(M)$ rounded up to the nearest integer, where M is the number of PEs supported by the interface plus 1.

I _{SRVVS}	In the GICv5 Stream Protocol, all commands have a length which is a multiple of 16 bits.
--------------------	--

I_{DWDQB}	<p>The xTKEEP and xTSTRB signals are defined as optional in <i>AMBA® AXI-Stream Protocol Specification</i>[13].</p> <p>If the AXI5-Stream interfaces used for a GICv5 Stream Protocol link have data path that is 16-bit or smaller, Arm recommends that xTKEEP is not implemented. Otherwise, Arm recommends that xTKEEP is implemented.</p> <p>If the xTKEEP signal is implemented, it is asserted for bytes required for the command and deasserted for all the other bytes. If the xTKEEP signal is not implemented, any bytes not required to transmit a command are set to 0b00.</p> <p>Arm recommends that the xTSTRB signal is not implemented. If xTSTRB is implemented, it is asserted for all bytes for which xTKEEP is asserted and deasserted for all the other bytes.</p>
R_{ZZDQV}	<p>The xTLAST signal is asserted to indicate the last transfer for a command.</p> <p>A sender may assert xTLAST before all bytes of a command have been transferred. Any untransferred bytes in the command are treated as having the value 0b00 by the receiver.</p> <p>A sender may assert xTLAST late, making additional transfers not needed for the command. The receiver ignores any bytes not needed to represent the command.</p>
I_{QDLHQ}	<p>If parity checking is required, Arm recommends implementing the support for parity in AXI Stream defined in <i>AMBA® AXI-Stream Protocol Specification</i>[13].</p>
I_{TNVVV}	<p>For further information about the signals used by the GICv5 Stream Protocol interface, and for details about handshaking, see <i>AMBA® AXI-Stream Protocol Specification</i>[13].</p>

A2.2 Channel identification

I _{NMPSG}	The AXI Stream connection can connect one IRS to multiple end-points. A connection between the IRS and an end-point is referred to as a channel. An Interrupt Handling channel connects an IRS to a CPU interface. An Interrupt Signaling channel connects an IRS to an interrupt source.
I _{PDQSW}	The xTID and xTDEST signals are used to route packets and identify the type of channel.
R _{BRZBR}	For downstream commands, IRSTID indicates the channel type: 0: Interrupt Handling channel IRSTDEST indicates the target CPU interface. Values are allocated contiguously from 0, in order of ascending affinity. 1: Interrupt Signaling channel IRSTDEST indicates the target interrupt source.
R _{RBSBN}	For upstream commands, ICCTDEST indicates the channel type: 0: Interrupt Handling channel. ICCTID indicates the originating CPU interface. 1: Interrupt Signaling channel. ICCTID indicates the originating interrupt source.
R _{TCSCP}	The order and encoding of end-points is the same for IRSTDEST and IDDTID.
R _{JCRRQ}	The values used to identify CPU interfaces are allocated contiguously from 0, in order of ascending affinity.
R _{WZJVS}	The values used to identify interrupt sources is IMPLEMENTATION DEFINED.
I _{ZJTKX}	Arm recommends that the values used to identify interrupt sources are allocated contiguously from 0.
R _{NDDWL}	Transfers on the same AXI5-Stream Interface for different commands are not interleaved. Once transmission of a command is started subsequent transfers continue the same command, with the same values of xTID and xTDEST, until xTLAST is asserted.
I _{RQBR}	The restriction on interleaving transfers does not prevent upstream and downstream commands being sent in parallel.
I _{DMDNC}	An example system might have two CPU interfaces and an interrupt source sharing an AXI Stream connection, as shown in the diagram below:

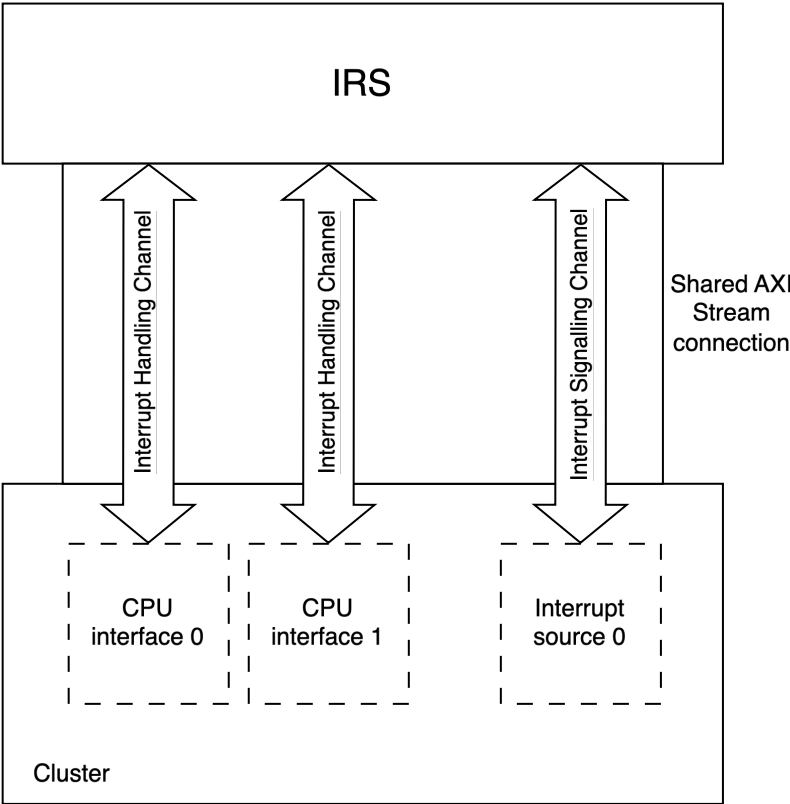


Figure A2.4: GICv5 Stream Protocol Interface

Downstream commands:

IRSTID	IRSTDEST	Channel type	Target
0	0	Interrupt Handling	CPU interface 0
0	1	Interrupt Handling	CPU interface 1
1	0	Interrupt Signaling	Interrupt source 0

Upstream commands:

ICCTDEST	ICCTID	Channel type	Source
0	0	Interrupt Handling	CPU interface 0
0	1	Interrupt Handling	CPU interface 1
1	0	Interrupt Signaling	Interrupt source 0

A2.3 Link status

$D_{WDBC\bar{G}}$	When the AXI Stream connection is described as <i>online</i> , data can be transferred across the upstream and downstream AXI Stream links. When a AXI Stream connection is described as <i>offline</i> , no data can be transferred across the AXI Stream links.
R_{SMLNR}	When one or more of the channels that share the AXI Stream connection is online the AXI Stream connection is online.
$R_{DG\bar{D}FT}$	When none of the channels that share the AXI Stream connection is online, it is IMPLEMENTATION DEFINED whether the AXI Stream connection is online or offline.
R_{NBCWM}	How the AXI Stream connection transitions from online to offline, or offline to online, is IMPLEMENTATION DEFINED.

Chapter A3

Common behaviors

I_{QHRMH}	This section defines behaviors common to the Interrupt Handling and Interrupt Signaling channels.
D_{PWLWC}	Some commands require an acknowledgement from the receiver. A command that requires an acknowledgement is defined as <i>outstanding</i> until the acknowledgement is received.
D_{TYKBT}	A command that requires acknowledgement is described as <i>complete</i> once the acknowledgement is received. All other commands are described as complete once received by the recipient.
R_{NQQBH}	If multiple transfers are required to transmit a command, the receiver does not acknowledge the command until all of the command has been received.

Chapter A4

Interrupt Handling channel

- \mathbb{I}_{HTMK} The Interrupt Handling channel defines the connection between one CPU interface and one IRS, it includes the following capabilities:
- Forwarding candidate HPPIs to the CPU interface.
 - Activating and deactivating interrupts as they are handled.
 - Managing the current resident VPE.
 - Managing interrupt configuration and state.
- \mathbb{I}_{LYNC} The rules in this section only apply to the Interrupt Handling channel.

A4.1 Command summary

I_{LHGV} The upstream commands that a CPU interface can send to the IRS are:

ID	Command	Description
0x0	UpstreamControl	Communicates control information from the CPU interface to the IRS.
0x1	-	-
0x2	Activate	Activates a previously sent interrupt.
0x3	Release	Releases a previously sent interrupt.
0x4	DownstreamControl Ack	Acknowledges a DownstreamControl command.
0x5	-	-
0x6	Deactivate	Deactivates an interrupt.
0x7	IMPLEMENTATION DEFINED	Reserved for IMPLEMENTATION DEFINED functionality.
0x8	SetEnabled	Sets the individual Enabled value of an interrupt.
0x9	SetTarget	Sets an LPI or SPI's target PE.
0xA	SetPriority	Sets an interrupt's priority.
0xB	SetHandling	Sets an interrupt's Handling mode.
0xC	RequestConfig	Requests an interrupt's configuration and state
0xD	Sync	Requests confirmation that the effects of previous commands are globally observable.
0xE	SetPending	Sets an LPI or SPI's Pending state.
0xF	SetResident	Sets the current resident VPE.

I_{PYPQZ} The downstream commands that an IRS can send to the CPU interface are:

ID	Command	Description
0x0	UpstreamControl Ack	Acknowledges an UpstreamControl command.
0x1	Forward	Sets an interrupt as pending on the CPU interface.
0x2	Activate Ack	Acknowledges an Activate command.
0x3	Recall	Recalls a previously sent pending interrupt.
0x4	DownstreamControl	Communicates control information from the IRS to the CPU interface.
0x5	WakeRequest	Requests the channel is brought online.
0x6	Deactivate Ack	Acknowledges a Deactivate command.
0x7	IMPLEMENTATION DEFINED	Reserved for IMPLEMENTATION DEFINED functionality.
0x8	Set Ack	Acknowledges a SetEnabled, SetPending, SetPriority, SetTarget, or SetHandling command.
0x9	-	-
0xA	-	-
0xB	-	-

ID	Command	Description
0xC	RequestConfig Ack	Acknowledges a RequestConfig command and returns the requested data.
0xD	Sync Ack	Acknowledges a Sync command, indicating that the effect of previous commands are globally observable.
0xE	-	-
0xF	SetResident Ack	Acknowledges a SetResident command.

R_{LDZYS} ID 0x7 in upstream and downstream is reserved for IMPLEMENTATION DEFINED functionality.

When supported, IMPLEMENTATION DEFINED commands do not impact the functionality of other commands. If not supported, IMPLEMENTATION DEFINED commands are ignored.

Arm strongly recommends that IMPLEMENTATION DEFINED commands are not used unless it can be guaranteed that both the sender and receiver agree on the usage.

A4.2 Outstanding commands

R_{DGPBP} Commands which require acknowledgement are acknowledged in finite time, apart from Forward commands.

I_{YYGKH} The following commands require acknowledgement:

Command	Direction	Acknowledged by
Activate	Upstream	Activate Ack or Forward with <code>ActivateAck</code> set to 1.
Deactivate	Upstream	Deactivate Ack
DownstreamControl	Downstream	DowstreamControl Ack.
RequestConfig	Upstream	RequestConfig Ack
Forward	Downstream	Activate or Release. SetResident, with <code>Valid</code> set to 0. UpstreamControl with Identifier set to 0b0001 (Quiesce)
SetEnabled	Upstream	Set Ack
SetPending	Upstream	Set Ack
SetHandling	Upstream	Set Ack
SetPriority	Upstream	Set Ack
SetResident	Upstream	SetResident Ack
SetTarget	Upstream	Set Ack
Sync	Upstream	Sync Ack
UpstreamControl	Upstream	UpstreamControl Ack

All other commands do not require acknowledgement.

I_{CWTZR} For commands that require acknowledgement, the architecture places limits on the number of outstanding commands that are permitted.

I_{YYPPQ} An UpstreamControl command with Identifier set to 0b0000 (Reset) initiates a reset of the channel. Resetting the channel returns pending interrupts to the IRS and cancels outstanding commands. Acknowledges to commands issued before the reset started might still be received during the reset process.

R_{YDBHS} For each Interrupt Handling channel, the IRS is permitted to have one DownstreamControl command outstanding.

I_{KGBV} The WakeRequest command does not have an acknowledge, meaning that is considered complete after being sent rather than requiring an explicit acknowledge. WakeRequests are implicitly acknowledged by the CPU interface sending UpstreamControl with Identifier set to 0b0000 (Reset) to bring the channel online.

I_{MJDMH} Arm recommends that once an IRS has issued a WakeRequest command it does not send further wake requests until the next time the channel is offline.

R_{PSRJH} For each Interrupt Handling channel, the CPU interface is permitted to have the following numbers of outstanding commands:

- One Activate command.
- One RequestConfig command.
- One SetResident command.
- One UpstreamControl command.
- One Sync command.

- One Deactivate command.
- In total, one of the following commands:
 - SetEnabled.
 - SetHandling.
 - SetPending.
 - SetPriority.
 - SetTarget.

A4.3 Connection management

D _{GRHNS}	When an Interrupt Handling channel is described as <i>online</i> , upstream and downstream commands can be sent between the IRS and CPU interface. When an Interrupt Handling channel is described as <i>offline</i> , the only commands that can be sent are those to bring the channel online.
R _{GYWDK}	When an Interrupt Handling channel is offline: <ul style="list-style-type: none"> • The IRS sends no commands other than WakeRequest. • The CPU interface sends no commands other than UpstreamControl with Identifier set to 0b0000 (Reset).
R _{WZVTW}	If an Interrupt Handling channel is offline, the IRS may send a WakeRequest command to request the channel to be brought online. Whether an IRS sends wake requests over GICv5 Stream is IMPLEMENTATION DEFINED and an IRS might use other IMPLEMENTATION DEFINED mechanisms as well as or instead of GICv5 Stream.
I _{TXLVK}	Arm expects an IRS to only send a wake request if there is a candidate HPPI for the PE.
R _{ZPMWQ}	The IRS does not send WakeRequest commands if the Interrupt Handling channel is online.
I _{TSRMX}	A WakeRequest is implicitly acknowledged by an UpstreamControl command with Identifier set to 0b0000 (Reset).
R _{BTJBS}	To bring an Interrupt Handling channel online, or reset an online channel: <ul style="list-style-type: none"> • The CPU interface sends an UpstreamControl command with Identifier set to 0b0000 (Reset). • The IRS responds with a DownstreamControl command with Identifier set to 0b0001 (Flush). • The CPU interface acknowledges the DownstreamControl command, with Flush set to 1. • The IRS acknowledges the UpstreamControl command. <p>Once the UpstreamControl command is acknowledged, the Interrupt Handling channel is online.</p>
R _{RNZYL}	When an UpstreamControl command with Identifier set to 0b0000 (Reset) is acknowledged: <ul style="list-style-type: none"> • There are no outstanding upstream or downstream commands. • There are no pending physical or virtual interrupts from the IRS pending on the CPU interface. • There is no resident VPE. • There are no 1ofN selection hints set.
R _{MSHYJ}	Between sending an UpstreamControl command with Identifier set to 0b0000 (Reset) and the command being acknowledged, the CPU interface sends no commands other than to acknowledge a DownstreamControl with Identifier set to 0b0001 (Flush). Between receiving an UpstreamControl command with Identifier set to 0b0000 (Reset) and sending the acknowledge, the IRS sends no other commands than DownstreamControl with Identifier set to 0b0001 (Flush).
I _{PQTNV}	If a reset request is received part way through a previous incomplete reset sequence, it is possible that operations are seen out of order or repeated.
R _{DLBTV}	If the IRS receives a DownstreamControl Acknowledge with Flush set to 1 before it has sent the DownstreamControl with Identifier set to 0b0001 (Flush) due the most recent UpstreamControl command with Identifier set to 0b0000 (Reset), it is IMPLEMENTATION DEFINED whether the IRS sends DownstreamControl. If the DownstreamControl command with Identifier set to 0b0001 (Flush) is sent, the IRS treats the command as already acknowledged. If the CPU interface receives a DownstreamControl with Identifier set to 0b0001 (Flush) after it has sent DownstreamControl Acknowledge with Flush sent to 1, the command is ignored and the DownstreamControl is treated as already acknowledged.
I _{GSQPK}	GICv5 uses IAFFIDs to identify the different connected PEs. The IRS informs the CPU interface of its IAFFID as part of bringing the Interrupt Handling Channel online. Some systems might require an IMPLEMENTATION DEFINED sequence before the correct IAFFIDs are known for the CPU interfaces. If the Interrupt Handling Channel is brought online before this sequence is completed the

IRS can re-communicate the correct IAFFID using DownstreamControl command with Identifier set to 0b0010 (IAFFID).

In the PE, the IAFFID is reported in ICC_IAFFIDR_EL1.

I_CDBHF If a system requires an IMPLEMENTATION DEFINED sequence to initialise the correct IAFFIDs, Arm recommends that this occurs no more than once between GIC resets.

R_GYHSP The CPU interface does not cache the IAFFID value across resets of the Interrupt Handling Channel.

I_SYLKT Under normal operation, an UpstreamControl command with Identifier set to 0b0000 (Reset) is sent only to bring the connection between the CPU interface and IRS online. In which case there are no other outstanding commands, outstanding pending interrupts or resident VPE.

However, the CPU interface may send an UpstreamControl command with Identifier set to 0b0000 when the Interrupt Handling channel is already online. Sending the command resets the connection, returning all pending interrupts and cancelling outstanding commands.

When an online connection is being reset, it is possible that the CPU interface receives downstream commands between sending the UpstreamControl and seeing the acknowledgement. The CPU interface does not respond to these commands and the IRS does not expect to receive responses.

R_HXGXY To take an Interrupt Handling channel offline:

- The CPU interface sends an UpstreamControl command with Identifier set to 0b0001 (Quiesce).
 - The CPU interface sends no further commands, other than acknowledgements to downstream commands, until UpstreamControl Ack is received.
- The IRS completes any outstanding requests from the CPU interface, then sends an UpstreamControl Ack.

Once the UpstreamControl command is acknowledged, the Interrupt Handling channel is offline.

I_QDJXY There is a single Interrupt Handling channel between the CPU interface and the IRS, shared by all implemented Interrupt Domains. Therefore, the UpstreamControl commands are independent of Interrupt Domain.

R_YZFJG When the Interrupt Handling channel is offline, the IRS and CPU interface treat all previously forwarded physical or virtual interrupts having been recalled.

I_QNSGW Taking the Interrupt Handling channel offline acts as an implicit Release for any pending physical or virtual interrupts of any Interrupt Domain, including Forwards received after the Quiesce is sent. If there is an outstanding acknowledge for an interrupt, the CPU interface must send the Activate command before sending the UpstreamControl command with Identifier set to 0b0001 (Quiesce).

I_GWCSK The UpstreamControl command provides a mechanism for the CPU interface to provide hints to the 1ofN selection algorithm in the IRS. The algorithm for 1ofN selection is IMPLEMENTATION DEFINED and might ignore any hints provided by the CPU interface. Whether a CPU interface sends 1ofN hints is IMPLEMENTATION DEFINED.

R_FMGVS When an Interrupt Handling channel is online, the CPU interface can set hints for 1ofN target selection.

The CPU interface sends an UpstreamControl command with:

- Identifier set to 0b0010 to set 1ofN hints.
- Identifier set to 0b0011 to clear 1ofN hints.

The Data field in the UpstreamControl command specifies which hints are being set or cleared. The state of a 1ofN hint not specified in Data is unaffected by set or clear operations.

Unsupported or unknown hints are ignored by the IRS.

I_MLPDY The CPU interface is permitted to clear hints that it has not previously set.

I_YCCYG Each hint defines which Interrupt Domain, or Domains, it applies to.

A4.4 Managing the resident VPE

<code>I_{XNGZK}</code>	When an Interrupt Handling channel transitions from offline to online, there is no resident VPE.
<code>R_{TPKML}</code>	<p>To make a VPE resident, the CPU interface sends a SetResident command with <code>Valid</code> set to 1. The VPE being made resident is identified by the following fields:</p> <ul style="list-style-type: none"> • <code>Domain</code> - the Physical Interrupt Domain the VPE is associated with. • <code>VM</code> - the virtual machine the VPE is associated with. • <code>VPE</code> - the VPE within the specified virtual machine. <p>The IRS acknowledges the SetResident command with a SetResident Ack command with <code>Fault</code> set to indicate whether specified VPE was made resident.</p> <p>The IRS does not generate a SetResident Ack until one of the following is true:</p> <ul style="list-style-type: none"> • The specified VPE was made resident and one of the following is true: <ul style="list-style-type: none"> – The IRS has forwarded the first virtual interrupt for the VPE. – The IRS has determined that there is no virtual interrupt to forward. • The IRS has determined that it is unable to make the specified VPE resident. <p>Once the SetResident Ack command with <code>Fault</code> set to 0 is sent, the specified VPE is resident.</p>
<code>I_{HSHXS}</code>	An IRS might be unable to make the requested VPE resident, for example because the VPE is invalid or the VPE is already resident on a different CPU interface.
<code>I_{VSJLW}</code>	There might be no virtual interrupt to forward because there is no candidate HPPI for the VPE.
<code>R_{HLJNT}</code>	<p>To clear the resident VPE, the CPU interface sends a SetResident command with <code>Valid</code> set to 0. When the IRS acknowledges the SetResident command by sending a SetResident Ack, there is no resident VPE.</p> <p>When the IRS acknowledges the SetResident command by sending a SetResident Ack, the IRS and CPU interface treat all previously forwarded virtual interrupts for the VPE as having been recalled.</p>
<code>I_{QTCTP}</code>	<p>Clearing the resident VPE acts as an implicit Release for any outstanding pending virtual interrupts, returning the interrupts to the IRS.</p> <p>If there is any outstanding Activate for a virtual interrupt, the CPU interface must send the Activate command before sending SetResident.</p>
<code>I_{LDFMM}</code>	The CPU interface might send a SetResident command with <code>Valid</code> set to 0 while earlier virtual commands remain outstanding. The IRS is not required to delay sending SetResident Ack until it has acknowledged those earlier virtual commands. However, the IRS must process commands with the context that was valid at the time the command was received.
<code>R_{VCZPC}</code>	The CPU interface does not send a SetResident command with <code>Valid</code> set to 1 when there is a resident VPE.
<code>R_{RZRYM}</code>	Between sending SetResident and receiving SetResident Ack, the CPU interface sends no commands with <code>Virtual</code> set to 1 other than SetPending, and commands that acknowledge downstream commands.

A4.4.1 Interrupt Handling Channel behaviors when there is a resident VPE

<code>R_{MGZZN}</code>	All downstream and upstream commands with <code>Virtual</code> set to 1, other than SetPending, are for the resident VPE. The VM targeted by SetPending is specified as part of the command's payload.
<code>R_{PYWKM}</code>	<p>If the IRS receives a command with <code>Virtual</code> set to 1 where <code>Domain</code> does not correspond to the IRS's record of the resident VPE, the behavior depends on the type of the command as follows:</p> <ul style="list-style-type: none"> • RequestConfig Ack commands return <code>F</code> as 1 with zeros for all configuration fields. • SetResident commands with <code>Valid</code> set to 0 clear the resident VPE, and <code>DB</code> is treated as being 0. • All other commands are acknowledged but otherwise have no effect.

R_{QMVYG}

If the CPU interface receives a command with `Virtual` set to 1 and where `Domain` does not correspond to the current value of `SCR_EL3.{NSE,NS}`, the behavior depends on the type of the command as follows:

- Forward command: The command replaces any previously received virtual interrupts, but the new interrupt is not considered when determining the virtual HPPI.
 - If the Forward command results in an implicit Recall, the CPU interface sets `Domain` in the Release or Activate command to the Interrupt Domain of the interrupt that is being Recalled.
- Recall command: Any resulting Activate or Release command is issued with `Domain` set to the Interrupt Domain of the interrupt that is being Recalled.

A4.4.2 Interrupt Handling Channel behaviors when there is no resident VPE

R_{JRPVW}

If there is no resident VPE:

- The IRS:
 - Does not issue Forward commands with `Virtual` set to 1, except as part of the sequence to make a VPE resident.
 - Acknowledges `RequestConfig` commands that have `Virtual` set to 1 with `RequestConfig Ack` commands that have `F` set to 1.
 - Acknowledges, but otherwise ignores, all other upstream commands with `Virtual` set to 1.
- The CPU interface:
 - Does not issue upstream commands with `Virtual` set to 1, other than `SetResident` and `SetPending`.
 - Downstream Forward commands with `Virtual` set to 1 are not considered when determining the virtual HPPI until a VPE is resident.

A4.5 Forwarding, recalling, and releasing interrupts

<code>I_{GLBSF}</code>	<p>The IRS forwards a candidate HPPI for an Interrupt Domain, or for the resident VPE, using a Forward command. An interrupt can be retrieved from the CPU interface using a Recall command, or by sending another Forward command for the same Interrupt Domain.</p> <p>For each Physical Interrupt Domain, and for the resident VPE, the IRS presents at most one candidate HPPI at a time. If the selected candidate HPPI changes, the IRS sends a replacement Forward command, which acts an implicit Recall if there is an outstanding Forward for the same Interrupt Domain.</p> <p>A CPU interface might issue an Activate instead of a Release for the previous pending interrupt.</p> <p>The protocol does not require that the new Forward is for an interrupt with higher priority than the outstanding Forward it replaces. A lower priority interrupt might replace a higher priority interrupt. For example, if an LPI is being re-targeted while it is pending, it must be retrieved from the old target CPU interface before being presented to the new target CPU interface. If there are other pending interrupts for the old target, one of these might now be presented to the CPU interface and this interrupt could be lower priority than the interrupt being re-targeted.</p>
<code>I_{TZBHV}</code>	There is no limit on the number of Forward commands that can be outstanding.
<code>R_{KNCCF}</code>	For a given Interrupt Domain, there can never be multiple outstanding Forward commands for the same INTID.
<code>I_{CLNKP}</code>	If the IRS needs to re-issue a pending interrupt with new properties, for example a different priority value, it must first fully retrieve the previous pending interrupt.
<code>R_{GGHVV}</code>	<p>When issuing a Forward command, the IRS sets <code>Returnable</code> based on the routing mode of the pending interrupt being forwarded:</p> <ul style="list-style-type: none"> • Targeted: <code>Returnable</code> is set to 0. • 1ofN: It is IMPLEMENTATION DEFINED whether <code>Returnable</code> is set to 0 or 1. <ul style="list-style-type: none"> – Arm recommends that an implementation does not set <code>Returnable</code> to 1 unless there are other valid targets for the interrupt.
<code>R_{HWRFH}</code>	<p>Forward commands are acknowledged by either an Activate command or Release command. For a given Interrupt Domain, or for the resident VPE, Forward commands are acknowledged in order.</p> <p>A single Release command can acknowledge multiple Forward commands. The Release command acknowledges the INTID specified in the command, and all earlier unacknowledged Forward commands for the same Interrupt Domain or resident VPE.</p>
<code>I_{HJMSR}</code>	For example, the CPU interface will acknowledge all the Forward commands for physical interrupts for the Non-secure Domain in the order it received them. But Forward commands for different Domains might be acknowledged out of order with respect to each other.
<code>R_{MYXCZ}</code>	For edge-triggered interrupts, receiving an ActivateAck command from the IRS in response to an Activate command guarantees that an edge generated by an interrupt source as a direct or indirect result of an instruction executed on the PE after observing the ActivateAck will not be merged into the acknowledged instance of the interrupt.
<code>R_{XSDST}</code>	<p>The CPU interface issues Release commands in response to:</p> <ul style="list-style-type: none"> • A Recall command. • A Forward that has been replaced by a subsequent Forward for the same Interrupt Domain. • The most recent Forward for an Interrupt Domain had <code>Returnable</code> set to 1 and the CPU interface has for IMPLEMENTATION DEFINED reasons decided to return the interrupt.
<code>R_{MPTHB}</code>	For a given Interrupt Domain, or for the resident VPE, a Recall command applies to the most recent Forward command.
<code>R_{PKSKS}</code>	If a Recall command is received when there is no outstanding Forward command, the command is ignored.
<code>R_{GMQMJ}</code>	An IRS is able to accept an Activate or Release command for any outstanding Forward.

I_{TJkMB}

The GICv5 Stream Protocol has no way for an IRS to signal that a Release or Activate has failed. Therefore, an IRS only forwards an interrupt to a CPU interface if it can accept either an Activate or Release for that interrupt.

A4.6 INTID configuration

<code>R_{DHHRZ}</code>	<p>For any of the following the commands, when the IRS acknowledges the command, the IRS has observed the command and the effects of the command will be globally observable in finite time:</p> <ul style="list-style-type: none"> • Deactivate • SetEnabled • SetHandling • SetPending • SetPriority • SetTarget
<code>R_{DTMNP}</code>	<p>The CPU interface issues an upstream Sync command to synchronize the effects of earlier commands.</p> <p>The IRS responds with a Sync Ack once all the following are true:</p> <ul style="list-style-type: none"> • Any Activate, Deactivate, Disable, SetEnabled, SetPending, SetPriority, or SetTarget command received before the Sync has been acknowledged. • The effects of any Activate, Deactivate, Disable, SetEnabled, SetHandling, SetPending, SetPriority, or SetTarget command received before the Sync are globally observable.
<code>R_{RXYSV}</code>	<p>Between sending Sync and receiving Sync Ack, the CPU interface sends no commands other than to acknowledge downstream commands.</p>
<code>I_{PLQBY}</code>	<p>The CPU interface might send an UpstreamControl command with <code>Identifier</code> set to Reset between sending a Sync and receiving the Sync Ack.</p> <p>Arm expects resetting an online channel to be rare and be due to the PE resetting. If an online channel is reset, it is possible that some stale downstream commands are received from the IRS. The CPU interface ignores such commands.</p>
<code>I_{MDVSS}</code>	<p>The CPU interface can request the current configuration and state of an INTID by sending a RequestConfig command. The IRS returns the requested INTID's configuration when acknowledging the command with RequestConfig Ack.</p>
<code>R_{DBJFG}</code>	<p>Between sending RequestConfig and receiving RequestConfig Ack, the CPU interface sends no commands other than to acknowledge downstream commands.</p>
<code>I_{GSGBS}</code>	<p>The CPU interface might send an UpstreamControl command with <code>Identifier</code> set to Reset between sending a RequestConfig and receiving the RequestConfig Ack.</p> <p>Arm expects resetting an online channel to be rare and be due to the PE resetting. If an online channel is reset, it is possible that some stale downstream commands are received from the IRS. The CPU interface ignores such commands.</p>

A4.7 IRS and CPU interface capabilities

I _{MNVFZ}	The GICv5 Stream interface supports independent development of an IRS and a PE containing a CPU interface. This could lead to mismatches between the capabilities of the IRS and the connected PEs.
R _{CQTBC}	The CPU interface and IRS connected by an Interrupt Handling channel implement the same set of Interrupt Domains. Neither the CPU interface nor IRS generate any command with the <code>Domain</code> field set to an unimplemented Interrupt Domain.
I _{MTKNY}	Support for virtualization is reported in the IRS by <code>IRS_IDR0.VIRT</code> and in the PE by <code>ID_AA64PFR0_EL1.EL2</code> .
I _{TGHLW}	In GICv5, either both the IRS and CPU interface support virtualization or neither support virtualization.
I _{PDCJG}	<p>GICv5 supports up to 5 bits of interrupts priority. The number of supported priority bits is reported by the IRS in <code>IRS_IDR1.PRI_BITS</code> and by the CPU interface in <code>ICC_IDR0_EL1.PRI_BITS</code>.</p> <p>In GICv5 Stream, if the component receiving a command implements fewer than 5 bits of priority, the unimplemented least significant priority bits in the command are ignored and treated as being 0.</p> <p>If the CPU interface and IRS implement different number of priority bits, Arm recommends that software restricts its usage to only those bits supported by both components.</p>
I _{RJJYX}	<p>GICv5 supports up to 24 bits of <code>INTID.ID</code> namespace for each interrupt type. The number of supported <code>INTID.ID</code> bits is reported by the IRS in <code>IRS_IDR2.ID_bits</code> and by the CPU interface in <code>ICC_IDR0_EL1.ID_Bits</code>. When fewer than 24 bits of identifier bits are supported, the unimplemented bits are the most significant bits.</p> <p>Arm recommends that all components in a GICv5 system implement the same number of <code>INTID</code> identifier bits. If different components support different number of <code>INTID.ID</code> bits, Arm recommends that software restricts its usage to only those bits supported by all components.</p>
R _{YGCKH}	If the IRS receives a command with <code>INTID.ID</code> beyond the implemented range, the command is acknowledged but otherwise ignored.
R _{LQZNQ}	<p>If the CPU interface receives a command with <code>INTID.ID</code> beyond the implemented range, the behavior depends on the type of command as follows:</p> <ul style="list-style-type: none"> Forward command: The command is treated as an implicit Recall of any outstanding Forward for the specified Interrupt Domain or resident VPE, but is otherwise ignored. Recall command: The command is ignored.

Chapter A5

Interrupt Signaling channel

I _{TJQNR}	Interrupt Signaling channels allows interrupt sources to be connected to an IRS, and are intended as an alternative to using dedicated interrupt signals or MSIs.
I _{HRNDT}	The rules in this section only apply to the Interrupt Signaling channel.
I _{GSNZT}	How interrupts delivered by an Interrupt Signaling channel map to SPIs or IWB inputs is IMPLEMENTATION SPECIFIC.
I _{QBKQL}	Each Interrupt Source connected via an Interrupt Signaling channel can have one or more distinct interrupts.

A5.1 Command summary

I_{WSFCN}

The upstream commands that the Interrupt Source can send to the IRS are:

ID	Command	Description
0x0	Reset	Resets the channel and deasserts all interrupts.
0x1	INT	Communicates the current state of an interrupt.
0x2	Resample Ack	Acknowledges a Resample command.
0x3	Quiesce	Requests the channel is taken offline.
0x4	Flush Ack	Acknowledges a Flush command.
0x7	IMPLEMENTATION DEFINED	Reserved for IMPLEMENTATION DEFINED functionality.

All other IDs are reserved.

I_{YMHEW}

The downstream commands that an IRS can send to the Interrupt Source are:

ID	Command	Description
0x0	Reset Ack	Acknowledges a Reset command.
0x1	-	-
0x2	Resample	Requests current state of interrupt or interrupt levels.
0x3	Quiesce Ack	Acknowledges a Quiesce command.
0x4	Flush	Flushes commands from the IRS as part of a reset sequence.
0x7	IMPLEMENTATION DEFINED	Reserved for IMPLEMENTATION DEFINED functionality.

All other IDs are reserved.

R_{DNTTB}

ID 0x7 in upstream and downstream is reserved for IMPLEMENTATION DEFINED functionality.

When supported, IMPLEMENTATION DEFINED commands do not impact the functionality of other commands. If not supported, IMPLEMENTATION DEFINED commands are ignored.

Arm strongly recommends that IMPLEMENTATION DEFINED commands are not used unless it can be guaranteed that both the sender and receiver agree on the usage.

A5.2 Outstanding commands

I_{QBJBL} The following commands require acknowledgement:

Command	Direction	Acknowledged by
Reset	Upstream	Reset Ack
Resample	Downstream	Resample Ack or Reset
Quiesce	Upstream	Quiesce Ack
Flush	Downstream	Flush Ack

- R_{FFYPH} All other commands do not require acknowledgement.
- I_{BBGBS} Commands which require acknowledgement are acknowledged in finite time.
- R_{SWNBP} For commands that require acknowledgement, the architecture places limits on the number of outstanding commands that are permitted.
- R_{HKNHF} For each Interrupt Signaling channel, the IRS can have:
- Up to one Resample command outstanding.
 - Up to one Flush command outstanding.
- I_{QNLWZ} For each Interrupt Signaling channel, the CPU interface can have:
- Up to one Quiesce command outstanding.
- See [A5.4 Connection management](#) for more information about outstanding Reset commands.

A5.3 Signaling interrupts to the IRS

R_{NVLNT} An Interrupt Source can support up to 256 interrupts, numbered 0 to 255.

R_{GYFBQ} An Interrupt Source sends an INT command to indicate the state of an interrupt:

- INT.ID specifies the interrupt whose state is being communicated.
- INT.Edge and INT.Level indicate the current state of the interrupt.

Edge	Level	Meaning	Usage
0	0	The interrupt is de-asserted.	Used to indicate a level-sensitive interrupt has become de-asserted.
0	1	The interrupt is asserted.	Used to respond to Resample commands.
1	0	The interrupt was asserted and became de-asserted.	Used for edge-triggered interrupts to indicate a pulse.
1	1	The interrupt is asserted and remains asserted.	Used to indicate a level-sensitive interrupt has become asserted.

I_{CDSFZ} INT.Edge is used to identify positive edges, it is not used to indicate negative edges.

R_{DBMY} When an interrupt changes state, the Interrupt Source sends an INT command in finite time.

I_{BMDPL} When the state of an interrupt changes multiple times before an INT command can be sent, and the interrupt is level-sensitive, it is permissible for the Interrupt Source to only present the final state of the interrupt.

I_{MPJMH} Arm recommends that an Interrupt Source does not send INT commands unless the interrupt has changed state or it has received a Resample command.

R_{TVKMG} It is IMPLEMENTATION DEFINED which values of INT.ID that a CPU interface may send.

R_{TLDXR} The IRS can request the current state of an interrupt or of all interrupts by issuing a Resample command.

On receiving a Resample command:

- If Resample.IDV is 0:
 - The Interrupt Source sends an INT command for any interrupt source that is currently asserted, then sends a Resample Acknowledge command.
 - Any interrupt, for which the Interrupt Source does not send an INT command, between the Resample and Resample Acknowledge, is treated by the IRS as being de-asserted.
- If Resample.IDV is 1:
 - If the interrupt identified by Resample.ID is invalid, the Interrupt Source sends a Resample Acknowledge command.
 - If the interrupt identified by Resample.ID is valid and asserted, the Interrupt Source sends an INT command, then sends a Resample Acknowledge command.
 - If the interrupt identified by Resample.ID is valid and not asserted, the Interrupt Source sends a Resample Acknowledge command.

A5.4 Connection management

R _{XSJTH}	The Interrupt Source sending a Reset command sets all interrupts to de-asserted and acknowledges any outstanding downstream commands.
R _{VRGRT}	<p>To bring an Interrupt Signaling channel online, or reset an online channel:</p> <ul style="list-style-type: none">• The Interrupt Source sends a Reset command.• The IRS responds with a Flush command.• The Interrupt Source acknowledges the Flush command.• The IRS acknowledges the Reset command. <p>Once the Reset command is acknowledged, the Interrupt Handling channel is online.</p>
R _{FWSZJ}	<p>When a Reset command is acknowledged:</p> <ul style="list-style-type: none">• There are no outstanding upstream or downstream commands.• All interrupts are treated as de-asserted.
R _{CSVKH}	<p>Between sending a Reset command and the command being acknowledged, the CPU interface sends no commands other than to acknowledge a Flush command.</p> <p>Between receiving Reset command and sending the acknowledge, the IRS sends no other commands than Flush.</p>
R _{TMFQN}	The IRS generates a Flush command at most once per channel reset. The Interrupt Source generates a Flush Acknowledge command at most once per channel reset.
I _{LFQPR}	If multiple Reset commands are issued in quick succession, the Interrupt Source might not be able to determine whether a received Flush is due to the most recent Reset or the previous Reset. To avoid the Interrupt Source sending more Flush Acknowledge commands than there are Flush commands, the Interrupt Source is restricted to sending Flush Acknowledge once per-Reset.
R _{ZNKNG}	<p>If the IRS receives Flush Acknowledge before it has sent Flush in the channel reset flow, it is IMPLEMENTATION DEFINED whether the IRS sends a Flush command. If the Flush command is sent after receiving a Flush Acknowledge, the IRS treats the command as already acknowledged.</p> <p>If the Interrupt source receives a Flush command after it has sent a Flush Acknowledge, the Flush command is ignored and treated as acknowledged.</p>
R _{GDBXW}	<p>To take an Interrupt Signaling channel offline:</p> <ul style="list-style-type: none">• The Interrupt Source sends a Quiesce command.<ul style="list-style-type: none">– The Interrupt Source sends no further commands, other than acknowledgements to downstream commands, until Quiesce Ack is received.• The IRS completes any outstanding commands and then sends Quiesce Ack.<ul style="list-style-type: none">– Once the IRS has received an upstream Quiesce command it does not send further Resample commands. <p>Once the IRS responds with Quiesce Acknowledge, the channel is offline.</p>
I _{WMGK}	<p>It is possible that the Interrupt Source receives a downstream Resample command after the upstream Quiesce command is sent. The Interrupt Source is required to send Resample Ack before the IRS can respond with Quiesce Ack.</p> <p>Taking an Interrupt Source channel offline implicitly deasserts all the interrupts. Therefore, if there is an outstanding Resample when Quiesce is issued, Arm recommends that the Interrupt Source sends Resample Ack without sending any INT commands, which indicates that all interrupts are deasserted. It is, however, permitted for the Interrupt Source to send INT commands and then Resample Ack.</p>
R _{YWLGB}	<p>If an Interrupt Signaling channel is offline, all of the following are true:</p> <ul style="list-style-type: none">• The IRS sends no commands.• The Interrupt Source sends no commands, other than Reset to bring the channel online.• All interrupts are treated as de-asserted.

Chapter A6

Alphabetical list of commands

This section contains the definition of each command for the GICv5 Stream Protocol.

A6.1 Interrupt Handling channel

This section contains the definition of each command for the Interrupt Handling channel for the GICv5 Stream Protocol.

A6.1.1 Activate, Activate command (CPUIF -> IRS)


The Activate characteristics are:

Attributes

Activate is a 2-byte structure.

Field descriptions

The Activate bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0010	Activate

Virtual, bit [4]

Whether the operation is on a physical or virtual interrupt.

Virtual	Meaning
0b0	Physical interrupt.
0b1	Virtual interrupt.

Bit [5]

Reserved, RES0.

Domain, bits [7:6]

The Domain the operation affects.

Domain	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

Bits [15:8]

Reserved, RES0.

A6.1.2 ActivateAck, Activate Acknowledge command (IRS -> CPUIF)


The ActivateAck characteristics are:

Attributes

ActivateAck is a 2-byte structure.

Field descriptions

The ActivateAck bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0010	Activate Ack

Bits [15:4]

Reserved, RES0.

A6.1.3 Deactivate, Deactivate interrupt command (CPUIF -> IRS)


The Deactivate characteristics are:

Attributes

Deactivate is a 6-byte structure.

Field descriptions

The Deactivate bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0110	Deactivate

Virtual, bit [4]

Whether the operation is on a physical or virtual interrupt.

Virtual	Meaning
0b0	Physical interrupt.
0b1	Virtual interrupt.

TYPE, bits [7:5]

The type of the interrupt.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Domain, bits [9:8]

The Domain the operation affects.

Domain	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

Bits [15:10]

Reserved, RES0.

ID, bits [39:16]

The ID of the interrupt.

Bits [47:40]

Reserved, RES0.

A6.1.4 DeactivateAck, Deactivate interrupt Acknowledge command (IRS -> CPUIF)


The DeactivateAck characteristics are:

Attributes

DeactivateAck is a 2-byte structure.

Field descriptions

The DeactivateAck bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0110	Deactivate Ack

Bits [15:4]

Reserved, RES0.

A6.1.5 DownstreamControl, Downstream Control (IRS -> CPUIF)


The DownstreamControl characteristics are:

Attributes

DownstreamControl is a 4-byte structure.

Field descriptions

The DownstreamControl bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0100	DownstreamControl

Identifier, bits [7:4]

Specifies the type of Downstream Control operation.

Identifier	Meaning
0b0001	Flush
0b0010	IAFFID

Bits [15:8]

Reserved, RES0.

IAFFID, bits [31:16]

The PE interrupt affinity.

This field is RES0 when Identifier is not 0b0010 (IAFFID).

A6.1.6 DownstreamControlAck, Downstream Control Acknowledge command (CPUIF -> IRS)


The DownstreamControlAck characteristics are:

Attributes

DownstreamControlAck is a 2-byte structure.

Field descriptions

The DownstreamControlAck bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0100	DownstreamControl Ack

Flush, bit [4]

Whether the operation being acknowledged is a Flush.

Flush	Meaning
0b0	Not acknowledging a Flush.
0b1	Acknowledging a Flush.

Bits [15:5]

Reserved, RES0.

A6.1.7 Forward, Forward command (IRS -> CPUIF)


The Forward characteristics are:

Attributes

Forward is a 6-byte structure.

Field descriptions

The Forward bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0001	Forward

Virtual, bit [4]

Whether the operation is on a physical or virtual interrupt.

Virtual	Meaning
0b0	Physical interrupt.
0b1	Virtual interrupt.

TYPE, bits [7:5]

The type of the interrupt.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Domain, bits [9:8]

The Domain the operation affects.

Domain	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

ActivateAck, bit [10]

Whether the Forward also Acknowledges a previous Activate command.

ActivateAck	Meaning
0b0	The Forward does not acknowledge any previous Activate.
0b1	The Forward acknowledges a previous Activate.

When the Forward command acknowledges a previous Activate, it applies to an Activate for the same Domain and setting of Virtual.

Priority, bits [15:11]

The priority of the interrupt.

Lower numeric values are higher priority, meaning that 0b00000 is the highest supported priority and 0b11111 is the lowest priority.

The CPU interface treats priority bits it does not implement as RES0.

The IRS sets priority bits it does not implement to 0b0.

ID, bits [39:16]

The ID of the interrupt.

Bits [46:40]

Reserved, RES0.

Returnable, bit [47]

Indicates whether the CPU interface can choose to proactively return the interrupt.

Returnable	Meaning
0b0	Interrupt is not returnable.
0b1	Interrupt is returnable.

A6.1.8 Recall, Recall command (IRS -> CPUIF)


The Recall characteristics are:

Attributes

Recall is a 2-byte structure.

Field descriptions

The Recall bit assignments are:

***CMD, bits [3:0]***

Indicates the command type.

CMD	Meaning
0b0011	Recall

Virtual, bit [4]

Whether the operation is on a physical or virtual interrupt.

Virtual	Meaning
0b0	Physical interrupt.
0b1	Virtual interrupt.

Bit [5]

Reserved, RES0.

Domain, bits [7:6]

The Domain the operation affects.

Domain	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

Bits [15:8]

Reserved, RES0.

A6.1.9 Release, Release command (CPUIF -> IRS)


The Release characteristics are:

Attributes

Release is a 6-byte structure.

Field descriptions

The Release bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0011	Release

Virtual, bit [4]

Whether the operation is on a physical or virtual interrupt.

Virtual	Meaning
0b0	Physical interrupt.
0b1	Virtual interrupt.

TYPE, bits [7:5]

The type of the interrupt.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Domain, bits [9:8]

The Domain the operation affects.

Domain	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

Bits [15:10]

Reserved, RES0.

ID, bits [39:16]

The ID of the interrupt.

Bits [47:40]

Reserved, RES0.

A6.1.10 RequestConfig, Request Interrupt Configuration command (CPUIF -> IRS)


The RequestConfig characteristics are:

Attributes

RequestConfig is a 6-byte structure.

Field descriptions

The RequestConfig bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b1100	RequestConfig

Virtual, bit [4]

Whether the operation is on a physical or virtual interrupt.

Virtual	Meaning
0b0	Physical interrupt.
0b1	Virtual interrupt.

TYPE, bits [7:5]

The type of the interrupt.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Domain, bits [9:8]

The Domain the operation affects.

Domain	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

Bits [15:10]

Reserved, RES0.

ID, bits [39:16]

The ID of the interrupt.

Bits [47:40]

Reserved, RES0.

A6.1.11 RequestConfigAck, Request Interrupt Configuration Acknowledge command (IRS -> CPUIF)


The RequestConfigAck characteristics are:

Attributes

RequestConfigAck is a 4-byte structure.

Field descriptions

The RequestConfigAck bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b1100	RequestConfigAck

Bit [4]

Reserved, RES0.

Active, bit [5]

Whether the interrupt is Active.

Active	Meaning
0b0	Inactive
0b1	Active

Fault, bit [6]

Whether the IRS was able to retrieve the requested data.

Fault	Meaning
0b0	Success
0b1	Fault

When this field is set to 1 it indicates the requested INTID was unreachable, the returned configuration is UNKNOWN.

Enable, bit [7]

Whether the interrupt is individually enabled.

Enable	Meaning
0b0	Disabled
0b1	Enabled

HM, bit [8]

Whether the interrupt is Level or Edge.

HM	Meaning
0b0	Edge
0b1	Level

IRM, bit [9]

Interrupt routing mode

IRM	Meaning
0b0	Interrupt Routing mode is Targeted.
0b1	Interrupt Routing mode is 1ofN.

Pending, bit [10]

Whether the interrupt is pending.

Pending	Meaning
0b0	Not pending.
0b1	Pending.

Priority, bits [15:11]

The priority of the interrupt.

IAFFID, bits [31:16]

The interrupt Affinity value.

When IRM is 1, this field is IMPLEMENTATION SPECIFIC.

A6.1.12 SetAck, Set interrupt configuration acknowledge command (IRS -> CPUIF)


The SetAck characteristics are:

Attributes

SetAck is a 2-byte structure.

Field descriptions

The SetAck bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b1000	Set Ack

Bits [15:4]

Reserved, RES0.

A6.1.13 SetEnabled, Set interrupt Enabled command (CPUIF -> IRS)


The SetEnabled characteristics are:

Attributes

SetEnabled is a 6-byte structure.

Field descriptions

The SetEnabled bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b1000	SetEnabled

Virtual, bit [4]

Whether the operation is on a physical or virtual interrupt.

Virtual	Meaning
0b0	Physical interrupt.
0b1	Virtual interrupt.

TYPE, bits [7:5]

The type of the interrupt.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Domain, bits [9:8]

The Domain the operation affects.

Domain	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

EN, bit [10]

Whether the interrupt is Enabled.

EN	Meaning
0b0	Disabled
0b1	Enabled

Bits [15:11]

Reserved, RES0.

ID, bits [39:16]

The ID of the interrupt.

Bits [47:40]

Reserved, RES0.

A6.1.14 SetHandling, Set Interrupt Handling Mode command (CPUIF -> IRS)


The SetHandling characteristics are:

Attributes

SetHandling is a 6-byte structure.

Field descriptions

The SetHandling bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b1011	SetHandling

Virtual, bit [4]

Whether the operation is on a physical or virtual interrupt.

Virtual	Meaning
0b0	Physical interrupt.
0b1	Virtual interrupt.

TYPE, bits [7:5]

The type of the interrupt.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Domain, bits [9:8]

The Domain the operation affects.

Domain	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

HM, bit [10]

Interrupt Handling mode.

HM	Meaning
0b0	Edge
0b1	Level

Bits [15:11]

Reserved, RES0.

ID, bits [39:16]

The ID of the interrupt.

Bits [47:40]

Reserved, RES0.

A6.1.15 SetPending, Set interrupt Pending command (CPUIF -> IRS)


The SetPending characteristics are:

Attributes

SetPending is a 8-byte structure.

Field descriptions

The SetPending bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b1110	SetPending

Virtual, bit [4]

Whether the operation is on a physical or virtual interrupt.

Virtual	Meaning
0b0	Physical interrupt.
0b1	Virtual interrupt.

TYPE, bits [7:5]

The type of the interrupt.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Domain, bits [9:8]

The Domain the operation affects.

Domain	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

Pending, bit [10]

Whether the interrupt is Pending.

Pending	Meaning
0b0	Generate CLEAR event.
0b1	Generate SET event.

Bits [15:11]

Reserved, RES0.

ID, bits [39:16]

The ID of the interrupt.

Bits [47:40]

Reserved, RES0.

VM, bits [63:48]

The Virtual Machine identifier. This field is RES0 when Virtual is 0.

A6.1.16 SetPriority, Set Interrupt Priority command (CPUIF -> IRS)


The SetPriority characteristics are:

Attributes

SetPriority is a 6-byte structure.

Field descriptions

The SetPriority bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b1010	SetPriority

Virtual, bit [4]

Whether the operation is on a physical or virtual interrupt.

Virtual	Meaning
0b0	Physical interrupt.
0b1	Virtual interrupt.

TYPE, bits [7:5]

The type of the interrupt.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Domain, bits [9:8]

The Domain the operation affects.

Domain	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

Bit [10]

Reserved, RES0.

Priority, bits [15:11]

The priority of the interrupt.

The CPU interface sets priority bits it does not implement to 0b0

The IRS treats priority bits it does not implement as 0b0.

ID, bits [39:16]

The ID of the interrupt.

Bits [47:40]

Reserved, RES0.

A6.1.17 SetResident, Set Resident command (CPUIF -> IRS)


The SetResident characteristics are:

Attributes

SetResident is a 6-byte structure.

Field descriptions

The SetResident bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b1111	SetResident

Virtual, bit [4]

Whether the operation is physical or virtual.

Virtual	Meaning
0b1	Virtual interrupt.

Bit [5]

Reserved, RES0.

Domain, bits [7:6]

The Interrupt Domain of the VPE being made resident or non-resident.

Domain	Meaning
0b00	Secure
0b01	Non-secure
0b11	Realm

Valid, bit [8]

Whether a VPE is resident or not.

Valid	Meaning
0b0	The currently resident VPE is being cleared.
0b1	The VPE indicated by the Domain and VPE fields is being made resident.

VPE, bits [23:9]

The VPE identifier. This field is RES0 when Valid is 0.

Bits [25:24]

Reserved, RES0.

DB, bit [26]

Doorbell Request.

DB	Meaning
0b0	No doorbell requested
0b1	Doorbell requested

This field is RES0 when Valid is 1.

DBPM, bits [31:27]

Doorbell Priority Mask.

This field is RES0 when DB is 0.

VM, bits [47:32]

The Virtual Machine identifier. This field is RES0 when Valid is 0.

A6.1.18 SetResidentAck, Set Resident acknowledge command (IRS -> CPUIF)


The SetResidentAck characteristics are:

Attributes

SetResidentAck is a 2-byte structure.

Field descriptions

The SetResidentAck bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b1111	SetResident Ack

Bits [5:4]

Reserved, RES0.

Fault, bit [6]

Whether the requested VPE was made resident.

Fault	Meaning
0b0	Requested VPE was made resident.
0b1	Requested VPE was not made resident.

When acknowledging a SetResident with Valid set to 0, this field is RES0.

Bits [15:7]

Reserved, RES0.

A6.1.19 SetTarget, Set Interrupt Target command (CPUIF -> IRS)


The SetTarget characteristics are:

Attributes

SetTarget is a 8-byte structure.

Field descriptions

The SetTarget bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b1001	SetTarget

Virtual, bit [4]

Whether the operation is on a physical or virtual interrupt.

Virtual	Meaning
0b0	Physical interrupt.
0b1	Virtual interrupt.

TYPE, bits [7:5]

The type of the interrupt.

TYPE	Meaning
0b010	LPI
0b011	SPI

Values not defined above are reserved.

Domain, bits [9:8]

The Domain the operation affects.

Domain	Meaning
0b00	Secure
0b01	Non-secure
0b10	EL3
0b11	Realm

IRM, bit [10]

Interrupt routing mode

IRM	Meaning
0b0	The interrupt Routing mode is Targeted.
0b1	The interrupt Routing mode is 1ofN.

Bits [15:11]

Reserved, RES0.

ID, bits [39:16]

The ID of the interrupt.

Bits [47:40]

Reserved, RES0.

IAFFID, bits [63:48]

Target PE.

If IRM is 0, this field indicates the PE interrupt affinity of the target PE.

If IRM is 1, this field provides an IMPLEMENTATION DEFINED hint to 1ofN selection algorithm. A value of 0 means that no hint is provided.

A6.1.20 Sync, synchronizes previously sent configuration changes (CPUIF -> IRS)


The Sync characteristics are:

Attributes

Sync is a 2-byte structure.

Field descriptions

The Sync bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b1101	Sync

Bits [15:4]

Reserved, RES0.

A6.1.21 SyncAck, synchronizes previously sent configuration changes (IRS -> CPUIF)


The SyncAck characteristics are:

Attributes

SyncAck is a 2-byte structure.

Field descriptions

The SyncAck bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b1101	Sync Ack

Bits [15:4]

Reserved, RES0.

A6.1.22 UpstreamControl, Upstream Control (CPUIF -> IRS)


The UpstreamControl characteristics are:

Attributes

UpstreamControl is a 2-byte structure.

Field descriptions

The UpstreamControl bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0000	UpstreamControl

Identifier, bits [7:4]

Specifies the type of Upstream Control operation.

Identifier	Meaning
0b0000	Reset
0b0001	Quiesce
0b0010	Set 1ofN hints, Data specifies which hints are being set.
0b0011	Clear 1ofN hints, Data specifies which hints are being cleared.

Data, bits [15:8]

Data associated with the Upstream Control operation.

When Identifier is 0b0010 or 0b0011, Data indicates which 1ofN hints are being set or cleared.

Data	Meaning
0bxxxxxxxx1	Selecting this PE is likely to cause exit from a low power state. Applies to all Domains.

Values not defined above are reserved.

For all other values of Identifier, this field is RES0.

A6.1.23 UpstreamControlAck, Upstream Control Acknowledge command (IRS -> CPUIF)

The UpstreamControlAck characteristics are:

Attributes

UpstreamControlAck is a 4-byte structure.

Field descriptions

The UpstreamControlAck bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0000	UpstreamControl Ack

Bits [14:4]

Reserved, RES0.

Reset, bit [15]

Whether a reset request is being acknowledged.

Reset	Meaning
0b0	Not acknowledging a reset request.
0b1	Acknowledging a reset request.

IAFFID, bits [31:16]

The PE interrupt affinity.

This field is RES0 when acknowledging an UpstreamControl with an Identifier that is not 0b0000 (Reset).

A6.1.24 WakeRequest, Wake request (IRS -> CPUIF)


The WakeRequest characteristics are:

Attributes

WakeRequest is a 2-byte structure.

Field descriptions

The WakeRequest bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0101	WakeRequest

Bits [15:4]

Reserved, RES0.

A6.2 Interrupt Signaling channel

This section contains the definition of each command for the Interrupt Signaling channel for the GICv5 Stream Protocol.

A6.2.1 INT, Interrupt command (Interrupt Source -> IRS)


The INT characteristics are:

Attributes

INT is a 2-byte structure.

Field descriptions

The INT bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0001	INT

Level, bit [4]

Level	Meaning
0b0	De-asserted.
0b1	Asserted.

Edge, bit [5]

Edge	Meaning
0b0	Falling-edge, or no edge.
0b1	Positive-edge.

Bits [7:6]

Reserved, RES0.

ID, bits [15:8]

The ID of the interrupt.

A6.2.2 Flush, Flush command (IRS -> Interrupt Source)


The Flush characteristics are:

Attributes

Flush is a 2-byte structure.

Field descriptions

The Flush bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0100	Flush

Bits [15:4]

Reserved, RES0.

A6.2.3 FlushAck, Flush acknowledge command (Interrupt Source -> IRS)


The FlushAck characteristics are:

Attributes

FlushAck is a 2-byte structure.

Field descriptions

The FlushAck bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0100	Flush Ack

Bits [15:4]

Reserved, RES0.

A6.2.4 Quiesce, Quiesce command (Interrupt Source -> IRS)


The Quiesce characteristics are:

Attributes

Quiesce is a 2-byte structure.

Field descriptions

The Quiesce bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0011	Quiesce

Bits [15:4]

Reserved, RES0.

A6.2.5 QuiesceAck, Quiesce acknowledge command (IRS -> Interrupt Source)


The QuiesceAck characteristics are:

Attributes

QuiesceAck is a 2-byte structure.

Field descriptions

The QuiesceAck bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0011	Quiesce Ack

Bits [15:4]

Reserved, RES0.

A6.2.6 Resample, Resample request command (IRS -> Interrupt Source)


The Resample characteristics are:

Attributes

Resample is a 2-byte structure.

Field descriptions

The Resample bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0010	Resample

IDV, bit [4]

IDV	Meaning
0b0	Resample request for all asserted interrupts.
0b1	Resample request for the interrupt identified by ID field.

Bits [7:5]

Reserved, RES0.

ID, bits [15:8]

The ID of the interrupt. This field is RES0 when IDV is 0.

A6.2.7 ResampleAck, Resample request acknowledge command (Interrupt Source -> IRS)

The ResampleAck characteristics are:

Attributes

ResampleAck is a 2-byte structure.

Field descriptions

The ResampleAck bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0010	Resample Ack

Bits [15:4]

Reserved, RES0.

A6.2.8 Reset, Reset command (Interrupt Source -> IRS)


The Reset characteristics are:

Attributes

Reset is a 2-byte structure.

Field descriptions

The Reset bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0000	Reset

Bits [15:4]

Reserved, RES0.

A6.2.9 ResetAck, Reset acknowledge command (IRS -> Interrupt Source)


The ResetAck characteristics are:

Attributes

ResetAck is a 2-byte structure.

Field descriptions

The ResetAck bit assignments are:



CMD, bits [3:0]

Indicates the command type.

CMD	Meaning
0b0000	Reset Ack

Bits [15:4]

Reserved, RES0.

Chapter A7

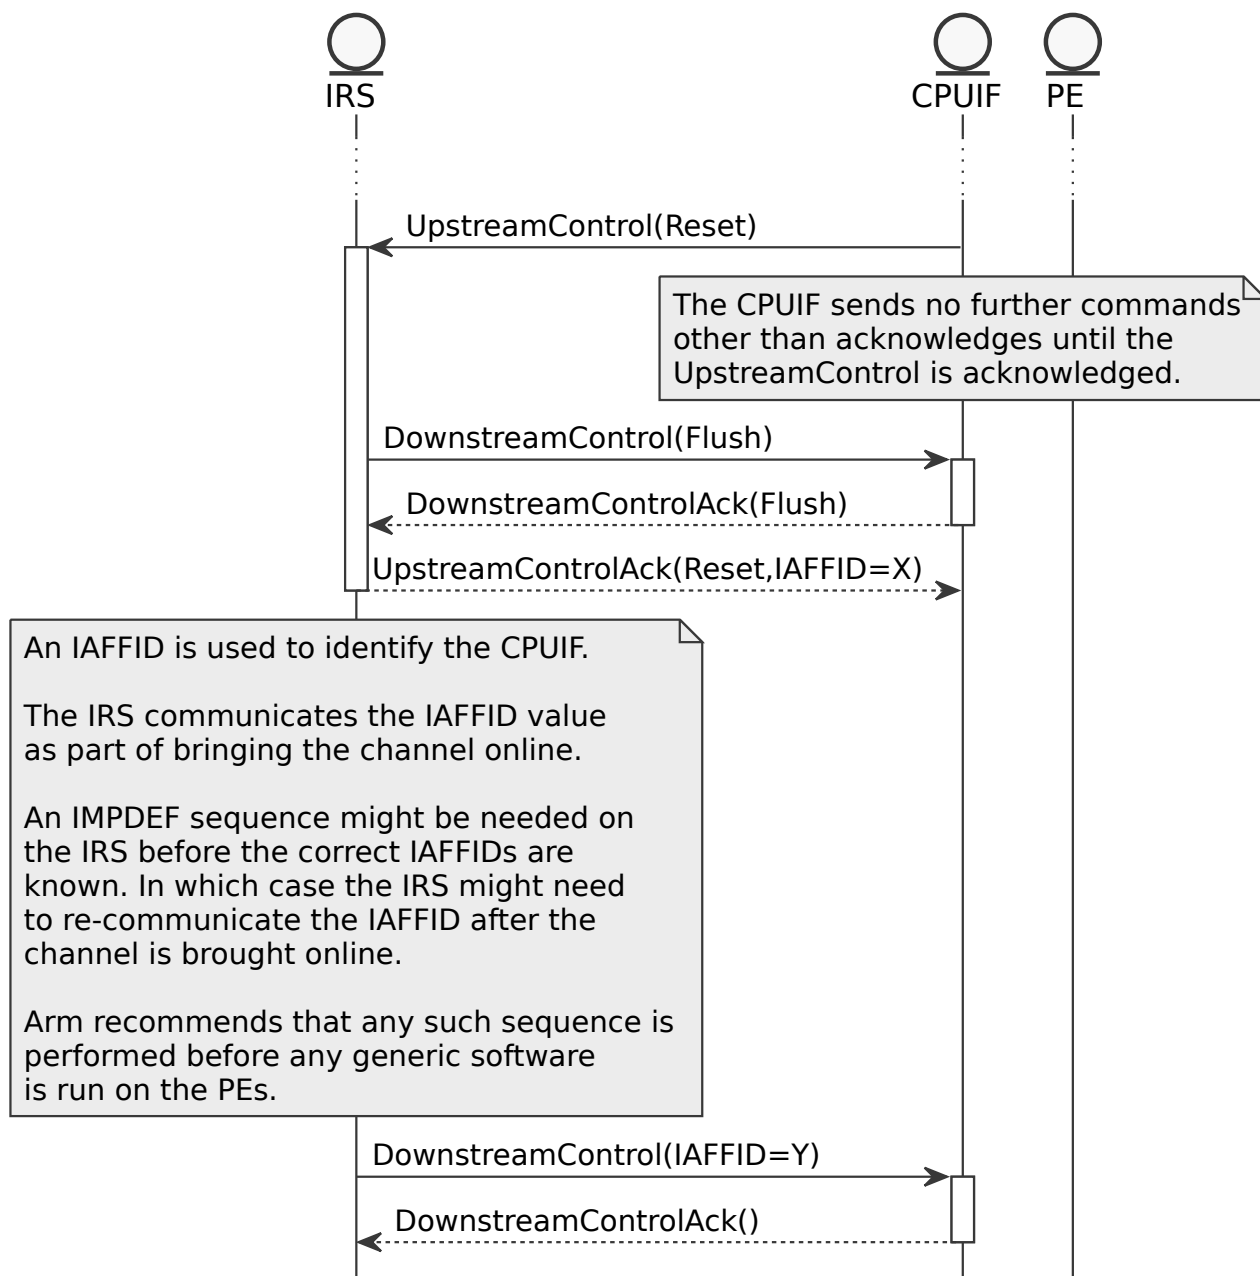
Example sequences

This section gives examples of typical sequences.

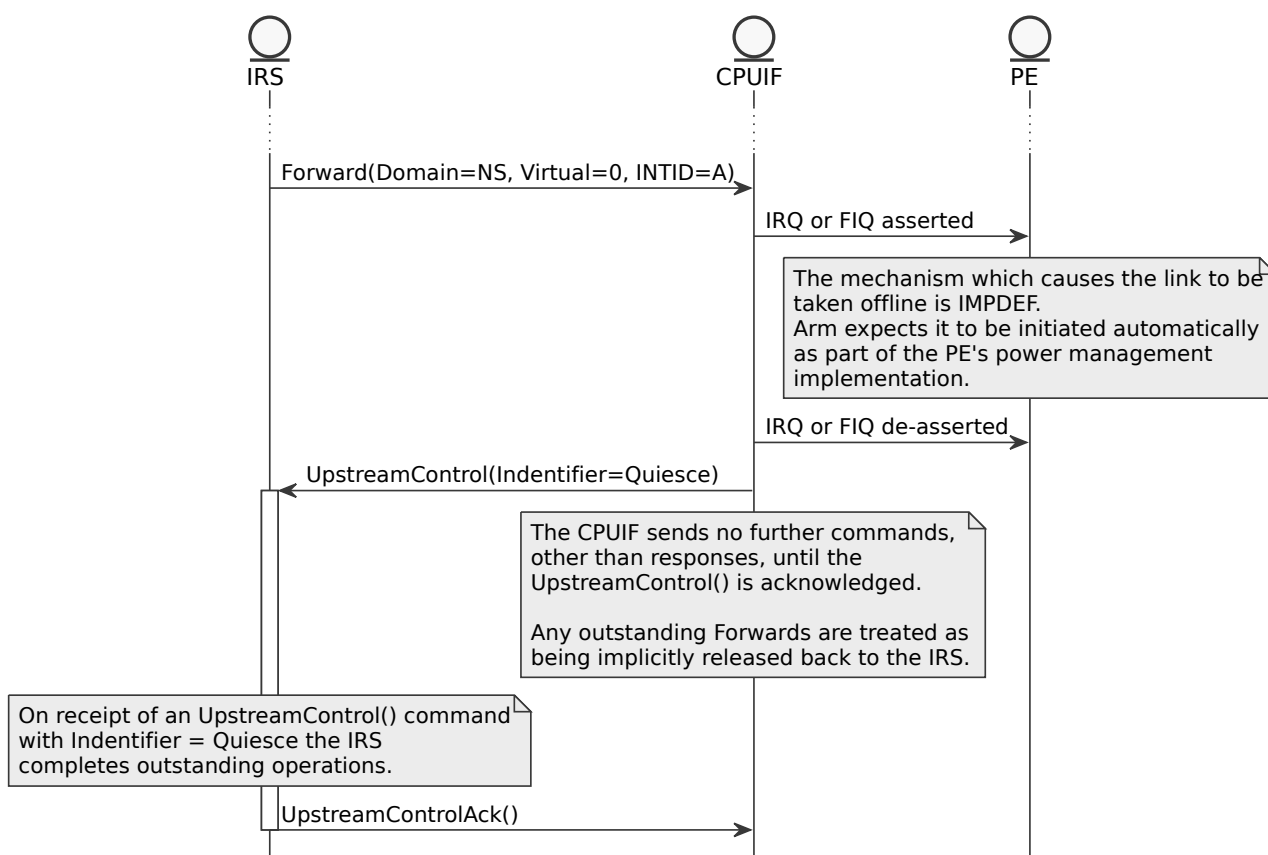
The flows given in this section are for illustration purposes only.

A7.1 Bringing the Interrupt Handling channel online and taking it offline

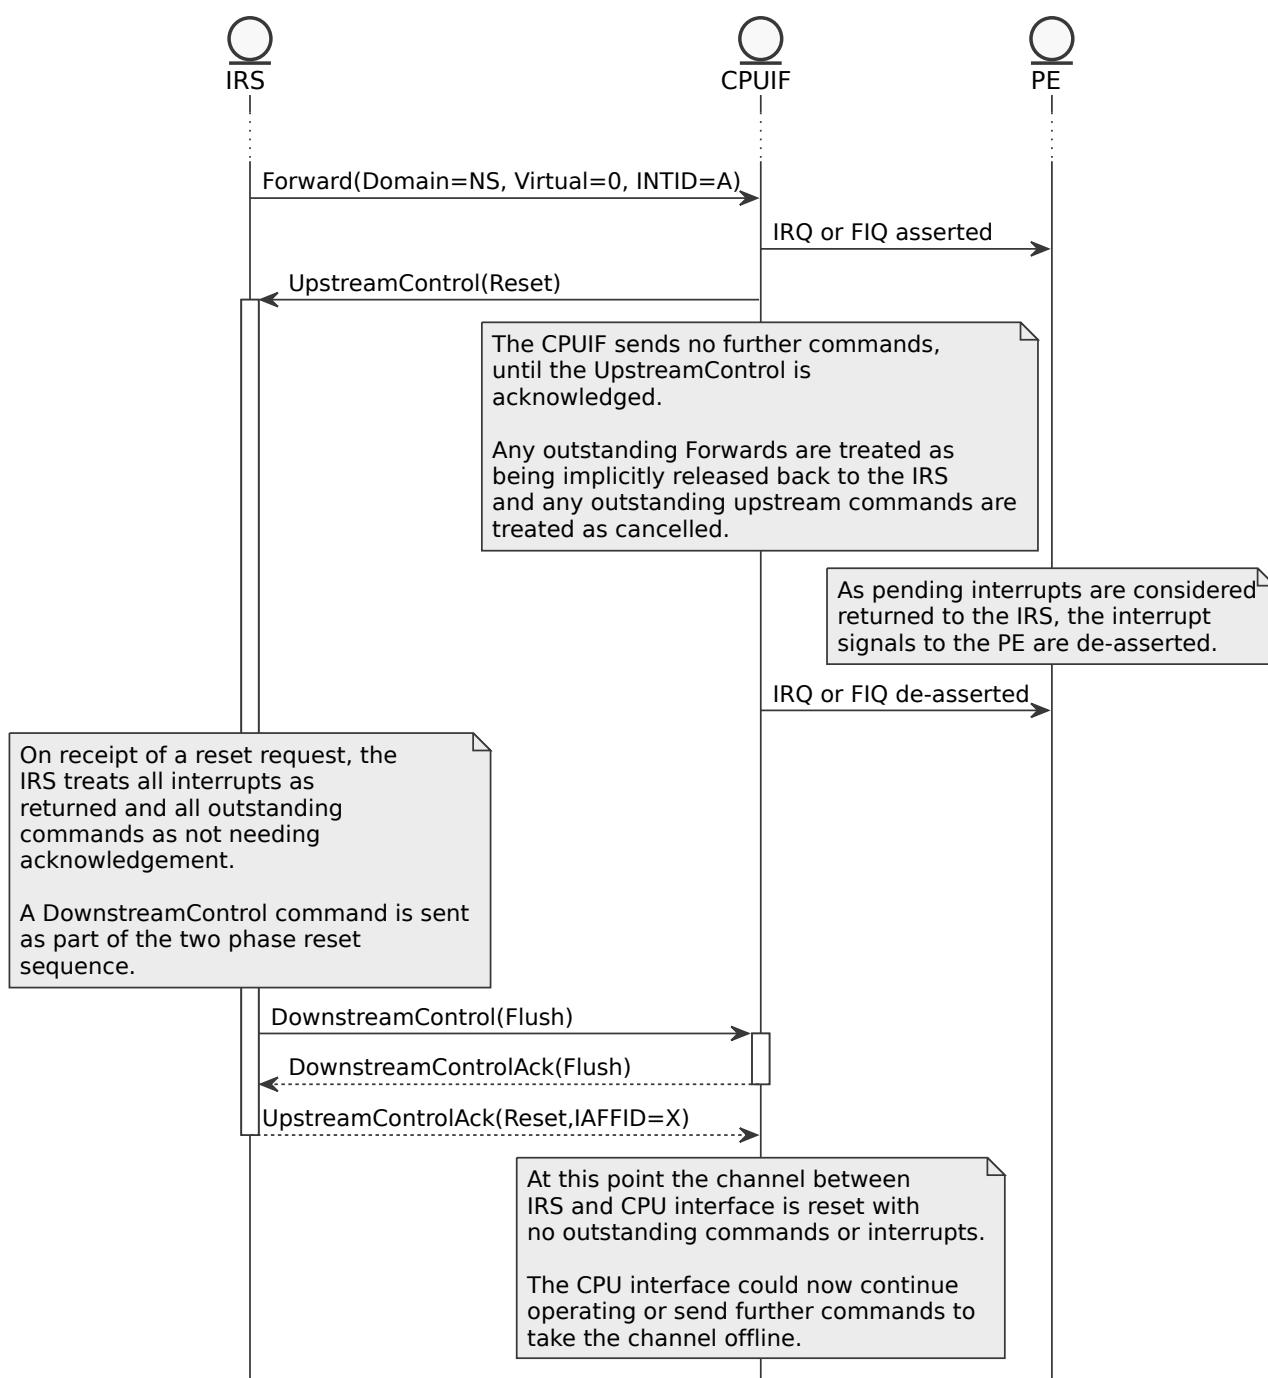
If the Interrupt Handling channel between the CPU interface and IRS is offline, the first command that is sent is an UpstreamControl command. A two-phase reset process is used, as illustrated below:



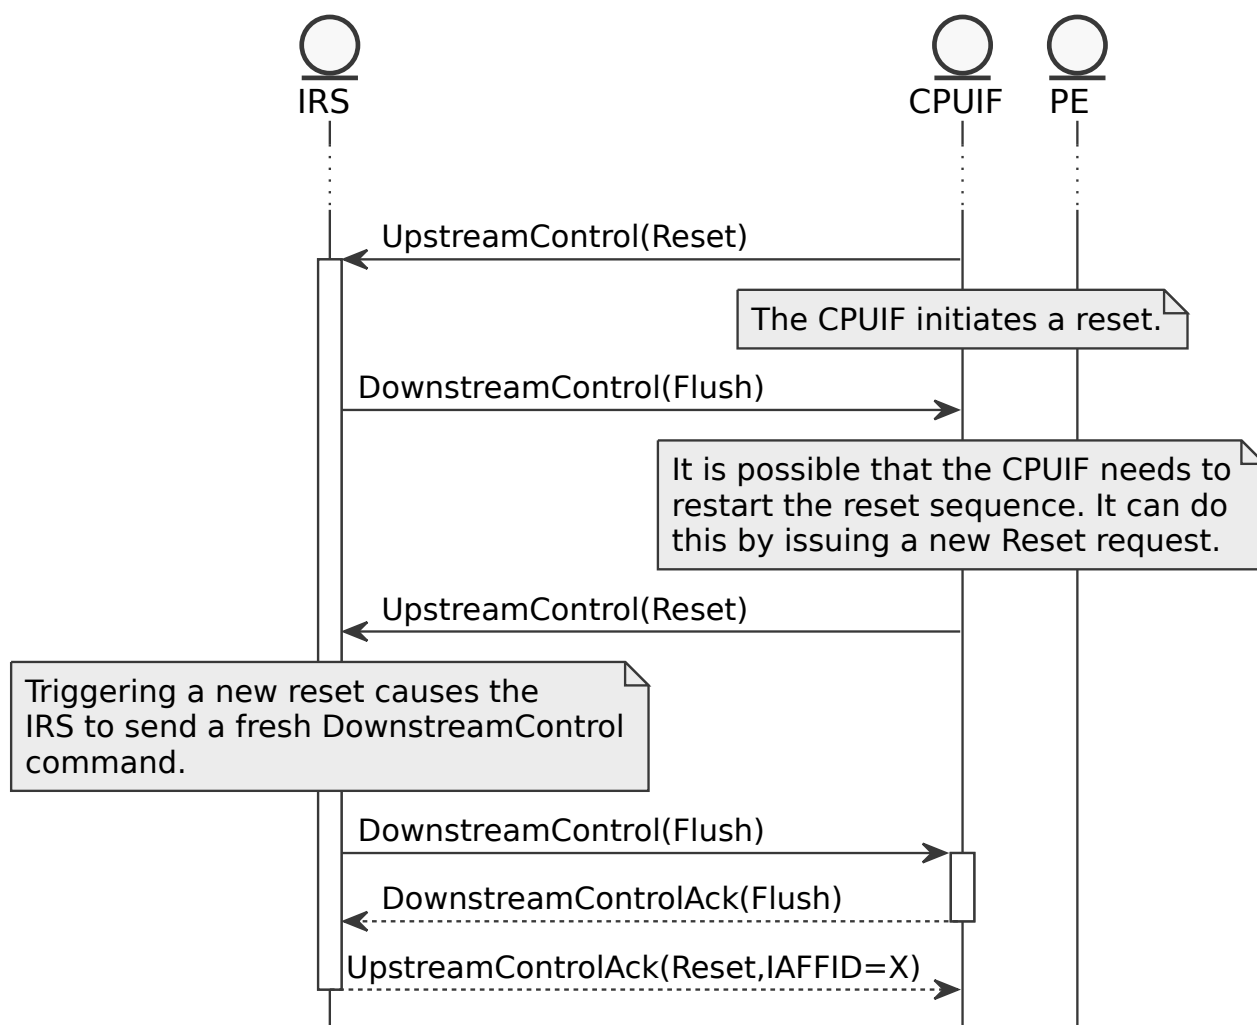
The Interrupt Handling channel is taken offline by the CPU interface sending an UpstreamControl command with Identifier set to Quiesce. On receipt of the UpstreamControl command the IRS completes all outstanding commands, before sending a UpstreamControlAck. The IRS treats all outstanding Forward commands as being implicitly returned by the link that is taken offline. Once the UpstreamControlAck is received by the PE the Interrupt Handling channel is offline.



There may be cases where the Interrupt Handling channel is classed as online but needs to be reset. This can be achieved by the CPU interface sending an `UpstreamControl` command with `Identifier` set to `Reset`. Sending the `UpstreamControl` command with `Identifier` set to `Reset` returns any pending interrupts to the IRS and cancels outstanding commands. This process is illustrated below:

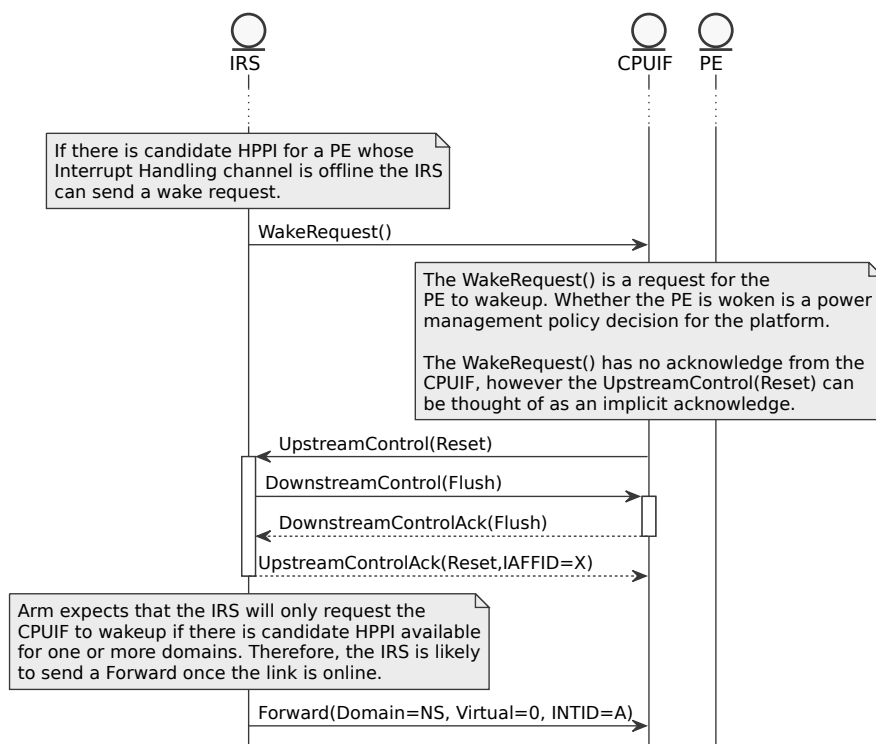


The CPU interface might need to restart the reset process part way through. To do this, the CPU interface sends a new Reset request.




Arm recommends that resetting an online connection is only used for recovery and not for normal operation.

If the Interrupt Handling channel is offline, the IRS can request that it is brought online by sending a wake request:



A7.2 Simple interrupt life-cycle

The diagram below shows a physical interrupt being presented to a PE by an IRS:

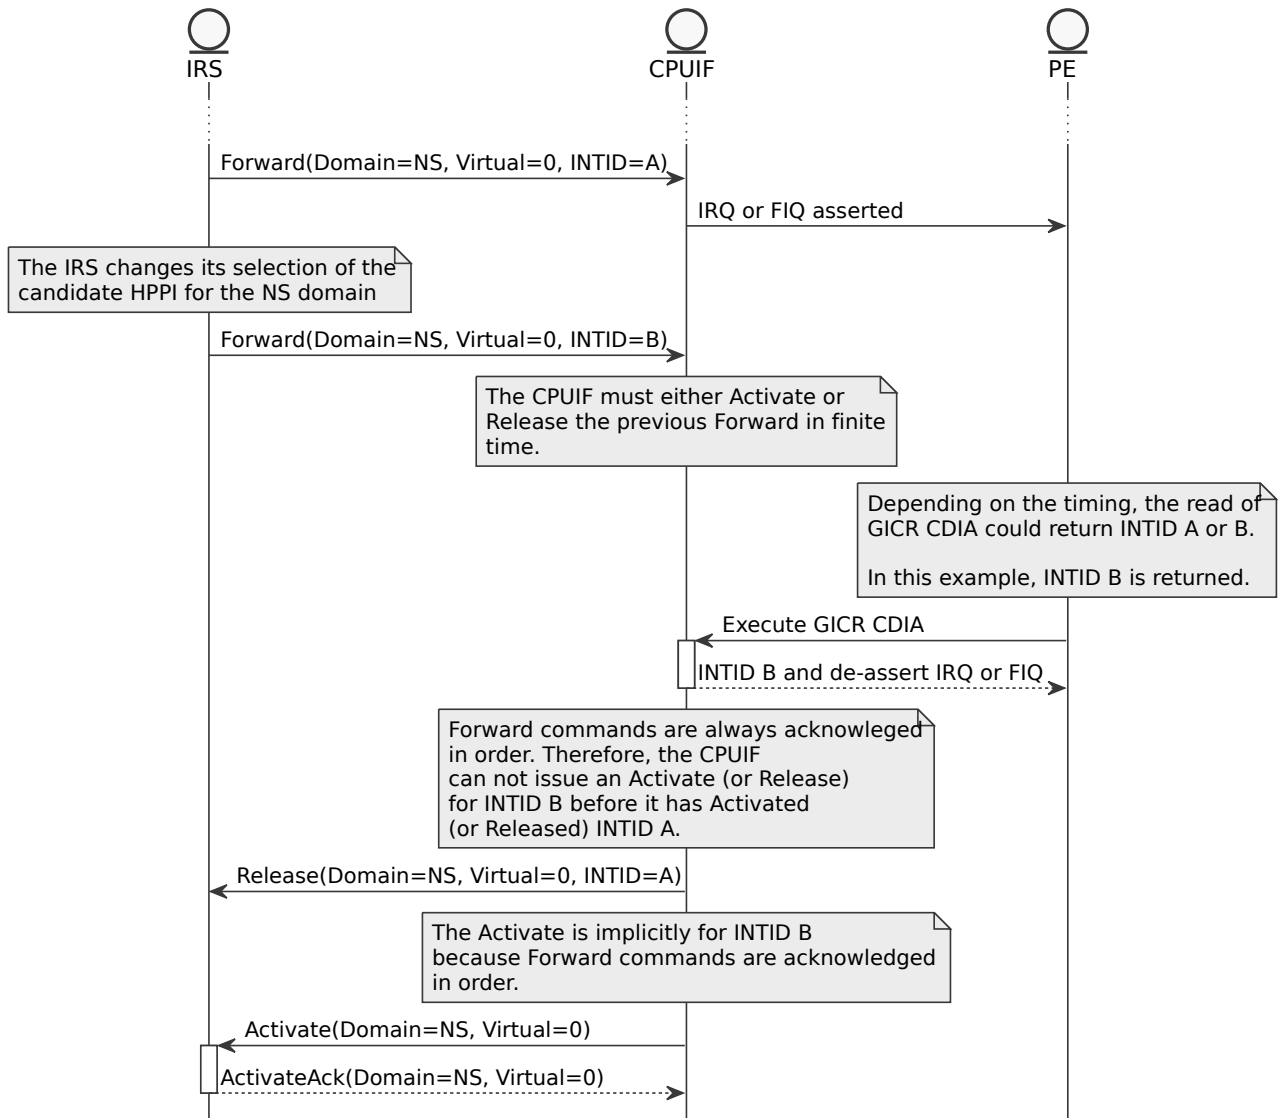


The effect of acknowledging an interrupt is not guaranteed to have been observed by the IRS until after a `GSB ACK` or `GSB SYS` instruction has completed. With GICv5 Stream, a PE can determine the activation has been observed by waiting for an `ActivateAck()` before completing the barrier.

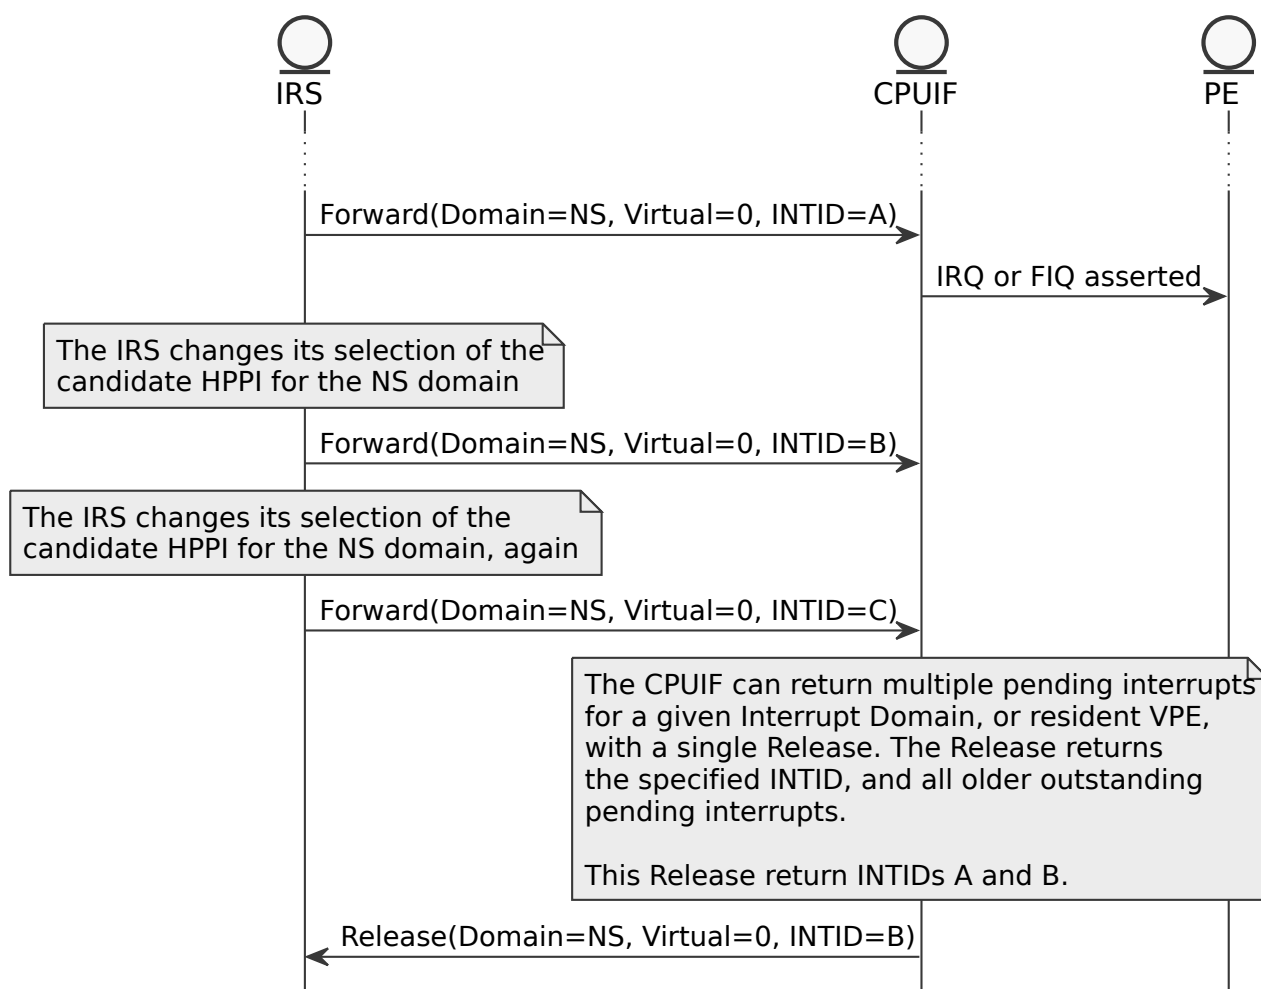
With edge-triggered interrupts, until the acknowledging of an instance of an interrupt is visible to the IRS, further edges from the interrupt source might continued to be merged.

A7.3 Replacing the candidate HPPI for an Interrupt Domain, or resident VPE

The IRS forwards the candidate HPPI for a given Interrupt Domain, or resident VPE, to the CPU interface. The chosen candidate HPPI might change, for example due to another interrupt becoming pending or a change to interrupt configuration. In this case, the IRS can replace the candidate HPPI for the Interrupt Domain, or resident VPE, by sending another Forward. The CPU interface is required to either release or activate the replaced interrupt in finite time.

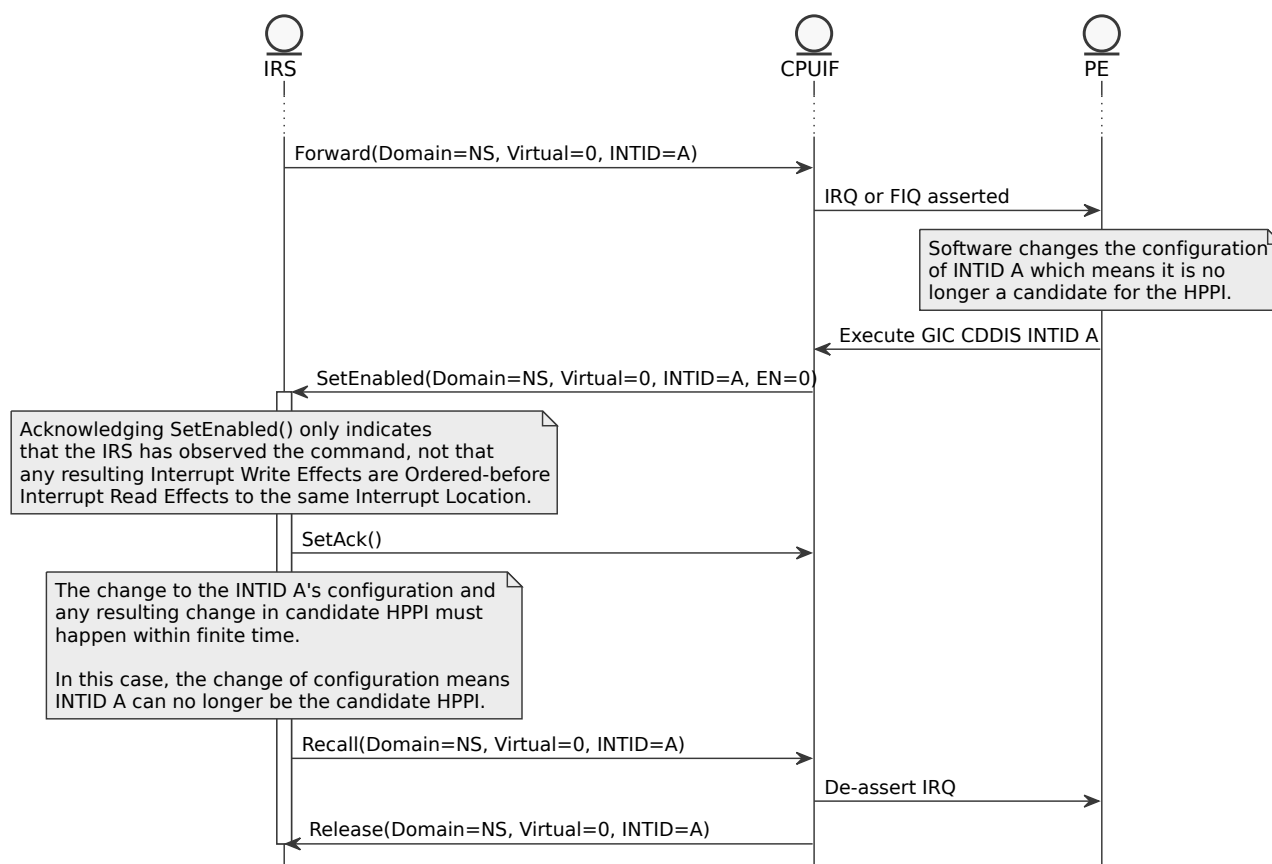


It is possible that the IRS again changes the candidate HPPI before the CPU interface has responded in finite time. The IRS can send further Forward commands, each replacing the previous. The CPU interface can return multiple pending interrupts for a given Interrupt Domain, or VPE, with a single Release command. The Release command releases the specified INTID, and all earlier outstanding Forward commands for that Interrupt Domain, or VPE.



For a given Interrupt Domain, or the resident VPE, Forward commands are always acknowledged in the order they are received. In the above example, the CPU interface cannot issue an Activate for INTID C before INTIDs A and B are released.

In some cases the candidate HPPI might change due to GIC instructions executed on the PE that the interrupt has been forwarded to. In these cases the CPU interface communicates the change in configuration to the IRS and the IRS is responsible for recalling the interrupt if required. The CPU interface does not proactively issue a Release.




In the above example, the IRS acknowledges the SetEnabled command before the Interrupt Write effect of clearing the INTID Enable is visible. Acknowledging a SetEnabled, SetPending, SetPriority or SetTarget indicates that the IRS has observed the command and that the effects of the command will become visible in finite time. To guarantee that the effects of previous operations are complete, software must issue a GSB SYS barrier which results in a Sync command being sent to the IRS.

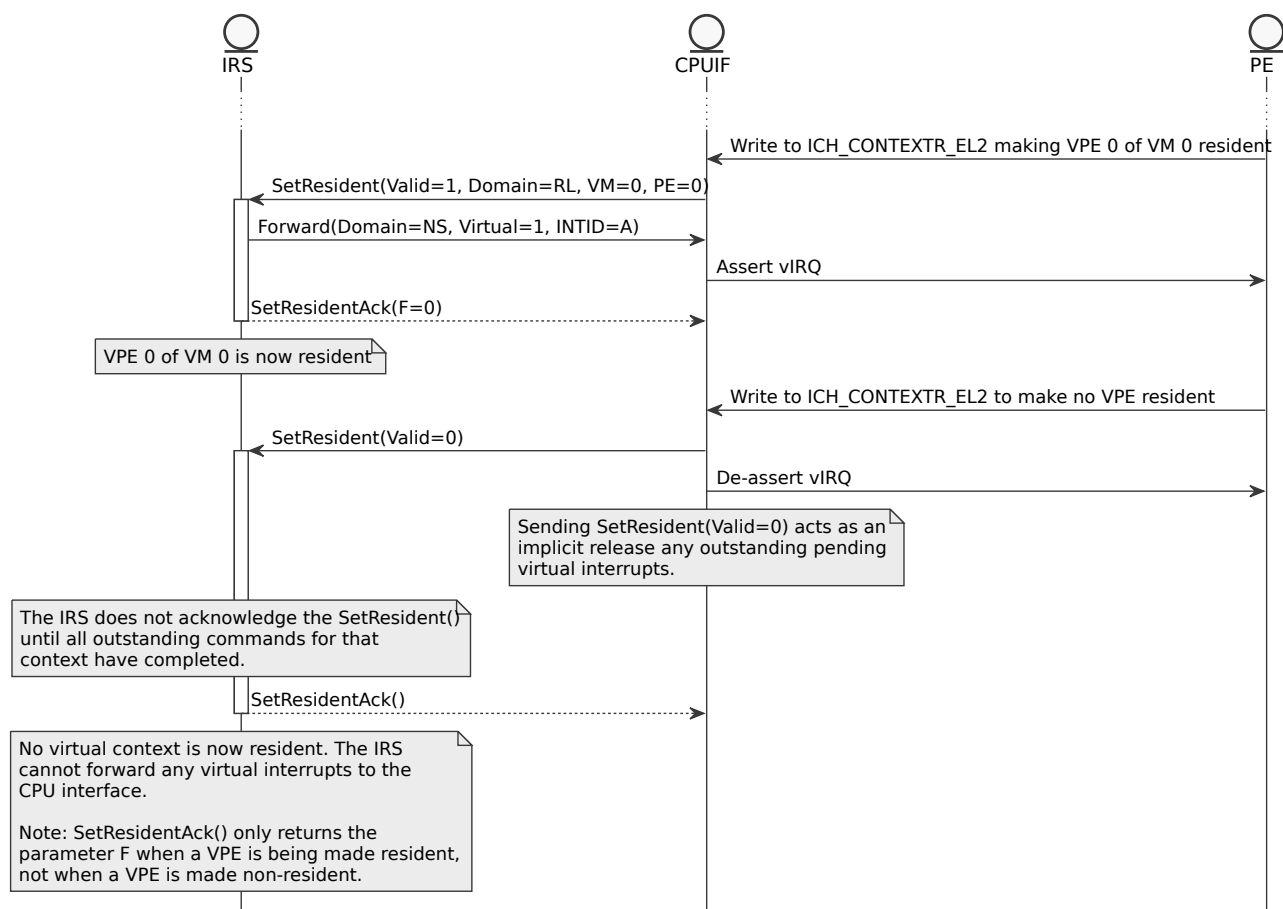
Another execution where the IRS sends the SetAck after the Interrupt Write effects are globally visible is also possible and permitted by the architecture.

A7.4 Making a VPE resident


To set the resident VPE, the CPU interface sends a `SetResident()` command. The IRS responds with a `SetResidentAck()` to show that it has completed the operation. This is illustrated here:



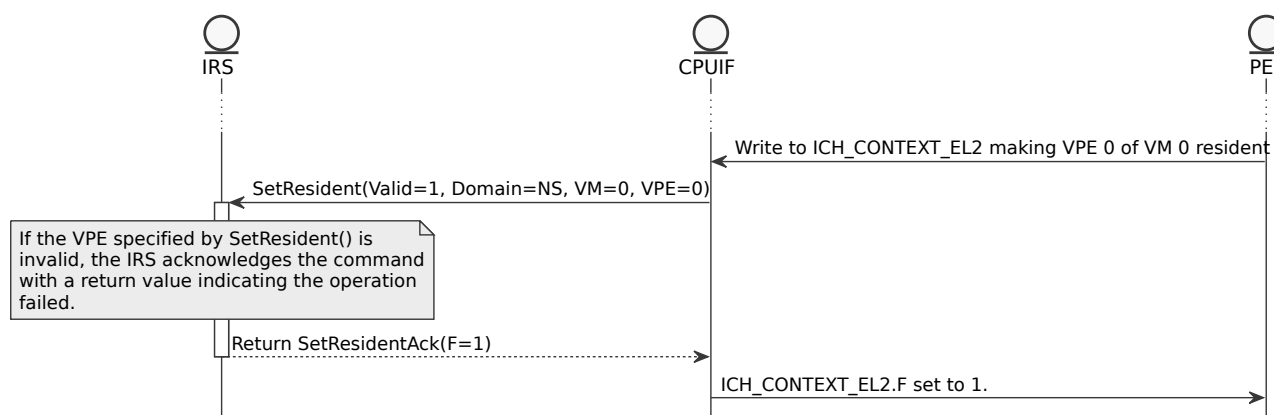
Software can also configure the GIC so that there is no resident VPE, as illustrated here:



To change the resident VPE, the current resident VPE must first be made non-resident:

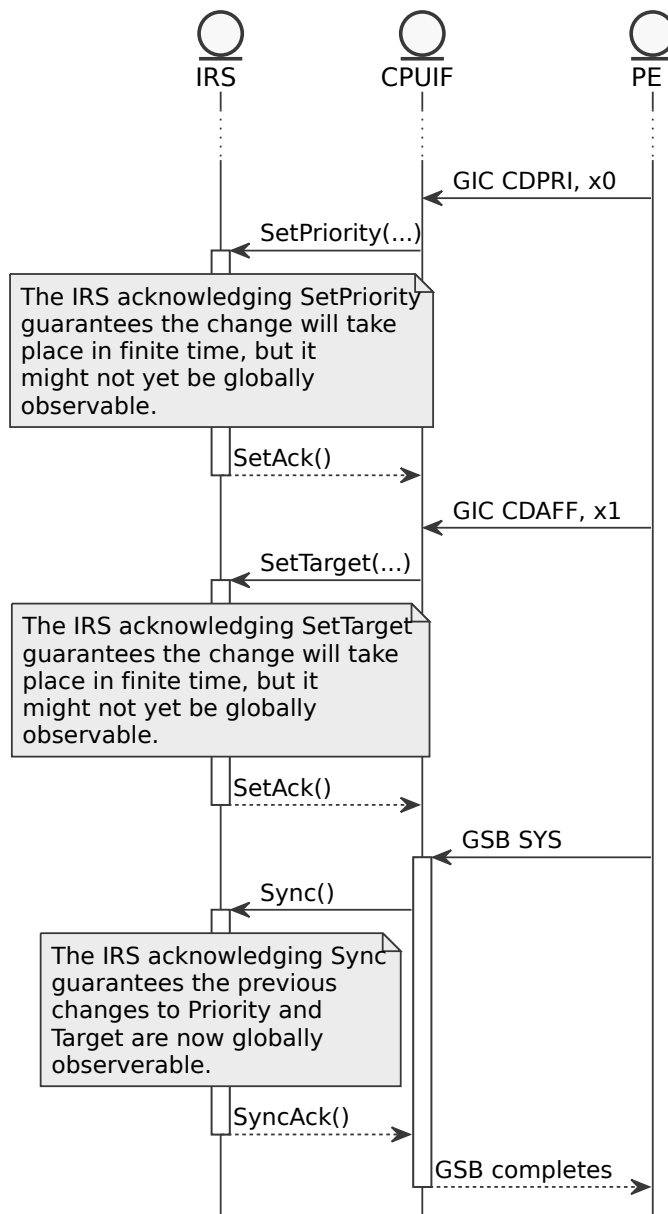


If the CPU interface attempts to make an invalid VPE resident, the IRS acknowledges the command indicating that operation failed:



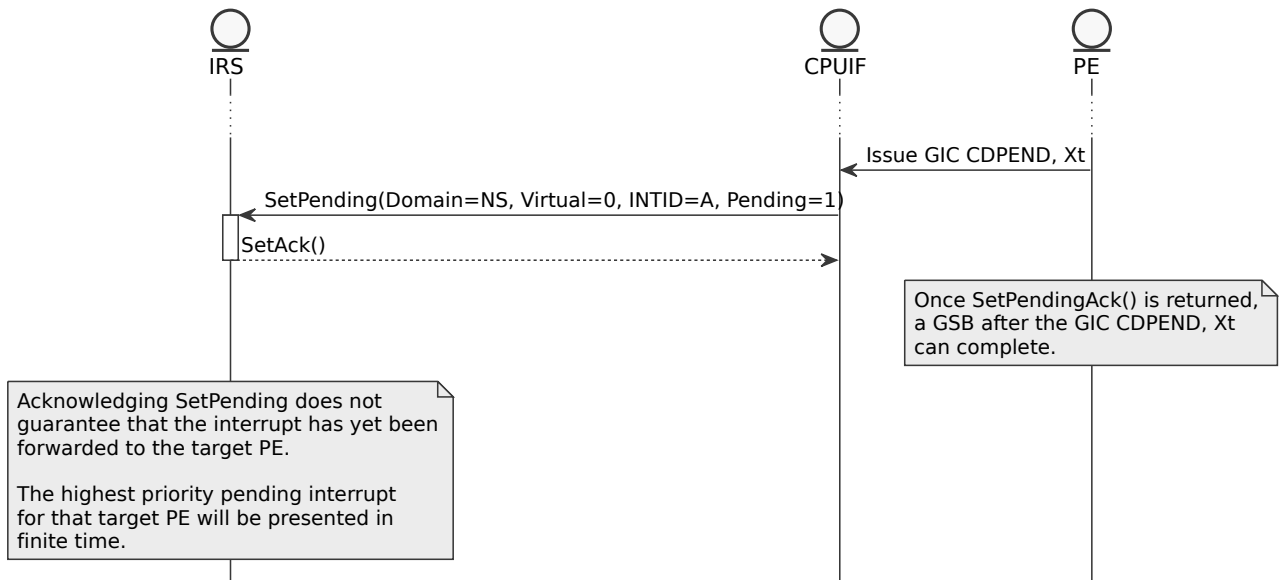
A7.5 Interrupt configuration

The flow for updating the configuration of an LPI or SPI is illustrated below:

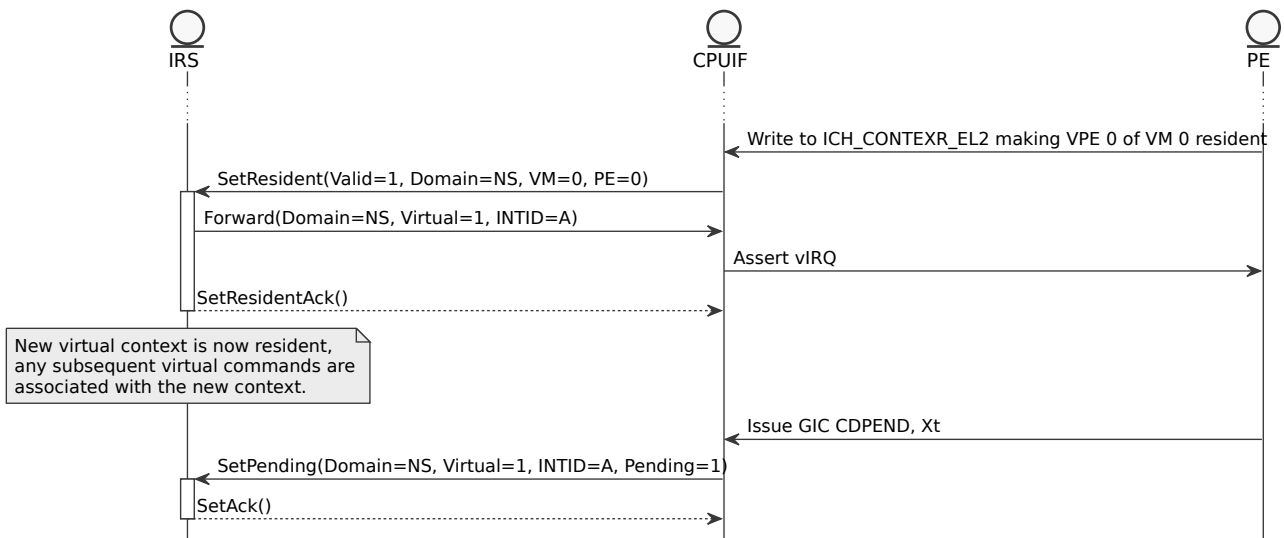


A7.6 Sending IPIs

The flow for generating a physical IPI is illustrated here:

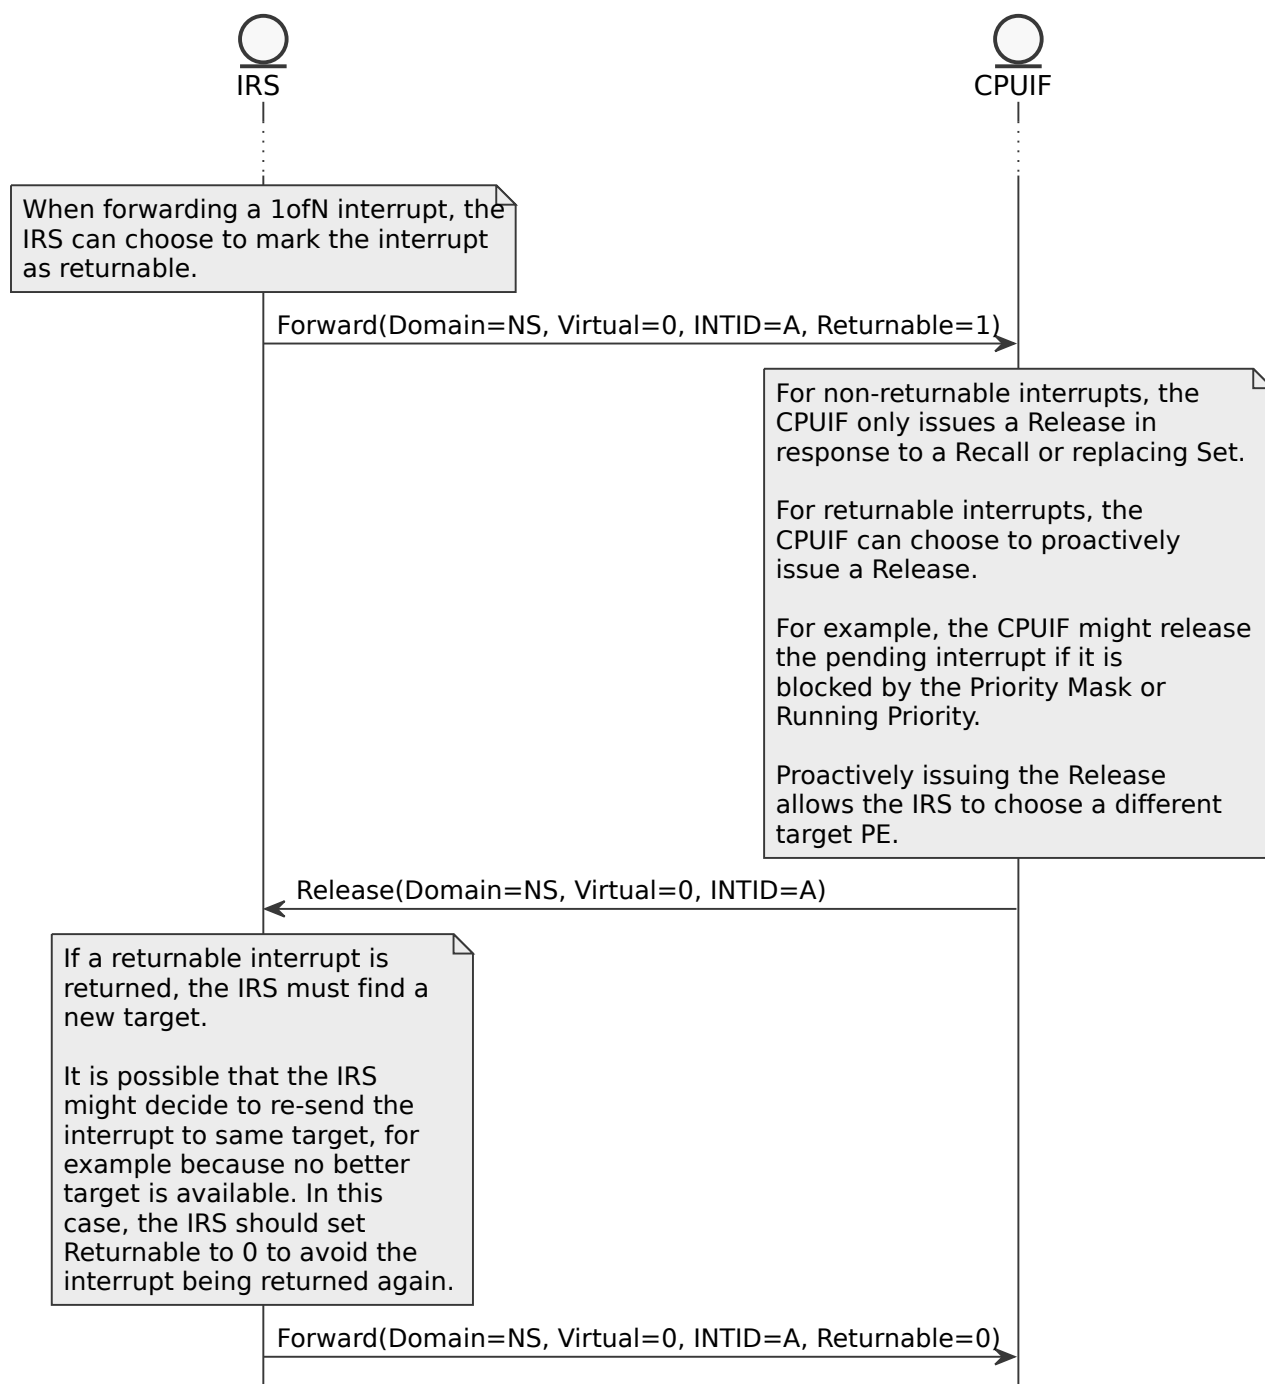


For virtual IPIs, the VPE is taken from the previous SetResident() command, as illustrated here:



A7.7 1 of N interrupts

The following flow demonstrates how the `Returnable` argument to a `Forward` command can be used for 1ofN interrupts:



Part B

Litmus tests

Chapter B1

Interrupt ordering litmus tests

This section gives examples of the use of GIC instructions with and without the use of the GIC specific barrier instructions introduced in this document.

The aim of these tests is to help programmers, hardware design engineers, and validation engineers understand the need for the different kinds of barriers and inherent ordering guarantees.

See also:

- [2.12 Interrupt ordering model and synchronization requirements](#)

B1.1 Interrupt litmus test assumptions

The litmus tests make the following assumptions related to GIC instructions, configuration of the GICv5 CPU interface, and interrupt state and configuration:

- We assume all interrupts are initialized with the following values:

```
[INTID(x)]=(pending:0, enabled:1, priority:1, target:P0, routing_mode:targeted,  
↪ handling_mode:edge)
```
- We assume no asynchronous events take place, including no interrupts become pending without explicit actions taken by one of the processes in the litmus test.
- We assume that ICC_PCR_EL1 and ICC_HAPR_EL1 are both set at the idle priority.
- We assume all GICv5 CPU interface and IRS controls are set to enable delivery of interrupts to PEs included in the litmus test.
- We assume that SCTLR_ELx.NMI is 1.
- We assume interrupt exceptions are masked using PSTATE including interrupts with Superpriority.
- We assume all INTIDs referenced from the litmus tests are reachable from the PEs.
- We assume all PEs are executing in the same Interrupt Domain.
- Interrupts are delivered in finite time.
- Some litmus tests rely on a peripheral with the following properties:
 - The peripheral generates interrupt events to INTID(A) as a result of writing 1 to an MMIO register on the device.
 - A mapping to the interrupt generating MMIO register is annotated with //PERIP in the test.
 - The register initially has the value 0 when read back.
 - The register preserves the values written to it.
 - The peripheral generates a single edge event when there is a write to its MMIO register.
 - The interrupt is made Pending in finite time after the write to the peripheral's MMIO register.
- The WAIT(<instr>; <cond>) construct executes the instruction in a loop enough times that either <cond> is true, or any number of finite time periods have completed.
 - For example, a WAIT(GICR Xt, CDIA; Xt=(valid:1)) will eventually acknowledge an interrupt if all of the following are true:
 - * The interrupt is Pending or becomes Pending in finite time.
 - * The interrupt is Inactive, targeted to the PE, and of sufficient Priority to be acknowledged.

See the documentation of the hertools7 toolsuite (<https://diy.inria.fr/>) for more information about the syntax and use of litmus tests.

B1.2 Atomicity of interrupt updates by GIC system instructions

B1.2.1 Notes

A GIC system instruction generates a [Successful Read-Modify-Write pair](#) of INTID Effects to update an interrupt. The purpose of this section is to show that the Successful Read-Modify-Write pair is required to be [atomic](#).

B1.2.2 Litmus test

```
Arch64 gic.pri+gic.en
{
  [INTID(A)]=(priority:0,enabled:0);
  0:X1=(intid:A, priority:1);
  1:X1=(intid:A);
}
P0          | P1          ;
GIC CDPRI, X1 | GIC CDEN, X1 ;

exists (INTID(A)=(priority:0,enabled:1) /\ INTID(A)=(priority:1,enabled:0))

kinds: gic.pre+gic.en Forbid
```

B1.2.2.1 Explanation

On P0, the GIC CDPRI, X1 system instruction updates the priority of the interrupt INTID(A):


- Generates an Interrupt Read Effect E₁ to INTID(A).
- Modifies the `priority` field of the value read.
- Generates an Interrupt Write Effect E₂ to INTID(A).

E₁ and E₂ are a Successful Read-Modify-Write pair of Interrupt Effects and have to be atomic.

On P1, the GIC CDEN, X1 system instruction enables the interrupt INTID(A):

- Generates an Interrupt Read Effect E₃ to INTID(A).
- Modifies the `enabled` field of the value read.
- Generates an Interrupt Write Effect E₄ to INTID(A).

E₃ and E₄ are a Successful Read-Modify-Write pair of Interrupt Effects and have to be atomic.

Therefore, an execution where E₁ is [Coherence-before](#) E₄, E₄ is [Coherence-before](#) E₂ and as a result INTID(A) has the final value (`priority:1, enabled:0`) is architecturally forbidden. Also, an execution where E₃ is [Coherence-before](#) E₂ and E₂ is [Coherence-before](#) E₄ and as a result INTID(A) has the final value (`priority:0,  enabled:1`) is architecturally forbidden.

B1.3 Multiple updates of the same Interrupt Location

B1.3.1 Notes

The purpose of this test is to show that multiple updates to the same Interrupt Location by GIC system instructions with commands other than `DI` are observed in program order.

B1.3.2 Litmus test with two configuration updates

```
AArch64 coWW-gic
{
  [INTID(A)]=(priority:0);
  0:X1=(intid:A, priority:1);
  0:X2=(intid:A, priority:2);
}
P0
GIC CDPRI, X1
GIC CDPRI, X2

exists(INTID(A)=(priority:0) /\ INTID(A)=(priority:1))


kinds: coWW-gic Forbid
```

B1.3.2.1 Explanation

On P0:

- The first `GIC CDPRI, X1` generates an Explicit Interrupt Write Effect to `INTID(A)` E_1 that updates its priority.
- The second `GIC CDPRI, X1` generates an Explicit Interrupt Write Effect to `INTID(A)` E_2 that updates its priority.

E_1 and E_2 are a Successful Read-Modify-Write pair of Interrupt Effects and are required to be [atomic](#).

Therefore, an execution where E_1 is [Coherence-before](#) E_4 , E_4 is [Coherence-before](#) E_2 and as a result `INTID(A)` has the final value `(priority:1, enabled:0)` is architecturally forbidden. Also, an execution where E_3 is [Coherence-before](#) E_2 and E_2 is [Coherence-before](#) E_4 and as a result `INTID(A)` has the final value `(priority:0,  enabled:1)` is architecturally forbidden.

B1.3.3 Litmus test with an interrupt disable and an interrupt deactivate

```
AArch64 coWR-gic+dis-di+rcfg
{
  [INTID(A)]=(enabled:1, active:1);
  0:X1=(intid:A);
  1:X1=(intid:A);
}
P0 | P1
GIC CDDIS,X1 | GIC CDRCFG, X1
GIC CDDI,X1 | ISB
| MRS X3, ICC_ICSR_EL1
exists(1:X3=(active:0,enabled:1))

kinds: coWR-gic+dis-di+rcfg Allowed
```

B1.3.3.1 Explanation

On P0:

- GIC CDDIS, X1 generates an Explicit Interrupt Write Effect to INTID(A) E₁ and an Explicit Interrupt Write Effect to INTID(A) E₂.
- GIC CDDI, X1 generates an Implicit Interrupt Read Effect to INTID(A) E₃ and an Implicit Interrupt Write Effect to INTID(A) E₄.

On P1:

- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₅.

In the execution that satisfies the postcondition, E₅ **Reads-from** E₄ and E₄ is **Coherence-before** E₁. This is architecturally allowed.

B1.3.4 Litmus test with an interrupt deactivate and an interrupt disable

```
AArch64 coWR-gic+dis-di+rcfg
{
  [INTID(A)]=(enabled:1, active:1);
  0:X1=(intid:A);
  1:X1=(intid:A);
}
P0          | P1
GIC CDDI,X1  | GIC CDRCFG, X1      ;
GIC CDDIS,X1 | ISB          ;
              | MRS X3, ICC_ICSR_EL1 ;
exists(1:X3=(active:1,enabled:0))

kinds: coWR-gic+di-dis+rcfg Allowed
```

B1.3.4.1 Explanation

On P0:

- GIC CDDI, X1 generates an Implicit Interrupt Write Effect to INTID(A) E₁ and an Implicit Interrupt Write Effect to INTID(A) E₂.
- GIC CDDIS, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₃ and an Explicit Interrupt Write Effect to INTID(A) E₄.

On P1:

- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₅.

In the execution that satisfies the postcondition, E₅ **Reads-from** E₄ and E₄ is **Coherence-before** E₁. This is architecturally allowed.

B1.4 Reading back interrupt writes on a single PE

B1.4.1 Notes

The purpose of these tests is to show that writing the configuration of an interrupt or acknowledging an interrupt and reading back the configuration on the same PE guarantees that the written values are observed without the use of explicit synchronization.

B1.4.2 Litmus test with a configuration update

```
AArch64 coWR-gic+en-rcfg
{
  [INTID(A)]=(enabled:0, affinity:P0);
  0:X1=(intid:A);
}
P0
GIC CDEN, X1
GIC CDRCFG, X1
ISB
MRS X3, ICC_ICSR_EL1

exists(0:X3=(affinity:P0,enabled:0))

kinds: coWR-gic+en-rcfg Forbid
```

B1.4.2.1 Explanation

On P0:

- GIC CDEN, X1 generates an Explicit Interrupt Write Effect to INTID(A) E_1 .
- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E_2 .

In the execution that satisfies the postcondition, E_2 [Coherence-before](#) E_1 . This violates the architectural requirements of the [Coherence-before](#) relation and as a result, this execution is architecturally forbidden.

B1.4.3 Litmus test with deactivate

```
AArch64 coWR-gic+di-rcfg
{
  [INTID(A)]=(active:1);
  0:X1=(intid:A);
}
P0
GIC CDDI, X1
GIC CDRCFG, X1
ISB
MRS X2, ICC_ICSR_EL1

exists(0:X2=(active:1))

kinds: coWR-gic+di-rcfg Forbid
```

B1.4.3.1 Explanation

On P0:

- GIC CDDI, X1 generates an Implicit Interrupt Write Effect to INTID(A) E_1 .
- GIC CDRCFG, X1 generates an Interrupt Read Effect to INTID(A) E_2 .

In the execution that satisfies the postcondition, E_2 [Coherence-before](#) E_1 . This violates the architectural requirements of the [Coherence-before](#) relation and as a result, this execution is architecturally forbidden.

B1.4.4 Litmus test with acknowledgement

```
AArch64 coWR-gic+ia-rcfg
{
  [INTID(A)]=(enabled:1, pending:1, affinity:P0);
  0:X1=(intid:A);
}
P0
GICR    X2, CDIA          ;
GIC     CDRCFG, X1        ;
ISB                                           ;
MRS     X3, ICC_ICSR_EL1  ;

exists(0:X2=(valid:1) /\ 0:X3=(pending:1))

kinds: coWR-gic+ia-rcfg Forbid
```

B1.4.4.1 Explanation

On P0:

- GIC X2, CDIA generates an Implicit Interrupt Write Effect to INTID(A) E_1 .
- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E_2 .

In the execution that satisfies the postcondition, E_2 is [Coherence-before](#) E_1 . This violates the architectural requirements of the [Coherence-before](#) relation and as a result, this execution is architecturally forbidden.

B1.4.5 Litmus test with two configuration updates

```
AArch64 coWWR-gic+aff-en-rcfg
{
  [INTID(A)]=(enabled:0, affinity:P0);
  0:X0=(intid:A, affinity:P1);
  0:X1=(intid:A);
}
P0
GIC CDAFF, X0          ;
GIC CDEN, X1           ;
GIC CDRCFG, X1         ;
ISB                    ;
MRS X3, ICC_ICSR_EL1  ;

exists(0:X3=(affinity:P0,enabled:0) /\
      0:X3=(affinity:P1,enabled:0) /\
      0:X3=(affinity:P0,enabled:1))

kinds: coWWR-gic+aff-en-rcfg Forbid
```

B1.4.5.1 Explanation

On P0:

- GIC CDAFF, X0 generates an Explicit Interrupt Write Effect to INTID(A) E_1 .
- GIC CDEN, X1 generates an Explicit Interrupt Read Effect to INTID(A) E_2 and an Explicit Interrupt Write Effect to INTID(A) E_3 .

- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₄.

In the executions that satisfy the postcondition, any of the following applies:

- E₄ is **Coherence-before** E₁. This violates the architectural requirements of the **Coherence-before** relation and as a result, this execution is architecturally forbidden.
- E₂ is **Coherence-before** E₁ and E₄~ **Reads-from** E₃. This violates the architectural requirements of the **Coherence-before** relation and as a result, this execution is architecturally forbidden.

B1.5 Reading interrupt configurations and subsequent updates

B1.5.1 Notes

The purpose of these tests is to show that when reading the configuration of an interrupt and then updating the configuration of that same interrupt, the value read is not affected by the updates.

B1.5.2 Litmus test with update to priority

```
AArch64 coRW-gic+rcfg-pri
{
  [INTID(A)]=(priority:0);
  0:X0=(intid:A);
  0:X1=(intid:A, priority:1);
}
P0
GIC CDRCFG, X0
GIC CDPRI, X1
ISB
MRS X3, ICC_ICSR_EL1

exists(0:X3=(priority:1))

kinds: coRW-gic+rcfg-pri Forbid
```

B1.5.2.1 Explanation

On P0:

- GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₁.
- GIC CDPRI, X1 generates an Explicit Interrupt Write Effect to INTID(A) E₂.

It is architecturally required the Interrupt Read Effect generated by an GIC system instruction with the RCFG command E₁ does not [Read-from](#) the Interrupt Write effect E₂.

Note that the ISB and MRS X3, ICC_ICSR_EL1 sequence can be placed before the GIC CDPRI, X1 instruction without changing the outcome of the test.

B1.5.3 Litmus test with deactivate

```
AArch64 coRW-gic+rcfg-di
{
  [INTID(A)]=(active:1);
  0:X0=(intid:A);
  0:X1=(intid:A);
}
P0
GIC CDRCFG, X0
GIC CDDI, X1
ISB
MRS X3, ICC_ICSR_EL1

exists(0:X3=(active:0))

kinds: coRW-gic+rcfg-di Forbid
```

B1.5.3.1 Explanation

On P0:

- GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₁.
- GIC CDDI, X1 generates an Implicit Interrupt Write Effect to INTID(A) E₂.

It is architecturally required the Interrupt Read Effect generated by an GIC system instruction with the RCFG command E₁ does not [Read-from](#) the Interrupt Write effect E₂.

Note that the ISB and MRS X3, ICC_ICSR_EL1 sequence can be placed before the GIC CDDI, X1 instruction without changing the outcome of the test.

B1.5.4 Litmus test with acknowledge

```
AArch64 coRW-gic+rcfg-ia
{
  [INTID(A)]=(priority:1,pending:1);
  0:X0=(intid:A);
}
P0
GIC CDRCFG, X0      ;
GICR X1,CDIA        ;
ISB                  ;
MRS X3, ICC_ICSR_EL1 ;

exists(0:X3=(priority:1,pending:0,active:1))

kinds: coRW-gic+rcfg-ia Forbid
```

B1.5.4.1 Explanation

On P0:

- GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₁.
- GICR X1,CDIA generates an Implicit Interrupt Write Effect to INTID(A) E₂.

It is architecturally required the Interrupt Read Effect generated by an GIC system instruction with the RCFG command E₁ does not [Read-from](#) the Interrupt Write effect E₂.

Note that the ISB and MRS X3, ICC_ICSR_EL1 sequence can be placed before the GICR X1,CDIA instruction without changing the outcome of the test.

B1.6 Configuration and acknowledgement

B1.6.1 Notes

The purpose of this test is to show that disabling an interrupt is only required to be observed by an acknowledge when using explicit synchronization.

Illustrates a common principle of interrupt acknowledgement not being guaranteed to observe a preceding write in program order without either using explicit synchronization or observing the write through an explicit read.

See also:

- [B1.11 IPI and acknowledgement](#)
- [B1.33 Reading interrupt configuration and exception status](#)

B1.6.2 Litmus test using disable without explicit synchronization

```
AArch64 coWR-gic.dis.ack
{
  [INTID(A)]=(enabled:1, pending:1);
  0:X1=(intid:A);
}
P0
GIC CDDIS, X1
GICR X0, CDIA

exists(0:X0=(valid:1, intid:A))

kinds: coWR-gic.dis.ack Allow
```

B1.6.2.1 Explanation

In the execution that satisfies the postcondition, on P0:

- GIC CDDIS, X1 generates an Explicit Interrupt Write Effect to INTID(A) E_1 .
- GICR X0, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E_2 .

In the execution that satisfies the postcondition, E_2 is [Coherence-before](#) E_1 . This is permitted by the architecture.

B1.6.3 Litmus test using disable with explicit synchronization

```
AArch64 coWR-gic.dis.ack+gsb.sys
{
  [INTID(A)]=(enabled:1, pending:1);
  0:X1=(intid:A);
}
P0
GIC CDDIS, X1
GSB SYS
GICR X0, CDIA

exists(0:X0=(valid:1, intid:A))

kinds: coWR-gic.dis.ack+gsb.sys Forbid
```

B1.6.3.1 Explanation

In the execution that satisfies the postcondition, on P0:

- GIC CDDIS, X1 generates an Explicit Interrupt Write Effect to INTID(A) E_1 .

- GSB SYS generates a GSB.SYS Effect E_2 .
- GICR X0, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E_3 .

Therefore, E_1 is GSB-Ordered-before E_3 . As per the postcondition, E_3 is Coherence-before E_1 . Consequently, this execution creates a cycle in the Ordered-before relation and as a result, violates the External visibility requirement.

B1.6.4 Litmus test with deactivate

```
AArch64 coWW-gic+di-ia+rcfg
{
  [INTID(A)]=(pending:1,enabled:1,affinity:P0);
  0:X0=(intid:A);
  1:X0=(intid:A);
}
P0          | P1
GIC CDDI,X0 | GIC CDRCFG,X0
GICR X1,CDIA | ISB
              | MRS X2,ICC_ICSR_EL1
exists(0:X1=(valid:1,intid:A) /\ 1:X2=(pending:0,active:0,enabled:1))

kinds: coWW-gic+di-ia+rcfg
```

B1.6.4.1 Explanation

In the execution that satisfies the postcondition, on P0:

- GIC CDDI, X0 generates an Implicit Interrupt Write Effect to INTID(A) E_1 and an Implicit Interrupt Write Effect to INTID(A) E_2 .
- GICR X1, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E_3 and an Implicit Interrupt Write Effect to INTID(A) E_4 .
- GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect to INTID(A) E_5 .

In the execution that satisfies the postcondition, E_4 is [Coherence-before](#) E_2 and E_5 [Reads-from](#) E_2 . This means that E_2 is [Coherence-before](#) E_3 and that violates the architectural requirements of the [Coherence-before relation](#). As a result, this execution is architecturally forbidden.

B1.6.5 Litmus test using priority without explicit synchronization

```
AArch64 coWR-gic+pcr-isb-pri-ia
{
  [INTID(A)]=(enabled:1, pending:1, priority:1);
  0:X1=(intid:A, priority:2);
}
P0
MOV X0, #1
MSR ICC_PCR_EL1, X0
ISB
GIC CDPRI, X1
GICR X0, CDIA

exists(0:X0=(valid:1, intid:A))

kinds: coWR-gic+pcr-isb-pri-ia Allow
~~~
```

B1.6.5.1 Explanation

In the execution that satisfies the postcondition, on P0:

- MSR ICC_PCR_EL1, X0 generates a Direct System Register Write Effect to ICC_PCR_EL1 E₁.
- ISB generates a Instruction Fetch Barrier Effect E₂.
- GIC CDPRI, X1 generates an Explicit Interrupt Write Effect to INTID(A) E₃.
- GICR X0, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₄ and an Indirect System Register Read Effect to ICC_PCR_EL1 E₅.

E₁ is IFB-Ordered-before E₅ and, consequently, E₅ Reads-from-register E₁. This means that GICR X0, CDIA ↪ cannot acknowledge any interrupt with priority lower than 1. As per the postcondition, GICR X0, CDIA acknowledges interrupt INTID(A) and since the priority of INTID(A) cannot have been lower than 1, E₄ is Coherence-before E₃. This is an architecturally allowed execution.

B1.6.6 Litmus test using priority with explicit synchronization.

```
AArch64 coWR-gic+pcr-isb-pri-gsb.sys-ia
{
  [INTID(A)]=(enabled:1, pending:1, priority:1);
  0:X1=(intid:A, priority:2);
}
P0
MOV  X0, #1
MSR  ICC_PCR_EL1, X0
ISB
GIC  CDPRI, X1
GSB  SYS
GICR X0, CDIA

exists(0:X0=(valid:1, intid:A))

kinds: coWR-gic+pcr-isb-pri-gsb.sys-ia Forbid
```

B1.6.6.1 Explanation

In the execution that satisfies the postcondition, on P0:

- MSR ICC_PCR_EL1, X0 generates a Direct System Register Write Effect to ICC_PCR_EL1 E₁.
- ISB generates a Instruction Fetch Barrier Effect E₂.
- GIC CDPRI, X1 generates an Explicit Interrupt Write Effect to INTID(A) E₃.
- GSB SYS generates a GSB.SYS Effect E₄.
- GICR X0, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₅ and an Indirect System Register Read Effect to ICC_PCR_EL1 E₆.

E₁ is IFB-Ordered-before E₆ and, consequently, E₆ Reads-from-register E₁. This means that GICR X0, CDIA ↪ cannot acknowledge any interrupt with priority lower than 1. As per the postcondition, GICR X0, CDIA acknowledges interrupt INTID(A) and since the priority of INTID(A) cannot have been lower than 1, E₅ is Coherence-before E₃. At the same time, E₃ is GSB-Ordered-before E₅. This creates a cycle in the Ordered-before relation which violates the External visibility requirement.

B1.7 Acknowledge followed by interrupt changes

B1.7.1 Notes

The update to an interrupt caused by an interrupt acknowledge is required to be ordered with respect to a deactivate but not with any other configuration update.

B1.7.2 Litmus test with deactivate

```
AArch64 ack-deactivate
{
  [INTID(A)]=(enabled:1,affinity:P0,pending:1,active:0);
  0:X0=(intid:A);
}
P0
GICR  X1, CDIA          ;
GIC   CDDI, X0          ;
GIC   CDRCFG, X0        ;
ISB
MRS   X2, ICC_ICSR_EL1  ;

exists(0:X1=(valid:1) /\ 0:X2=(pending:0, active:1))

kinds: ack-deactivate Forbid
```

B1.7.2.1 Explanation

In the execution that satisfies the postcondition, on P0:

- GICR X1, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₁. As per the postcondition, the GICR system instruction acknowledges INTID(A) and, as a result, generates an Implicit Interrupt Write Effect E₂ to INTID(A) that resets its pending state to 0 and sets its active state to 1.
- GIC CDDI, X0 generates an Implicit Interrupt Write Effect to INTID(A) E₃.
- GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₄.

As per the postcondition, the Explicit Interrupt Read Effect E₄ observes that INTID(A) is active, and as a result, E₄ [Reads-from](#) E₂. This means that either:

- E₄ is [Coherence-before relation](#) E₃ which violates the architectural requirements of the [Coherence-before relation](#), or
- E₃ is [Coherence-before relation](#) E₂ and consequently E₁ [Reads-from](#) E₃ which violates the architectural requirements of the [Reads-from relation](#).

B1.7.3 Litmus test with make pending

```
AArch64 edge-merging.ipi
{
  [INTID(A)]=(enabled:1, pending:1, affinity:P0);
  0:X0=(intid:A, pending:1); 0:X1=(intid:A);
}
P0
GICR  X2, CDIA          ;
GIC   CDPEND, X0        ;
GIC   CDRCFG, X1        ;
ISB
MRS   X3, ICC_ICSR_EL1  ;

exists(0:X2=(valid:1, intid:A) /\
```

```
0:X3=(pending:0, active:1))
```

```
kinds: edge-merging.ipi Forbid
```

B1.7.3.1 Explanation

On P0:

- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₁ and an Implicit Interrupt Write Effect to INTID(A) E₂ that resets the pending state to 0 and sets the active state to 1.
- GIC CDPEND, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₃ and an Explicit Interrupt Write Effect to INTID(A) E₄ that sets the pending state to 1.
- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect E₅ to INTID(A).

In the execution that satisfies the postcondition, E₅ [Reads-from](#) E₂, and any of the following applies:

- E₄ is [Coherence-before](#) E₂. This violates the requirements of the [Coherence-before](#) and, as a result, is architecturally forbidden.
- E₂ is [Coherence-before](#) E₄ and, consequently, E₅ is [Coherence-before](#) E₄. This violates the [requirements of the Coherence-before](#) and, as a result, is architecturally forbidden.

B1.8 Multiple updates with interleaved read

B1.8.1 Notes

The purpose of this test is to show that a write, followed by a read, followed by a write to the same interrupt location must result in the read observing the first write and not the second write.

B1.8.2 Litmus test

```
AArch64 coWRW-gic
{
  [INTID(A)]=(priority:0);
  0:X0=(intid:A);
  0:X1=(intid:A, priority:1);
  0:X2=(intid:A, priority:2);
}
P0
GIC CDPRI, X1
GIC CDRCFG, X0
GIC CDPRI, X2
ISB
MRS X3, ICC_ICSR_EL1

exists(0:X3=(priority:2) /\
      0:X3=(priority:0))

kinds: coWRW-gic Forbid
```

B1.8.2.1 Explanation

On P0:

- The first `GIC CDPRI, X1` generates an Explicit Interrupt Write Effect E_1 .
- `GIC CDRCFG, X1` generates an Explicit Interrupt Read Effect E_2 .
- The second `GIC CDPRI, X2` generates an Explicit Interrupt Write Effect E_3 .

It is architecturally required that:

- The Interrupt Read Effect generated by an `GIC` system instruction with the `RCFG` command E_2 is [Coherence-after](#) the Interrupt Write effect E_1 .
- The Interrupt Read Effect generated by an `GIC` system instruction with the `RCFG` command E_1 does not [Read-from](#) the Interrupt Write effect E_3 .

Note that the `ISB` and `MRS X3, ICC_ICSR_EL1` sequence can be placed before the second `GIC CDPRI, X2` instruction without changing the outcome of the test.

B1.9 Configuration write and IRQ unmask in PSTATE

B1.9.1 Notes

The purpose of these tests is to show the required synchronization of to ensure that after disabling an pending interrupt there cannot be a IRQ exception when possible.

B1.9.2 Litmus test

```
AArch64 gic.cddis-gsb.sys-mrs.isr_el1
{
  [INTID(A)]=(enabled:1, pending:1);
  0:X1=(intid:A);
}
P0
GIC CDDIS, X1
GSB SYS
MRS X0, ISR_EL1
exists(0:X0=(I=1))

kinds: gic.cddis-gsb.sys-mrs.isr_el1 Forbid
```

B1.9.2.1 Explanation

For the execution that validates the postcondition:

- GIC CDDIS, X1 generates an Interrupt Write Effect E_1 to Interrupt Location INTID(A).
- GSB SYS generates a GSB.SYS Effect E_2 .
- MRS X0, ISR_EL1 generates a Direct System Register Read to ISR_EL1 E_3 .

In the execution that satisfies the postcondition, E_3 indicates that INTID(A) is still the HPPI. However because of E_2 INTID(A) is disabled and cannot be the HPPI. As a result, this execution is architecturally forbidden.

B1.9.3 Litmus test with initially masked IRQs

```
AArch64 coWex+gic.cddis-gsb.sys-clr.i
{
  [INTID(A)]=(enabled:1, pending:1);
  0:X1=(intid:A);
}
P0
GIC CDDIS, X1
GSB SYS
DAIFClr PSTATE.I
exists(async(P0, IRQ))

kinds: coWex+gic.cddis-gsb.sys-clr.i Forbid
```

B1.9.3.1 Explanation

For the execution that validates the postcondition:

- GIC CDDIS, X1 generates an Interrupt Write Effect E_1 to Interrupt Location INTID(A).
- GSB SYS generates a GSB.SYS Effect E_2 .
- DAIFClr PSTATE.I generates a Direct System Register Write to PSTATE E_3 .

In the execution that satisfies the postcondition, `INTID(A)` has become a pending IRQ. This execution is architectural forbidden as the `GSB.SYS` is required to have made the `E1` observable before the IRQ was unmasked.

B1.9.4 Litmus test with disable of a PPI

```
AArch64 ppi-disable-isb-clr.i
{
  0:[ppi(27)]=(enabled:1,pending:1,active:0);
}
P0
MOV    X0, #0                ;
MSR    ICC_PPI_ENABLER, X0    ;
ISB                                ;
DAIFClr PSTATE.I              ;

exists(async(P0, IRQ)) Forbid
```

B1.9.4.1 Explanation

The ISB guarantees completion of the MSR instructions and the effects of disabling the PPI which can no longer be the HPPI. As a result, this execution is architecturally forbidden.

B1.10 Configuration write and exception status on a single PE

B1.10.1 Notes

The purpose of this test is to show that an interrupt exception as a result of enabling a pending interrupt is not required to be observed even when explicit synchronization is used.

B1.10.2 Litmus test without wait for IRQ exception to be signaled

```
AArch64 coWR+gic.cden-gsb.sys-isb-isr
{
  [INTID(A)]=(enabled:0, pending:1);
  0:X1=(intid:A);
}
P0
GIC CDEN, X1
GSB SYS
ISB
MRS X0, ISR_EL1

exists(0:X0=(I=0))

kinds: coWR+gic.cden-gsb.sys-isb-isr Allow
```

B1.10.2.1 Explanation

On P0:

- GIC CDEN, X1 generates an Explicit Interrupt Write Effect to INTID(A) E₁.
- GSB SYS generates a GSB.SYS Effect E₂.
- ISB generates a Instruction Fetch Barrier Effect E₃.
- MRS X0, ISR_EL1 generates a Direct System Register Read Effect to ISR_EL1 E₄.

E₁ is IFB-Ordered-before E₁. As a result, INTID(A) is already a candidate HPPI when P0 generates E₄. However, the architecture does not guarantee that INTID(A) will be presented as the HPPI immediately. And as a result, the execution where reading ISR_EL1 indicates there is no IRQ pending is architecturally allowed.

B1.10.3 Litmus test with wait for IRQ exception to be signaled

```
AArch64 coWR+gic.cden-gsb.sys-isb-isr+wait
{
  [INTID(A)]=(enabled:0, pending:1);
  0:X1=(intid:A);
}
P0
GIC CDEN, X1
GSB SYS
ISB
WAIT (MRS X0, ISR_EL1; X0=(I=1))

exists(0:X0=(I=0))

kinds: coWR+gic.cden-gsb.sys-isb-isr+wait Forbid
```

B1.10.3.1 Explanation

On P0:

- `GIC_CDEN, X1` generates an Explicit Interrupt Write Effect to `INTID(A)` E_1 .
- `GSB_SYS` generates a `GSB.SYS` Effect E_2 .
- `ISB` generates a Instruction Fetch Barrier Effect E_3 .
- The execution of `MRS X0, ISR_EL1` in a loop generates a Direct System Register Read Effect to `ISR_EL1` E_4, E_5, \dots, E_n .

E_1 is IFB-Ordered-before E_1 . As a result, `INTID(A)` is already a candidate HPPI when `P0` generates E_4 . The architecture does not guarantee that `INTID(A)` will be presented as the HPPI immediately. However, in finite time the GIC will present `INTID(A)` as the HPPI, as a result, E_n where n is a finite number reads that the `I` field of `ISR_EL1` is set to 1 and an IRQ is pending.

B1.11 IPI and acknowledgement

B1.11.1 Notes

The purpose of these tests is to show that sending an IPI is only required to be observed by an acknowledge when using explicit synchronization.

Illustrates a common principle of interrupt acknowledgement not being guaranteed to observe a preceding write in program order without either using explicit synchronization or observing the write through an explicit read.

See also:

- [B1.6 Configuration and acknowledgement](#)

B1.11.2 Litmus test without explicit synchronization

```
AArch64 coWR-gic.ipi.ack
{
  [INTID(A)]=(enabled:1, pending:1, affinity:P0);
  0:X0=(intid:A, pending:1); 0:X1=(intid:A);
}
P0
GIC      CDPEND, X0          ;
GICR     X2, CDIA           ;
GICR     X2, CDIA           ;
GIC      CDRCFG, X1         ;
ISB                               ;
MRS      X3, ICC_ICSR_EL1   ;

exists(0:X2=(valid:1) /\ 0:X3=(pending:1, active:1))

kinds: coWR-gic.ipi.ack Allow
```

B1.11.2.1 Explanation

On P0:

- GIC CDPEND, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₁ and an Explicit Interrupt Write Effect to INTID(A) E₂ that sets the pending state to 1.
- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₄ and an Implicit Interrupt Write Effect to INTID(A) E₅ that resets the pending state to 0 and sets the active state to 1.
- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₆.

In the execution that satisfies the postcondition, GICR X2, CDIA acknowledges INTID(A) and consequently, either:

- E₆ [Reads-from](#) the initial value of INTID(A) and therefore E₆ is [Coherence-before](#) E₂. This violates [architectural requirements of the Coherence-before relation](#).
- E₆ [Reads-from](#) E₂. This means that either:
 - E₂ is [Coherence-before](#) E₅ and E₆ is [Coherence-before](#) E₅. This violates [architectural requirements of the Coherence-before relation](#).
 - E₅ is [Coherence-before](#) E₂. This execution is architecturally allowed.

B1.11.3 Litmus test with explicit synchronization

```
AArch64 coWR-gic.ipi.ack+gsb.sys
{
  [INTID(A)]=(enabled:1, pending:1, affinity:P0);
```



```

0:X0=(intid:A, pending:1); 0:X1=(intid:A);
}
P0
GIC      CDPEND, X0      ;
GSB      SYS             ;
GICR     X2, CDIA        ;
GIC      CDRCFG, X1      ;
ISB                      ;
MRS      X3, ICC_ICSR_EL1 ;

exists(0:X2=(valid:1) /\ 0:X3=(pending:1))

kinds: coWR-gic.ipi.ack+gsb.sys Forbid

```

B1.11.3.1 Explanation

On P0:

- GIC CDPEND, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₁ and an Explicit Interrupt Write Effect to INTID(A) E₂ that sets the pending state to 1.
- GSB SYS generates a GSB.SYS Effect E₃.
- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₄ and an Implicit Interrupt Write Effect to INTID(A) E₅ that resets the pending state to 0 and sets the active state to 1.
- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₆.

In the execution that satisfies the postcondition, GICR X2, CDIA acknowledges INTID(A) and consequently, either:

- E₆ Reads-from the initial value of INTID(A) and therefore E₆ is Coherence-before E₂. This violates architectural requirements of the Coherence-before relation.
- E₆ Reads-from E₂. This means that either:
 - E₂ is Coherence-before E₅ and E₆ is Coherence-before E₅. This violates architectural requirements of the Coherence-before relation.
 - E₅ is Coherence-before E₂ which would mean that E₁ Reads-from E₅ and as a result E₅ is GIC-Observed-by E₁. At the same time, E₁ is GSB-ordered-before E₅. This creates a cycle in the Ordered-before relation and, as a result, the execution violates the External visibility requirement.

B1.12 Observing multiple writes on a different PE

B1.12.1 Notes

The purpose of this test is to show that the external observations of updates to the same interrupt location cannot contradict the Coherence order.

B1.12.2 Litmus test

```
AArch64 coWW-gic+R
{
  [INTID(A)]=(enabled:0, affinity:P0);
  0:X1=(intid:A, affinity:P1);
  0:X2=(intid:A);
  1:X1=(intid:A);
}
P0          | P1
GIC CDAFF, X1 | GIC CDRCFG, X1 ;
GIC CDEN, X2  | ISB ;
              | MRS X0, ICC_ICSR_EL1 ;

exists(1:X0=(affinity:P0, enabled:1))

kinds: coWW-gic+R Forbid
```

B1.12.2.1 Explanation

On P0:

- GIC CDAFF, X1 generates an Explicit Interrupt Write Effect to INTID(A) E_1 .
- GIC CDEN, X2 generates an Explicit Interrupt Read Effect to INTID(A) E_2 and an Explicit Interrupt Write Effect to INTID(A) E_3 .

On P1:

- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E_4 .

For the execution that satisfies the post-condition, E_4 Reads-from E_3 and E_2 is Coherence-before E_1 . This violates the [architectural requirement of the Reads-from relation](#).

B1.13 Read of the configuration of an interrupt

B1.13.1 Notes

The purpose of this section is to demonstrate the ordering properties of the Indirect System Register Write to ICC_ICSR_EL1 generated by the execution of a GIC system instruction with the RCFG command.

B1.13.2 Litmus test without ISB

```
AArch64 gic.rcfg-mrs.icc_icsr_el1
{
  [INTID(A)]=(priority:5);
  0:X0=(intid:A);
}
P0
GIC CDRCFG, X0
MRS X1, ICC_ICSR_EL1 ;

exists(1:X1!=(priority:5))

kinds: gic.rcfg-mrs.icc_icsr_el1 Allow
```

B1.13.2.1 Explanation

On P0:

- GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect E₁ and an Indirect System Register Write Effect E₂.
- MRS X1, ICC_ICSR_EL1 generates a Direct System Register Read Effect to ICC_ICSR_EL1 E₃ and a Register Write Effect to X1 E₄.

In the execution that satisfies the postcondition, E₁ Reads-from the initial configuration of INTID(A) and consequently E₃ is Coherence-before E₁. This is architecturally allowed.

B1.13.3 Litmus test with ISB

```
AArch64 gic.rcfg-isb-mrs.icc_icsr_el1
{
  [INTID(A)]=(priority:5);
  0:X0=(intid:A);
}
P0
GIC CDRCFG, X0
ISB
MRS X1, ICC_ICSR_EL1 ;

exists(1:X1!=(priority:5))

kinds: gic.rcfg-isb-mrs.icc_icsr_el1 Forbid
```

B1.13.3.1 Explanation

On P0:

- GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect E₁ and an Indirect System Register Write Effect E₂.
- ISB generates an Instruction Fetch Barrier Effect E₂ which is a Context Synchronization Event.

- `MRS X1, ICC_ICSR_EL1` generates a Direct System Register Read Effect to `ICC_ICSR_EL1` E_3 and a Register Write Effect to `X1` E_4 .

In the execution that satisfies the postcondition, E_1 Reads-from the initial configuration of `INTID(A)` and consequently E_3 is Coherence-before E_1 . This is architecturally forbidden. E_3 is required to Read-from E_1 because E_2 is a Context Synchronization Event.

See also *Arm® Architecture Reference Manual, for A-profile architecture*[\[1\]](#) Table D24-1 Synchronization requirements' for more information.

B1.14 Multiple reads of the same config

B1.14.1 Notes

The purpose of this test is to show that if a read has observed a write to the same location, a second read must also have observed the same write.

The first litmus test uses an ISB after both GIC CDRCFG, X0 instructions to ensure that the implicit writes to ICC_ICSR_EL1 are observed by each MRS instruction after the ISB. The subsequent litmus tests illustrate the impact of omitting one or both ISBs.

B1.14.2 Litmus test with ISBs

```
AArch64 coRR-gic+cdpri+cfg-isb-cfg-isb
{
  [INTID(A)]=(priority:0);
  0:X1=(intid:A, priority:1);
  1:X0=(intid:A);
}

P0          | P1
GIC CDPRI, X1 | GIC CDRCFG, X0
              | ISB
              | MRS X1, ICC_ICSR_EL1
              | GIC CDRCFG, X0
              | ISB
              | MRS X2, ICC_ICSR_EL1

exists(1:X1=(priority:1) /\ 1:X2=(priority:0))

kinds: coRR-gic+cdpri+cfg-isb-cfg-isb Forbid
```

B1.14.2.1 Explanation

On P0:

- GIC CDPRI, X1 generates an Explicit Interrupt Write Effect E₁.

On P1:

- The first GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect E₂ and an Indirect System Register Write Effect to ICC_ICSR_EL1 E₃.
- The first ISB generates an IFB Effect E₄.
- MSR C1, ICC_ICSR_EL1 generates a Direct System Register Read Effect E₅.
- The second GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect E₆ and an Indirect System Register Write Effect to ICC_ICSR_EL1 E₇.
- The second ISB generates an IFB Effect E₈.
- MSR C1, ICC_ICSR_EL1 generates a Direct System Register Read Effect E₉.
- The first GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect E₁₀ and an Indirect Write to System Register ICC_ICSR_EL1 E₁₁.

In the execution that satisfies the postcondition, E₂ Reads-from E₁ and as a result E₂, E₁ is Observer-by E₂ and one of the following applies:

E₃ is Coherence-before E₁ and as a result, E₂ is [GIC-Hazard-Ordered-before](#) E₁. This creates a cycle in the Ordered-before relation which violates the External visibility requirement.

B1.15 GIC coherence order

B1.15.1 Notes

The purpose of this test is to show that due to the total order on writes to an interrupt location, two reads of the same location by two separate observers must observe the same order on the writes.

B1.15.2 Litmus test

```
AArch64 IRIW-loc+gic.cdpr
{
  [INTID(A)]=(priority:0);
  0:X1=(intid:A, priority:1);
  1:X1=(intid:A, priority:2);
  2:X0=(intid:A);
  3:X0=(intid:A);
}

P0          | P1          | P2          | P3          ;
GIC CDPRI, X1 | GIC CDPRI, X1 | GIC CDRCFG, X0 | GIC CDRCFG, X0 ;
              |              | ISB          | ISB          ;
              |              | MRS X1, ICC_ICSR_EL1 | MRS X1, ICC_ICSR_EL1 ;
              |              | GIC CDRCFG, X0 | GIC CDRCFG, X0 ;
              |              | ISB          | ISB          ;
              |              | MRS X2, ICC_ICSR_EL1 | MRS X2, ICC_ICSR_EL1 ;

exists((2:X1=(priority:1) /\ 2:X2=(priority:2) /\
        3:X1=(priority:2) /\ 3:X2=(priority:1))

kinds: IRIW-loc+gic.cdpr Forbid
```

B1.15.2.1 Explanation

In the execution that satisfies the postcondition:

- On P0, GIC CDPRI, X1 generates an Explicit Interrupt Write Effect to INTID(A) E₁ that sets its priority to 1.
- On P1, GIC CDPRI, X1 generates an Explicit Interrupt Write Effect to INTID(A) E₂ that sets its priority to 2.
- On P2:
 - The first GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₃ that reads the priority as 1.
 - The second GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₄ that reads the priority as 2.
- On P3:
 - The first GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₅ that reads the priority as 2.
 - The second GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₆ that reads the priority as 1.

Any of the following applies:

- E₁ is [Coherence-before](#) E₂. In which case E₆ is [Coherence-before](#) E₂ and consequently E₅ is [GIC-hazard-ordered-before](#) E₂. At the same time, E₅ [Reads-from](#def:cpuif:ordering_model:rf) E₂ and as a result E₂ is [GIC-observed-by](#def:cpuif:ordering_model:GIC-obs) E₅. This creates a cycle in the Ordered-before relation and, as a result, the execution violates the External visibility requirement.

- E_2 is **Coherence-before** E_1 . In which case E_4 is **Coherence-before** E_1 and consequently E_3 is **GIC-hazard-ordered-before** E_1 . At the same time, E_3 [Reads-from](#def:cpuif:ordering_model:rf) E_1 and as a result E_1 is [GIC-observed-by](#def:cpuif:ordering_model:GIC-obs) E_3 . This creates a cycle in the Ordered-before relation and, as a result, the execution violates the External visibility requirement.

B1.16 Independent reads of independent writes

B1.16.1 Notes

The purpose of this test is to show that writes which can be observed by any observer must be observable by all observers.

In other words, the test shows that the interrupt ordering model has other multicopy atomicity.

B1.16.2 Litmus test

```
AArch64 IRIW-2loc+gic.cdpri
{
  [INTID(A)]=(priority:0);
  [INTID(B)]=(priority:0);
  0:X1=(intid:A, priority:1);
  1:X1=(intid:B, priority:1);
  2:X0=(intid:A); 2:X1=(intid:B);
  3:X0=(intid:A); 3:X1=(intid:B);
}
P0          | P1          | P2          | P3          ;
GIC CDPRI, X1 | GIC CDPRI, X1 | GIC CDRCFG, X0 | GIC CDRCFG, X1 ;
              |              | ISB           | ISB          ;
              |              | MRS X2, ICC_ICSR_EL1 | MRS X2, ICC_ICSR_EL1 ;
              |              | GSB SYS       | GSB SYS      ;
              |              | GIC CDRCFG, X1 | GIC CDRCFG, X0 ;
              |              | ISB           | ISB          ;
              |              | MRS X3, ICC_ICSR_EL1 | MRS X3, ICC_ICSR_EL1 ;

exists(2:X2=(priority:1) /\ 2:X3=(priority:0) /\
      3:X2=(priority:1) /\ 3:X3=(priority:0))

kinds: IRIW-2loc+gic.cdpri Forbid
```

B1.16.2.1 Explanation

In the execution that satisfies the postcondition:

- On P0, GIC CDPRI, X1 generates an Explicit Interrupt Write Effect to INTID (A) E₁.
- On P1, GIC CDPRI, X1 generates an Explicit Interrupt Write Effect to INTID (B) E₂.
- On P2:
 - GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect to INTID (A) E₃.
 - GSB SYS generates a GSB SYS Effect E₄.
 - GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID (B) E₅.
- On P3:
 - GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID (B) E₆.
 - GSB SYS generates a GSB SYS Effect E₇.
 - GIC CDRCFG, X0 generates an Explicit Interrupt Read Effect to INTID (A) E₈.

E₁ is **GIC-observed-by** E₃ as E₃ **Reads-from** E₁. E₃ is **GSB-ordered-before** E₅. E₅ is **GIC-observed-by** E₂ as E₅ is **Coherence-before** E₂. E₂ is **GIC-observed-by** E₆ as E₆ **Reads-from** E₂. E₆ is **GSB-ordered-before** E₈. E₈ is **GIC-observed-by** E₁ as E₈ is **Coherence-before** E₁. As a result, this execution creates a cycle in the Ordered-before relation and violates the External visibility requirement.

B1.17 Message passing via flag in memory

B1.17.1 Notes

The purpose of this test is to show the required barrier use when using a flag in memory to inform a different PE that an interrupt configuration has been updated.

B1.17.2 Litmus test

```
AArch64 MP+gsb.sys+dsb.ld
{
  [INTID(A)]=(affinity:P0);
  0:X1=x;
  0:X2=(intid:A, affinity:P1);
  0:X3=(intid:A);
  1:X1=x;
  1:X2=(intid:A);
}

P0      | P1
GIC CDAFF, X2 | LDR X0, [X1]
GSB SYS      | DSB LD
MOV X0, #1    | GIC CDRCFG, X2
STR X0, [X1]  | ISB
            | MRS X3, ICC_ICSR_EL1

exists(1:X0=1 /\ 1:X3=(affinity:P0))

kinds: MP+gsb.sys+dsb.ld Forbid
```

B1.17.2.1 Explanation

In the execution that satisfies the postcondition:

- On P0:
 - GIC CDAFF, X2 generates an Explicit Interrupt Write Effect to INTID(A) E_1 that sets its affinity to P1.
 - GSB SYS generates a GSB SYS Effect E_2 .
 - STR X0, [X1] generates an Explicit Write Memory Effect to x E_3 .
- On P1:
 - LDR X0, [X1] generates an Explicit Memory Read Effect to x E_4 .
 - DSB LD generates a DSB LD Effect E_5 .
 - GIC RCFG, X2 generates an Explicit Interrupt Effect E_6 .

In this execution:

- E_4 Reads-from-memory E_3 and as a result E_3 is Explicit-observed-by E_4 .
- E_4 is DSB-ordered-before E_6 .
- E_6 is [Coherence-before](#) E_1 .
- E_1 is [GSB-ordered-before](#) E_4 .

This creates a cycle in the Ordered-before relation and, as a result, the execution violates the External visibility requirement.

B1.18 Message passing via interrupt priority configuration

B1.18.1 Notes

The purpose of this test is to show that writes to an interrupt configuration can be ordered with respect to an update of the interrupt Pending state such that observing the interrupt as Pending guarantees that the writes to the interrupt configuration are observed.

B1.18.2 Litmus test

```
AArch64 MP-pri+gsb.sys+ack
{
  [INTID(A)]=(enabled:1,affinity:P1,priority:1,pending:0);
  [PTE(x)]=(oa:PA(x),attrs:(device-nGRE)); // PERIP
  0:X1=x;
  0:X2=(intid:A, priority:5);
  1:X1=(intid:A);
}

P0          | P1
GIC CDPRI, X2 | GICR X0, CDIA
GSB SYS      | GIC CDRCFG, X1
STR X0, [X1]  | ISB
              | MRS X1, ICC_ICSR_EL1 ;

exists(1:X0=(valid:1, intid:A) /\
      1:X1=(enabled:1,affinity:P1,priority:1,active:1))

kinds: MP-pri+gsb.sys+ack Forbid
```

B1.18.2.1 Explanation

In the execution that satisfies the postcondition:

- On P0:
 - GIC CDPRI, X2 generates an Explicit Interrupt Write Effect to INTID(A) E_1 that sets its priority to 5.
 - GSB SYS generates a GSB SYS Effect E_2 .
 - STR X0, [X1] generates an Explicit Write Memory Effect to x . When the Explicit Write Effect to Location x reaches its endpoint, the Peripheral device signals an interrupt event to the IRI which generates an Interrupt Read Effect E_3 to INTID(A) and an Interrupt Write Effect to INTID(A) E_4 which sets the pending state to 1.
- On P1:
 - GICR X0, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E_5 and an Implicit Interrupt Write Effect to INTID(A) E_6 that resets the pending state to 0 and sets the active state to 1.
 - GIC CDRCFG, X2 generates an Explicit Interrupt Read Effect E_7 .

In this execution:

- E_7 [Reads-from](#) E_6 and as a result E_6 and E_5 are [Coherence-before](#) E_1 .
- E_1 is [GSB-ordered-before](#) E_4 .
- E_5 [Reads-from](#) E_4 and as a result, E_4 is [GIC-observed-by](#) E_5

This creates a cycle in the Ordered-before relation and, as a result, the execution violates the External visibility requirement.

B1.19 Message passing with an LPI and a device read

B1.19.1 Notes

The purpose of this test is to show that if a write to a peripheral has the side effect of making an interrupt pending, then observing that pending state on a different PE means that the write to the peripheral can also be observed.

B1.19.2 Litmus test

```
AArch64  MP-fLPI+imp+gsb.ack
{
  [INTID(A)]=(enabled:1,affinity:P1,priority:1,pending:0);
  [PTE(x)]=(oa:PA(x),attrs:(device-nGRE)); // PERIP
  0:X3=x;
  1:X1=x;
}
P0          | P1          ;
MOV  W2, #1  | GICR X0, CDIA  ;
STR  W2, [X3] | GSB  ACK       ;
          | LDR  W2, [X1]  ;

exists(1:X0=(valid:1, intid:A) /\ 1:X2=0);

kinds: MP-fLPI+imp+gsb.ack Forbid
```

B1.19.2.1 Explanation

In the execution that satisfies the postcondition:

- On P0:
 - STR W2, [X3] generates an Explicit Write Memory Effect to x E_1 . When the Explicit Write Effect to Location x reaches its endpoint, the Peripheral device signals an interrupt event to the IRI which generates Interrupt Read Effect E_2 to $\text{INTID}(A)$ and an Interrupt Write Effect to $\text{INTID}(A)$ E_3 which sets the pending state to 1.
- On P1:
 - GICR X0, CDIA generates an Implicit Interrupt Read Effect to $\text{INTID}(A)$ E_4 and an Implicit Interrupt Write Effect to $\text{INTID}(A)$ E_5 that resets the pending state to 0 and sets the active state to 1.
 - GSB SYS generates a GSB SYS Effect E_6 .
 - LDR W2, [X1] generates an Explicit Memory Read Effect E_7 .

In this execution:

- E_4 **Reads-from** E_3 and as a result E_3 is **GIC-observed-by** E_4
- E_4 is **GSB-ordered-before** E_7 .
- E_7 is Coherence-before E_1 and as a result E_7 is Explicit-observed-by E_1 .
- E_1 is Ordered-before E_3 .

This creates a cycle in the Ordered-before relation and, as a result, the execution violates the External visibility requirement.

B1.19.3 Litmus test with address dependency

The GSB ACK can be omitted in favor of creating an address dependency on the result of the acknowledge.

This may provide performance benefits in some cases, however, an address dependency cannot always be used a replacement for the GSB ACK. For example, a store following LDR X2, [X1, X5] without any dependency on previous instructions can be re-ordered before the GICR X0, CDIA and the effects of the store can be observed before the effects of the GICR X0, CDIA on another PE.

```
AArch64 MP-fLPI+imp+addrpt
{
  [INTID(A)]=(enabled:1,affinity:P1,priority:1,pending:0);
  [PTE(x)]=(oa:PA(x),attrs:(device-nGRE)); // PERIP
  0:X1=x;
  1:X1=x;
}
P0          | P1          ;
MOV W2, #1   | GICR X0, CDIA   ;
STR W2, [X1] | EOR X5, X0, X0  ;
              | LDR W2, [X1, X5] ;

exists(1:X0=(valid:1, intid:A) /\ 1:X2=0);

kinds: MP-fLPI+imp+addrpt Forbid
```

B1.19.3.1 Explanation

In the execution that satisfies the postcondition:

- On P0:
 - STR W2, [X1] generates an Explicit Write Memory Effect to x E_1 . When the Explicit Write Effect to Location x reaches its endpoint, the Peripheral device signals an interrupt event to the IRI which generates an Interrupt Read Effect E_2 to INTID (A) and an Interrupt Write Effect to INTID (A) E_3 which sets the pending state to 1.
- On P1:
 - GICR X0, CDIA generates an Implicit Interrupt Read Effect to INTID (A) E_4 and an Implicit Interrupt Write Effect to INTID (A) E_5 that resets the pending state to 0 and sets the active state to 1.
 - LDR W2, [X1, X5] generates an Explicit Memory Read Effect E_7 .

In this execution:

- E_4 Reads-from E_3 and as a result E_3 is GIC-observed-by E_4
- There is an address dependency from E_4 to E_7 and, as a result, E_4 is Dependency-ordered-before E_7 .
- E_7 is Coherence-before E_1 and as a result E_7 is Explicit-observed-by E_1 .
- E_1 is Ordered-before E_3 .

This creates a cycle in the Ordered-before relation and, as a result, the execution violates the External visibility requirement.

B1.20 Message passing with an LPI and a GSB

B1.20.1 Notes

The purpose of this test is to show that when a write to memory is ordered before a write which has the side effect of making an interrupt pending, then observing the interrupt being pending through acknowledgment on a different PE guarantees that the write is also observed.

If the peripheral generating an interrupt exists in the shared address space of P0 and P1, a the `DSB ST` on P0 can be relaxed to a `DMB ST`, but because the test makes no such assumptions, we must use a `DSB ST`.

B1.20.2 Litmus test

```
AArch64 MP-fSPI+dsb.st+gsb.ack
{
  [INTID(A)]=(enabled:1,affinity:P1,priority:1,pending:0);
  [PTE(y)]=(oa:PA(y),attrs:(device-nGRE)); // PERIP
  0:X1=x; 0:X3=y;
  1:X1=x;
}
P0          | P1          ;
MOV  W0, #1  | GICR X0, CDIA ;
STR  W0, [X1] | GSB  ACK      ;
DSB  ST      | LDR  X2, [X1]  ;
MOV  W2, #1  |              ;
STR  W2, [X3] |              ;

exists(1:X0=(valid:1, intid:A) /\ 1:X2=0);

kinds: MP-fSPI+dsb.st+gsb.ack Forbid
```

B1.20.2.1 Explanation

In the execution that satisfies the postcondition:

- On P0:
 - `STR W0, [X1]` generates an Explicit Memory Write Effect to `INTID(A)` E_1 .
 - `DSB ST` generates a `DSB ST` Effect E_2 .
 - `STR W2, [X3]` generates an Explicit Write Memory Effect to x . When the Explicit Write Effect to Location x reaches its endpoint, the Peripheral device signals an interrupt event to the IRI which generates an Interrupt Read Effect E_3 to `INTID(A)` and an Interrupt Write Effect to `INTID(A)` E_4 which sets the pending state to 1.
- On P1:
 - `GICR X0, CDIA` generates an Implicit Interrupt Read Effect to `INTID(A)` E_5 and an Implicit Interrupt Write Effect to `INTID(A)` E_6 that resets the pending state to 0 and sets the active state to 1.
 - `GSB ACK` generates a `GSB ACK` Effect E_7 .
 - `LDR X2, [X1]` generates an Explicit Memory Read Effect E_8 .

In this execution:

- E_5 [Reads-from](#) E_4 and, as a result, E_4 is [GIC-observed-by](#) E_5 .
- E_5 is [GSB-ordered-before](#) E_8 .
- E_8 is Coherence-before E_1 and, as a result, E_8 is Explicit-observed-by E_1 .

- E_1 is DSB-ordered-before E_4 .

This creates a cycle in the Ordered-before relation and, as a result, the execution violates the External visibility requirement.

B1.21 Message passing with an IPI and a GSB

B1.21.1 Notes

The purpose of this test is to show that a write to memory ordered before sending an IPI guarantees that the memory write is observed when the IPI is observed on another PE.

This is illustrating a commonly used pattern for cross-CPU communication by software.

B1.21.2 Litmus test

```
AArch64 MP-fIPI+dsb.st+gsb.ack
{
  [INTID(A)]=(enabled:1, pending:0, affinity:P1);
  0:X1=x; 0:X2=(intid:A, pending:1);
  1:X1=x;
}
P0          | P1          ;
MOV W0, #1   | GICR X0, CDIA   ;
STR W0, [X1] | GSB ACK         ;
DSB ST       | LDR W2, [X1]    ;
GIC CDPEND, X2 |              ;

exists(1:X0=(valid:1, intid:A) /\ 1:X2=0)

kinds: MP-fIPI+dsb.st+gsb.ack Forbid
```

B1.21.2.1 Explanation

In the execution that satisfies the postcondition:

- On P0:
 - STR W0, [X1] generates an Explicit Memory Write Effect to INTID(A) E₁.
 - DSB ST generates a DSB ST Effect E₂.
 - GIC CDPEND, X2 generates an Explicit Interrupt Read Effect E₃ to INTID(A) and an Explicit Interrupt Write Effect to INTID(A) E₄ which sets the pending state to 1.
- On P1:
 - GICR X0, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₅ and an Implicit Interrupt Write Effect to INTID(A) E₆ that resets the pending state to 0 and sets the active state to 1.
 - GSB ACK generates a GSB ACK Effect E₇.
 - LDR X2, [X1] generates an Explicit Memory Read Effect E₈.

In this execution:

- E₅ [Reads-from](#) E₄ and, as a result, E₄ is [GIC-observed-by](#) E₅.
- E₅ is [GSB-ordered-before](#) E₈.
- E₈ is Coherence-before E₁ and, as a result, E₈ is Explicit-observed-by E₁.
- E₁ is DSB-ordered-before E₄.

This creates a cycle in the Ordered-before relation and, as a result, the execution violates the External visibility requirement.

B1.22 Message passing using deactivate

B1.22.1 Notes

The purpose of these tests is to show that without explicit synchronization, or a combination of explicit synchronization and an address dependency, there is no ordering between writing to memory and deactivating an interrupt, or between reading interrupt configuration and reading from memory.

This test illustrates part of a real-world example in the following way. P0 illustrates the end of an interrupt handler sequence, which performs memory accesses as part of the handler logic and then finishes the handler by deactivating the interrupt. P1 reads the interrupt configuration and state and determines that the interrupt is no longer actively being handled, and therefore P1 expects to observe all memory accesses performed as part of the handler on P0.

B1.22.2 Litmus test with explicit synchronization

```
AArch64 MP-fDI+W-dsb-W+R-gsb-R
{
  [INTID(A)]=(active:1);
  0:X0=x; 0:X1=(intid:A);
  1:X0=x; 1:X1=(intid:A);
}

P0                                | P1                                ;
MOV      X3, #1                   | GIC CDRCFG, X1                   ;
STR      X3, [X0]                 | ISB                              ;
DSB      ST                       | MRS X2, ICC_ICSR_EL1            ;
GIC      CDDI, X1                 | GSB SYS                          ;
                                | LDR X3, [X0]                    ;

exists(1:X2=(active:0) /\ 1:X3=0)

kinds: MP-fDI+W-dsb-W+R-gsb-R Forbid
```

B1.22.2.1 Explanation

In the execution that satisfies the postcondition:

- On P0:
 - STR X3, [X0] generates an Explicit Memory Write Effect to INTID(A) E₁.
 - DSB ST generates a DSB ST Effect E₂.
 - GIC CDDI, X1 generates an Implicit Interrupt Read Effect E₃ to INTID(A) and an Implicit Interrupt Write Effect to INTID(A) E₄ which resets the active state to 0.
- On P1:
 - GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₅.
 - GSB SYS generates a GSB ACK Effect E₆.
 - LDR X3, [X0] generates an Explicit Memory Read Effect E₇.

In this execution:

- E₅ [Reads-from](#) E₄ and, as a result, E₄ is [GIC-observed-by](#) E₅.
- E₅ is [GSB-ordered-before](#) E₇.
- E₇ is Coherence-before E₁ and, as a result, E₇ is Explicit-observed-by E₁.
- E₁ is DSB-ordered-before E₄.

This creates a cycle in the Ordered-before relation and, as a result, the execution violates the External visibility requirement.

B1.22.3 Litmus test with address dependency

```
AArch64 MP-fDI+W-dsb-W+R-addrpt-R
{
  [INTID(A)]=(active:1);
  0:X0=x; 0:X1=(intid:A);
  1:X0=x; 1:X1=(intid:A);
}

P0          | P1
MOV    X3, #1      | GIC CDRCFG, X1      ;
STR    X3, [X0]    | ISB                  ;
DSB    ST          | MRS X2, ICC_ICSR_EL1 ;
GIC    CDDI, X1    | EOR X5, X2, X2      ;
          | LDR X3, [X0, X5]    ;

exists(1:X2=(active:0) /\ 1:X3=0)

kinds: MP-fDI+W-dsb-W+R-addrpt-R Forbid
```

B1.22.3.1 Explanation

In the execution that satisfies the postcondition:

- On P0:
 - STR X3, [X0] generates an Explicit Memory Write Effect to INTID(A) E₁.
 - DSB ST generates a DSB ST Effect E₂.
 - GIC CDDI, X1 generates an Implicit Interrupt Read Effect E₃ to INTID(A) and an Implicit Interrupt Write Effect to INTID(A) E₄ which resets the active state to 0.
- On P1:
 - GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₅.
 - LDR X3, [X0, X5] generates an Explicit Memory Read Effect E₇.

In this execution:

- E₅ Reads-from E₄ and, as a result, E₄ is GIC-observed-by E₅.
- There is an address dependency from E₅ to E₇ and, as a result, E₅ is Dependency-ordered-before E₇.
- E₇ is Coherence-before E₁ and, as a result, E₇ is Explicit-observed-by E₁.
- E₁ is DSB-ordered-before E₄.

This creates a cycle in the Ordered-before relation and, as a result, the execution violates the External visibility requirement.

B1.22.4 Litmus test with control dependency

```
AArch64 MP-fDI+W-dsb-W+R-ctrlisb-R
{
  [INTID(A)]=(active:1);
  0:X0=x; 0:X1=(intid:A);
  1:X0=x; 1:X1=(intid:A); 1:X5=(active:1);
}

P0          | P1
MOV    X3, #1      | GIC CDRCFG, X1      ;
```

```

STR      X3, [X0]      | ISB                               ;
DSB      ST             | MRS X2, ICC_ICSR_EL1      ;
GIC      CDDI, X1       | TST X2, X5                ;
                        | B.NE L0                   ;
                        | ISB                       ;
                        | LDR X3, [X0]              ;
                        | L0:                       ;
                        | NOP                       ;

```

```
exists(1:X2=(active:0) /\ 1:X3=0)
```

```
kinds: MP-fDI+W-dsb-W+R-ctrlisb-R Forbid
```

B1.22.4.1 Explanation

In the execution that satisfies the postcondition:

- On P0:
 - STR X3, [X0] generates an Explicit Memory Write Effect to `INTID(A)` E_1 .
 - DSB ST generates a DSB ST Effect E_2 .
 - GIC CDDI, X1 generates an Implicit Interrupt Read Effect E_3 to `INTID(A)` and an Implicit Interrupt Write Effect to `INTID(A)` E_4 which resets the active state to 0.
- On P1:
 - GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to `INTID(A)` E_5 .
 - ISB (the second instance) generates an Instruction Fetch Barrier Effect E_6 .
 - LDR X3, [X0] generates an Explicit Memory Read Effect E_7 .

In this execution:

- E_5 Reads-from E_4 and, as a result, E_4 is GIC-observed-by E_5 .
- There is a control dependency from E_5 to E_6 and as a result, because E_6 is an Instruction Fetch Barrier Effect, E_5 is IFB-ordered-before E_7 .
- E_7 is Coherence-before E_1 and, as a result, E_7 is Explicit-observed-by E_1 .
- E_1 is DSB-ordered-before E_4 .

This creates a cycle in the Ordered-before relation and, as a result, the execution violates the External visibility requirement.

B1.22.5 Litmus test without explicit synchronization

```

AArch64 MP-fDI+RR
{
  [INTID(A)]=(active:1);
  0:X0=x; 0:X1=(intid:A);
  1:X0=x; 1:X1=(intid:A);
}
P0      | P1
MOV      X3, #1      | GIC CDRCFG, X1          ;
STR      X3, [X0]     | ISB                     ;
GIC      CDDI, X1     | MRS X2, ICC_ICSR_EL1   ;
                        | LDR X3, [X0]           ;

exists(1:X2=(active:0) /\ 1:X3=0)

```

kinds: MP-fDI+RR Allow

B1.22.5.1 Explanation

In the execution that satisfies the postcondition:

- On P0:
 - STR X3, [X0] generates an Explicit Memory Write Effect to INTID(A) E₁.
 - GIC CDDI, X1 generates an Implicit Interrupt Read Effect E₃ to INTID(A) and an Implicit Interrupt Write Effect to INTID(A) E₄ which resets the active state to 0.
- On P1:
 - GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₅.
 - LDR X3, [X0] generates an Explicit Memory Read Effect E₇.

In this execution:

- E₅ Reads-from E₄ and, as a result, E₄ is GIC-observed-by E₅.
- E₇ is Coherence-before E₁ and, as a result, E₇ is Explicit-observed-by E₁.

This means that there is no cycle in the Ordered-before relation and, as a result, this execution satisfies the External visibility requirement. This execution is architecturally allowed.

B1.22.6 Litmus test with a DSB but without a GSB

AArch64 MP-fDI+W-dsb-W+RR

```
{
  [INTID(A)]=(active:1);
  0:X0=x; 0:X1=(intid:A);
  1:X0=x; 1:X1=(intid:A);
}

P0          | P1          ;
MOV    X3, #1      | GIC CDRCFG, X1      ;
STR    X3, [X0]    | ISB                  ;
DSB    ST          | MRS X2, ICC_ICSR_EL1 ;
GIC    CDDI, X1    |                      ;
                | LDR X3, [X0]        ;
```

exists(1:X2=(active:0) /\ 1:X3=0)

kinds: MP-fDI+W-dsb-W+RR Allow

B1.22.6.1 Explanation

In the execution that satisfies the postcondition:

- On P0:
 - STR X3, [X0] generates an Explicit Memory Write Effect to INTID(A) E₁.
 - DSB ST generates a DSB ST Effect E₂.
 - GIC CDDI, X1 generates an Implicit Interrupt Read Effect E₃ to INTID(A) and an Implicit Interrupt Write Effect to INTID(A) E₄ which resets the active state to 0.
- On P1:
 - GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₅.
 - LDR X3, [X0] generates an Explicit Memory Read Effect E₇.

In this execution:

- E₅ **Reads-from** E₄ and, as a result, E₄ is **GIC-observed-by** E₅.
- E₇ is Coherence-before E₁ and, as a result, E₇ is Explicit-observed-by E₁.
- E₁ is DSB-ordered-before E₄.

This means that there is no cycle in the Ordered-before relation and, as a result, this execution satisfies the External visibility requirement. This execution is architecturally allowed.

B1.22.7 Litmus test without a DSB but with a GSB

```
AArch64 MP-fDI+WW+R-gsb-R
```

```
{
  [INTID(A)] = (active:1);
  0:X0=x; 0:X1=(intid:A);
  1:X0=x; 1:X1=(intid:A);
}

P0          | P1
MOV    X3, #1      | GIC CDRCFG, X1      ;
STR     X3, [X0]    | ISB                  ;
GIC     CDDI, X1    | MRS X2, ICC_ICSR_EL1 ;
          | GSB SYS              ;
          | LDR X3, [X0]          ;
```

```
exists(1:X2=(active:0) /\ 1:X3=0)
```

```
kinds: MP-fDI+WW+R-gsb-R Allow
```

B1.22.7.1 Explanation

In the execution that satisfies the postcondition:

- On P0:
 - STR X3, [X0] generates an Explicit Memory Write Effect to INTID(A) E₁.
 - GIC CDDI, X1 generates an Implicit Interrupt Read Effect E₃ to INTID(A) and an Implicit Interrupt Write Effect to INTID(A) E₄ which resets the active state to 0.
- On P1:
 - GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₅.
 - GSB SYS generates a GSB SYS Effect E₆.
 - LDR X3, [X0] generates an Explicit Memory Read Effect E₇.

In this execution:

- E₅ **Reads-from** E₄ and, as a result, E₄ is **GIC-observed-by** E₅.
- E₅ is **GSB-ordered-before** E₇.
- E₇ is Coherence-before E₁ and, as a result, E₇ is Explicit-observed-by E₁.

This means that there is no cycle in the Ordered-before relation and, as a result, this execution satisfies the External visibility requirement. This execution is architecturally allowed.

B1.23 IPI edge merging and message passing

B1.23.1 Notes

The purpose of this test is to show that it is possible to inform a different PE that a GIC operation has completed by using a GSB SYS and a flag.

B1.23.2 Litmus test

```
AArch64 edge-merging.ipi.MP
{
  [INTID(A)]=(enabled:1, pending:1, affinity:P1);
  0:X0=(intid:A, pending:1);
  0:X1=x;
  1:X1=x;
  1:X3=(intid:A);
}
P0                                | P1                                ;
GIC    CDPEND, X0                | LDR  W0, [X1]                    ;
GSB     SYS                      | DSB  LD                          ;
MOV     X2, #1                   | GICR X2, CDIA                    ;
STR     W2, [X1]                 | GIC CDRCFG, X3                   ;
                                | ISB                              ;
                                | MRS  X4, ICC_ICSR_EL1           ;

exists(1:X0=1 /\ 1:X2=(valid:1, intid:A) /\
      1:X4=(pending:1, active:1, affinity:P1))

kinds: edge-merging.ipi.MP Forbid
```

B1.23.2.1 Explanation

In the execution that satisfies the postcondition:

- On P0:
 - GIC CDPEND, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₁ and an Explicit Interrupt Write Effect to INTID(A) E₂ that sets the pending state to 1.
 - GSB SYS generates a GSB ACK Effect E₃.
 - STR W2, [X1] generates an Explicit Memory Write Effect E₄.
- On P1:
 - LDR W0, [X1] generates an Explicit Memory Read Effect to INTID(A) E₅.
 - DSB LD generates a DSB LD Effect E₆.
 - GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₇ and an Implicit Interrupt Write Effect to INTID(A) E₈ that resets the pending state to 0 and sets the active state to 1.
 - GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect E₉ to INTID(A).

In the execution that satisfies the postcondition:

- E₉ Reads-from E₂ and, as a result, E₇ is Coherence-before E₂ and E₇ is GIC-observed-by E₂.
- E₂ GSB-ordered-before E₄
- E₅ Reads-from-memory E₄ and, as a result, E₄ is Explicit-observed-by E₅.
- E₅ is DSB-ordered-before E₇.

This creates a cycle in the Ordered-before relation and, as a result, the execution violates the External visibility requirement.

B1.24 Device edge merging with GSB ACK

B1.24.1 Notes

The purpose of this test is to show that two separate instances of an LPI being Pending are not allowed to be merged across an acknowledgement.

The LPI is initially Pending, it is then successfully acknowledged by GICR X1, CDIA which writes Pending = 0.

The STR has an implicit effect of writing Pending = 1, which happens in finite time, and the GIC will forward the Pending interrupt in finite time so that the second GICR X2, CDIA will return valid == 1 if not merged with the first instance.

Note that this test does not depend on the interrupt becoming Pending or being forwarded to the PE as a direct result of the write to the peripheral. The test relies solely on the finite time guarantees which are covered by the WAIT construct.

See also:

- [B1.1 Interrupt litmus test assumptions](#)

B1.24.2 Litmus test

```
AArch64 edge-merging.LPI+gsb.ack
{
  [INTID(A)]=(enabled:1,pending:1,affinity:P0);
  [PTE(x)]=(oa:PA(x),attrs:(device-nGnRnE)); // PERIP
  0:X0=x;
}
P0
    GICR  X1, CDIA                ;
    GSB   ACK                      ;
    MOV   X3, #1                   ;
    STR   X3, [X0]                 ; trigger intid A via device
    GIC   CDDI, X1                 ;
    WAIT  (GICR  X2, CDIA; X2=(valid:1)) ;

exists(0:X1=(valid:1) /\ 0:X2=(valid:0))

kinds: edge-merging.LPI+gsb.ack Forbid
```

B1.24.2.1 Explanation

In the execution that satisfies the postcondition, on P0:

- GICR X1, CDIA generates an Implicit Interrupt Write Effect to INTID(A) E₁ and an Implicit Interrupt Write Effect to INTID(A) E₂ which resets the pending state to 0 and sets the active state to 1.
- GSB ACK generates a GSB ACK Effect E₃.
- STR X3, [X0] generates an Explicit Write Effect E₄ to Location x. When the Explicit Write Effect to Location x reaches its endpoint, the Peripheral device signals an interrupt event to the IRI which generates Interrupt Read Effect E₄ to INTID(A) and an Interrupt Write Effect to INTID(A) E₅ which sets the pending state to 1.
- GIC CDDI, X1 generates Implicit Interrupt Write Effect to INTID(A) E₆ and an Implicit Interrupt Write Effect to INTID(A) E₇ which resets the pending state to 0.
- The repeated execution of GICR X2, CDIA generates Implicit Interrupt Read Effects E₈, E₉, ..., E_n where n is a finite number.

All Interrupt Write Effects are required to complete in finite time and consequently, Interrupt Write Effects E_2 , E_5 and E_7 are **Coherence-before** E_n . In the execution that satisfies the postcondition, no interrupt is acknowledged and therefore $\text{INTID}(A)$ does not become the HPPI. As a result one of the following has to be true:

- E_5 is **Coherence-before** E_2 and consequently, E_1 **Reads-from** $E_{\sim 5}$ and E_5 is **GIC-observed-by** E_1 . E_1 is **GSB-ordered-before** E_5 . As a result, this execution violates the external visibility requirement as it creates a cycle in the Ordered-before relation.
- E_7 is **Coherence-before** E_2 and consequently, E_1 **Reads-from** $E_{\sim 7}$ and E_7 is **GIC-observed-by** E_1 . E_1 is **GSB-ordered-before** E_7 . As a result, this execution violates the external visibility requirement as it creates a cycle in the Ordered-before relation.

B1.25 Configuration read and interrupt acknowledge

B1.25.1 Notes

The purpose of these tests is to show that if an interrupt is disabled, and that disable is observed by a configuration read, then the interrupt cannot be acknowledged after observing the disable.

We show both a single-threaded and multi-threaded version of the test.

Illustrates a common principle of an interrupt acknowledge observing side effects of a write if that write can be observed via an explicit read.

B1.25.2 Single-threaded litmus test

```
AArch64 coRR.ack
{
  [INTID(A)]=(enabled:1,affinity:P0,pending:1);
  0:X1=(intid:A);
}
P0
GIC CDDIS, X1
GIC CDRCFG, X1
ISB
MRS X2, ICC_ICSR_EL1
GICR X0, CDIA

exists(0:X0=(valid:1, intid:A))

kinds: coRR.ack Forbid
```

B1.25.2.1 Explanation

On P0:

- GIC CDDIS, X1 generates an Explicit Interrupt Write Effect to INTID(A) E₁ and an Explicit Interrupt Write Effect to INTID(A) E₂ which sets the enabled state to 0
- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₃.
- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₄ and an Implicit Interrupt Write Effect to INTID(A) E₅.

In the execution that satisfies the postcondition:

- E₄ [Reads-from](#) the initial value of INTID(A) and, as a result, due to the [atomicity requirement](#) for successful pairs of Read Interrupt and Write Interrupt Effects, E₅ is [Coherence-before](#) E₂.
- Any of the following applies:
 - E₃ [Reads-from](#) the initial state of INTID and, as a result, E₃ is [Coherence-before](#) E₂. This violates the [requirements of the Coherence-before](#) relation. As a result, this execution is architecturally forbidden.
 - E₃ [Reads-from](#) E₅. This violates the [Reads-from](#) relation. As a result, this execution is architecturally forbidden.

B1.25.3 Multi-threaded litmus test

```
AArch64 coRR.ack+dis+Rack
{
  [INTID(A)]=(enabled:1,affinity:P1,pending:1);
  0:X1=(intid:A);
  1:X1=(intid:A);
}
```

```

}
P0                                | P1                                ;
GIC CDDIS, X1                    | GIC CDRCFG, X1                    ;
                                | ISB                               ;
                                | MRS X2, ICC_ICSR_EL1             ;
                                | GICR X0, CDIA                    ;

exists(1:X2=(enabled:0, affinity:P1, pending:1) /\
      1:X0=(valid:1, intid:A))

kinds: coRR.ack+dis+Rack Forbid

```

B1.25.3.1 Explanation

- GIC CDDIS, X1 generates an Explicit Interrupt Write Effect to INTID(A) E₁ and an Explicit Interrupt Write Effect to INTID(A) E₂ which sets the enabled state to 0

On P1:

- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₃.
- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₄ and an Implicit Interrupt Write Effect to INTID(A) E₅.

In the execution that satisfies the postcondition:

- E₃ is in program order before E₄.
- E₃ and E₄ are to the same Interrupt Location INTID(A).
- E₄ [Reads-from](#) the initial state of INTID(A) and as a result E₄ is [Coherence-before](#) E₂. This means that E₃ is [GIC-Hazard-Ordered-before](#) E₂
- E₃ [Reads-from](#) E₂ and as a result, E₂ is [GIC-Observed-by](#) E₃.

This creates a cycle in the Ordered-before relation and as a result, the execution violates the External visibility requirement.

B1.26 Atomicity of interrupt acknowledge

B1.26.1 Notes

The purpose of this test is to show that the two interrupt write effects of acknowledging an interrupt are atomic with respect to any read effect to the interrupt.

B1.26.2 Litmus test

```
AArch64 atomic-gic.ack
{
  [INTID(A)]=(enabled:1,affinity:P0,active:0,pending:1,handling_mode:edge);
  1:X0=(intid:A)
}
P0          | P1
GICR X0, CDIA | GIC CDRCFG, X0
              | ISB
              | MRS X1, ICC_ICSR_EL1 ;

exists(1:X1=(pending:1, active:1) \ /
      1:X1=(pending:0, active:0))

kinds: atomic-gic.ack Forbid
```

B1.26.2.1 Explanation

On P0:

- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E_1 and an Implicit Interrupt Write Effect to INTID(A) E_2 and atomically sets the pending state to 0 and the active state to 1.

On P1:

- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E_3 .

In the execution that satisfies the postcondition, E_3 reads that either the pending state is 1 and the active state is 1, or the pending state is 0 and the active state is 0. Both postconditions violate the atomicity requirement of E_2 .

B1.27 Atomicity of interrupt disable and acknowledge

B1.27.1 Notes

This test demonstrates the interaction between disabling (clearing Enabled) and acknowledging an INTID from different PEs.

Illustrates a common principle of interrupt acknowledgement observing side effects of a write if that write can be observed via an explicit read.

B1.27.2 Litmus test

```
AArch64 atomic.ack.dis.gic+W+R
{
  [INTID(A)]=(enabled:1,affinity:P0,pending:1,active:0);
  1:X1=(intid:A); 2:X1=(intid:A);
}
P0          | P1          | P2          ;
GICR X2, CDIA | GIC CDDIS, X1 | GIC CDRCFG, X1 ;
              |              | ISB          ;
              |              | MRS X2, ICC_ICSR_EL1 ;

exists(0:X2=(valid:1,intid:A) /\
      2:X2=(enabled:0, pending:1, active:0))

kinds: atomic.ack.dis.gic+W+R Forbid
```

B1.27.2.1 Explanation

On P0:

- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₁ and an Implicit Interrupt Write Effect to INTID(A) E₂ and atomically sets the pending state to 0 and the active state to 1.

On P1:

- GIC CDDIS, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₃ and an Explicit Interrupt Write Effect to INTID(A) E₄ and atomically sets the enabled state to 0.

On P2:

- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₅.

In the execution that satisfies the postcondition, E₅ Reads-from E₄ and E₄ is GIC-observed-by E₅. The postcondition also requires that E₃ Reads-from the initial state of INTID(A). At the same time, E₁ Reads-from the initial state of INTID(A). This violates the [atomicity requirement](#) for successful pairs of Read Interrupt and Write Interrupt Effects which should apply for the pairs of E₃, E_{~4} and E₁, E₂.

B1.28 Interrupt handler completion

B1.28.1 Notes

The purpose of this test is to show that when an interrupt can be acknowledged on a PE, then disabling the interrupt on a different PE followed by reading back the interrupt configuration and state, exactly one of the following is true:

- The interrupt is observed as Active and P0 has acknowledged the interrupt.
- The interrupt is not observed as Active and P0 will not acknowledge the interrupt.

This illustrates a concept that operating systems would rely on, which is that if they've disabled the interrupt and observe an interrupt is not active, then that interrupt cannot be taken and become active without the interrupt being enabled again.

B1.28.2 Litmus test

```
AArch64 atomic.ack.gic-ib+W-R-si-R
{
  [INTID(A)]=(enabled:1,affinity:P0,pending:1,active:0);
  1:X1=(intid:A);
}
P0                                | P1                                ;
GICR X2, CDIA                    | GIC CDDIS, X1                      ;
                                | GIC CDRCFG, X1                     ;
                                | ISB                               ;
                                | MRS X2, ICC_ICSR_EL1              ;

exists(0:X2=(valid:1, intid:A) /\
      1:X2=(enabled:0, affinity:P0, pending:1, active:0))

kinds: atomic.ack.gic-ib+W-R-si-R Forbid
```

B1.28.2.1 Explanation

On P0:

- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₁ and an Implicit Interrupt Write Effect to INTID(A) E₂ and atomically sets the pending state to 0 and the active state to 1.

On P1:

- GIC CDDIS, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₃ and an Explicit Interrupt Write Effect to INTID(A) E₄ and atomically sets the enabled state to 0.
- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₅.

In the execution that satisfies the postcondition:

- E₁ [Reads-from](#) the initial value of INTID(A) and, as a result, E₁ is [Coherence-before](#) E₄.
- E₅~ [Reads-from](#) E₃.
- E₅ is [Coherence-before](#) E₂ and, as a result, E₄ is also [Coherence-before](#) E₂.

This execution violates the [atomicity requirement](#) for the successful pair of Read Interrupt Effect E₁ and Write Interrupt Effect E₂, as E₁ is [Coherence-before](#) E₄ and E₄ is [Coherence-before](#) E₂.

B1.29 Configuration update while disabled

B1.29.1 Notes

The purpose of this test is to show that if multiple writes occur to the same interrupt while the interrupt is disabled, the interrupt will not be acknowledged and observe a partially updated interrupt.

B1.29.2 Litmus test

```
AArch64 coWWW-gic.dis-en+ack+ack
{
  [INTID(A)]=(enabled:1, affinity:P1, active:0, pending:1, priority:5);
  0:X0=(intid:A);
  0:X1=(intid:A, priority:4);
  0:X2=(intid:A, affinity:P2);

  1:X1=(intid:A);
}
P0          | P1          | P2          |
GIC CDDIS, X0 | GICR X2, CDIA | GICR X2, CDIA |
GIC CDPRI, X1 | ISB          | ISB          |
GIC CDAFF, X2 | MRS X0, ICC_HAPR_EL1 | MRS X0, ICC_HAPR_EL1 |
GIC CDEN, X0  |              |              |

exists((1:X2=(valid:1, intid:A) /\ 1:X0=(priority:4)) /\
      (2:X2=(valid:1, intid:A) /\ 2:X0=(priority:5)))

kinds: coWWW-gic.dis-en+ack+ack Forbid
```

B1.29.2.1 Explanation

On P0:

- GIC CDDIS, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₁ and an Explicit Interrupt Write Effect to INTID(A) E₂ and sets its *enabled* state to 0.
- GIC CDPRI, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₃ and an Explicit Interrupt Write Effect to INTID(A) E₄ and sets its *priority* to 4.
- GIC CDAFF, X2 generates an Explicit Interrupt Read Effect to INTID(A) E₅ and an Explicit Interrupt Write Effect to INTID(A) E₆ and sets its *affinity* to P2.
- GIC CDEN, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₇ and an Explicit Interrupt Write Effect to INTID(A) E₈ and sets its *enabled* state to 1.

On P1:

- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₉.

On P2:

- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₁₀.

In the execution that satisfies the postcondition, any of the following applies:

- E₄ is [Coherence-before](#) E₂, E₂ is [Coherence-before](#) E₆ and E~8. E₉ [Reads-from](#) E₄. In this execution, E₃ is [Coherence-before](#) E₂ which violates the [requirements of the Coherence-before](#) relation.
- E₂ is [Coherence-before](#) E₄, E₄ is [Coherence-before](#) E₈, E₈ is [Coherence-before](#) E₆. E₉ [Reads-from](#) E₈. In this execution, E₇ is [Coherence-before](#) E₆ which violates the [requirements of the Coherence-before](#) relation.
- E₆ is [Coherence-before](#) E₂, E₂ is [Coherence-before](#) E₄ and E~8. E₁₀ [Reads-from](#) E₆. In this execution, E₅ is [Coherence-before](#) E₂ which violates the [requirements of the Coherence-before](#) relation.

- E_2 is **Coherence-before** E_6 , E_6 is **Coherence-before** E_8 , E_8 is **Coherence-before** E_4 . E_{10} **Reads-from** E_8 . In this execution, E_7 is **Coherence-before** E_4 which violates the **requirements of the Coherence-before** relation.

B1.30 Atomicity of interrupt acknowledge and retarget

B1.30.1 Notes

The purpose of this test is to show that even though an interrupt can be acknowledged and subsequently made pending, it cannot be acknowledged a second time without a deactivate of the interrupt taking place.

B1.30.2 Litmus test

```
AArch64 co.ack-gic+WP+ack+ack
{
  [INTID(A)]=(enabled:1,affinity:P1,pending:1,active:0);
  0:X0=(intid:A, pending:1);
  0:X1=(intid:A, affinity:P2);
}
P0          | P1          | P2          ;
GIC CDAFF, X1 | GICR X2, CDIA | GICR X2, CDIA ;
GIC CDPEND, X0 |          |          ;

exists(1:X2=(valid:1,intid:A) /\ 2:X2=(valid:1,intid:A))

kinds: co.ack-gic+WP+ack+ack Forbid
```

B1.30.2.1 Explanation

On P0:

- GIC CDAFF, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₁ and an Explicit Interrupt Write Effect to INTID(A) E₂ that sets its affinity to P2.
- GIC CDPEND, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₃ and an Explicit Interrupt Write Effect to INTID(A) E₄ that set its pending state to 1.

On P1:

- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₅ and an Implicit Interrupt Read Effect to INTID(A) E₆ that resets its pending state to 0 and its active state to 1.

On P3:

- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₇ and an Implicit Interrupt Read Effect to INTID(A) E₈ that resets its pending state to 0 and its active state to 1.

In one of the possible executions that satisfies the postcondition:

- E₅ Reads-from the initial state of INTID(A) and as result E₅ is Coherence-before E₂.
- E₇ Reads-from E₂.
- E₂ is Coherence-before E₆.

This execution violates the [atomicity requirement](#) for the successful pair of Read Interrupt Effect E₅ and Write Interrupt Effect E₆, as E₅ is Coherence-before E₂ and E₂ is Coherence-before E₆.

There are other executions that satisfy the postcondition all of which violate the [atomicity requirement](#) for a successful pair of Read Interrupt and Write Interrupt Effects.

B1.31 1ofN interrupt acknowledge

B1.31.1 Notes

The purpose of this test is to show that a 1ofN interrupt can only be acknowledged on a single PE.

B1.31.2 Litmus test

```
AArch64 co.ack.1ofN-gic
{
  [INTID(A)]=(enabled:1,affinity:1ofN,pending:1,active:0);
}
P0                                | P1                                ;
GICR X2, CDIA                    | GICR X2, CDIA                    ;

exists(0:X2=(valid:1,intid:A) /\ 1:X2=(valid:1,intid:A))

kinds: co.ack.1ofN-gic Forbid
```

B1.31.2.1 Explanation

- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₁ and an Implicit Interrupt Read Effect to INTID(A) E₂ that resets its pending state to 0 and its active state to 1.

On P3:

- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₃ and an Implicit Interrupt Read Effect to INTID(A) E₄ that resets its pending state to 0 and its active state to 1.

In one of the execution that satisfies the postcondition:

- E₁ [Reads-from](#) the initial state of INTID(A) and as result E₁ and E₂ are [Coherence-before](#) E₄.
- E₃ [Reads-from](#) the initial state of INTID(A) and as result E₃ and E₄ are [Coherence-before](#) E₂.

This execution violates the [atomicity requirement](#) for the successful pair of Read Interrupt Effect E₁ and Write Interrupt Effect E₂, as E₁ is [Coherence-before](#) E₄ and E₄ is [Coherence-before](#) E₂.

B1.32 Retargeting interrupts without synchronization

B1.32.1 Notes

The purpose of this test is to show that due to the total order on writes to the same interrupt location, a write to change an interrupt's Affinity followed by a write to make the interrupt Pending, results in an interrupt not being taken on the old PE.

B1.32.2 Litmus test

```
AArch64 coWW.ack-gic+WP+ack
{
  [INTID(A)]=(enabled:1,affinity:P1,pending:0,active:0);
  0:X0=(intid:A, pending:1); 0:X1=(intid:A, affinity:P0);
  1:X0=x; 1:X1=(intid:A);
}
P0          | P1          ;
GIC CDAFF, X1 | GICR X2, CDIA ;
GIC CDPEND, X0 |          ;

exists(1:X2=(valid:1,intid:A))

kinds: coWW.ack-gic+WP+ack Forbid
```

B1.32.2.1 Explanation

On P0:

- GIC CDAFF, X1 generates an Explicit Interrupt Read Effect to INTID(A) E₁ and an Explicit Interrupt Write Effect to INTID(A) E₂ that sets its affinity to P0.
- GIC CDPEND, X0 generates an Explicit Interrupt Read Effect to INTID(A) E₃ and an Explicit Interrupt Write Effect to INTID(A) E₄ that set its pending state to 0.

On P1:

- GICR X2, CDIA generates an Implicit Interrupt Read Effect to INTID(A) E₅.

In the execution that satisfies the postcondition:

- E₅ [Reads-from](#) E₄.
- E₄ is [Coherence-before](#) E₂ and, as a result, E₃ is also [Coherence-before](#) E₂. This violates the [requirements of the Coherence-before](#) relation and, as a result, this is an architecturally forbidden execution.

B1.33 Reading interrupt configuration and exception status

B1.33.1 Notes

The purpose of these tests is to show the interaction between observing a change to an interrupt configuration and observing a corresponding change in the interrupt exception pending status.

See also:

- [B1.6 Configuration and acknowledgement](#)

B1.33.2 Litmus test with ISB before reading IRQ pending status

```
AArch64 coRex-gic+dis+RisbR
{
  [INTID(A)]=(enabled:1, pending:1, affinity:P1);
  0:X1=(intid:A);
  1:X1=(intid:A);
}
P0                                | P1                                ;
GIC  CDDIS, X1                    | GIC  CDRCFG, X1                    ;
                                | ISB                                ;
                                | MRS  X2, ICC_ICSR_EL1              ;
                                | MRS  X0, ISR_EL1                  ;

exists(1:X2=(enabled:0) /\ 1:X0=(I:1))

kinds: coRex-gic+dis+RisbR Forbid
```

B1.33.2.1 Explanation

For the execution that validates the postcondition:

On P0:

- GIC CDDIS, X1 generates an Explicit Interrupt Write Effect E_1 to INTID(A).

On P1:

- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect E_2 to INTID(A).
- ISB generates an Instruction Fetch Barrier Effect E_3 .
- MRS X0, ISR_EL1 generates a Direct System Register Read Effect to ISR_EL1 E_4 .

In the execution that satisfies the postcondition, E_2 [Reads-from](#) E_1 . But E_4 indicates that INTID(A) is still the HPPI. However because of E_3 INTID(A) is disabled and cannot be the HPPI. As a result, this execution is architecturally forbidden.

B1.33.3 Litmus test with ISB after reading IRQ pending status

```
AArch64 coRex-gic+dis+RR
{
  [INTID(A)]=(enabled:1, pending:1, affinity:P1);
  0:X1=(intid:A);
  1:X1=(intid:A);
}
P0                                | P1                                ;
GIC  CDDIS, X1                    | GIC  CDRCFG, X1                    ;
                                | MRS  X0, ISR_EL1                  ;
                                | ISB                                ;
                                | MRS  X2, ICC_ICSR_EL1              ;
```

```
exists(1:X2=(enabled:0) /\ 1:X0=(I:1))  
  
kinds: coRex-gic+dis+RR Allow
```

B1.33.3.1 Explanation

For the execution that validates the postcondition:

On P0:

- GIC CDDIS, X1 generates an Explicit Interrupt Write Effect E_1 to `INTID(A)`.

On P1:

- GIC CDRCFG, X1 generates an Explicit Interrupt Read Effect E_2 to `INTID(A)`.
- MRS X0, ISR_EL1 generates a Direct System Register Read Effect to `ISR_EL1` E_4 .

In the execution that satisfies the postcondition, E_2 [Reads-from](#) E_1 . But E_4 indicates that `INTID(A)` is still the HPPI. This execution is architecturally allowed.

B1.34 Reading interrupt configuration and IRQ unmask in PSTATE

B1.34.1 Notes

The purpose of these tests is to show the interaction between observing a change to an interrupt configuration and observing the corresponding interrupt exception.

See also:

- [B1.9 Configuration write and IRQ unmask in PSTATE](#)

B1.34.2 Explanation

The interrupt read effect R1 of GIC CDRCFG, X1 on P1 may or may not observe the interrupt write effect W2 of GIC CDDIS, X1 on P0, but if it does, it means that the interrupt has been recalled from P1 on completion of W2.

If R1 from GIC CDRCFG, X1 observes W2 from GIC CDDIS, X1, then effects of W2 on the execution context (whether an IRQ exception is pending) must be observed by the instructions following the ISB. This means that when DAIFClr PSTATE.I is executed, the IRQ interrupt exception will not be taken.

B1.34.3 Litmus test with ISB before unmasking IRQ in PSTATE

```
AArch64 coRR+gic.cddis+gic.cdrcfg-isb-clr.i
{
  [INTID(A)]=(enabled:1, pending:1, affinity:P1);
  0:X1=(intid:A);
  1:X1=(intid:A);
}
P0                                | P1                                ;
GIC CDDIS, X1                    | GIC CDRCFG, X1                    ;
                                | ISB                                ;
                                | MRS X2, ICC_ICSR_EL1             ;
                                | DAIFClr PSTATE.I                 ;

exists(1:X2=(enabled:0) /\ async(P1, IRQ)

kinds: coRR+gic.cddis+gic.cdrcfg-isb-clr.i Forbid
```

B1.34.4 Litmus test with ISB after unmasking IRQ in PSTATE

```
AArch64 coRR+gic.cddis+gic.cdrcfg-clr.i-isb
{
  [INTID(A)]=(enabled:1, pending:1, affinity:P1);
  0:X1=(intid:A);
  1:X1=(intid:A);
}
P0                                | P1                                ;
GIC CDDIS, X1                    | GIC CDRCFG, X1                    ;
                                | DAIFClr PSTATE.I                 ;
                                | ISB                                ;
                                | MRS X2, ICC_ICSR_EL1             ;

exists(1:X2=(enabled:0) /\ async(P1, IRQ)

kinds: coRR+gic.cddis+gic.cdrcfg-clr.i-isb Allow
```

B1.35 PPI activate and system register read

B1.35.1 Notes

The PPI active state is stored in system registers.

These tests show that an ISB is required to guarantee the updates to the system registers from acknowledge interrupts are visible to a system register read.

Illustrates a common principle that any instruction with effects to state held in system registers, including GICR and GIC instructions affecting PPI state, can be observed following an ISB.

B1.35.2 Litmus test without ISB

```
AArch64 coWmrs.ack
{
  0:[ppi(0)]=(enabled:1,pending:1,active:0);
}
P0
GICR X0, CDIA ;
MRS X1, ICC_PPI_SACTIVER0_EL1 ;

exists(0:X0=(valid:1, initd:ppi0) /\ 0:X1=0)

kinds: coWmrs.ack Allow
```

B1.35.2.1 Explanation

On P0:

- GICR X0, CDIA acknowledges ppi(0) and as a result generates an Indirect System Register Write Effect to ICC_PPI_SACTIVER0_EL1 E₁.
- MRS X1, ICC_PPI_SACTIVER0_EL1 generates a Direct System Register Read Effect to ICC_PPI_SACTIVER0_EL1 E₂.

As per the postcondition, E₂ is Coherence-before E₁ which is architecturally allowed. See ‘Table D24-1 Synchronization requirements’ in the Architecture Reference Manual for A-profile [1] for more information.

B1.35.3 Litmus test with GSB

```
AArch64 coWmrs.ack+gsb
{
  0:[ppi(0)]=(enabled:1,pending:1,active:0);
}
P0
GICR X0, CDIA ;
GSB ACK ;
MRS X1, ICC_PPI_SACTIVER0_EL1 ;

exists(0:X0=(valid:1, initd:ppi0) /\ 0:X1=0)

kinds: coWmrs.ack+gsb Allow
```

B1.35.3.1 Explanation

On P0:

- GICR X0, CDIA acknowledges ppi(0) and as a result generates an Indirect System Register Write Effect to ICC_PPI_SACTIVER0_EL1 E₁.

- GSB ACK generates a GSB ACK Effect E_2 .
- MRS $X1, ICC_PPI_SACTIVER0_EL1$ generates a Direct System Register Read Effect to $ICC_PPI_SACTIVER0_EL1$ E_3 .

As per the postcondition, E_3 is Coherence-before E_1 which is architecturally allowed. See ‘Table D24-1 Synchronization requirements’ in the Architecture Reference Manual for A-profile [1] for more information.

B1.35.4 Litmus test with ISB

```
AArch64 coWmrs.ack+isb
{
  0:[ppi(0)]=(enabled:1,pending:1,active:0);
}
P0
GICR X0, CDIA ;
ISB ;
MRS X1, ICC_PPI_SACTIVER0_EL1 ;

exists(0:X0=(valid:1, initd:ppi0) /\ 0:X1=0)

kinds: coWmrs.ack+isb Forbid
```

B1.35.4.1 Explanation

On P0:

- GICR $X0, CDIA$ acknowledges $ppi(0)$ and as a result generates an Indirect System Register Write Effect to $ICC_PPI_SACTIVER0_EL1$ E_1 .
- ISB generates an Instruction Fetch Barrier Effect E_2 which is a Context Synchronization Event.
- MRS $X1, ICC_PPI_SACTIVER0_EL1$ generates a Direct System Register Read Effect to $ICC_PPI_SACTIVER0_EL1$ E_3 .

As per the postcondition, E_3 is Coherence-before E_1 . This violates the requirements of the Coherence-before definition. As a result this execution is architecturally forbidden.

See also ‘Table D24-1 Synchronization requirements’ in the Architecture Reference Manual for A-profile [1] for more information.

B1.36 PPI disable and acknowledge

B1.36.1 Notes

The PPI configuration is stored in system registers.

The purpose of the test is to show that the effects of updates to the execution context are completed by the ISB, including which PPIs can be the HPPI.

Illustrates a common principle that any instruction with effects to state held in system registers, including GICR and GIC instructions affecting PPI state, can be observed following an ISB.

B1.36.2 Litmus test with ISB

```
AArch64 comsrR.ack+isb
{
  0:[ppi(0)]=(enabled:1,pending:1,active:0);
  0:X1=(ppi0:1);
}
P0
MSR    ICC_PPI_CENABLER0_EL1, X1 ;
ISB                                         ;
GICR   X0, CDIA                          ;

exists(0:X0=(valid:1, initd:ppi0))

kinds: comsrR.ack+isb Forbid
```

B1.36.2.1 Explanation

On P0:

- MSR ICC_PPI_CENABLER0_EL1, X1 generates a Direct System Register Write Effect to ICC_PPI_CENABLER0_EL1 E₁ and disables ppi(0).
- ISB generates an Instruction Fetch Barrier Effect E₂ which is a Context Synchronization Event.
- GICR X0, CDIA acknowledges ppi(0) and as a result generates an Indirect System Register Read Effect to ICC_PPI_CENABLER0_EL1 E₃.

As per the postcondition, E₃ is Coherence-before E₁ as E₃ finds ppi(0) still enabled. This violates the requirements of the Coherence-before definition. As a result this execution is architecturally forbidden.

See also ‘Table D24-1 Synchronization requirements’ in the Architecture Reference Manual for A-profile [1] for more information.

B1.37 PPI acknowledgement

B1.37.1 Notes

The PPI active state is stored in system registers.

The purpose of the test is to show that a GICR instruction that acknowledges a PPI generates an Indirect System Register Write to ICC_PPI_SPEND<n>_EL1.

B1.37.2 Litmus test

```
AArch64 coWR-ppi+ia-isb-mrs.sprendr1
{
  0:[ppi(64)]=(enabled:1,pending:1,active:0); // Edge-triggered IMPDEF PPI 64
  0:X1=1;
}
P0
GICR  X0, CDIA                ;
ISB                                     ;
MRS   ICC_PPI_SPENDR1_EL1, X2    ;

exists(0:X2=1)

kinds: coWR-ppi+ia-isb-mrs.sprendr1 Forbid
```

B1.37.2.1 Explanation

In the execution that satisfies the postcondition, on P0:

- GICR X1, CDIA acknowledges ppi(64) and generates an Indirect System Register Write to ICC_PPI_SPEND1_EL1 E₁ and sets Pend0 to 0.
- ISB generates an Instruction Fetch Barrier Effect E₂ which is a Context Synchronization Event.
- MRS ICC_PPI_SPEND1_EL1, X2 generates a Direct System Register Read to ICC_PPI_SPEND1_EL1 E₃.

As per the postcondition, E₃ is Coherence-before E₁. This violates the requirements of the Coherence-before definition. As a result this execution is architecturally forbidden. See ‘Table D24-1 Synchronization requirements’ in the Architecture Reference Manual for A-profile [1] for more information.

B1.37.3 Litmus test

```
AArch64 edge-merging.ppi
{
  0:[ppi(64)]=(enabled:1,pending:1,active:0); // Edge-triggered IMPDEF PPI 64
  0:X1=1;
}
P0
GICR  X0, CDIA                ;
MSR   ICC_PPI_SPENDR1_EL1, X1    ;
MRS   ICC_PPI_SPENDR1_EL1, X2    ;

exists(0:X2=0)

kinds: edge-merging.ppi Forbid
```

B1.37.3.1 Explanation

In the execution that satisfies the postcondition, on P0:

- GICR X_1 , CDIA acknowledges `ppi(64)` and generates an Indirect System Register Write to ICC_PPI_SPEND1_EL1 E_1 and sets `Pend0` to 0.
- MSR ICC_PPI_SPEND1_EL1, X_1 generates a Direct System Register Write to ICC_PPI_SPEND1_EL1 E_2 and sets `Pend0` to 1.
- MRS ICC_PPI_SPEND1_EL1, X_2 generates a Direct System Register Read to ICC_PPI_SPEND1_EL1 E_3 .

As per the postcondition, E_3 Reads-from E_1 . This is architecturally forbidden as E_1 and E_2 are ordered without the need of synchronization, and E_3 has to Read-from E_2 as both are Direct System Register Effects.

See also *Arm® Architecture Reference Manual, for A-profile architecture*[\[1\]](#) Table D24-1 Synchronization requirements' for more information.

B1.38 Write after changing resident VM

B1.38.1 Notes

The purpose of this test is to show that an ISB is required to change the resident VM on a PE so that interrupt write effects are made to the expected interrupts.

Illustrates a common principle that any instruction with effects to state held in system registers, including GICR and GIC instructions affecting PPI state, can be observed following an ISB.

B1.38.2 Litmus test

```
AArch64 vmW-gic+R
{
  [VM0:INTID(A)]=(enabled:0);
  [VM1:INTID(A)]=(enabled:0);
  0:ICH_CONTEXTR_EL2=(VM1, VPE0, VALID);
  0:X0=(intid:A); 0:X1=(VM0, VPE0, VALID);
  1:ICH_CONTEXTR_EL2=(VM1, VPE1, VALID);
  1:X0=(intid:A);
}
P0                                | P1                                ;
MSR ICH_CONTEXTR_EL2, XZR | GIC VDRCFG, X0                ;
MSR ICH_CONTEXTR_EL2, X1  | ISB                               ;
ISB                          | MRS X2, ICC_ICSR_EL1      ;
GIC VDEN, X0                |                             ;

exists(1:X2=(enabled:1))

kinds: vmW-gic+R Forbid
```

B1.38.2.1 Explanation

On P0:

- MSR ICH_CONTEXTR_EL2, XZR generates a Direct System Register Write Effect to ICH_CONTEXTR_EL2 E₁ and ensures that there is no resident VPE.
- MSR ICH_CONTEXTR_EL2, X1 generates a Direct System Register Write Effect to ICH_CONTEXTR_EL2 E₂ and makes VPE 0 of VM 0 resident.
- ISB generates an Instruction Fetch Barrier Effect E₃ which is a Context Synchronization Event.
- GIC VDEN, X0 generates an Indirect System Register Read to ICH_CONTEXTR_EL2 E₄ and an Explicit Interrupt Write Effect to INTID(A) E₅ enabling INTID(A).

On P1:

- GIC VDRCFG, X0 generates an Indirect System Register Read to ICH_CONTEXTR_EL2 E₆ and an Explicit Interrupt Read Effect to INTID(A) E₇.

As per the postcondition, E₇ reads that INTID(A) is enabled which could only have happened if E₇ [Reads-from](#) E₅. This in turn means that E₅ was to VM1:INTID(A) and as a result, E₄ Reads-from the initial value of ICH_CONTEXTR_EL2 and E₄ is Coherence-before E₂. This violates the requirements of the Coherence-before definition and this execution is architecturally forbidden.

See ‘Table D24-1 Synchronization requirements’ in the Architecture Reference Manual for A-profile [1] for more information.

B1.39 Write before changing resident VM

B1.39.1 Notes

The purpose of this test is to show that an ISB is not needed to ensure that interrupt write effects are made to the resident VM when an instruction to change the resident VM occurs in program order after the instruction causing the interrupt write effect.

B1.39.2 Litmus test

```
AArch64 Wvm-gic+R
{
  [VM0:INTID(A)]=(enabled:0);
  [VM1:INTID(A)]=(enabled:0);
  0:ICH_CONTEXTR_EL2=(VM0, VPE0);
  0:X0=(intid:A); 0:X1=(VM1, VPE0);
  1:ICH_CONTEXTR_EL2=(VM1, VPE0);
  1:X0=(intid:A);
}
P0                                | P1                                ;
GIC VDEN, X0                     | GIC VDRCFG, X0                     ;
MSR ICH_CONTEXTR_EL2, XZR | ISB                                           ;
MSR ICH_CONTEXTR_EL2, X1 | MRS X2, ICC_ICSR_EL1                           ;

exists(1:X2=(enabled:1))

kinds: Wvm-gic+R Forbid
```

B1.39.2.1 Explanation

On P0:

- GIC VDEN, X0 generates an Indirect System Register Read to ICH_CONTEXTR_EL2 E₁ and an Explicit Interrupt Write Effect to INTID(A) E₂ enabling INTID(A).
- MSR ICH_CONTEXTR_EL2, XZR generates a Direct System Register Write Effect to ICH_CONTEXTR_EL2 E₃ and ensures that there is no resident VPE.
- MSR ICH_CONTEXTR_EL2, X1 generates a Direct System Register Write Effect to ICH_CONTEXTR_EL2 E₄ and makes VPE 0 of VM 0 resident.

On P1:

- GIC VDRCFG, X0 generates an Indirect System Register Read to ICH_CONTEXTR_EL2 E₅ and an Explicit Interrupt Read Effect to INTID(A) E₆.

As per the postcondition, E₆ reads that INTID(A) is enabled which could only have happened if E₆ [Reads-from](#) E₂. This in turn means that E₂ was to VM1:INTID(A) and as a result, E₁ Reads-from E₄. This violates the requirements of the Reads-from definition and this execution is architecturally forbidden.

See ‘Table D24-1 Synchronization requirements’ in the Architecture Reference Manual for A-profile [1] for more information.

B1.40 Completion of GIC and GICR instructions in finite time

B1.40.1 Notes

The purpose of this test is to show that writes to Interrupt Locations are required to become visible to observers in finite time.

B1.40.2 Litmus test with a priority update observed by a configuration read

```
AArch64 gic.pri+gic.rcfg-wait
{
  [INTID(A)]=(priority:0);
  0:X1=(intid:A, priority:1);
  1:X0=(intid:A); 1:X2=(priority:1);
}
P0          | P1          ;
GIC CDPRI, X1 |L0:          ;
              | GIC CDRCFG, X0    ;
              | ISB              ;
              | MRS X1, ICC_ICSR_EL1 ;
              | CMP X1, X2          ;
              | B.NE L0           ;
existsN(1:X1=(priority:1))

kinds: gic.pri+gic.rcfg-wait Require
```

B1.40.2.1 Explanation

On P0:

- GIC CDPRI, X1 generates an Explicit Interrupt Write Effect E_1 to INTID(A).

On P1:

- GIC CDRCFG, X0 generates Explicit Interrupt Read Effects $E_2, E_3, \dots E_n$ to INTID(A).

It is architecturally required that there is a finite number n for which E_n Reads-from E_1 .

B1.40.3 Litmus test with an interrupt acknowledge observed by a configuration read

```
AArch64 gic.ia+gic.rcfg-wait
{
  [INTID(A)]=(enabled:1, pending:1, active:0, priority:1, affinity:P0);
  1:X0=gicarg:(intid:A);
  1:X2=(enabled:1, pending:0, active:1, priority:1, affinity:P0);
}
P0          | P1          ;
GICR X1,CDIA |L0:          ;
              | GIC CDRCFG, X0    ;
              | ISB              ;
              | MRS X1, ICC_ICSR_EL1 ;
              | CMP X1, X2          ;
              | B.NE L0           ;
existsN(0:X1=(valid:1, intid:A) /\ 1:X1=(priority:1))

kinds: gic.ia+gic.rcfg-wait Require
```

B1.40.3.1 Explanation

On P0:

- GICR X1, CDIA generates an Implicit Interrupt Write Effect E_1 to INTID(A).

On P1:

- GIC CDRCFG, X0 generates Explicit Interrupt Read Effects $E_2, E_3, \dots E_n$ to INTID(A).

It is architecturally required that there is a finite number n for which E_n Reads-from E_1 .

B1.40.4 Litmus test with an interrupt becoming pending and then acknowledged

```
AArch64 gic.pend1+gic.ia-wait
{
  [INTID(A)]=(enabled:1, pending:0, active:0, priority:1, affinity:P1);
  0:X0=(intid:A,pending:1);
  1:X2=(valid:1,intid:A);
}
P0          | P1          ;
GIC CDPEND,X0 | L0:          ;
              | GICR X1,CDIA ;
              | CMP X1, X2   ;
              | B.NE L0      ;
existsN(1:X1=(valid:1,intid:A))

kinds: gic.pend1+gic.ia-wait Require
```

B1.40.4.1 Explanation

On P0:

- GIC CDPEND, X0 generates an Explicit Interrupt Write Effect E_1 to INTID(A).

On P1:

- GICR X1, CDIA generates Implicit Interrupt Read Effects $E_2, E_3, \dots E_n$ to INTID(A).

It is architecturally required that there is a finite number n for which E_n Reads-from E_1 .

Chapter B2

Effects of disabling a PPI source on the PPI Pending state

B2.0.1 Notes

The purpose of this test is to show that updates to a PPI Pending state can not be synchronized by any instruction.

B2.0.2 Litmus test with disable of the timer state

```
AArch64 timer-disable-isb-ppi-read
{
  0: [ppi(27)] = (enabled:1, pending:1, active:0);
}
P0
MRS  X0, CNTV_CTL_EL0          ;
MOV  X1, #0                    ;
MSR  CNTV_CTL_EL0, X1          ;
ISB                               ;
MRS  X2, ICC_PPI_SPENDR0       ;

exists(0:X0=(ISTATUS=1, IMASK=0, ENABLE=1) /\ 0:X2=0x00020000)

kinds: timer-disable-isb-ppi-read Allow
```

B2.0.2.1 Explanation

Updates from a PPI source to the PPI Pending state is an autonomous asynchronous event. Even though the PPI Pending state can be observed by reading the ICC_PPI_SPENDCR0 System register, the update to the PPI Pending state is not an indirect System register write to ICC_PPI_SPENDRR0 caused by the direct System register write to CNTV_CTL_EL0. Instead, the write to CNTV_CTL_EL0 is observed by the autonomous asynchronous event that causes the EL1 Virtual Timer to evaluate its state and update the Pending state of the PPI.

Part C

Model

Chapter C1

Operational model

I_QZJCC

The following code is an operational model of the GIC, written in Arm Specification Language (ASL). This section is at Alpha quality. Alpha quality means that most major features of the specification are described in this release, but some features and details might be missing.

See also:

- Arm Specification Language Reference Manual [14].

```
// exceptions.asl
//
// This is the pseudocode implementing the exceptions signaling logic.

INTID GetHPPI(bits(2) domain)
    return DOMAIN_HPPIS[UInt(domain)].intid;

boolean HasHPPI(bits(2) domain)
    return DOMAIN_HPPIS[UInt(domain)].valid;

boolean DomainEnabled(bits(2) domain)
    if domain == DOM_EL3 then
        return ICC_CR0_EL3.EN == '1';
    else
        return BankedICC_CR0_EL1(domain).EN == '1';

bits(2) PhysicalIRQTarget()
    route_to_el2 = (PSTATE.EL IN {EL0, EL1} && EL2Enabled() &&
        (HCR_EL2.TGE == '1' || HCR_EL2.IMO == '1'));
```

```

    if PSTATE.EL == EL2 || route_to_el2 then
        assert PSTATE.EL != EL3;
        return EL2;
    else
        assert PSTATE.EL IN {EL0, EL1};
        return EL1;

boolean NMIEntered()
    if PhysicalIRQTarget() == EL2 then
        return SCTL_EL2.NMI == '1';
    if PhysicalIRQTarget() == EL1 then
        return SCTL_EL1.NMI == '1';
    return FALSE;

integer GetRunningPriority(bits(2) domain)
    return 31; // Should return the running priority based on the value in the
    ↪banked ICC_APR_EL1 or ICC_APR_EL3

integer GetPriorityMask(bits(2) domain)
    case domain of
        when DOM_EL3 return UInt(ICC_PCR_EL3.Priority);
        when DOM_S return UInt(ICC_PCR_EL1[0].Priority);
        when DOM_NS return UInt(ICC_PCR_EL1[1].Priority);
        when DOM_RL return UInt(ICC_PCR_EL1[2].Priority);

boolean HPPIAvailable(bits(2) domain)
    if !DomainEnabled(domain) then
        return FALSE;

    hppi = GetHPPI(domain);
    hppi_priority = UInt(hppi.Priority);
    running_priority = GetRunningPriority(domain);
    priority_mask = GetPriorityMask(domain);

    return hppi_priority < running_priority && hppi_priority <= priority_mask;

bits(2) PreemptiveDomain()
    return ICC_CR0_EL3.PID;

boolean PreemptiveHPPIAvailable(bits(2) domain)
    if domain == DOM_EL3 then
        return FALSE; // 0b10 in ICC_CR0_EL3.PID encodes no preemptive domain

    if PreemptiveDomain() != domain then
        return FALSE;

    return HPPIAvailable(domain) && UInt(GetHPPI(domain).Priority) < UInt(
    ↪BankedICC_CR0_EL1(domain).IPPT);

// GICSignalException()
// =====
// Signal CPU interrupt exception based on HPPI availability in the Interrupt
// Domains.

GICPhysicalExceptions()
    if CurrentDomain() == DOM_EL3 then
        if HPPIAvailable(DOM_EL3) ||
            HPPIAvailable(DOM_S) ||
            HPPIAvailable(DOM_NS) ||

```

```

        HPPIAvailable(DOM_RL) then
            FIQ();
    else // Non-EL3 Domain
        if HPPIAvailable(EL3) || PreemptiveHPPIAvailable(PreemptiveDomain()) then
            FIQ();
        elsif HPPIAvailable(CurrentDomain()) then
            IRQ();

    return;

// gic_shared_functions.asl

sysreg_ICC_CR0_EL1 BankedICC_CR0_EL1(bits(2) domain)
    case domain of
        when DOM_EL3 assert FALSE;
        when DOM_S return ICC_CR0_EL1[0];
        when DOM_NS return ICC_CR0_EL1[1];
        when DOM_RL return ICC_CR0_EL1[2];

bits(2) CurrentDomain()
    if PSTATE.EL == EL3 then
        return DOM_EL3;
    else
        bits(2) state_bits;
        if HaveRME() then
            state_bits = SCR_EL3.<NSE,NS>;
        else
            state_bits = '0' : SCR_EL3.NS;

        case state_bits of
            when '00' return DOM_S;
            when '01' return DOM_NS;
            when '10' Unreachable();
            when '11' return DOM_RL;

// HavePreemptiveDomain()
// =====
// Returns TRUE when the Domain selected by ICC_CR0_EL3 is implemented, FALSE
// ↪otherwise

// MostPrivilegedSecurityState()
// =====
// Returns the most privileged Security state

SecurityState MostPrivilegedSecurityState()
    // We can always ask SecurityStateAtEL with 'EL3' to get the right answer even
    // ↪if EL3 is not implemented
    return SecurityStateAtEL(EL3)

// IsAccessMPPAS()
// =====
// Returns TRUE when an access is made with the Most Privileged PAS

boolean IsAccessMPPAS();
    mpss = MostPrivilegedSecurityState()

    case mpss of
        when SS_Root return IsAccessRoot();
        when SS_Secure return IsAccessSecure();

```

```

        when SS_NonSecure return IsAccessNonSecure();
        when SS_Realm Unreachable();

// IsAccessIWBMPAS()
// =====
// Returns TRUE when an access is made with the Most Privileged PAS of the IWB

boolean IsAccessIWBMPAS();
    iwb_mppas = IWBMPAS()

    case iwb_mppas of
        when PAS_Root return IsAccessRoot();
        when PAS_Secure return IsAccessSecure();
        when PAS_NonSecure return IsAccessNonSecure();
        otherwise Unreachable();

bits(2) IWBMPAS()
    domains = IWB_IDR0.INT_DOMS
    case domains of
        when '0001' return PAS_Secure;
        when '0010' return PAS_NonSecure;
        when '0111' return PAS_Secure;
        when '1110' return PAS_Root;
        when '1111' return PAS_Root;
        otherwise Unreachable();

// IsWireAccessible()
// =====
// Returns TRUE when a wire can be configured on an IWB for a given access PAS

boolean IsWireAccessible(bits(11) n, bits(5) x)
    // IWB MPPAS can access state of all wires
    if IsAccessIWBMPAS() then
        return TRUE;

    reg_index = (n * 2) + (x / 16);
    wire_index = x MOD 16;
    wire_domain = IWB_WDOMAINR[reg_index].WDOM[wire_index];

    if IsAccessRoot() then
        // If PAS is Root and it is not the IWB MPPAS, then either EL3 Domain is
        // not supported or it is associated with the Secure PAS.
        return FALSE;

    if IsAccessSecure() then
        // Wire cannot be accessed if Secure Interrupt Domain is not implemented.
        if IWB_IDR0.INT_DOMS[0] == 0 then
            return FALSE;
        if wire_domain == DOM_S then
            return TRUE;

    if IsAccessRealm() then
        // Wire cannot be accessed if Realm Interrupt Domain is not implemented.
        if IWB_IDR0.INT_DOMS[3] == 0 then
            return FALSE;
        if wire_domain == DOM_RL then
            return TRUE;

```

```

    if IsAccessNonSecure() then
        // Wire cannot be accessed if Non-secure Interrupt Domain is not
        // ↪implemented.
        if IWB_IDR0.INT_DOMS[1] == 0 then
            return FALSE;
        if wire_domain == DOM_NS then
            return TRUE;

    Unreachable();

// IsWireDomainRO()
// =====
// Returns TRUE if the Interrupt Domain assignment of an IWB input wire is fixed

boolean IsWireDomainRO(bits(16) wire_idx)
    return IMPLEMENTATION_DEFINED "Whether the Interrupt Domain assignment of a
    // ↪wire is fixed.";

// IsWireConfigRO()
// =====
// Returns TRUE if the Trigger mode of an IWB input wire is read-only

boolean IsWireConfigRO(bits(16) wire_idx)
    return IMPLEMENTATION_DEFINED "Whether an IWB input wire Trigger mode is read-
    // ↪only";

// IsPPIAssignedToCurrentDomain()
// =====
// Returns TRUE when a physial PPI is assigned to the Current Physical Interrupt
// ↪Domain

boolean IsPPIAssignedToCurrentDomain(bits(7) PPI_ID)
    n = PPI_ID DIV 32;
    x = PPI_ID MOD 32;

    if !HaveEL(EL3) then
        return TRUE;

    return ICC_PPI_DOMAINNR_EL3[n]<x + 1:x> == CurrentDomain();

// IsSPIAssignedToDomain()
// =====
// Returns TRUE if the SPI is assigned to the specified domain.
// This would be correctly modeled by accessing the internal SPI Domain assignment
// ↪state that is exposed via IRS_SPI_DOMAINNR.

boolean IsSPIDVISupported(bits(29) spi_id, bits(2) domain);

// IsSPIDVISupported()
// =====
// Returns TRUE if the SPI supports being assigned to a VM.
// Arm strongly recommends that an SPI that can be driven by a signal external to
// ↪the IRS supports assignment to a VM.

```

```
boolean IsSPIDVISupported(bits(29) spi_id)
    return IMPLEMENTATION_DEFINED "Whether the SPI can be assigned to a VM.";
```

Glossary

Abort

Any situation where the memory system returns an error response to a memory access by the GIC. Examples of error responses could include GPF faults returned by an SMMU[5] or DECERR in AMBA AXI[7] systems.

Accepted

An event has been Accepted by a GICv5 system component when it has sent a reply to the sender of the event. See [3.2 Communication between GIC system components](#) for more information.

Application PE

A PE used by the operating system or hypervisor to execute user application or kernel threads.

ASL

Arm Specification Language
Language used to express pseudocode implementations. Formal language definition can be found in the *Arm® Specification Language Reference Manual*[14].

Candidate HPPI

An interrupt identified among a subset of all interrupts as a candidate for the HPPI. See [2.7 Interrupt Prioritization](#) for more information.

Doorbell PPI

A PPI used to signal to an Interrupt Domain that there is an interrupt pending for another Interrupt Domain. See [2.9.6 Doorbell PPIs](#) for more information.

Downstream

An interrupt signal flows *downstream* from the interrupt source through the ITS, IRS, and is presented to PEs.

Where two components are involved, *upstream* is the component farthest away from the PE, and *downstream* is the component closest to the PE.

This also applies to the GICv5 Stream protocol where downstream packets flow from the IRS to the PE, and upstream packets flow from the PE to the IRS.

GPC

Granule Protection Check. The process of checking whether an access to a Location is permitted by the Granule Protection Table. This term is also used to refer to an MMU-attached or SMMU-attached Granular PAS filter that implements the Granule Protection Check. See *Arm® Architecture Reference Manual, for A-profile architecture*[1] for more information.

HPPI

Highest Priority Pending Interrupt. See [2.7 Interrupt Prioritization](#) for more information.

IPI

Inter-processor Interrupt. See [2.5 Inter-Processor Interrupts](#) for more information.

IRI

[Interrupt Routing Infrastructure](#). See [Chapter 1 Introduction](#) and [Chapter 3 GICv5 system architecture](#) for more information.

IRS

Interrupt Routing Service. See [Chapter 4 *Interrupt routing service \(IRS\)*](#) for more information.

ITS

Interrupt Translation Service. See [Chapter 5 *Interrupt translation service \(ITS\)*](#) for more information.

IWB

Interrupt Wire Bridge. See [Chapter 6 *Interrupt Wire Bridge \(IWB\)*](#) for more information.

LPI

Logical Peripheral Interrupt. See [2.4 *Interrupt types and identifiers*](#) for more information.

MPPAS

Most Privileged PAS. See [3.1 *Interrupt Domains*](#) for more information.

MPSS

Most Privileged Security State. See [3.1 *Interrupt Domains*](#) for more information.

NMI

Non-maskable Interrupt. See [2.9 *The physical CPU interface*](#) for more information.

PAS

Physical Address Space. See *Arm® Architecture Reference Manual, for A-profile architecture*[1] for more information.

PPI

PE-Private Peripheral Interrupt. See [2.4 *Interrupt types and identifiers*](#) for more information.

Processed

An interrupt event has been processed by an IRS when the interrupt effects of an interrupt event can be observed by PEs. See [4.5 *IRS synchronization requests*](#) for more information.

RAS

Reliability, Availability, and Serviceability. See *Arm® Architecture Reference Manual, for A-profile architecture*[1] and *Arm® Reliability, Availability, and Serviceability (RAS) System Architecture for A-profile architecture*[15] for more information.

Reachable

An INTID is reachable from a PE if the PE can access the configuration and state of the interrupt. See [2.6 *GIC System instructions*](#) and [4.6 *Interrupt configuration and state*](#) for more information.

Running priority

The highest active priority on an Interrupt Domain. See [2.8 *Interrupt handling*](#) for more information.

SGI

Software Generated Interrupt. A dedicated interrupt type for IPIs in GIC version 3 and version 4. See *Arm® Generic Interrupt Controller Architecture Specification, GIC architecture version 3 and version 4*[3] for more information.

SPI

Shared Peripheral Interrupt. See [2.4 *Interrupt types and identifiers*](#) for more information.

Unreachable

See Reachable.

Upstream

See Downstream

VM

Virtual machine.

VMSA

Virtual Memory System Architecture defined in *Arm[®] Architecture Reference Manual, for A-profile architecture*[\[1\]](#).

VPE

Virtual PE.

VPE doorbell

An LPI which is signaled when a non-resident VPE receives an interrupt. See [4.10.7 VPE doorbells](#) for more information.

Write-One-to-Clear (W1C)

Writes of 0 to the bit are ignored. A write of 1 clears the bit to 0. See *Arm[®] Architecture Reference Manual, for A-profile architecture*[\[1\]](#) for more information.