



Arm[®] Architecture Reference Manual for A-profile architecture

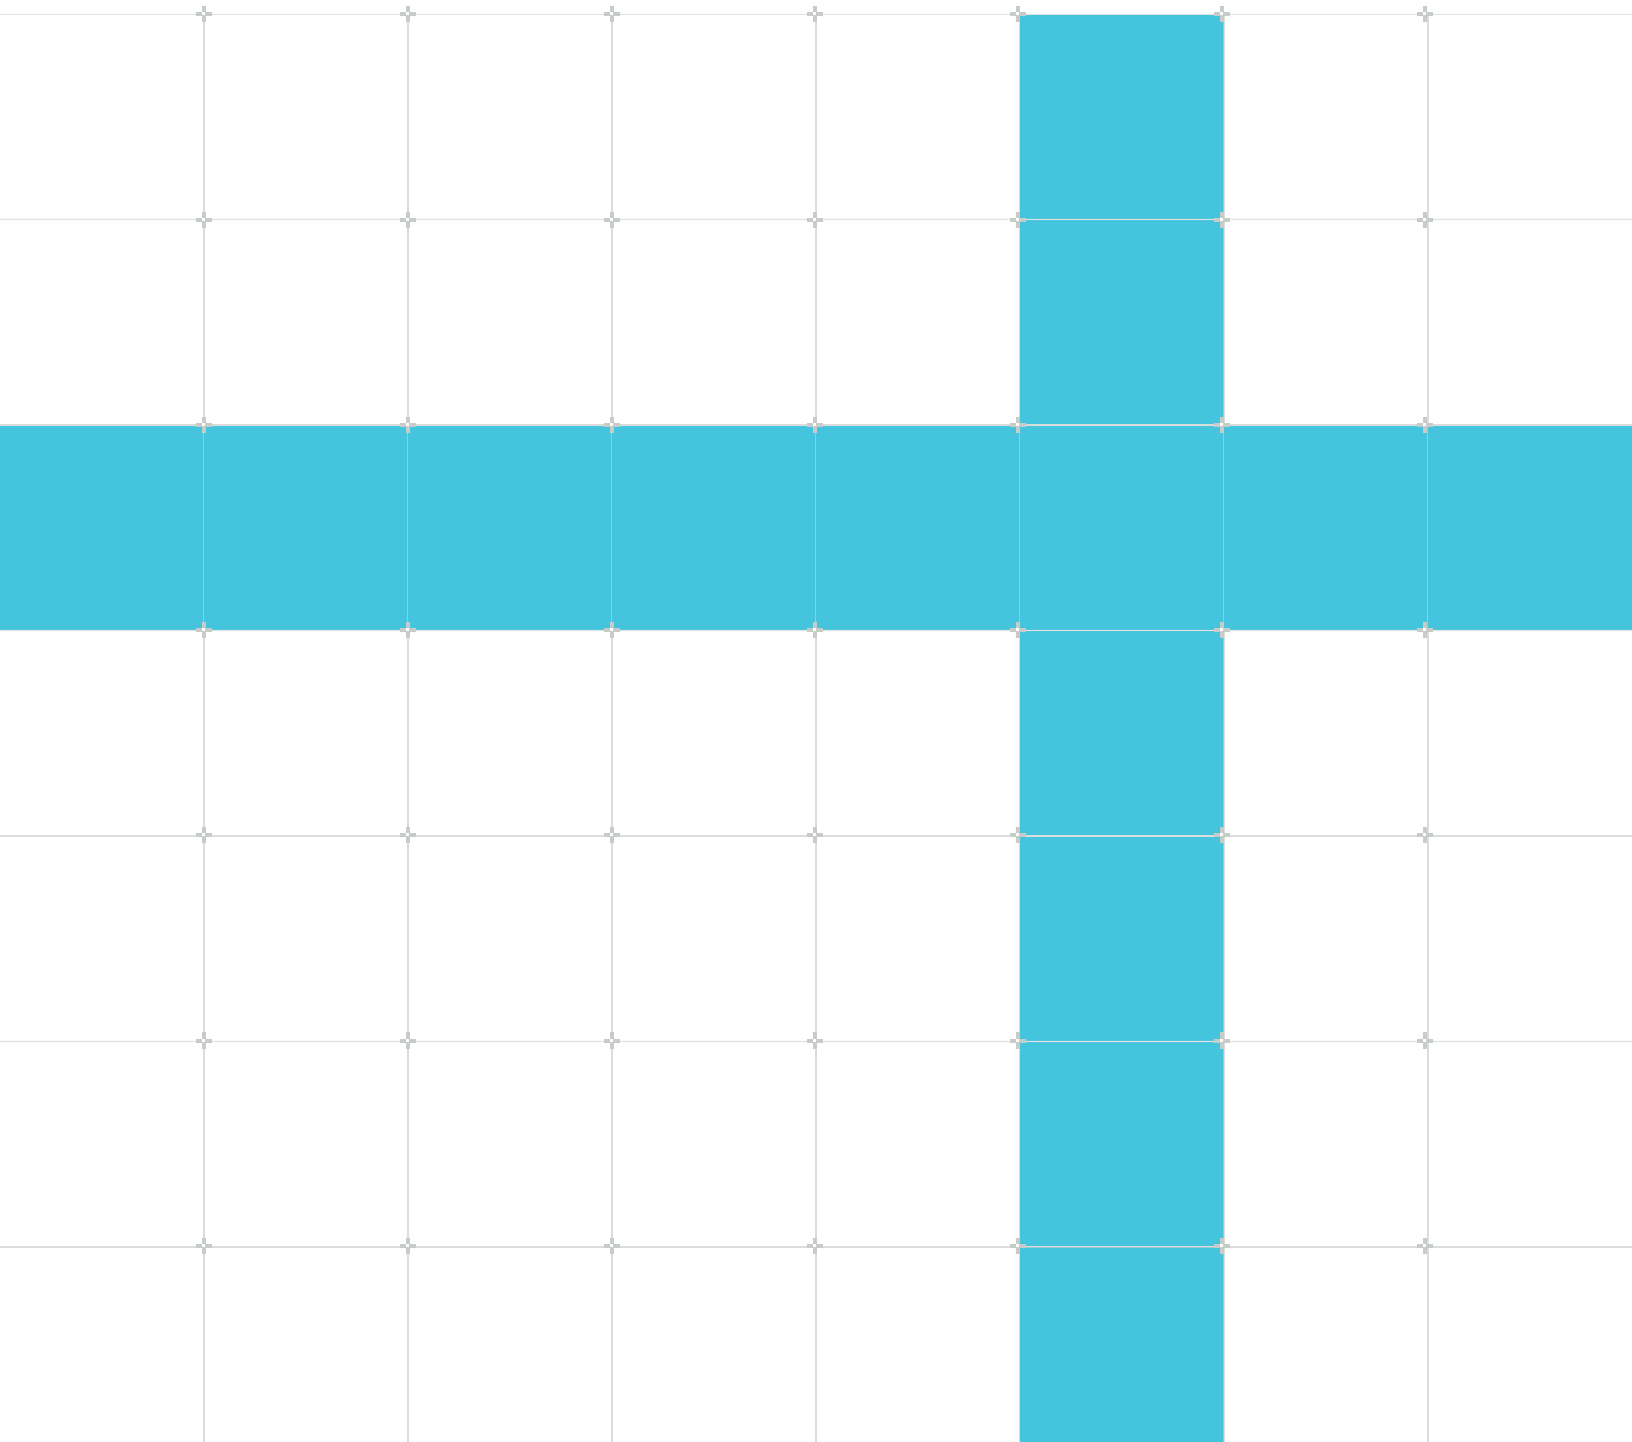
Known issues in Issue L.b

Non-Confidential

Copyright © 2020, 2022–2025 Arm Limited (or its affiliates).
All rights reserved.

Issue 02

102105_L.b_02_en



Arm® Architecture Reference Manual for A-profile architecture Known issues in Issue L.b

Copyright © 2020, 2022–2025 Arm Limited (or its affiliates). All rights reserved.

Release information

Document history

Issue	Date	Confidentiality	Change
L.b-02	4 August 2025	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue L.b, as of 4 August 2025
L.b-01	3 July 2025	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue L.b, as of 1 July 2025
L.b-00	30 May 2025	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue L.b, as of 28 May 2025
L.a-05	6 May 2025	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue L.a, as of 6 May 2025
L.a-04	1 April 2025	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue L.a, as of 1 April 2025
L.a-03	6 March 2025	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue L.a, as of 6 March 2025
L.a-02	7 February 2025	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue L.a, as of 7 February 2025
L.a-01	7 January 2025	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue L.a, as of 7 January 2025
L.a-00	2 December 2024	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue L.a, as of 2 December 2024
K.a-08	30 November 2024	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue K.a, as of 30 November 2024
J.a-08	4 March 2024	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue J.a, as of 4 March 2024
I.a-06	21 April 2023	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue I.a, as of 31 March 2023

Issue	Date	Confidentiality	Change
H.a-06	22 July 2022	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue H.a, as of 22 July 2022
G.b-05	31 January 2022	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue G.b, as of 7 January 2022
F.c-04	18 December 2020	Non-Confidential	Known Issues in Arm® Architecture Reference Manual, Issue F.c, as of 18 December 2020

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.

Arm Limited. Company 02557590 registered in England.

110 Fulbourn Road, Cambridge, England CB1 9NJ.

PRE-1121-V1.0

Confidentiality Status

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

Unrestricted Access is an Arm internal classification.

Product Status

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

Feedback

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

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

Inclusive language commitment

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

We believe that this document contains no offensive language. To report offensive language in this document, email terms@arm.com.

Contents

1. Introduction.....	11
1.1 Conventions.....	11
1.2 Useful resources.....	12
1.3 Other information.....	13
2. Known issues.....	14
2.1 D17119.....	14
2.2 D17268.....	14
2.3 C17296.....	17
2.4 C17536.....	18
2.5 D18847.....	20
2.6 D19403.....	21
2.7 D19798.....	21
2.8 C19822.....	22
2.9 D20234.....	23
2.10 C20429.....	24
2.11 D20732.....	24
2.12 D20796.....	25
2.13 D21388.....	27
2.14 D22004.....	28
2.15 D22031.....	29
2.16 E22039.....	30
2.17 D22105.....	31
2.18 D22207.....	35
2.19 C22264.....	37
2.20 C22369.....	38
2.21 C22430.....	39
2.22 D22446.....	41
2.23 C22566.....	41
2.24 C22584.....	41
2.25 D22595.....	42
2.26 E22689.....	42

2.27 D22719.....	44
2.28 D22806.....	47
2.29 R22882.....	48
2.30 D22979.....	48
2.31 D23002.....	49
2.32 R23026.....	49
2.33 D23122.....	51
2.34 R23144.....	53
2.35 D23154.....	54
2.36 D23168.....	55
2.37 C23225.....	56
2.38 D23239.....	56
2.39 C23271.....	57
2.40 D23403.....	58
2.41 D23404.....	59
2.42 C23431.....	60
2.43 D23442.....	61
2.44 D23462.....	62
2.45 C23479.....	63
2.46 D23518.....	70
2.47 D23543.....	71
2.48 R23545.....	72
2.49 D23551.....	73
2.50 C23552.....	74
2.51 C23586.....	75
2.52 C23587.....	79
2.53 D23647.....	80
2.54 D23659.....	80
2.55 D23680.....	82
2.56 D23683.....	83
2.57 D23687.....	84
2.58 D23689.....	84
2.59 C23701.....	85
2.60 D23713.....	86
2.61 C23721.....	88
2.62 D23733.....	90

2.63 D23744.....	91
2.64 C23753.....	92
2.65 C23766.....	93
2.66 D23768.....	94
2.67 D23772.....	95
2.68 C23775.....	102
2.69 C23782.....	104
2.70 D23794.....	107
2.71 D23796.....	108
2.72 D23797.....	108
2.73 D23807.....	110
2.74 C23808.....	110
2.75 D23812.....	111
2.76 E23814.....	111
2.77 R23822.....	111
2.78 D23823.....	112
2.79 C23830.....	113
2.80 C23842.....	114
2.81 C23844.....	116
2.82 C23845.....	116
2.83 R23848.....	117
2.84 D23852.....	118
2.85 D23854.....	118
2.86 C23859.....	119
2.87 C23866.....	119
2.88 D23869.....	120
2.89 C23900.....	120
2.90 R23907.....	122
2.91 R23909.....	123
2.92 R23917.....	124
2.93 R23928.....	124
2.94 D23933.....	127
2.95 D23934.....	130
2.96 C23935.....	131
2.97 D23936.....	131
2.98 D23947.....	132

2.99 C23956.....	132
2.100 D23962.....	133
2.101 D23966.....	134
2.102 C23967.....	135
2.103 C23982.....	138
2.104 D23988.....	139
2.105 D23995.....	139
2.106 C23997.....	139
2.107 D24015.....	141
2.108 D24019.....	142
2.109 D24020.....	142
2.110 C24031.....	143
2.111 C24054.....	143
2.112 D24057.....	144
2.113 C24061.....	144
2.114 D24062.....	145
2.115 C24064.....	146
2.116 C24066.....	146
2.117 D24086.....	147
2.118 D24093.....	148
2.119 C24096.....	148
2.120 D24098.....	149
2.121 D24099.....	149
2.122 D24101.....	150
2.123 D24112.....	150
2.124 C24124.....	151
2.125 D24129.....	153
2.126 D24146.....	153
2.127 C24218.....	155
2.128 D24149.....	155
2.129 C24151.....	156
2.130 D24153.....	156
2.131 D24156.....	156
2.132 C24161.....	157
2.133 D24162.....	157
2.134 D24166.....	158

2.135 R24169.....	158
2.136 C24177.....	159
2.137 D24180.....	160
2.138 C24186.....	162
2.139 C24190.....	163
2.140 C24218.....	163
2.141 D24224.....	163
2.142 C24236.....	164
2.143 C24254.....	165
2.144 D24257.....	166
2.145 R24624.....	166
2.146 D24274.....	166
2.147 D24282.....	167
2.148 D24321.....	168
2.149 C24337.....	168
2.150 D24341.....	168
2.151 D24352.....	169

1. Introduction

1.1 Conventions

The following subsections describe conventions used in Arm documents.

Glossary

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

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

Typographic conventions

Arm documentation uses typographical conventions to convey specific meaning.

Convention	Use
italic	Citations.
bold	Interface elements, such as menu names. Terms in descriptive lists, where appropriate.
monospace	Text that you can enter at the keyboard, such as commands, file and program names, and source code.
monospace <u>underline</u>	A permitted abbreviation for a command or option. You can enter the underlined text instead of the full command or option name.
<and>	Encloses replaceable terms for assembler syntax where they appear in code or code fragments. For example: <pre>MRC p15, 0, <Rd>, <CRn>, <CRm>, <Opcode_2></pre>
SMALL CAPITALS	Terms that have specific technical meanings as defined in the <i>Arm® Glossary</i> . For example, IMPLEMENTATION DEFINED , IMPLEMENTATION SPECIFIC , UNKNOWN , and UNPREDICTABLE .



We recommend the following. If you do not follow these recommendations your system might not work.



Your system requires the following. If you do not follow these requirements your system will not work.



You are at risk of causing permanent damage to your system or your equipment, or harming yourself.



This information is important and needs your attention.



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



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

1.2 Useful resources

This document contains information that is specific to this product. See the following resources for other useful information.

Arm documents are available on developer.arm.com/documentation.

Confidential documents are only available to licensees, when logged in. Each document link in the tables below provides direct access to the online version of the document.

Arm product resources	Document ID	Confidentiality
Arm® Architecture Reference Manual for A-profile architecture, Issue L.b	DDI 0487L.b	Non-Confidential

1.3 Other information

See the Arm website for other relevant information.

- [Arm® Developer](#).
- [Arm® Documentation](#).
- [Technical Support](#).
- [Arm® Glossary](#).

2. Known issues

This document records known issues in the Arm Architecture Reference Manual for A-profile architecture (DDI 0487), Issue L.b.

Key

- C = Clarification.
- D = Defect.
- R = Relaxation.
- E = Enhancement.

2.1 D17119

In sections F3.1.10.2 “Advanced SIMD two registers and shift amount” and F4.1.22.2 “Advanced SIMD two registers and shift amount”, the following constraints are added to VMOVL:

- 'L' must be '0'.
- 'imm3H' cannot be '000'.

2.2 D17268

In section D24.2.182 “TCR_EL1, Translation Control Register (EL1)”, in the description for field “IPS, bits [34:32]”, the text that reads:

The value 0b110 represents the following output address sizes:

- For the 64KB translation granule size, if FEAT_LPA is implemented, the value 0b110 represents 52 bits.
- For the 4KB and 16KB granule sizes, only if both FEAT_LPA and the effective value of TCR_EL1.DS is 0b1, the value 0b110 represents 52 bits.
- Otherwise, the value 0b110 behaves as the 0b101 value and represents 48 bits.

If the value of ID_AA64MMFRO_EL1.PARange is 0b0111, and the value of this field is not 0b111 or a value treated as 0b111, then bits[55:52] of every translation table base address are 0b0000 for the stage of translation controlled by TCR_EL1.

is changed to read:

The values 0b110 and 0b111 represent different output address sizes depending on implementation choices and translation configuration.

The following table captures the output address size represented by the value 0b110:

Descriptor Format	ID_AA64MMFR0_EL1.PARange	Translation Granule	TCR_EL1.DS	Represented OA size
Any	0b0101	Any	Any	48 bits, 256TB
VMSAv8-64	0b011x	4KB or 16KB	0	48 bits, 256TB
VMSAv8-64	0b011x	4KB or 16KB	1	52 bits, 4PB
VMSAv8-64	0b011x	64KB	N/A	52 bits, 4PB
VMSAv9-128	0b011x	Any	N/A	52 bits, 4PB

If 52-bit PA is supported, and the translation table descriptors cannot express an OA larger than 48 bits, then bits[51:48] of every translation table base address are treated as 0b0000 for the stage of translation controlled by TCR_EL1.

The following table captures the output address size represented by the value 0b111:

Descriptor Format	ID_AA64MMFR0_EL1.PARange	Translation Granule	TCR_EL1.DS	Represented OA size
Any	0b0110	Any	Any	OA size represented by 0b110
VMSAv8-64	0b0111	Any	Any	OA size represented by 0b110
VMSAv9-128	0b0111	Any	N/A	56 bits, 64PB

If 56-bit PA is supported, and the translation table descriptors cannot express an OA larger than 52 bits, then bits[55:52] of every translation table base address are treated as 0b0000 for the stage of translation controlled by TCR_EL1.

If the output address size represented by this field is larger than the supported PA size expressed in ID_AA64MMFR0_EL1.PARange, then the output address size is treated as being the same as the supported PA size. Arm strongly recommends that software avoids configuring this field to a value representing an output address size larger than the supported PA size.

The equivalent changes are made in the following sections:

- D24.2.183 “TCR_EL2, Translation Control Register (EL2)”.
- D24.2.184 “TCR_EL3, Translation Control Register (EL3)”.
- D24.2.210 “VTCR_EL2, Virtualization Translation Control Register”.

In section D24.2.195 “TTBRO_EL1, Translation Table Base Register 0 (EL1)”, under the heading “When FEAT_D128 is not implemented or TCR2_EL1.D128 == 0:”, in field “BADDR[47:1], bits [47:1]”, the text that reads:

The BADDR field represents a 52-bit address if one of the following applies:

- FEAT_LPA is implemented, the 64KB granule size is in use, and the value of TCR_EL1.IPS is 0b110.
- FEAT_LPA2 is implemented, the 4KB or 16KB granule size is in use, and the Effective value of TCR_EL1.DS is 1.
- FEAT_D128 is implemented, 56-bit PAs are supported, the 64KB granule size is in use, and the value of TCR2_EL1.D128 is 0.

is changed to read:

The BADDR field represents a 52-bit address if the value of TCR_EL1.IPS represents an output address size of 52 bits.

If 52-bit PA is supported, and BADDR field does not represent a 52-bit address, then bits[51:48] of the base address are treated as 0b0000.

Additionally, the note in the same field that reads:

For the 64KB granule, if FEAT_LPA is not implemented, and the value of TCR_EL1.IPS is 0b110, one the following **IMPLEMENTATION DEFINED** behaviors occur:

is changed to read:

For the 64KB granule, if 52-bit PA is not supported, and the value of TCR_EL1.IPS is 0b110 or 0b111, one of the following **IMPLEMENTATION DEFINED** behaviors occur:

In section D8.1.7 “Output address size configuration”, under the rule R_{CLPQT} , the text that reads:

The OA size from a translation stage cannot be larger than the implemented PA size. If the value of TCR_ELx.{}PS is larger than the implemented PA size, then {}PS is treated as being the same as the implemented PA size.

is changed to read:

The OA size from a translation stage cannot be larger than the implemented PA size. If the OA size represented by TCR_ELx.{}PS is larger than the implemented PA size, then the OA size is treated as being the same as the implemented PA size.

Under the rule I_{HGDWP} , the text that reads:

If a translation stage is enabled and the translation stage OA size is larger than the implemented PA size, then an Address size fault is generated at the lookup level in the translation stage that generated the OA.

is changed to read:

If a translation stage is enabled and the translation stage OA is larger than the implemented PA size, then an Address size fault is generated at the lookup level in the translation stage that generated the OA.

Under the rule R_{BZHGM} , the text that reads:

If all of the following apply, then a stage 2 Translation fault is generated:

- Two address translation stages are used.
- Stage 2 address translation is enabled.
- The stage 1 OA size does not generate a stage 1 Address size fault.
- The stage 1 OA size is larger than the specified stage 2 translation IA size.

is changed to read:

If all of the following apply, then a stage 2 Translation fault is generated:

- Two address translation stages are used.
- Stage 2 address translation is enabled.
- The stage 1 OA does not generate a stage 1 Address size fault.
- The stage 1 OA is larger than the specified stage 2 translation IA size.

2.3 C17296

In section D14.2 “The PMU event number space and common events”, the text that reads:

0x0005, L1D_TLB_REFILL, Level 1 data TLB refill

The counter does not count the access if any of the following are true:

- The access misses in the TLB and generates a translation table walk but does not cause a refill of the TLB.
- The access is due to a TLB maintenance instruction.
- The access generates a Translation fault because the applicable TCR_ELx.EPDy bit is 1.
- FEAT_EOPD is implemented and the access is an unprivileged access that generates a Translation fault because the applicable TCR_ELx.EOPDy bit is 1.
- FEAT_SVE is implemented and the access is a non-fault access that fails because the applicable TCR_ELx.NFDy bit is 1.

It is **IMPLEMENTATION DEFINED** whether the counter counts the access if any of the following are true:

- The refill is not allocated in the TLB.
- The access generates a Translation fault for any other reason.

is changed to read:

0x0005, L1D_TLB_REFILL, Level 1 data TLB refill

The counter does not count the access if any of the following are true:

- The access is due to a TLB maintenance instruction.
- The access generates a Translation fault because the applicable TCR_ELx.EPDy bit is 1.
- FEAT_EOPD is implemented and the access is an unprivileged access that generates a Translation fault because the applicable TCR_ELx.EOPDy bit is 1.
- FEAT_SVE is implemented and the access is a non-fault access that fails because the applicable TCR_ELx.NFDy bit is 1.

It is **IMPLEMENTATION DEFINED** whether the counter counts the access if any of the following are true:

- The access generates a Translation fault for any other reason.
- The access misses in the TLB and generates a translation table walk, but the result is not allocated into the TLB for any reason other than a Translation fault.

The equivalent changes are made in the following events: L1I_TLB_REFILL, L2D_TLB_REFILL, and L2I_TLB_REFILL.

2.4 C17536

In section D14.3.2 “Common microarchitectural events”, the description of DP_SPEC that reads:

0x0073, DP_SPEC, Operation speculatively executed, integer data processing

The counter counts each operation counted by INST_SPEC that is an integer data processing operation.

Operations due to the following instructions are counted as integer data-processing operation:

- In AArch64 state instructions from the following sections:
 - “Data processing - immediate”
 - “Data processing - register”
 - “System register instructions”
 - “System instructions” other than Memory-writing instructions
 - “Hint instructions”
 - When FEAT_SVE is implemented and the SVE_SPEC event is implemented, non-SIMD SVE instructions.
- In AArch32 state instructions from the following sections:
 - “Data-processing instructions”.
 - “PSTATE and banked register access instructions”.
 - “Banked register access instructions”.
 - “Miscellaneous instructions other than ISB and prefetches”.
 - “System register access instructions other than LDC and STC instructions”.

This includes MOV and MVN instructions.

It is **IMPLEMENTATION DEFINED** whether the preload instructions PRDM, PLD, PLDW, and PLI count as integer data-processing operations or load operations. Arm recommends that if the instructions are not implemented as a **NOP** then it is counted as a load operation.

When FEAT_PMUv3p9 is not implemented, it is **IMPLEMENTATION DEFINED** whether ISB is counted as an integer data-processing operation or a Software change of the PC.

When FEAT_PMUv3p9 is implemented, ISB is counted as an integer data-processing operation.

It is **IMPLEMENTATION DEFINED** whether the following instructions are counted as integer data-processing operations, SIMD operations, or floating-point operations, but Arm recommends that the instructions are counted as integer data-processing operations:

- In AArch64 state:
 - Instructions from Floating-point move (register) that transfers data between a general-purpose register and a SIMD&FP register without conversion: FMOV (general).
 - Instructions from “SIMD Move” that transfers data between a general-purpose register and an element or elements in a SIMD&FP register: DUP (general), SMOV, UMOV, and INS (general). This includes the aliases MOV (from general) and MOV (to general).
 - When FEAT_SVE is implemented and the SVE_SPEC event is not implemented, non-SIMD SVE instructions.
- In AArch32 state:
 - VDUP (general-purpose register).
 - All VMOV instructions that transfer data between a general-purpose register and a SIMD&FP register.
 - VMRS and VMSR.

When FEAT_PMUv3p8 is not implemented, this is an **IMPLEMENTATION DEFINED** event.

is changed to read:

0x0073, DP_SPEC, Operation speculatively executed, integer data processing

The counter counts each operation counted by INST_SPEC and not counted as any of:

- A load or store operation, counted by LDST_SPEC.
- A Software change of the PC operation, counted by PC_WRITE_SPEC.
- A scalar floating-point data processing operation, counted by VFP_SPEC.
- An Advanced SIMD, SVE, or SME data processing operation, counted by SE_SPEC.
- A cryptographic operation, counted by CRYPTO_SPEC.

This includes operations due to the following instructions, which are counted as integer data-processing operations

- In AArch64 state instructions from the following sections:
 - “Data processing - immediate”
 - “Data processing - register”
 - “System register instructions”
 - “Instructions with register argument”
 - “System instructions” other than Memory-writing instructions
 - “Hint instructions”
 - When FEAT_SVE is implemented and the SVE_SPEC event is implemented, non-SIMD SVE instructions.

- In AArch32 state instructions from the following sections:
 - “Data-processing instructions”.
 - “PSTATE and banked register access instructions”.
 - “Banked register access instructions”.
 - “Miscellaneous instructions other than ISB and prefetches”.
 - “System register access instructions other than LDC and STC instructions”.

When FEAT_PMUv3p8 is not implemented, this is an **IMPLEMENTATION DEFINED** event.

2.5 D18847

In section J1.1.3 “aarch64/functions”, in the pseudocode function AArch64.MemSingleRead(), the code segment that reads:

```
(bits(size*8), AddressDescriptor, PhysMemRetStatus) AArch64.MemSingleRead(bits(64)
address,
                                                    integer
size,
AccessDescriptor accdesc_in,
                                                    boolean
aligned)
    assert size IN {1, 2, 4, 8, 16};
    bits(size*8) value = bits(size*8) UNKNOWN;
    PhysMemRetStatus memstatus = PhysMemRetStatus UNKNOWN;
    AccessDescriptor accdesc = accdesc_in;
    if IsFeatureImplemented(FEAT_LSE2) then
        assert AllInAlignedQuantity(address, size, 16);
    else
        assert IsAligned(address, size);
        if IsFeatureImplemented(FEAT_MTE2) && accdesc.tagchecked then
            accdesc.tagchecked = AArch64.AccessIsTagChecked(address, accdesc);
        AddressDescriptor memaddrdesc;
        memaddrdesc = AArch64.TranslateAddress(address, accdesc, aligned, size);
        if IsFault(memaddrdesc) then
            return (value, memaddrdesc, memstatus);
```

is changed to read:

```
(bits(size*8), AddressDescriptor, PhysMemRetStatus) AArch64.MemSingleRead(bits(64)
address,
                                                    integer
size,
AccessDescriptor accdesc_in,
                                                    boolean
aligned)
    assert size IN {1, 2, 4, 8, 16};
    bits(size*8) value = bits(size*8) UNKNOWN;
    PhysMemRetStatus memstatus = PhysMemRetStatus UNKNOWN;
    AccessDescriptor accdesc = accdesc_in;
    if IsFeatureImplemented(FEAT_LSE2) then
        assert AllInAlignedQuantity(address, size, 16);
    else
        assert IsAligned(address, size);
        if IsFeatureImplemented(FEAT_MTE2) && accdesc.tagchecked then
```

```
    accdesc.tagchecked = AArch64.AccessIsTagChecked(address, accdesc);  
    AddressDescriptor memaddrdesc;  
    memaddrdesc = AArch64.TranslateAddress(address, accdesc, aligned, size);  
    if !IsFault(memstatus) && accdesc.acctype == AccessType_IFETCH then  
        memaddrdesc.fault = AArch64.CheckDebug(memaddrdesc.fault.vaddress, accdesc,  
size);  
        if IsFault(memaddrdesc) then  
            return (value, memaddrdesc, memstatus);
```

2.6 D19403

In section D2.12 “Synchronization and debug exceptions” and section G2.11 “Synchronization and debug exceptions”, the following text is deleted:

Some register updates are self-synchronizing. Others require an explicit Context Synchronization event.

In section D24.10.20 “CNTPCTSS_ELO, Counter-timer Self-Synchronized Physical Count Register”, the text that reads:

This register is a self-synchronised view of the CNTPCT_ELO counter, and cannot be read speculatively.

is changed to read:

This register is a view of the CNTPCT_ELO register for which reads appear to occur in program order relative to other instructions, without the need for any explicit synchronization. Reads of this register return a value consistent with the counter not being read until the read instruction is known to be non-speculative.

Equivalent changes are made in the following sections:

- D24.10.29 “CNTVCTSS_ELO, Counter-timer Self-Synchronized Virtual Count Register”.
- G8.7.20 “CNTPCTSS, Counter-timer Self-Synchronized Physical Count register”.
- G8.7.25 “CNTVCTSS, Counter-timer Self-Synchronized Virtual Count register”.

2.7 D19798

In section G8.2.26 “CFPRCTX, Control Flow Prediction Restriction by Context”, in the accessibility for EL1, the following code is deleted:

```
    elsif EffectiveHCR_EL2_NVx() IN {'xx1'} then  
        AArch64.SystemAccessTrap(EL2, 0x03);
```

Equivalent changes are made to:

- G8.2.29 “COSPRCTX, Clear Other Speculative Prediction Restriction by Context”.
- G8.2.34 “CPPRCTX, Cache Prefetch Prediction Restriction by Context”.

- G8.2.51 “DVPRCTX, Data Value Prediction Restriction by Context”.

2.8 C19822

In sections D24.2.40 “ESR_EL1, Exception Syndrome Register (EL1)”, D24.2.41 “ESR_EL2, Exception Syndrome Register (EL2)” and D24.2.42 “ESR_EL3, Exception Syndrome Register (EL3)”, under the heading “ISS encoding for an exception from the Memory Copy and Memory Set instructions”, the text that reads:

WrongOption, bit [17]

Algorithm option.

0b0	WrongOption is false.
0b1	WrongOption is true.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

OptionA, bit [16]

Algorithm type indicated by the PSTATE.C bit.

0b0	OptionB indicated by PSTATE.C is 0.
0b1	OptionA indicated by PSTATE.C is 1.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

is changed to read:

FormatOption, bit [17:16]

Reports the Option used to encode the initial Xs, Xd and Xn register values provided to the instruction that generated the exception.

FormatOption	Meaning
0b00	Option B
0b01	Option A
0b10	Option A
0b11	Option B

For more information, see D1.3.5.7 “Memory Copy and Memory Set exceptions”.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

Note:

This field was previously presented as two separate bits, WrongOption (bit [17]) and OptionA (bit [16]), which were already expected to be used together and not as individual bits.

In section D1.3.5.7 Memory Copy and Memory Set exceptions, under rule I_{CNTMJ}, the code that reads:

```
if OptionA && WrongOption || OptionB && !WrongOption then
    // format is from option B
    ...
```

is changed to read:

```
if FormatOption == 0b00 || FormatOption == 0b11 then
    // format is from option B
    ...
```

In the same section, under rule I_{MWFQH}, the code that reads:

```
if !OptionA && WrongOption || (OptionA && !WrongOption) then
    // format is from option A
    ...
```

is changed to read:

```
if FormatOption == 0b01 || FormatOption == 0b10 then
    // format is from option A
    ...
```

2.9 D20234

In section D14.2 “The PMU event number space and common events”, table D14-2 “Allocation of the PMU event number space”, the text that reads:

Event numbers	Allocation
...	...
0x8000 - 0x80FF	Common architectural and microarchitectural events.
0x8100 - 0x81FF	From Armv8.6, common architectural and microarchitectural events.
0x8200 - 0xC0BF	Reserved.
...	...

is changed to read:

Event numbers	Allocation
...	...

Event numbers	Allocation
0x8000 - 0x8FFF	Common architectural and microarchitectural events.
0x9000 - 0xC0BF	Reserved.
...	...

2.10 C20429

In the section D17.8.2 “Buffer Full event”, the text that reads:

If, after writing a sample record, there is not sufficient space in the Profiling Buffer for a sample record of the size indicated by PMSIDR_EL1.MaxSize, and PMBSR_EL1.S is 0, a Profiling Buffer management event is generated:

- ...
- PMBPTR_EL1 is set to the first byte after the last complete sample record. PMBSR_EL1.DL is unchanged
- ...

is changed to read:

If, after writing a sample record, there is not sufficient space in the Profiling Buffer for a sample record of the size indicated by PMSIDR_EL1.MaxSize, and PMBSR_EL1.S is 0, a Profiling Buffer management event is generated:

- ...
- PMBPTR_EL1 is set to the first byte after the last complete sample record, which can include any padding that the implementation uses to force alignment. PMBSR_EL1.DL is unchanged.
- ...

Equivalent changes are made in the following sections: - D17.8.3 “Faults and watchpoints”. - D17.8 “Profiling Buffer management”.

2.11 D20732

In section G3.2.1 “Controls to prohibit trace at Exception levels”, the text that reads:

If SelfHostedTraceEnabled() == TRUE, no events are exported to the trace unit when tracing is prohibited.

If SelfHostedTraceEnabled() == FALSE, no events are exported to the trace unit when the PE is in Secure state and counting in Secure state is prohibited.

If FEAT_ETE is not implemented, when PMCR_ELO.X==0 or PMCR.X==0, no PMU events are exported to the trace unit.

Otherwise, PMU events are exported to the trace unit.

is changed to read:

If FEAT_ETE is not implemented, when `SelfHostedTraceEnabled() == TRUE`, no events are exported to the trace unit when tracing is prohibited.

If FEAT_ETE is not implemented, when `SelfHostedTraceEnabled() == FALSE`, no events are exported to the trace unit when the PE is in Secure state and counting in Secure state is prohibited.

If FEAT_ETE is not implemented, when `PMCR_ELO.X==0` or `PMCR.X==0`, no PMU events are exported to the trace unit.

If FEAT_ETE is implemented, see External inputs for the rules governing export of events to the trace unit.

Otherwise, PMU events are exported to the trace unit.

2.12 D20796

In section B2.6.9 Data Synchronization Barrier, the text that reads:

A DSB instruction executed by a PE, PEE, completes when all of the following apply:

- All explicit memory effects of the required access types appearing in program order before the DSB are complete for the set of observers in the required shareability domain.
- If the required access types of the DSB is reads, or reads and writes, all implicit tag memory read effects appearing in program order before the DSB are complete.
- If the required access types of the DSB is reads and writes, the following instructions issued by PEE before the DSB are complete for the required shareability domain:
 - All cache maintenance instructions.
 - All TLB maintenance instructions.
 - All PSB instructions.
- When FEAT_XS is implemented, if the required access types of the DSB is reads and writes, completion of the DSB instruction with the nXS qualifier executed by a PE, PEE, ensures that:
 - All previous TLBInXS maintenance operations generated by AArch64 TLB maintenance instructions with the nXS qualifier executed by PEE are finished for all PEs in the shareability domain of the DSB instruction.
 - All previous TLBInXS maintenance operations generated by AArch32 or AArch64 TLB maintenance instructions executed at EL1 by PEE when `HCRX_EL2.FnXS` is 1 are finished for all PEs in the shareability domain of the DSB instruction.
- Completion of the DSB instruction with the nXS qualifier executed by a PE, PEE, does not ensure that:

- All previous TLB maintenance operations generated by AArch32 or AArch64 TLB maintenance instructions executed at EL1 by PEe when HCRX_EL2.FnXS is 0 are finished for all PEs in the shareability domain of the DSB instruction.
- All previous TLB maintenance operations generated by AArch32 or AArch64 TLB maintenance instructions executed at EL2 or EL3 by PEe are finished for all PEs in the shareability domain of the DSB instruction.

is changed to read:

A DSB instruction executed by a PE completes when all of the following are complete for the set of observers in the required shareability domain:

- If the required access types of the DSB is reads, all of the following:
 - All Explicit Memory Read Effects generated by instructions in program order before the DSB.
 - If FEAT_GCS is implemented, all GCS Memory Read Effects that appear in program order before a GCSB instruction that appears in program order before the DSB.
 - All Implicit Tag Memory Read Effects generated by instructions in program order before the DSB.
- If the required access types of the DSB is writes, all of the following:
 - All Explicit Memory Write Effects generated by instructions in program order before the DSB.
 - If FEAT_GCS is implemented, all GCS Memory Write Effects that appear in program order before a GCSB instruction that appears in program order before the DSB.
- If the required access types of the DSB is reads and writes, all of the following:
 - All Explicit Memory Effects generated by instructions in program order before the DSB.
 - All Implicit Tag Memory Read Effects generated by instructions in program order before the DSB.
 - All Implicit TTD Memory Effects generated by instructions in program order before the DSB.
 - All Implicit Instruction Memory Read Effects generated by instructions in program order before the DSB.
 - All cache maintenance instructions in program order before the DSB.
 - If FEAT_SPMU is implemented, all Direct System Register Write Effects to a System PMU register generated by instructions in program order before the DSB.
 - If FEAT_SPE is implemented, all PSB CSYNC instructions in program order before the DSB.
 - If FEAT_XS is implemented and the DSB has the nXS qualifier, all of the following:
 - All TLBInXS maintenance operations generated by AArch64 TLB maintenance instructions with the nXS qualifier in program order before the DSB.
 - If HCRX_EL2.FnXS is 1, all TLBInXS maintenance operations generated by AArch32 or AArch64 TLB maintenance instructions executed at EL1 in program order before the DSB.

- If FEAT_XS is not implemented or the DSB does not have the nXS qualifier, all of the following:
 - All TLBI maintenance operations generated by TLB maintenance instructions in program order before the DSB.
- If FEAT_HDBSS is implemented and the DSB is executed at EL2 or EL3, all HDBSS Memory Write Effects generated by instructions in program order before the DSB.
- If FEAT_TRBE is implemented, all TSB CSYNC instructions in program order before the DSB.
- All CFP RCTX, COSP RCTX, CPP RCTX, DVP RCTX instructions in program order before the DSB.
- If FEAT_MTE_ASYNC is implemented, all Indirect System Register Write Effects to the Tag Fault Status Register accessible at ELx generated by instructions in program order before the DSB.
- If FEAT_GCS is implemented, all GCS Memory Effects that appear in program order before a GCSB instruction that appears in program order before the DSB.

2.13 D21388

In section D24.2.40 “ESR_EL1, Exception Syndrome Register (EL1)”, the text that reads:

If the Armv8-A architecture is implemented with an ETMv4 implementation, MCRR and MRRC accesses to trace registers are **UNDEFINED** and the resulting exception is higher priority than an exception due to these traps.

is changed to read:

If FEAT_ETMv4 or FEAT_ETE are implemented, MCRR and MRRC accesses to trace registers are **UNDEFINED** and the resulting exception is higher priority than an exception due to these traps.

Equivalent changes are made to:

- D24.2.41 “ESR_EL2, Exception Syndrome Register (EL2)”.
- D24.2.42 “ESR_EL3, Exception Syndrome Register (EL3)”.

In D24.2.34 “CPTR_EL2, Architectural Feature Trap Register (EL2)” in the description of the TTA field, the text that reads:

The ETMv4 architecture does not permit ELO to access the trace registers. If the trace unit implements FEAT_ETMv4, ELO accesses to the trace registers are **UNDEFINED**, and any resulting exception is higher priority than an exception that would be generated because the value of CPTR_EL2.TTA is 1.

is changed to read:

FEAT_ETMv4 and FEAT_ETE do not permit ELO to access the trace registers. ELO accesses to the trace System registers are **UNDEFINED**, and any resulting exception is higher priority than an exception that would be generated because the value of CPTR_EL2.TTA is 1.

2.14 D22004

The following text is deleted from section “Glossary”:

Explicit memory effect

A read from memory, or a write to memory, generated by a load or store instruction executed by the PE. Reads and writes generated by hardware translation table accesses, as well as instruction fetches and SPE writes to the Profiling Buffer, are not explicit memory effects.

The following text is added to section B2.3.1 “Basic definitions”:

Explicit Memory Effect

A Memory Read Effect or a Memory Write Effect that is generated by one of the following:

- a load instruction, store instruction, atomic instruction, Tag and Data store instruction, DC ZVA, DC GZVA, or STZGM; to access the Memory Location(s) addressed by any of the following from the instruction:
 - the Xn|SP operand, plus any applicable offset
 - the PC plus the offset specified by the <label> argument
- a Tag load instruction, Tag store instruction, Tag and data store instruction, Bulk Allocation tag access instruction, DC GZVA, or DC GVA; to access the Tag Location(s) addressed by the Xn|SP operand, plus any applicable offset, of the instruction.
- a Memory Copy or Memory Set instruction; to access the Memory Location(s) addressed by the Xs or Xd operand of the instruction, including all of the bytes within the range specified by the Xn operand of the instruction.
- a Memory Set with tag setting instruction; to access the Tag Location(s) addressed by the Xs or Xd operand of the instruction, including all of the Allocation tags within the range specified by the Xn operand of the instruction.
- an MRS or MSR instruction that is transformed into a Memory Effect due to FEAT_NV2; to access the Memory Locations addressed by VNCR_EL2 plus the offset for the specified register.

Note: Prefetch memory instructions do not generate any Explicit Memory Effects.

2.15 D22031

In section C6.2.287 “PRFUM”, under the “Assembler Symbols” heading, the table for the “<prfop>” symbol that reads:

Rt	<prfop>
00000	PLDL1KEEP
00001	PLDL1STRM
00010	PLDL2KEEP
00011	PLDL2STRM
00100	PLDL3KEEP
00101	PLDL3STRM
01000	PLIL1KEEP
01001	PLIL1STRM
01010	PLIL2KEEP
01011	PLIL2STRM
01100	PLIL3KEEP
01101	PLIL3STRM
10000	PSTL1KEEP
10001	PSTL1STRM
10010	PSTL2KEEP
10011	PSTL2STRM
10100	PSTL3KEEP
10101	PSTL3STRM

is changed to read:

Rt	<prfop>	Architectural Feature
00000	PLDL1KEEP	-
00001	PLDL1STRM	-
00010	PLDL2KEEP	-
00011	PLDL2STRM	-
00100	PLDL3KEEP	-
00101	PLDL3STRM	-
00110	PLDSLCKEEP	FEAT_PRFMSLC
00111	PLDSLCSTRM	FEAT_PRFMSLC
01000	PLIL1KEEP	-
01001	PLIL1STRM	-
01010	PLIL2KEEP	-
01011	PLIL2STRM	-
01100	PLIL3KEEP	-
01101	PLIL3STRM	-

Rt	<prfop>	Architectural Feature
01110	PLISLCKEEP	FEAT_PRFM_SLC
01111	PLISLCSTRM	FEAT_PRFM_SLC
10000	PSTL1KEEP	-
10001	PSTL1STRM	-
10010	PSTL2KEEP	-
10011	PSTL2STRM	-
10100	PSTL3KEEP	-
10101	PSTL3STRM	-
10110	PSTSLCKEEP	FEAT_PRFM_SLC
10111	PSTSLCSTRM	FEAT_PRFM_SLC

In section D24.2.80 “ID_AA64ISAR2_EL1, AArch64 Instruction Set Attribute Register 2”, the field “PRFM_SLC, bits [43:40]” that reads:

Indicates whether the PRFM instructions support a system level cache option. The value of this field is an **IMPLEMENTATION DEFINED** choice of:

PRFM_SLC	Meaning
0b0000	The PRFM instructions do not support the SLC target.
0b0001	The PRFM instructions support the SLC target.

is changed to read:

Indicates whether the PRFM and PRFUM instructions support a system level cache option. The value of this field is an **IMPLEMENTATION DEFINED** choice of:

PRFM_SLC	Meaning
0b0000	The PRFM and PRFUM instructions do not support the SLC target.
0b0001	The PRFM and PRFUM instructions support the SLC target.

Equivalent changes made in section A2.2.10 “The Armv8.9 architecture extension”, under the heading “FEAT_PRFM_SLC, SLC target support for PRFM instructions”.

2.16 E22039

In section C6.2.80 “CPYFP, CPYFM, CPYFE”, the text that reads:

The memory copy performed by these instructions is in the forward direction only, so the instructions are suitable for a memory copy only where there is no overlap between the source and destination locations, or where the source address is greater than the destination address.

is changed to read:

The memory copy performed by these instructions is in the forward direction only, so the instructions are suitable for a memory copy only where there is no overlap between the source and destination locations, or where the source address is greater than or equal to the destination address.

2.17 D22105

In section D24.2.56 “HCR_EL2, Hypervisor Configuration Register”, in the field description of ‘IMO, bit [4]’, the text that reads:

IMO	Meaning
0b0	When executing at Exception levels below EL2, and EL2 is enabled in the current Security state: <ul style="list-style-type: none"> When the value of HCR_EL2.TGE is 0, Physical IRQ interrupts are not taken to EL2. When the value of HCR_EL2.TGE is 1, Physical IRQ interrupts are taken to EL2 unless they are routed to EL3. Virtual IRQ interrupts are disabled.
0b1	When executing at any Exception level, and EL2 is enabled in the current Security state: <ul style="list-style-type: none"> Physical IRQ interrupts are taken to EL2, unless they are routed to EL3. When the value of HCR_EL2.TGE is 0, then Virtual IRQ interrupts are enabled.

If EL2 is enabled in the current Security state and the value of HCR_EL2.TGE is 1:

- Regardless of the value of the FMO bit, physical FIQ Interrupts target EL2 unless they are routed to EL3.
- When the Effective value of HCR_EL2.E2H is not 1, the Effective value of this field is 1.
- When the Effective value of HCR_EL2.E2H is 1, the Effective value of this field is 0.

For more information, see ‘Asynchronous exception routing’.

is changed to read:

IMO	Meaning
0b0	When executing at Exception levels below EL3 and the Effective value of HCR_EL2.TGE is 0: <ul style="list-style-type: none"> Physical IRQ interrupts are not taken to EL2. Virtual IRQ interrupts are disabled.
0b1	When executing at Exception levels below EL3, EL2 is enabled in the current Security state, and the Effective value of HCR_EL2.TGE is 0: <ul style="list-style-type: none"> Physical IRQ exceptions are taken to EL2, unless they are routed to EL3. Virtual IRQ interrupts are enabled.

When executing at EL3, the Effective value of HCR_EL2.IMO has no impact on the behavior of the PE.

When executing at Exception levels below EL3, and EL2 is not enabled in the current Security state, the Effective value of HCR_EL2.IMO is 0.

When the Effective value of HCR_EL2.TGE is 1, regardless of the value of the IMO bit, all of the following are true:

- Physical IRQ Interrupts target EL2 unless they are routed to EL3.
- Virtual IRQ interrupts are disabled.

When executing at EL2 and the Effective value of HCR_EL2.{E2H, TGE} is {1, 0}, it is **IMPLEMENTATION DEFINED** whether the Effective value of HCR_EL2.IMO is 1 or the value programmed.

For more information, see 'Asynchronous exception types'.

In the same section, in the field description of 'FMO, bit [3]', the text that reads:

FMO	Meaning
0b0	When executing at Exception levels below EL2, and EL2 is enabled in the current Security state: <ul style="list-style-type: none"> • When the value of HCR_EL2.TGE is 0, Physical FIQ interrupts are not taken to EL2. • When the value of HCR_EL2.TGE is 1, Physical FIQ interrupts are taken to EL2 unless they are routed to EL3. • Virtual FIQ interrupts are disabled.
0b1	When executing at any Exception level, and EL2 is enabled in the current Security state: <ul style="list-style-type: none"> • Physical FIQ interrupts are taken to EL2, unless they are routed to EL3. • When HCR_EL2.TGE is 0, then Virtual FIQ interrupts are enabled.

If EL2 is enabled in the current Security state and the value of HCR_EL2.TGE is 1:

- Regardless of the value of the FMO bit, physical FIQ Interrupts target EL2 unless they are routed to EL3.
- When the Effective value of HCR_EL2.E2H is not 1, the Effective value of this field is 1.
- When the Effective value of HCR_EL2.E2H is 1, the Effective value of this field is 0.

For more information, see 'Asynchronous exception routing'.

is changed to read:

FMO	Meaning
0b0	When executing at Exception levels below EL3 and the Effective value of HCR_EL2.TGE is 0: <ul style="list-style-type: none"> • Physical FIQ interrupts are not taken to EL2. • Virtual FIQ interrupts are disabled.
0b1	When executing at Exception levels below EL3, EL2 is enabled in the current Security state, and the Effective value of HCR_EL2.TGE is 0: <ul style="list-style-type: none"> • Physical FIQ exceptions are taken to EL2, unless they are routed to EL3. • Virtual FIQ interrupts are enabled.

When executing at EL3, the Effective value of HCR_EL2.FMO has no impact on the behavior of the PE.

When executing at Exception levels below EL3, and EL2 is not enabled in the current Security state, the Effective value of HCR_EL2.FMO is 0.

When the Effective value of HCR_EL2.TGE is 1, regardless of the value of the FMO bit, all of the following are true:

- Physical FIQ Interrupts target EL2 unless they are routed to EL3.
- Virtual FIQ interrupts are disabled.

When executing at EL2 and the Effective value of HCR_EL2.{E2H, TGE} is {1, 0}, it is **IMPLEMENTATION DEFINED** whether the Effective value of HCR_EL2.FMO is 1 or the value programmed.

For more information, see 'Asynchronous exception types'.

In the same section, in the field description of 'AMO, bit [5]', the text that reads:

AMO	Meaning
0b0	Physical SError exceptions are unaffected by this mechanism. That is, physical SError exceptions are not taken to EL2 unless routed to EL2 by another control. Virtual SError exceptions are not enabled by this mechanism.
0b1	When executing at any Exception level, and EL2 is enabled in the current Security state, all of the following apply: <ul style="list-style-type: none"> • Physical SError exceptions are taken to EL2, unless they are routed to EL3. • If FEAT_E3DSE is implemented then delegated SError exceptions enabled by SCR_EL3.DSE are taken to EL2. • If HCR_EL2.TGE is 0 then virtual SError exceptions are enabled.

If EL2 is enabled in the current Security state and the value of HCR_EL2.TGE is 1:

- Regardless of the value of HCR_EL2.AMO, physical SError exceptions target EL2 unless they are routed to EL3.
- If FEAT_E3DSE is implemented then regardless of the value of HCR_EL2.AMO, delegated SError exceptions target EL2.
- When the Effective value of HCR_EL2.E2H is not 1, the Effective value of this field is 1.
- When the Effective value of HCR_EL2.E2H is 1, the Effective value of this field is 0.

For more information, see 'Asynchronous exception routing'.

is changed to read:

AMO	Meaning
0b0	When executing at Exception levels below EL3 and the Effective value of HCR_EL2.TGE is 0: <ul style="list-style-type: none"> • The routing of physical SError exceptions is unaffected by this mechanism. • If FEAT_E3DSE is implemented then delegated SError exceptions enabled by SCR_EL3.DSE are not taken to EL2. • Virtual SError exceptions are not enabled by this mechanism.

AMO	Meaning
0b1	<p>When executing at Exception levels below EL3, EL2 is enabled in the current Security state, and the Effective value of HCR_EL2.TGE is 0:</p> <ul style="list-style-type: none"> Physical SError exceptions are taken to EL2, unless they are routed to EL3. If FEAT_E3DSE is implemented then delegated SError exceptions enabled by SCR_EL3.DSE are taken to EL2. Virtual SError exceptions are enabled.

When executing at EL3, the value of HCR_EL2.AMO has no impact on the behavior of the PE.

When executing at Exception levels below EL3, and EL2 is not enabled in the current Security state, the Effective value of HCR_EL2.AMO is 0.

When the Effective value of HCR_EL2.TGE is 1, regardless of the value of the AMO bit, all of the following are true:

- Regardless of the value of HCR_EL2.AMO, physical SError exceptions target EL2 unless they are routed to EL3.
- If FEAT_E3DSE is implemented then regardless of the value of HCR_EL2.AMO, delegated SError exceptions target EL2.
- Virtual SError exceptions are disabled.

When executing at EL2 and the Effective value of HCR_EL2.{E2H, TGE} is {1, 0}, it is **IMPLEMENTATION DEFINED** whether the Effective value of HCR_EL2.AMO is 1 or the value programmed.

For more information, see 'Asynchronous exception types'.

In section D24.2.165 "SCTLR2_EL2, System Control Register (EL2)", in the field description of 'NMEA, bit [2]', the text that reads:

NMEA	Meaning
0b0	SError exceptions are not taken at EL2 if PSTATE.A == 1, unless routed to a higher Exception level.
0b1	SError exceptions are taken at EL2 regardless of the value of PSTATE.A, unless routed to a higher Exception level.

is changed to read:

NMEA	Meaning
0b0	The masking of SError exceptions at EL2 is unaffected by this mechanism.
0b1	SError exceptions are taken at EL2 regardless of whether they are masked by any mechanism, unless they are routed to EL3.

For more information, see 'Asynchronous exception types'.

In section D1.3.6.1 "Virtual interrupts", the rule R_{BCXJB} that reads:

If HCR_EL2.TGE is 0, setting an HCR_EL2.{FMO, IMO} routing control bit to 1 enables the corresponding virtual interrupt. If HCR_EL2.TGE is 1, all virtual interrupts are disabled and the Effective values of HCR_EL2.{FMO, IMO} are 0.

is changed to read:

If HCR_EL2.TGE is 0, setting an HCR_EL2.{FMO, IMO} routing control bit to 1 enables the corresponding virtual interrupt.

If HCR_EL2.TGE is 1, virtual interrupts are disabled.

2.18 D22207

In section D6.3.5 “The owning translation regime”, the text that reads:

R_{RRCNN}

When any of the following is true, the translation of addresses generated by the Trace Buffer Unit is **CONSTRAINED UNPREDICTABLE**:

- The owning Security state is Secure and SCR_EL3.NS is 1.
- When FEAT_RME is not implemented and the owning Security state is Non-secure and SCR_EL3.NS is 0.
- When FEAT_RME is implemented and the owning Security state is Non-secure and SCR_EL3.{NSE,NS} is not {0,1}.
- When FEAT_RME is implemented and the owning Security state is Realm and SCR_EL3.NSE is 0.

For these translations, the PE behaves as if one of the following is true:

- The owning Security state is Secure and SCR_EL3.NS is 0.
- When FEAT_RME is not implemented, the owning Security state is Non-secure and SCR_EL3.NS is 1.
- When FEAT_RME is not implemented, the owning Security state is Non-secure and SCR_EL3.{NSE,NS} is {0,1}.
- When FEAT_RME is not implemented, the owning Security state is Realm and SCR_EL3.{NSE,NS} is {1,1}.

Note:

The behavior might differ within the same translation.

is changed to read:

R_{RRCNN}

If the Effective value of SCR_EL3.{NSE, NS} does not match the owning Security state, it is **CONSTRAINED UNPREDICTABLE** whether any trace data:

- Is written to memory using a virtual address translated using the owning translation regime.

- Is written to memory using a virtual address translated using the translation regime formed from the owning Exception level and the Security state selected by the Effective value of SCR_EL3.{NSE, NS}.
- Is silently discarded and not written to memory.
- Is discarded and not written to memory and a trace buffer management event is generated as follows:
 - TRBSR_EL1.IRQ is set to 1.
 - If TRBSR_EL1.S is 0, then all of the following occur:
 - TRBSR_EL1.S is set to 1, collection is stopped.
 - TRBSR_EL1.EC is set to 0x00, other buffer management event.
 - TRBSR_EL1.BSC is set to 0b000000, access not allowed.
 - The other fields in TRBSR_EL1 are unchanged.

A corresponding update is made in section D17.9.1 “**UNPREDICTABLE** behavior”, where the text that reads:

If SCR_EL3.{NSE, NS} does not match the owning Security state, then it is **CONSTRAINED UNPREDICTABLE** whether the sample record or records:

- Are written to memory using a virtual address formed using the identity of the owning translation regime.
- Are written to memory using a virtual address formed using the value of SCR_EL3.{NSE, NS}.
- Are silently discarded and not written to memory.
- Are discarded and not written to memory, and an Access Not Allowed Profiling Buffer management event is generated.

is changed to read:

If the Effective value of SCR_EL3.{NSE, NS} does not match the owning Security state, then it is **CONSTRAINED UNPREDICTABLE** whether the profiling data:

- Is written to memory using a virtual address translated using the owning translation regime.
- Is written to memory using a virtual address translated using the translation regime formed from the owning Exception level and the Security state selected by the Effective value of SCR_EL3.{NSE, NS}.
- Is silently discarded and not written to memory.
- Is discarded and not written to memory, and an Access Not Allowed Profiling Buffer management event is generated.

2.19 C22264

In section C5.2.9 “FPMR, Floating-point Mode Register”, the field description for F8D, bits [8:6], that reads:

F8D, bits [8:6]

Destination result format for instructions that convert other floating-point values to an FP8 format.

F8D	Meaning
0b000	OFP8 E5M2 format.
0b001	OFP8 E4M3 format.

All other values are reserved. Reserved values identify an unsupported format, causing convert instructions that generate an FP8 result to perform a **CONSTRAINED UNPREDICTABLE** choice of one of the following behaviors:

- Setting the result to 0x FF and signaling an Invalid Operation floating-point exception.
- Generating the expected result of any of the supported FP8 formats.

is changed to read:

F8D, bits [8:6]

Destination result format for instructions that convert other floating-point values to an FP8 format.

F8D	Meaning
0b000	OFP8 E5M2 format.
0b001	OFP8 E4M3 format.

All other values are Reserved. Reserved values identify an unsupported format and behave as described in K1.2.19 Reserved values in System and memory-mapped registers and translation table entries. Additionally, FP8 instructions are permitted to set an FP8 result with an unsupported format to 0x FF and signal an Invalid Operation floating-point exception.

The field description for F8S2, bits [5:3], that reads:

F8S2, bits [5:3]

Second FP8 input data stream format for multiplication instructions with FP8 operands, and the corresponding instructions that convert an FP8 format to other floating-point formats.

F8S2	Meaning
0b000	OFP8 E5M2 format.
0b001	OFP8 E4M3 format.

All other values are reserved. Reserved values identify an unsupported format, and FP8 instructions treat the corresponding input as a **CONSTRAINED UNPREDICTABLE** choice of one of the following:

- A signaling NaN.
- Any of the supported FP8 formats.

is changed to read:

F8S2, bits [5:3]

Second FP8 input data stream format for multiplication instructions with FP8 operands, and the corresponding instructions that convert an FP8 format to other floating-point formats.

F8S2	Meaning
0b000	OFP8 E5M2 format.
0b001	OFP8 E4M3 format.

All other values are reserved. Reserved values identify an unsupported format and behave as described in K1.2.19 Reserved values in System and memory-mapped registers and translation table entries. Additionally FP8 instructions are permitted to treat FP8 input values with an unsupported format as a signaling NaN.

An equivalent change is made to the field F8S1, bits [2:0].

2.20 C22369

In section D24.3.18 “MDCR_EL3, Monitor Debug Configuration Register (EL3)”, field “NSPB, bits [13:12]”, the text that reads:

If profiling is enabled and MDCR_EL3.{NSPB, NSPBE} selects a reserved value, then the behavior is **CONSTRAINED UNPREDICTABLE**, and the Statistical Profiling Unit does one of:

...

- Selects an implemented Security state as the owning Security state.

is changed to read:

If profiling is enabled and MDCR_EL3.{NSPB, NSPBE} selects a reserved value, then the behavior is **CONSTRAINED UNPREDICTABLE**, and the Statistical Profiling Unit does one of:

...

- Selects an implemented Security state as the owning Security state. This includes generating Trap exceptions on accesses to Statistical Profiling and Profiling Buffer control registers in other Security states.

In field “NSTB, bits [25:24]”, the text that reads:

If the Trace Buffer Unit is enabled and using Self-hosted mode, and MDCR_EL3.{NSTB, NSTBE} selects a reserved value, then the behavior is **CONSTRAINED UNPREDICTABLE**, and the Trace Buffer Unit does one of:

...

- Selects an implemented Security state as the owning Security state.

is changed to read:

If the Trace Buffer Unit is enabled and using Self-hosted mode, and MDCR_EL3.{NSTB, NSTBE} selects a reserved value, then the behavior is **CONSTRAINED UNPREDICTABLE**, and the Trace Buffer Unit does one of:

...

- Selects an implemented Security state as the owning Security state. This includes generating Trap exceptions on accesses to Trace Buffer System registers in other Security states.

2.21 C22430

In section D.1.3.2.1 “Synchronous exception entry”, the rule that reads:

R_{BEJVV}

For GPC exceptions, a PA that characterizes the exception is captured in MFAR_EL3.

is changed to read:

R_{BEJVV}

For GPC exceptions, a PA and physical address space that characterizes the exception are captured in MFAR_EL3.

In section D.1.3.2.1 “Synchronous exception entry”, the rule that reads:

R_{TGQRC}

For Instruction Abort or Data Abort exceptions caused by an External abort, when FEAT_P FAR is implemented:

- If all of the following apply, then ESR_ELx.PFV is set to 0:
 - The exception is taken to EL1.
 - EL2 is implemented and enabled in the current Security state.
 - The Effective value of HCR_EL2.VM is 1
- Otherwise, ESR_ELx.PFV is set to an **IMPLEMENTATION DEFINED** value of 0 or 1.

For all other Instruction Aborts and Data Aborts, ESR_ELx.PFV is set to 0.

On taking a synchronous External Abort, if ESR_ELx.PFV is set to 1 by the PE then:

- An address within the same naturally-aligned fault granule as the faulting physical address is written to PFAR_ELx.PA or MFAR_EL3.PA as applicable. The fault granule size is defined by DLVGRB.
- The faulting physical address space is written to PFAR_ELx.{NSE,NS} or MFAR_EL3.{NSE,NS} as applicable.

Note:

PFAR_ELx never records the Intermediate Physical Address (IPA). PFAR_ELx might reveal a faulting physical addresses to a guest operating system if stage 2 translation is not being used and some other method is used to hide physical addresses from the guest (such as shadow page tables).

is changed to read:

R_{TGQRC}

For Instruction Abort or Data Abort exceptions caused by an External abort, when FEAT_P FAR is implemented:

- If all of the following apply, then ESR_ELx.PFV is set to 0:
 - The exception is taken to EL1.
 - EL2 is implemented and enabled in the current Security state.
 - The Effective value of HCR_EL2.VM is 1.
- Otherwise, ESR_ELx.PFV is set to an **IMPLEMENTATION DEFINED** value of 0 or 1.

For all other Instruction Abort and Data Abort exceptions, ESR_ELx.PFV is set to 0.

On taking an Instruction Abort or Data Abort exception:

- If ESR_ELx.PFV is set to 1 by the PE, then a PA and physical address space that characterizes the exception are captured in PFAR_ELx or MFAR_EL3 as applicable.
- If ESR_ELx.PFV is set to 0 by the PE, then the applicable PFAR_ELx or MFAR_EL3 register is set to an **UNKNOWN** value.

In section D.1.3.2.1 “Synchronous exception entry”, the rule that reads:

R_{ZTNQN}

On taking a synchronous External Abort, the following registers are **UNKNOWN** based on the ESR_ELx.PFV value:

- If ESR_EL1.PFV is set to 0, PFAR_EL1 is **UNKNOWN**.
- If ESR_EL2.PFV is set to 0, PFAR_EL2 is **UNKNOWN**.
- If ESR_EL3.PFV is set to 0, MFAR_EL3 is **UNKNOWN**.

is changed to read:

R_{X0001}

For all exceptions other than Instruction Abort, Data Abort, and GPC exceptions, the applicable PFAR_ELx or MFAR_EL3 register is set to an **UNKNOWN** value.

2.22 D22446

In section B2.5 “Restrictions on the effects of speculation”, the following text is added:

The Arm architecture places certain restrictions on the effects of speculation. These are:

...

- A speculative Explicit Memory Read Effect generated by an instruction (I_2) running in a hardware-defined context ($context_2$) will not speculatively read data from a speculative Explicit Memory Write Effect generated by an instruction (I_1) running in a different hardware-defined context ($context_1$) if I_1 generates an MMU Fault Effect instead of the Explicit Memory Write Effect.

2.23 C22566

In section A1.5.9.4 “Underflow exceptions”, the text that reads:

If the result of a floating-point operation is a denormalized number that is flushed to zero, then the Underflow floating-point exception is not generated.

is changed to read:

If the result of a floating-point operation is a denormalized number that is flushed to zero, then the Underflow floating-point exception is generated, however this floating-point exception is not trapped regardless of the value of FPCR.UFE.

2.24 C22584

In section J1.3 “Shared pseudocode”, the function IncrementInstructionCounter() with the code that reads:

```
if old_value<64> != new_value<64> then
    PMOVSSET_ELO.F0 = '1';
    PMOVSLR_ELO.F0 = '1';
```

is changed to read:

```
if old_value<64> != new_value<64> then
    PMOVSSET_ELO.F0 = '1';
```

Equivalent changes are made to functions CheckForPMUException(), CheckPMUOverflowCondition(), and IncrementEventCounter().

2.25 D22595

In section D14.3.2 “Common microarchitectural events”, the text for the PMU event INT_MUL64_SPEC that reads:

0x804c, INT_MUL64_SPEC, Integer operation speculatively executed, 64x64 multiply

The counter counts each Speculatively executed 64x64 integer multiply operation counted by INT_SPEC due to any of the following A64 instructions:

- Scalar: MADD, MSUB, MUL, SMULH, or UMULH.
- SVE: MAD, MLA, MLS, MSB, MUL, SMULH, or UMULH.
- SVE2: SVE2: CMLA (vectors), MLA, MLS, MUL, SMLALB, SMLALT, SMLSBLB, SMLSBLT, SMULH, SMULLB, SMULLT, SQDMLALB, SQDMLALBT, SQDMLALT, SQDMLSLB, SQDMLSLBT, SQDMLSLT, SQDMULH, SQDMULLB, SQDMULLT, SQRDCMLAH (vectors), SQRDCMLAH, SQRDMLSH, SQRDMULH, UMLALB, UMLALT, UMLSBLB, UMLSBLT, UMULH, UMULLB, or UMULLT.

is changed to read:

0x804c, INT_MUL64_SPEC, Integer operation speculatively executed, 64x64 multiply

The counter counts each Speculatively executed 64x64 integer multiply operation counted by INT_SPEC due to any of the following A64 instructions:

- Scalar: MADD, MSUB, MUL, SMULH, or UMULH.
- SVE: MAD, MLA, MLS, MSB, MUL, SMULH, or UMULH.
- SVE2: CMLA (vectors), MLA, MLS, MUL, SMULH, SQDMULH, SQRDCMLAH (vectors), SQRDCMLAH, SQRDMLSH, SQRDMULH, or UMULH.

The equivalent changes are made in event SVE_INT_MUL64_SPEC.

2.26 E22689

In D24.2.74 “ID_AA64DFR0_EL1, AArch64 Debug Feature Register 0”, in the field description of ‘BRPs, bits [15:12]’, the text that reads:

Number of breakpoints, minus 1. The value of this field is an **IMPLEMENTATION DEFINED** choice of:

0b0001...0b1111	The number of breakpoints, minus 1.
-----------------	-------------------------------------

If FEAT_Debugv8p9 is implemented and 16 or more breakpoints are implemented, then this field reads as 0b1111 and ID_AA64DFR1_EL1.BRPs indicates the number of breakpoints.

is changed to read:

Index of the highest numbered context-aware breakpoint. Identifies which of the implemented breakpoints are context-aware breakpoints.

The value of this field is an **IMPLEMENTATION DEFINED** choice of:

0b0001...0b1110	Index of the highest numbered context-aware breakpoint.
0b1111	If ID_AA64DFR0_EL1.CTX_CMPs and ID_AA64DFR1_EL1.CTX_CMPs indicate that 16 or fewer breakpoints are context-aware, then the index of the highest context-aware breakpoint is 15. Otherwise, the context-aware breakpoints are the lowest numbered breakpoints.

If ID_AA64DFR1_EL1.BRPs is zero, then this is also the number of implemented breakpoints, minus 1, meaning the context-aware breakpoints are the highest numbered breakpoints.

The value of this field is always less than the number of implemented breakpoints.

In the same section, in the field description of 'CTX_CMPs, bits [31:28]', the text that reads:

Number of context-aware breakpoints, minus 1.

The value of this field is an **IMPLEMENTATION DEFINED** choice of:

0b0000...0b1111	The number of context-aware breakpoints, minus 1.
-----------------	---

The value of this field is never greater than ID_AA64DFR0_EL1.BRPs. If FEAT_Debugv8p9 is implemented and 16 or more context-aware breakpoints are implemented, then this field reads as 0b1111 and ID_AA64DFR1_EL1.CTX_CMPs indicates the number of context-aware breakpoints.

is changed to read:

Number of context-aware breakpoints, minus 1.

The value of this field is an **IMPLEMENTATION DEFINED** choice of:

0b0000...0b1110	The number of context-aware breakpoints, minus 1.
0b1111	If ID_AA64DFR1_EL1.CTX_CMPs is zero, then 16 context-aware breakpoints are implemented. Otherwise, 16 or more context-aware breakpoints are implemented and ID_AA64DFR1_EL1.CTX_CMPs indicates how many.

Values greater than ID_AA64DFR0_EL1.BRPs are not permitted.

In D24.2.75 "ID_AA64DFR1_EL1, AArch64 Debug Feature Register 1", in the field description of 'BRPs, bits [15:8]', the following text is added:

If more than 16 breakpoints are implemented, or the index of the highest numbered context-aware breakpoint is not the number of breakpoints minus 1, then the value 0x00 is not permitted.

Nonzero values less than ID_AA64DFR0_EL1.BRPs are not permitted.

In the same section, in the field description of 'CTX_CMPs, bits [31:24]', the text that reads:

The value of this field is never greater than ID_AA64DFR1_EL1.BRPs.

is changed to read:

If 16 or fewer context-aware breakpoints are implemented, then the value of this field is either 0x00 or the same as ID_AA64DFR0_EL1.CTX_CMPs. If more than 16 context-aware breakpoints are implemented, then the value 0x00 is not permitted.

Values greater than ID_AA64DFR1_EL1.BRPs are not permitted.

2.27 D22719

In C5.5.92 "TLBIP RIPAS2E1, TLBIP RIPAS2E1NXS, TLB Range Invalidate by Intermediate Physical Address, Stage 2, EL1", under the heading 'TTL, bits [38:37]', the following text is removed:

If FEAT_LPA2 is not implemented, when using a 16KB translation granule, this value is reserved and hardware should treat this field as 0b00.

The equivalent change is applied to all other TLBIP instructions that apply to a range of addresses.

In C5.5.86 "TLBIP IPAS2E1, TLBIP IPAS2E1NXS, TLB Invalidate Pair by Intermediate Physical Address, Stage 2, EL1", under heading 'TTL, bits [47:44]', the text that reads:

- 0b01xx The entry comes from a 4KB translation granule. The level of walk for the leaf level 0bxx is encoded as:
 - 0b00 : If FEAT_LPA2 is implemented, level 0. Otherwise, treat as if TTL[3:2] is 0b00.
 - ...
- 0b10xx The entry comes from a 16KB translation granule. The level of walk for the leaf level 0bxx is encoded as:
 - 0b00 : Reserved. Treat as if TTL[3:2] is 0b00.
 - 0b01 : If FEAT_LPA2 is implemented, level 1. Otherwise, treat as if TTL[3:2] is 0b00.
 - ...

is changed to read:

- 0b01xx The entry comes from a 4KB translation granule. The level of walk for the leaf level 0bxx is encoded as:
 - 0b00 : Level 0.
 - ...
- 0b10xx The entry comes from a 16KB translation granule. The level of walk for the leaf level 0bxx is encoded as:

- 0b00 : Reserved. Treat as if TTL[3:2] is 0b00.
- 0b01 : Level 1.
- ...

The equivalent change is applied to all other TLBIP instructions that do not apply to a range of addresses.

In C5.5.59 “TLBI VAE1, TLBI VAE1NXS, TLB Invalidate by VA, EL1”, under the heading ‘Purpose’, the text that reads:

- The entry is one of the following:
 - A 64-bit stage 1 translation table entry.
 - If FEAT_D128 is implemented, a 128-bit stage 1 translation table entry, if TTL[3:2] is 0b00.
- The entry would be used to translate the specified VA, and one of the following applies:
 - The entry is from a level of lookup above the final level and matches the specified ASID.
 - The entry is a global entry from the final level of lookup.
 - The entry is a non-global entry from the final level of lookup that matches the specified ASID.

is changed to read:

- The entry would be used to translate the specified VA, and one of the following applies:
 - The entry is a Table entry and all the following apply:
 - The entry matches the specified ASID.
 - If FEAT_TTL is implemented, then one of the following applies:
 - TTL[3:2] is 0b00.
 - The entry is in VMSAv8-32 LPAE or VMSAv8-64 format and leads to a Page or Block entry translating the specified VA that matches the granule and level specified in TTL.
- The entry is a Page or Block entry and all of the following apply:
 - For a non-global entry, the entry matches the specified ASID.
 - If FEAT_TTL is implemented, then one of the following applies:
 - TTL[3:2] is 0b00.
 - The entry is in VMSAv8-32 LPAE or VMSAv8-64 format and matches the granule and level specified in TTL.

The equivalent changes are applied to other TLBI instructions that have a TTL hint.

In C5.5.128 “TLBIP VAE1, TLBIP VAE1NXS, TLB Invalidate Pair by VA, EL1”, under the heading ‘Purpose’, the text that reads:

- The entry is one of the following:
 - A 128-bit stage 1 translation table entry.
 - A 64-bit stage 1 translation table entry, if TTL[3:2] is 0b00.
- The entry would be used to translate the specified VA, and one of the following applies:
 - The entry is from a level of lookup above the final level and matches the specified ASID.
 - The entry is a global entry from the final level of lookup.
 - The entry is a non-global entry from the final level of lookup that matches the specified ASID.

is changed to read:

- The entry would be used to translate the specified VA, and one of the following applies:
 - The entry is a Table entry and all the following apply:
 - The entry matches the specified ASID.
 - If FEAT_TTL is implemented, then one of the following applies:
 - TTL[3:2] is 0b00.
 - The entry is in VMSAv9-128 format and leads to a Page or Block entry translating the specified VA that matches the granule and level specified in TTL.
- The entry is a Page or Block entry and all of the following apply:
 - For a non-global entry, the entry matches the specified ASID.
 - If FEAT_TTL is implemented, then one of the following applies:
 - TTL[3:2] is 0b00.
 - The entry is in VMSAv9-128 format and matches the granule and level specified in TTL.

The equivalent changes are applied to other TLBIP instructions that have a TTL hint.

In D8.17.5.3 “Translation table level hint”, the following new information statement is added:

For non-final levels of walk, a translation using VMSAv9-128 descriptors resolves a different number of IA bits from a translation using VMSAv8-64 descriptors. Because of this mismatch, the TTL hint applies differently between TLBI and TLBIP instructions:

- For TLBI instructions, the TTL hint applies for 64-bit descriptors, including VMSAv8-64 and VMSAv8-32LPAE descriptors.
- For TLBIP instructions, the TTL hint applies for 128-bit descriptors.

2.28 D22806

In section D24.2.137 “PFAR_EL1, Physical Fault Address Register (EL1)” and D24.2.138 “PFAR_EL2, Physical Fault Address Register (EL2)”, the description of the field NS, bit [63], that reads:

NS, bit [63]

...

When EL3 is implemented:

Non-secure. Reports the physical address space of the access that triggered the exception.

NS	Meaning
0b0	Secure physical address space.
0b1	Non-secure physical address space.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

Otherwise:

Reserved, **RES0**.

is changed to read:

NS, bit [63]

...

When EL3 is implemented or FEAT_Secure is implemented:

Non-secure. Reports the physical address space of the access that triggered the exception.

NS	Meaning
0b0	Secure physical address space.
0b1	Non-secure physical address space.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

Otherwise:

Reserved, **RES0**.

2.29 R22882

In section B2.12.5 “Load-Exclusive and Store-Exclusive instruction usage restrictions”, the text that reads:

LoadExcl/StoreExcl loops are guaranteed to make forward progress only if, for any LoadExcl/StoreExcl loop within a single thread of execution, the software meets all of the following conditions:

- ...

is changed to read:

If FEAT_LSE is not implemented, LoadExcl/StoreExcl loops are guaranteed to make forward progress only if, for any LoadExcl/StoreExcl loop within a single thread of execution, the software meets all of the following conditions.

If FEAT_LSE is implemented, software can maximize the likelihood that LoadExcl/StoreExcl loops make forward progress, for any LoadExcl/StoreExcl loop within a single thread of execution, by meeting all of the following conditions:

- ...

2.30 D22979

In section D8.5.2.2 “Implications of enabling the dirty state management mechanism”, under the rule R_{RKMHW} , the text that reads:

For the following instructions that require write permission, if the address specified in the instruction is translated by a writable-clean descriptor, then the descriptor is considered to grant write access and the hardware does not update the Dirty state of that descriptor:

- Data cache invalidation instruction, DC IVAC.
- Data cache and instruction cache maintenance instructions that are affected by FEAT_CMOW.
- Address translation instructions, AT S12E0W, AT S12E1W, AT S1E0W, AT S1E1W, AT S1E2W, AT S1E3W.

is changed to read:

For the following instructions that require write permission, if the address specified in the instruction is translated by a writable-clean descriptor and hardware management of Dirty state is enabled for the corresponding stage of translation, then the descriptor is considered to grant write access and the hardware does not update the Dirty state of that descriptor:

- Data cache invalidation instruction, DC IVAC.
- Data cache and instruction cache maintenance instructions that are affected by FEAT_CMOW.

- Address translation instructions, AT S12E0W, AT S12E1W, AT S1E0W, AT S1E1W, AT S1E2W, AT S1E3W.

2.31 D23002

In section D13.12 “PMU events and event numbers”, the following instructions are added to the event definitions:

Event	Instructions added
FPASE_LDST_REG_SPEC	LDAP1, LDAPUR, LDNP, LDP, LDR, LDUR, STL1, STLUR, STNP, STP, STR, STUR
FPASE_LD_REG_SPEC	LDAP1, LDAPUR, LDNP, LDP, LDR, LDUR
FPASE_ST_REG_SPEC	STL1, STLUR, STNP, STP, STR, STUR

2.32 R23026

In section A2.3.5 “The Armv9.4 architecture extension”, under the heading ‘FEAT_D128, 128-bit Translation Tables, 56 bit PA’, the text that reads:

FEAT_D128, 128-bit Translation Tables, 56 bit PA

FEAT_D128 adds support for the VMSAv9-128 translation system, comprising the following:

- 128-bit translation table descriptors.
- 56-bit physical addresses.
- 56-bit virtual addresses.
- 128-bit System registers.
- 128-bit atomic instructions.
- TLBIP VA*, TLBIP RVA*, TLBIP IPA*, TLBIP RIPA* instructions that can take 128-bit inputs.
- **IMPLEMENTATION DEFINED** System instructions that can take 128-bit inputs.

This feature is supported in AArch64 state only.

FEAT_D128 is OPTIONAL from Armv9.3.

If FEAT_D128 is implemented, then FEAT_SYSREG128 is implemented.

If FEAT_D128 is implemented, then FEAT_SYSINSTR128 is implemented.

If FEAT_D128 is implemented, then FEAT_LSE128 is implemented.

If FEAT_D128 is implemented, then FEAT_S1PIE is implemented.

If FEAT_D128 and EL2 are implemented, then FEAT_S2PIE is implemented.

If FEAT_D128 is implemented, then FEAT_AIE is implemented.

If FEAT_D128 is implemented, then FEAT_TCR2 is implemented.

If FEAT_D128 is implemented, then FEAT_LVA is implemented.

If FEAT_D128 is implemented, then FEAT_LPA2 is implemented.

The following field identifies the presence of FEAT_D128:

- ID_AA64MMFR3_EL1.D128.

is changed to read:

FEAT_D128, 128-bit Translation Tables

FEAT_D128 adds support for the VMSAv9-128 translation system, comprising the following:

- 128-bit translation table descriptors.
- Support for encoding up to 56-bit physical addresses in translation table descriptors.
- If FEAT_LVA or FEAT_LVA3 are implemented, support for translating up to 56-bit virtual addresses.
- TLBIP VA*, TLBIP RVA*, TLBIP IPA*, TLBIP RIPA* instructions that can take 128-bit inputs.

This feature is supported in AArch64 state only.

FEAT_D128 is OPTIONAL from Armv9.3.

If FEAT_D128 is implemented, then FEAT_SYSREG128 is implemented.

If FEAT_D128 is implemented, then FEAT_SYSINSTR128 is implemented.

If FEAT_D128 is implemented, then FEAT_LSE128 is implemented.

If FEAT_D128 is implemented, then FEAT_S1PIE is implemented.

If FEAT_D128 and EL2 are implemented, then FEAT_S2PIE is implemented.

If FEAT_D128 is implemented, then FEAT_AIE is implemented.

If FEAT_D128 is implemented, then FEAT_TCR2 is implemented.

The following field identifies the presence of FEAT_D128:

- ID_AA64MMFR3_EL1.D128.

2.33 D23122

In section D14.3.2 “Common microarchitectural events”, in the event description of ‘0x0015, L1D_CACHE_WB, Level 1 data cache write-back’ the text that reads:

The counter counts each write-back of data from the Level 1 data or unified cache to outside of the Level 1. For example:

- A write-back of a dirty cache line to a Level 2 cache or memory.
- A write-back of a recently fetched cache line that has not been allocated to the Level 1 data cache.

is changed to read:

The counter counts each write-back of data from the Level 1 data or unified cache to outside of the Level 1. For example:

- A write-back of a cache line to a Level 2 cache or memory.
- A write-back of a recently fetched cache line that has not been allocated to the Level 1 data cache.

In the same event description, the text that reads:

A write-back is attributable to the agent that generated the request that caused the write-back. This might not be the same agent that caused the data being written back to be allocated into the cache.

An Unattributable write-back event occurs when a requestor outside of the PE makes a coherency request that results in write-back. If the cache is shared, then an Unattributable write-back event is not counted. If the cache is not shared, then the event is counted.

is changed to read:

A write-back is attributable to one of the following, and it is **IMPLEMENTATION DEFINED** which:

- The PE or other agent that generated the request that caused the write-back. This might not be the same agent that caused the data to be allocated into the cache.
- If the cache is shared, the PE or other agent that caused the data to be allocated into the cache.
- If the cache is shared, the PE or other agent that last accessed the data when it was in the cache. If the data has not been accessed since being allocated, then the PE or other agent that caused the data to be allocated into the cache.

If the cache is shared, then the write-back event is counted only if it is attributable to the PE counting the event.

If the cache is not shared, then the event is counted by the PE that the cache is attached to, even if the write-back is not attributable to that PE. For example, a requestor outside of the PE made a coherency request that resulted in write-back.

In the same section, in the event description of '0x0046, L1D_CACHE_WB_VICTIM, Level 1 data cache write-back, victim', the text that reads:

The counter counts each write-back counted by L1D_CACHE_WB that occurs because the line is allocated for an access made by the PE.

is changed to read:

The counter counts each write-back counted by L1D_CACHE_WB that occurs because of a capacity eviction due to a line being allocated into the cache.

It is **IMPLEMENTATION DEFINED** whether this includes capacity evictions due to an agent other the PE counting the event. For example, a stashing request from outside of the PE, or an allocation due to another PE that shares the cache.

In all cases, the event is counted only if the eviction is also an L1D_CACHE_WB event. This means that, if the cache is shared, the eviction is counted only if it is attributable to the PE counting the event, as defined by L1D_CACHE_WB.

In the same section, in the event description of '0x0047, L1D_CACHE_WB_CLEAN, Level 1 data cache write-back, cleaning and coherency', the text that reads:

The counter counts each write-back counted by L1D_CACHE_WB that occurs because of a coherency operation made by another PE or, optionally, the execution of a cache maintenance instruction.

Whether write-backs that are caused by the execution of a cache maintenance instruction are counted is **IMPLEMENTATION DEFINED**.

Note:

The transfer of a dirty cache line from the Level 1 data cache of this PE to the data cache of another PE due to a hardware coherency operation is not counted unless the dirty cache line is also written back to a Level 2 cache or memory. If a coherency request from a requestor outside of the PE results in a write-back, it is an Unattributable event.

is changed to read:

The counter counts each write-back counted by L1D_CACHE_WB that occurs because of a coherency operation made by another PE or, optionally, the execution of a cache maintenance instruction.

It is **IMPLEMENTATION DEFINED** whether the transfer of a dirty cache line from the Level 1 data cache of this PE to the data cache of another PE due to a hardware coherency operation is counted when the dirty cache line is not also written back to a Level 2 cache or memory.

In all cases, the event is counted only if the write-back is also an L1D_CACHE_WB event. This means that, if the cache is shared, the write-back is counted only if it is attributable to the PE counting the event, as defined by L1D_CACHE_WB.

Similar updates are made for the following events:

- L2D_CACHE_WB (0x0018)
- L2D_CACHE_VICTIM (0x0056)
- L2D_CACHE_CLEAN (0x0057)
- L3D_CACHE_WB (0x002c)
- L3D_CACHE_VICTIM (0x00A6)
- L3D_CACHE_CLEAN (0x00A7)

2.34 R23144

In section D24.2.87 “ID_AA64PFR0_EL1, AArch64 Processor Feature Register 0”, in the field description of ‘RAS, bits [31:28]’, the text that reads:

0b0011	<p>FEAT_RASv2 implemented. As 0b0010 and adds support for:</p> <ul style="list-style-type: none"> • ERXGSR_EL1, to support System RAS agents. • Additional fine-grained EL2 traps for additional error record System registers. • The SCR_EL3.TWERR write control for error record System registers. <p>Error records accessed through System registers conform to RAS System Architecture v2.</p>
--------	---

is changed to read:

0b0011	<p>FEAT_RASv2 implemented. As 0b0010 and adds support for:</p> <ul style="list-style-type: none"> • The error group status register, ERXGSR_EL1. • The SCR_EL3.TWERR write trap control for error record System registers. • Additional syndrome in ESR_ELx for error exceptions. <p>Error records accessed through System registers conform to either RAS System Architecture v1.1 or RAS System Architecture v2.</p>
--------	---

In section A2.2.10 “The Armv8.9 architecture extension”, under the heading ‘FEAT_RASv2, RAS Extension v2’, the text that reads:

FEAT_RASv2 adds the following features to the Reliability, Availability, and Serviceability Extension:

- Adds the features defined by FEAT_RASSAv2 to System register error records.
- Defines the ERXGSR_EL1 register.
- Adds a Trap exception to EL3 for writes to RAS System registers.
- Adds additional syndrome to ESR_ELx on an error exception, to give information on whether a location being accessed has been updated.

is changed to read:

FEAT_RASv2 adds the following features to the Reliability, Availability, and Serviceability Extension:

- Allows the features defined by FEAT_RASSAv2 to System register error records.
- Defines the error group status register, ERXGSR_EL1.
- Adds a control to trap writes to RAS error record System registers to EL3.
- Adds additional syndrome to ESR_ELx for error exceptions.

2.35 D23154

In section J1.1.3 “aarch64/functions”, a new function is added:

```
boolean IsPFARValid(FaultRecord fault);
```

In the same section, in function AArch64.FaultSyndrome(...), the code that reads:

```
IssType AArch64.FaultSyndrome(...)
    if d_side then
        ...
    else
        ...
        isstype.iss2<11> = if fault.hdbssf then '1' else '0';
        if IsExternalAbort(fault) then
            isstype.iss<9> = fault.extflag;
```

is changed to read:

```
IssType AArch64.FaultSyndrome(...)
    if d_side then
        ...
    else
        ...
        isstype.iss2<11> = if fault.hdbssf then '1' else '0';
        if (IsFeatureImplemented(FEAT_PFAR) && IsExternalSyncAbort(fault) &&
            !(EL2Enabled() && (HCR_EL2.VM == '1' || HCR_EL2.DC == '1') &&
            target_el == EL1)) then
            isstype.iss<14> = if IsPFARValid(fault) then '1' else '0';
        if IsExternalAbort(fault) then
            isstype.iss<9> = fault.extflag;
```

An equivalent change is made in AArch64.PhysicalErrorSyndrome().

Additionally, in the function AArch64.TakePhysicalErrorException(...), the code that reads:

```
AArch64.TakePhysicalErrorException(...)
    ...
    constant bits(25) syndrome = AArch64.PhysicalErrorSyndrome(implicit_esb);
    ...
```

is changed to read:

```
AArch64.TakePhysicalSErrorException(...)
...
constant bits(25) syndrome = AArch64.PhysicalSErrorSyndrome(is_esb, implicit_esb);
...
if except.syndrome.iss<14> == '1' then
    except.pavalid = TRUE;
    except.paddress = fault.paddress;
else
    except.pavalid = FALSE;
    except.paddress = FullAddress UNKNOWN;
...
```

Equivalent changes are made in the following functions:

- AArch64.AbortSyndrome().
- HandleExternalTTWAbort().
- HandleExternalAbort().

2.36 D23168

In section D24.2.78 “ID_AA64ISAR0_EL1, AArch64 Instruction Set Attribute Register 0”, in the field description for ‘DP, bits [47:44]’, the text that reads:

From Armv8.2, the permitted values are 0b0000 and 0b0001.

is changed to read:

From Armv8.4, if FEAT_AdvSIMD is implemented, the value 0b0000 is not permitted.

In this same register, the following changes are also made:

In the field description for ‘RNDR, bits [63:60]’, the following text is removed:

From Armv8.5, the permitted values are 0b0000 and 0b0001.

In the field description for ‘FHM, bits [51:48]’, the text that reads:

From Armv8.2, the permitted values are 0b0000 and 0b0001.

is changed to read:

From Armv8.4, if FEAT_FP16 is implemented, the value 0b0000 is not permitted.

In the field description for ‘SM3, bits [39:36]’, the following text is removed:

From Armv8.2, the permitted values are 0b0000 and 0b0001.

In the field description for ‘SHA3, bits [35:32]’, the following text is removed:

From Armv8.2, the permitted values are 0b0000 and 0b0001.

In the field description for 'SM4, bits [43:40]', the text that reads:

From Armv8.2, the permitted values are 0b0000 and 0b0001.

is changed to read:

FEAT_SM4 implements the functionality identified by the value 0b0001.

2.37 C23225

In section B1.2.1 “FFR, First Fault Register” the text that reads:

I_{IXLZQY}

After a sequence of one or more SVE First-fault or Non-fault loads that follow a SETFFR instruction, the FFR contains a sequence of zero or more TRUE elements, followed by zero or more FALSE elements.

is changed to read:

I_{IXLZQY}

After a sequence of one or more SVE First-fault or Non-fault loads that follow a SETFFR instruction, without any intervening WRFFR instruction, the FFR contains a monotonic predicate value.

And the following text is added to the same section:

D_{X0001}

A monotonic predicate value is one that consists of, starting from predicate element 0, a sequence of zero or more TRUE elements, followed only by zero or more FALSE elements.

I_{X0001}

The WRFFR instruction requires that its source predicate contains a monotonic predicate value. If the source is not a monotonic predicate value, then the resulting value in the FFR is **UNKNOWN**.

2.38 D23239

In F6.1.32 “VAND (immediate)” and F6.1.157 “VORN (immediate)” the text that reads:

I_{I32}

is changed to read:

I16

and the text that reads:

I16

is changed to read:

I32

2.39 C23271

In section D14.3 “Common event numbers”, the description of event “0x0072, LDST_SPEC, Operation speculatively executed, load or store” that reads:

0x0072, LDST_SPEC, Operation speculatively executed, load or store

The counter counts each operation counted by INST_SPEC that is a load operation or a store operation. See LD_SPEC and ST_SPEC for these classifications.

When FEAT_PMUv3p8 is not implemented, this is an **IMPLEMENTATION DEFINED** event.

is changed to read:

0x0072, LDST_SPEC, Operation speculatively executed, load or store

The counter counts each operation counted by INST_SPEC that is a load operation or a store operation. See LD_SPEC, PRF_SPEC and ST_SPEC for these classifications.

If the PRF_SPEC event is implemented and the prefetch operations counted by PRF_SPEC are not counted by LD_SPEC, then it is **IMPLEMENTATION DEFINED** whether these operations are counted by LDST_SPEC. Arm recommends that these operations are counted by LDST_SPEC.

When FEAT_PMUv3p8 is not implemented, this is an **IMPLEMENTATION DEFINED** event.

In the same section, the description of event “0x0070, LD_SPEC, Operation speculatively executed, load” that reads:

0x0070, LD_SPEC, Operation speculatively executed, load

The counter counts each operation counted by LDST_SPEC that is a load operation. Operations due to Memory-reading instructions are counted as load operations.

It is **IMPLEMENTATION DEFINED** whether the prefetch instructions PRFM, PLD, PLDW, and PLI count as integer data-processing operations or load operations. Arm recommends that if the instruction is not implemented as a **NOP** then it is counted as a load operation.

When FEAT_PMUv3p8 is not implemented, this is an **IMPLEMENTATION DEFINED** event.

is changed to read:

0x0070, LD_SPEC, Operation speculatively executed, load

The counter counts each operation counted by LDST_SPEC that is a load operation. Operations due to Memory-reading instructions are counted as load operations.

It is **IMPLEMENTATION DEFINED** whether operations counted by PRF_SPEC count as load operations. Arm recommends that if the instruction is not implemented as a **NOP** and PRF_SPEC is not implemented, then the operation is counted as a load operation. If the operation is counted by neither LD_SPEC nor LDST_SPEC, then it is counted as a data-processing operation.

When FEAT_PMUv3p8 is not implemented, this is an **IMPLEMENTATION DEFINED** event.

The event “0x0080, LD_RETIRED, Operation architecturally executed, load” has the following text added:

If the prefetch instructions defined to be counted by PRF_SPEC are counted by LD_SPEC then they are also counted by LD_RETIRED.

2.40 D23403

In section D23.5.12 “PMEVTYPER<n>_ELO, Performance Monitors Event Type Registers, n = 0 - 30”, the following text is added to the description of field “MT, bit [25]”:

Unless otherwise stated:

- If the event counts PE cycles when a stall condition is true and a second condition is true, then the counter counts Processor cycles when the stall condition is true for all of these PEs, and the second condition is true for any of these PEs.
- If the event counts PE cycles when any other condition is true, then the counter counts Processor cycles when the condition is true for any of these PEs. Otherwise, the event counts by the sum of the count across all of these PEs.

For the stall events, the stall condition means the applicable condition described by the STALL, STALL_FRONTEND, or STALL_BACKEND event. The second condition is any condition in addition to this. For example, for the STALL_FRONTEND_L1I event, the stall condition is STALL_FRONTEND, and the second condition is when there is a demand instruction miss in the first level of instruction cache. For the STALL, STALL_FRONTEND, and STALL_BACKEND events themselves, the second condition is the null TRUE condition.

See also Cycle event counting.

Section D13.6 “Multithreaded implementations” is renamed to “Cycle event counting in multithreaded implementations” and the equivalent text is added.

More specific information is added to the affected STALL events.

2.41 D23404

In section D24.2.183 “TCR_EL2, Translation Control Register (EL2)”, the text in field ‘HWU062, bit[46]’ that reads:

For translations using TTBR0_EL1, bit[62] of each stage 1 translation table Block or Page entry cannot be used by hardware for an **IMPLEMENTATION DEFINED** purpose.

is changed to read:

For translations using TTBR0_EL2, bit[62] of each stage 1 translation table Block or Page entry cannot be used by hardware for an **IMPLEMENTATION DEFINED** purpose.

Equivalent changes are made for all references to TTBR0_EL1 in the HWU0xx fields.

In the same register, in the ‘TBI1, bit [38]’ field, the statement that reads:

- A branch or procedure return within ELO or EL1.
- An exception taken to EL1.
- An exception return to ELO or EL1.

is changed to read:

- A branch or procedure return within EL2.
- If the Effective value of HCR_EL2.{E2H, TGE} is {1,1}, a branch or procedure return within ELO.
- An exception taken to EL2.
- An exception return to EL2.
- If the Effective value of HCR_EL2.{E2H, TGE} is {1, 1}, then an exception return to ELO.

Equivalent change is made to the ‘TBI0, bit [37]’ field.

In the same register, the statement for the !EL2InHost() view of the register that reads:

It has an effect whether the EL2, or EL2&0, translation regime is enabled or not.

is changed to read:

It has an effect whether the EL2 translation regime is enabled or not.

In the same register, for the EL2InHost() view of the register, the text that reads:

It has an effect whether the EL2, or EL2&0, translation regime is enabled or not.

is changed to read:

It has an effect whether the EL2&O translation regime is enabled or not.

In section D24.2.168 “SCTLR_EL2, System Control Register (EL2)”, the text in the ‘EOE, bit[24]’ field that reads:

If an implementation only supports Little-endian accesses at ELO, then this bit is **RES0**. This option is not permitted when SCTLR_EL1.EE is **RES1**.

If an implementation only supports Big-endian accesses at ELO, then this bit is **RES1**. This option is not permitted when SCTLR_EL1.EE is **RES0**.

is changed to read:

If an implementation only supports Little-endian accesses at ELO, then this bit is **RES0**. This option is not permitted when SCTLR_EL2.EE is **RES1**.

If an implementation only supports Big-endian accesses at ELO, then this bit is **RES1**. This option is not permitted when SCTLR_EL2.EE is **RES0**.

In the same section, the text that reads:

This bit has no effect on the endianness of LDTR, LDTRH, LDTRSH, LDTRSW, STTR, and STTRH instructions executed at EL1.

is changed to read:

This bit has no effect on the endianness of LDTR, LDTRH, LDTRSH, LDTRSW, STTR, and STTRH instructions executed at EL2.

2.42 C23431

In section J1.1 “Pseudocode for AArch64 operation”, the pseudocode for function “AArch64.MemAtomicRCW()” that reads:

```
(bits(4), bits(size)) MemAtomicRCW(bits(64) address, bits(size) cmpoperand,  
bits(size) operand,  
AccessDescriptor accdesc_in)  
...  
return (nzcw, retvalue);
```

is changed to read:

```
(bits(4), bits(size)) MemAtomicRCW(bits(64) address, bits(size) cmpoperand,  
bits(size) operand,  
AccessDescriptor accdesc_in)  
...  
if SPESampleInFlight then  
    constant boolean is_load = TRUE;  
    SPESampleLoadStore(is_load, accdesc, memaddrdesc);  
return (nzcw, retvalue);
```

Similar changes are made in the following functions:

- AArch64.MemLoad64B().
- AArch64.MemStore64BWithRet().
- AArch64.MemStore64B().

2.43 D23442

In section D18.2.6 “Events packet”, under ‘Events packet payload’, in E[25], the text that reads:

E[25], byte 3 bit [1], when SZ == 0b10 or SZ == 0b11

SMCU or external coprocessor operation.

When FEAT_SPE_SME is implemented

E[25]	Description
0b0	Operation did not execute on an SMCU or external coprocessor.
0b1	Operation executed on an SMCU or external coprocessor.

When FEAT_SPEv1p4 is implemented

This bit reads-as-zero.

Otherwise

This bit reads as an **IMPLEMENTATION DEFINED** value.

is changed to read:

E[25], byte 3 bit [1], when SZ == 0b10 or SZ == 0b11

Streaming Mode Compute Unit (SMCU) or other shared resource operation.

When (FEAT_SPE_SME is implemented or FEAT_SPEv1p5 is implemented) and an SMCU or other shared resource is implemented

E[25]	Description
0b0	Operation did not execute on an SMCU or other shared resource.
0b1	Operation executed on an SMCU or other shared resource.

A shared resource means a resource that is shared between PEs for the execution of instructions. It is **IMPLEMENTATION DEFINED** which instructions can be executed on a shared resource.

When FEAT_SPEv1p4 is implemented or FEAT_SPE_SME is implemented

This bit reads-as-zero.

Otherwise

This bit reads as an **IMPLEMENTATION DEFINED** value.

In section D24.7.8 “PMSEVFR_EL1, Sampling Event Filter Register, in E[25], the text that reads:

E[25], bit [25] When FEAT_SPE_SME is implemented and event 25 is implemented: Filter on Streaming Mode Compute Unit (SMCU) or external coprocessor operation event.

E[25]	Meaning
0b0	SMCU or external coprocessor operation event is ignored.
0b1	Do not record samples that have the SMCU or external coprocessor operation event == 0.

is changed to read:

E[25], bit [25]

When (FEAT_SPE_SME is implemented or FEAT_SPEv1p5 is implemented) and event 25 is implemented:

Filter on Streaming Mode Compute Unit (SMCU) or other shared resource operation event.

E[25]	Meaning
0b0	SMCU or other shared resource operation event is ignored.
0b1	Do not record samples that have the SMCU or other shared resource operation event == 0.

The equivalent change is made in section D24.7.14 “PMSNEVFR_EL1, Sampling Inverted Event Filter Register”.

2.44 D23462

In section J1.3 “Shared Pseudocode”, in the pseudocode function IsDataAccess(), the code that reads:

```
boolean IsDataAccess(AccessType acctype)
    return ! acctype IN {AccessType_IFETCH,
                        AccessType_TTW,
                        AccessType_DC,
                        AccessType_IC,
                        AccessType_AT};
```

is changed to read:

```
boolean IsDataAccess(AccessType acctype)
    return ! acctype IN {AccessType_IFETCH,
                        AccessType_TTW,
                        AccessType_DC,
                        AccessType_IC,
                        AccessType_SPE,
                        AccessType_TRBE,
```

AccessType_AT);

2.45 C23479

In section A2.2.5 “The Armv8.4 architecture extension”, the text that reads:

FEAT_BBM, Translation table break-before-make levels

FEAT_BBM provides support to identify the requirements of hardware to have break-before-make sequences when changing between block size for a translation.

This feature is supported in AArch64 state only.

FEAT_BBM is OPTIONAL from Armv8.3.

FEAT_BBM is mandatory from Armv8.4.

The following field identifies the presence of FEAT_BBM:

- ID_AA64MMFR2_EL1.BBM.

For more information, see:

- Block descriptor and Page descriptor formats.
- Block translation entry.
- Support levels for changing block size..

is changed to read:

FEAT_BBML1, Translation table break-before-make level 1

FEAT_BBML1 provides support for reducing the requirements for following a break-before-make sequence when changing between block sizes for a translation.

This feature is supported in AArch64 state only.

FEAT_BBML1 is OPTIONAL from Armv8.3.

The following field identifies the presence of FEAT_BBML1:

- ID_AA64MMFR2_EL1.BBM.

For more information, see:

- Block descriptor and Page descriptor formats.
- Block translation entry.
- Support levels for changing block size.

FEAT_BBML2, Translation table break-before-make level 2

FEAT_BBML2 provides support for further reducing the requirements for following break-before-make sequences when changing between block sizes for a translation.

This feature is supported in AArch64 state only.

If FEAT_BBML2 is implemented, FEAT_BBML1 is implemented.

The following field identifies the presence of FEAT_BBML2:

- ID_AA64MMFR2_EL1.BBM.

For more information, see:

- FEAT_BBML1.

In section D8.3.1.2 “VMSAv8-64 Block descriptor and Page descriptor formats”, under rule R_{JJNHR}, the text in the table that reads:

Bit position	Field	Condition
Block[16]	RES0	FEAT_BBM is not implemented
Block[16]	nT	FEAT_BBM is implemented

is changed to read:

Bit position	Field	Condition
Block[16]	RES0	FEAT_BBML1 is not implemented
Block[16]	nT	FEAT_BBML1 is implemented

A similar table update is applied to rule R_{DMBGN}.

In section D8.7.3 “Block translation entry”, under rule R_{DFJNG}, the text that reads:

R_{DFJNG}

All statements in this section require implementation of FEAT_BBM.

is changed to read:

R_{DFJNG}

All statements in this section require implementation of FEAT_BBML1.

In the same section, under rule R_{MRRPW}, the text that reads:

R_{MRRPW}

If the implementation meets either level 1 or level 2 support requirements for changing table or block size, then when using a Table descriptor or Block descriptor with the nT bit set, the PE is permitted to do one of the following:

- Generate a Translation fault and not cache the entry in a TLB.
- If an entry that does not have the nT bit set is cached within a TLB and translates the same address to the same output address with consistent memory attributes and permissions, then the PE guarantees that accesses translated by the translation table entry with the nT bit set does not break coherency, ordering guarantees or uniprocessor semantics, or fail to clear the Exclusives monitors. For more information, see Support levels for changing table or block size.

is changed to read:

R_{MRRPW}

When using a Table descriptor or Block descriptor with the nT bit set, the PE is permitted to do one of the following:

- Generate a Translation fault and not cache the entry in a TLB.
- If an entry that does not have the nT bit set is cached within a TLB and translates the same address to the same output address with consistent memory attributes and permissions, then the PE guarantees that accesses translated by the translation table entry with the nT bit set does not break coherency, ordering guarantees or uniprocessor semantics, or fail to clear the Exclusives monitors. For more information, see Support levels for changing table or block size.

In section D8.15.1.1 “Translation fault”, under rule R_{FDQJL}, the text that reads:

R_{FDQJL}

If FEAT_BBML1 is implemented and supported at level 1 or higher, and the Block descriptor has the nT bit set, then the implementation is permitted to generate a Translation fault.

is changed to read:

R_{FDQJL}

If FEAT_BBML1 is implemented, and the Block descriptor has the nT bit set, then the implementation is permitted to generate a Translation fault.

In section D8.15.4 “MMU fault-checking sequence”, in the informational statement I_{PTYRT}, the text that reads:

1. If FEAT_BBML1 is implemented and supported at level 1 or higher, and the fetched descriptor is a block descriptor with the nT bit set, then the implementation can generate a Translation fault.

is changed to read:

1. If FEAT_BBML1 is implemented, and the fetched descriptor is a block descriptor with the nT bit set, then the implementation can generate a Translation fault.

In section D8.17.1 “Using break-before-make when updating translation table entries”, under rule R_{WHZWS}, the text that reads:

R_{WHZWS}

If the requirements enabled by FEAT_BBM level 1 or above cannot be followed, all of the following changes to the block size used by the translation system:

- Changing from a smaller size to a larger size, such as when a stage 2 Table descriptor is replaced with a Block descriptor.
- Changing from a larger size to a smaller size, such as when a stage 2 Block descriptor is replaced with a Table descriptor.

is changed to read:

R_{WHZWS}

If FEAT_BBML1 is not implemented, all of the following changes to the block size used by the translation system:

- Changing from a smaller size to a larger size, such as when a stage 2 Table descriptor is replaced with a Block descriptor.
- Changing from a larger size to a smaller size, such as when a stage 2 Block descriptor is replaced with a Table descriptor.

In section D8.17.2 “Support levels for changing table or block size”, the following rule is removed:

R_{KNVDX}

All statements in this section require implementation of FEAT_BBM.

In the same section, under rule R_{KFLJB} , the text that reads:

R_{KFLJB}

When a translation table entry is modified to change the table or block size, the hardware provides one of the following possible support levels affecting the break-before-make requirement to avoid breaking coherency, ordering guarantees or uniprocessor semantics, or failing to clear the Exclusives monitors:

- If level 0 is supported, then software is required to use break-before-make.
- If level 1 is supported, then software can use the level 0 approach or use the block translation entry bit, nT, in the Table descriptor or Block descriptor.
- If level 2 is supported, then all of the following apply:
 - Software can use the level 0 approach or the level 1 approach.
 - Changing table or block size does not break coherency, ordering guarantees or uniprocessor semantics, or fail to clear the Exclusives monitors. For more information, see Block translation entry.

is changed to read:

R_{KFLJB}

When a translation table entry is modified to change the table or block size, the hardware provides one of the following possible implementations affecting the break-before-make

requirement to avoid breaking coherency, ordering guarantees or uniprocessor semantics, or failing to clear the Exclusives monitors:

- If FEAT_BBML1 is not implemented, then software is required to use the break-before-make sequence.
- If FEAT_BBML1 is implemented, then software can use all of the following: - The break-before-make sequence. - The nT bit in the Table descriptor or Block descriptor.
- If FEAT_BBML2 is implemented, then software can change table or block size without breaking coherency, ordering guarantees or uniprocessor semantics, or failing to clear the Exclusives monitors.

In the same section, rule I_{HYQMB} the text that reads:

I_{HYQMB}

If any level is supported and the TLB entries are not invalidated after the writes that modified the translation table entries are completed, then a TLB conflict abort can be generated because in a TLB there might be multiple translation table entries that all translate the same IA. For Table descriptors, this also applies to intermediate TLB caching structures. For more information, see TLB conflict abort.

is changed to read:

I_{HYQMB}

If the TLB entries are not invalidated after the writes that modified the translation table entries are completed, then a TLB conflict abort can be generated because in a TLB there might be multiple translation table entries that all translate the same IA. For Table descriptors, this also applies to intermediate TLB caching structures. For more information, see TLB conflict abort.

In the same section, under rule R_{KHRBC}, the text that reads:

R_{KHRBC}

If level 1 or level 2 is supported, then changing the Contiguous bit in a set of Block descriptors or Page descriptors can be done without breaking coherency, ordering guarantees or uniprocessor semantics, or failing to clear the Exclusives monitors.

is changed to read:

R_{KHRBC}

If FEAT_BBML1 is implemented, then changing the Contiguous bit in a set of Block descriptors or Page descriptors can be done without breaking coherency, ordering guarantees or uniprocessor semantics, or failing to clear the Exclusives monitors.

In the same section, under rule R_{FCPSG}, the text that reads:

R_{FCPSG}

If level 1 or level 2 is supported and the Contiguous bit in a set of Block descriptors or Page descriptors is changed, then a TLB conflict abort can be generated because multiple translation table entries might exist within a TLB that translates the same IA.

is changed to read:

R_{FCPSG}

If FEAT_BBML1 is implemented and the Contiguous bit in a set of Block descriptors or Page descriptors is changed, then a TLB conflict abort can be generated because multiple translation table entries might exist within a TLB that translates the same IA.

In the same section, rule R_{FWRMB} that reads:

R_{FWRMB}

If all of the following apply, then a TLB conflict abort is reported to EL2:

- Level 1 or level 2 is supported.
- Stage 2 translations are enabled in the current translation regime.
- A TLB conflict abort is generated due to changing the block size or Contiguous bit.

is changed to read:

R_{FWRMB}

If all of the following apply, then a TLB conflict abort is reported to EL2:

- FEAT_BBML1 is implemented.
- Stage 2 translations are enabled in the current translation regime.
- A TLB conflict abort is generated due to changing the block size or Contiguous bit.

In the same section, in the informational statement I_{CFVVK}, the text that reads:

I_{CFVVK}

If level 1 or level 2 is supported and a TLB conflict abort is generated, then TLB maintenance is required to remove the multiple TLB entries that translate the same address. For Table descriptors, this also applies to intermediate TLB caching structures. For more information, see TLB maintenance due to TLB conflict.

is changed to read:

I_{CFVVK}

If FEAT_BBML1 is implemented and a TLB conflict abort is generated, then TLB maintenance is required to remove the multiple TLB entries that translate the same address. For Table descriptors, this also applies to intermediate TLB caching structures. For more information, see TLB maintenance due to TLB conflict.

In section D9.4.4 “GPT Contiguous descriptor”, under informational statement I_{JMVJS}, the text that reads:

I_{JMVJS}

This behavior is intended to be the same as the level 2 behavior that is specified in the FEAT_BBM feature, but with the option of TLB Conflict aborts removed. For more information, see Support levels for changing table or block size.

is changed to read:

I_{JMVJS}

This behavior is intended to be the same as FEAT_BBML2 behavior, but with the option of TLB Conflict aborts removed. For more information, see Support levels for changing table or block size.

In section D24.2.84 “ID_AA64MMFR2_EL1, AArch64 Memory Model Feature Register 2”, the text that reads:

BBM, bits [55:52]

Allows identification of the requirements of the hardware to have break-before-make sequences when changing block size for a translation.

BBM	Meaning
0b0000	Level 0 support for changing block size is supported.
0b0001	Level 1 support for changing block size is supported.
0b0010	Level 2 support for changing block size is supported.

All other values are reserved.

FEAT_BBM implements the functionality identified by the values 0b0000, 0b0001, and 0b0010.

From Armv8.4, the permitted values are 0b0000, 0b0001, and 0b0010.

Access to this field is RO.

is changed to read:

BBM, bits [55:52]

Allows identification of the requirements of the hardware to have break-before-make sequences when changing block or table size for a translation.

BBM	Meaning
0b0000	Break-before-make sequence must be used.
0b0001	Level 1 support for changing block size is supported.
0b0010	Level 2 support for changing block size is supported.

All other values are reserved.

FEAT_BBML1 implements the functionality identified by the value 0b0001.

FEAT_BBML2 implements the functionality identified by the value 0b0010.

Access to this field is RO.

2.46 D23518

In section A.2.2.5 “The Armv8.4 architecture extension”, the text that reads:

FEAT_MPAM, Memory Partitioning and Monitoring Extension

The MPAM Extension provides a framework for memory-system component controls that partition one or more of the performance resources of the component.

This feature is supported in AArch64 state only.

FEAT_MPAM is OPTIONAL from Armv8.2.

The following field identifies the presence of FEAT_MPAM:

- ID_AA64PFR0_EL1.MPAM.

is changed to read:

FEAT_MPAMv1p0, Memory Partitioning and Monitoring Extension version 1.0

FEAT_MPAMv1p0 introduces support for version 1.0 of the MPAM extension.

This feature is supported in AArch64 state only.

FEAT_MPAMv1p0 is OPTIONAL from Armv8.2.

The following field identifies the presence of FEAT_MPAMv1p0:

- ID_AA64PFR0_EL1.MPAM.
- ID_AA64PFR1_EL1.MPAM_frac.

The following text is added to the same section:

FEAT_MPAM, The Memory Partitioning and Monitoring Extension.

The MPAM Extension provides a framework for memory-system component controls that partition one or more of the performance resources of the component.

This feature is supported in AArch64 state only.

FEAT_MPAM is OPTIONAL from Armv8.2.

There are three versions of the MPAM extension: v1p0, v0p1, and v1p1.

The following fields identifies the presence of FEAT_MPAM:

- ID_AA64PFR0_EL1.MPAM.
- ID_AA64PFR1_EL1.MPAM_frac.

2.47 D23543

In section J1.1 “Pseudocode for AArch64 operation”, in function AArch64.MemSingleRead(), the code that reads:

```
(bits(size*8), AddressDescriptor, PhysMemRetStatus)
AArch64.MemSingleRead(bits(64) address,
                      integer size,
                      AccessDescriptor accdesc_in,
                      boolean aligned)
...
elseif accdesc.acctype == AccessType_ASIMD && size == 32 && accdesc.ispair
then
...
  for i = 0 to 3 do
    (memstatus, value<i*64+:64>) = PhysMemRead(memaddrdesc, 8, accdesc);
    if IsFault(memstatus) then
      return (value, memaddrdesc, memstatus);
    memaddrdesc.paddress.address = memaddrdesc.paddress.address + 8;
    memaddrdesc.vaddress = memaddrdesc.vaddress + 8;
  elseif aligned && accdesc.ispair then
    ...
    memaddrdesc.paddress.address = memaddrdesc.paddress.address + halfsize;
    memaddrdesc.vaddress = memaddrdesc.vaddress + halfsize;
    (memstatus, highhalf) = PhysMemRead(memaddrdesc, halfsize, accdesc);
    if IsFault(memstatus) then
      return (value, memaddrdesc, memstatus);
    value = highhalf:lowhalf;
  else
    for i = 0 to size-1 do
      (memstatus, Elem[value, i, 8]) = PhysMemRead(memaddrdesc, 1, accdesc);
      if IsFault(memstatus) then
        return (value, memaddrdesc, memstatus);
      memaddrdesc.paddress.address = memaddrdesc.paddress.address + 1;
      memaddrdesc.vaddress = memaddrdesc.vaddress + 1;
```

is changed to read:

```
(bits(size*8), AddressDescriptor, PhysMemRetStatus)
AArch64.MemSingleRead(bits(64) address,
                      integer size,
                      AccessDescriptor accdesc_in,
                      boolean aligned)
...
elseif accdesc.acctype == AccessType_ASIMD && size == 32 && accdesc.ispair
then
...
  for i = 0 to 3 do
    (memstatus, value<i*64+:64>) = PhysMemRead(memaddrdesc, 8, accdesc);
    if IsFault(memstatus) then
```

```

    return (value, memaddrdesc, memstatus);
    memaddrdesc.paddress.address = memaddrdesc.paddress.address + 8;
    memaddrdesc.vaddress = memaddrdesc.vaddress + 8;
elseif aligned && accdesc.ispair then
    ...
    memaddrdesc.paddress.address = memaddrdesc.paddress.address + halfsize;
    memaddrdesc.vaddress = memaddrdesc.vaddress + halfsize;
    (memstatus, highhalf) = PhysMemRead(memaddrdesc, halfsize, accdesc);
    if IsFault(memstatus) then
        return (value, memaddrdesc, memstatus);
    value = highhalf:lowhalf;
else
    for i = 0 to size-1 do
        (memstatus, Elem[value, i, 8]) = PhysMemRead(memaddrdesc, 1, accdesc);
        if IsFault(memstatus) then
            return (value, memaddrdesc, memstatus);
        memaddrdesc.paddress.address = memaddrdesc.paddress.address + 1;
        memaddrdesc.vaddress = memaddrdesc.vaddress + 1;

```

Equivalent changes are made in the following functions:

- AArch64.MemSingleWrite().
- MemStore64B().
- MemLoad64B().
- AArch64.WriteTagMem().

2.48 R23545

In section D14.3.2 “Common microarchitectural events”, the following text in the description of ‘0x8074, SVE_PRED_SPEC, Operation speculatively executed, SVE predicated’:

0x8074, SVE_PRED_SPEC, Operation speculatively executed, SVE predicated

The counter counts each speculatively executed SIMD data-processing, load, or store operation due to an instruction with a single Governing predicate operand that determines the Active elements.

When FEAT_SME is implemented, both operations due to SVE instructions and operations due to SME instructions operating on the SVE Z vectors with at least one Governing predicate operand are counted.

Note:

For outer product instructions which are widening, predication is considered with respect to the input element size.

Is changed to read:

0x8074, SVE_PRED_SPEC, Operation speculatively executed, SVE predicated

The counter counts each speculatively executed predicated SIMD data-processing, load, or store operation. This includes all of the following:

- A data processing or load operation that writes to one or more SVE Z vector destination registers under a Governing predicate using either zeroing or merging predication.
- A predicated store of one or more SVE Z vector registers.

It is **IMPLEMENTATION DEFINED** whether data processing operations due to instructions with a single Governing predicate operand that determines the Active elements that do not write to any SVE Z vector destination registers using either zeroing or merging predication are counted. For example, INCP. Arm recommends these operations are not counted when FEAT_SPE is implemented, for consistency with the Operation Type packet PRED field.

When FEAT_SME is implemented, both operations due to SVE instructions and operations due to SME instructions operating on the SVE Z vectors with at least one Governing predicate operand are counted.

Note:

For outer product instructions which are widening, predication is considered with respect to the input element size.

2.49 D23551

In section C6.2.96 “CPYP, CPYM, CPYE”, the text that reads:

For CPYP:

- After saturation is performed, the direction of the memory copy is based on the following:
 - If $(X_s > X_d)$ and $(X_d + \text{saturated copy size}) > X_s$, then the direction is forward.
 - If $(X_s < X_d)$ and $(X_s + \text{saturated copy size}) > X_d$, then the direction is backward.
 - Otherwise, the direction is an **IMPLEMENTATION DEFINED** choice between forward and backward.

is changed to read:

For CPYP:

- After saturation is performed, the direction of the memory copy is based on the following:
 - If $(X_s \langle 55:0 \rangle > X_d \langle 55:0 \rangle)$ and $(X_d \langle 55:0 \rangle + \text{saturated copy size}) > X_s \langle 55:0 \rangle$, then the direction is forward.
 - If $(X_s \langle 55:0 \rangle < X_d \langle 55:0 \rangle)$ and $(X_s \langle 55:0 \rangle + \text{saturated copy size}) > X_d \langle 55:0 \rangle$, then the direction is backward.
 - Otherwise, the direction is an **IMPLEMENTATION DEFINED** choice between forward and backward.

The equivalent changes are made in the following sections:

- C6.2.97 “CPYPN, CPYMN, CPYEN”.

- C6.2.98 “CPYPRN, CPYMRN, CPYERN”.
- C6.2.99 “CPYPRT, CPYMRT, CPYERT”.
- C6.2.100 “CPYPRTN, CPYMRTN, CPYERTN”.
- C6.2.101 “CPYPRTRN, CPYMRTRN, CPYERTRN”.
- C6.2.102 “CPYPRTWN, CPYMRTWN, CPYERTWN”.
- C6.2.103 “CPYPT, CPYMT, CPYET”.
- C6.2.104 “CPYPTN, CPYMTN, CPYETN”.
- C6.2.105 “CPYPTRN, CPYMTRN, CPYETRN”.
- C6.2.106 “CPYPTWN, CPYMTWN, CPYETWN”.
- C6.2.107 “CPYPWN, CPYMWN, CPYEWN”.
- C6.2.108 “CPYPWT, CPYMWT, CPYEWT”.
- C6.2.109 “CPYPWTN, CPYMWTN, CPYEWTN”.
- C6.2.110 “CPYPWTRN, CPYMWTRN, CPYEWTRN”.
- C6.2.111 “CPYPWTWN, CPYMWWTWN, CPYEWWTWN”.

2.50 C23552

In section C6.2.96 “CPYP, CPYM, CPYE”, the text that reads:

For the CPYE option B, when PSTATE.C = 1:

-
-
- If the copy is in the backward direction (PSTATE.N == 1), then: - Xs holds the highest address to be copied from. - Xd holds the highest address to be copied to. - On completion of the instruction:

is changed to read:

For the CPYE option B, when PSTATE.C = 1:

-
-
- If the copy is in the backward direction (PSTATE.N == 1), then: - Xs holds the highest address to be copied from + 1. - Xd holds the highest address to be copied to + 1. - On completion of the instruction:

Equivalent changes are made in the following sections:

- C6.2.97 “CPYPN, CPYMN, CPYEN”.
- C6.2.98 “CPYPRN, CPYMRN, CPYERN”.

- C6.2.99 “CPYPRT, CPYMRT, CPYERT”.
- C6.2.100 “CPYPRTN, CPYMRTN, CPYERTN”.
- C6.2.101 “CPYPRTN, CPYMRTN, CPYERTN”.
- C6.2.102 “CPYPRTWN, CPYMRTWN, CPYERTWN”.
- C6.2.103 “CPYPT, CPYMT, CPYET”.
- C6.2.104 “CPYPTN, CPYMTN, CPYETN”.
- C6.2.105 “CPYPTRN, CPYMTRN, CPYETRN”.
- C6.2.106 “CPYPTWN, CPYMTWN, CPYETWN”.
- C6.2.107 “CPYPWN, CPYMWN, CPYEWN”.
- C6.2.108 “CPYPWT, CPYMWT, CPYETWT”.
- C6.2.109 “CPYPWTN, CPYMWTN, CPYETWTN”.
- C6.2.110 “CPYPWTRN, CPYMWTRN, CPYEWTRN”.
- C6.2.111 “CPYPWTWN, CPYMWWTWN, CPYEWWTWN”.

2.51 C23586

In section B2.3.1 “Basic definitions”, the definition ‘Fault Effects, Exception Entry and Exception Return Effects’ that reads:

Fault Effects, Exception Entry, and Exception Return Effects

An instruction that encounters a fault generates a Fault Effect. An instruction that encounters a fault leading to a synchronous exception generates a Fault Effect that is also an Exception Entry Effect, for example as a result of reading an invalid TTD or fetching an illegal instruction.

An ERET instruction generates an Exception Return Effect.

A Fault Effect generated by the MMU is called an MMU Fault Effect.

A Fault Effect generated by a failed Tag Check is called a TagCheck Fault Effect.

is changed to read:

Fault Effects, Exception Entry, and Exception Return Effects

An instruction that encounters a fault generates a Fault Effect. An instruction that encounters a fault leading to a synchronous exception generates a Fault Effect that is also an Exception Entry Effect, for example as a result of reading an invalid TTD or fetching an illegal instruction.

An ERET instruction generates an Exception Return Effect.

A Fault Effect generated by the MMU is called an MMU Fault Effect.

A Fault Effect generated by a failed Tag Check is called a TagCheck Fault Effect.

Instructions that write to memory with Release semantics generate Fault Effects with Release semantics when they encounter a fault.

In section B2.3.7 “Ordering relations”, the definition of ‘Barrier-ordered-before’ that reads:

Barrier-ordered-before

An Effect E_1 is Barrier-ordered-before an Effect E_2 if one of the following applies:

- All of the following apply:
 - E_1 is an Explicit Memory Write Effect and is generated by an atomic instruction with both Acquire and Release semantics.
 - E_1 appears in program order before E_2 .
 - One of the following applies:
 - E_2 is an Explicit Memory Effect.
 - E_2 is an Implicit Tag Memory Read Effect.
 - E_2 is an MMU Fault Effect.
- All of the following apply:
 - E_1 is an Explicit Memory Effect or a Fault Effect and generated by an instruction with Release semantics.
 - E_1 appears in program order before E_2 .
 - E_2 is an Explicit Memory Effect generated by an instruction with Acquire semantics.
- All of the following apply:
 - One of the following applies:
 - E_1 is an Explicit Memory Effect generated by an instruction with Acquire semantics.
 - E_1 is an Explicit Memory Effect generated by an instruction with AcquirePC semantics.
 - E_1 appears in program order before E_2 .
 - One of the following applies:
 - E_2 is an Explicit Memory Effect.
 - E_2 is an Implicit Tag Memory Read Effect.
 - E_2 is an MMU Fault Effect.
- All of the following apply:
 - One of the following applies:
 - E_1 is an Explicit Memory Effect generated by an instruction with Acquire semantics.
 - E_1 is an Explicit Memory Effect generated by an instruction with AcquirePC semantics.
 - There is an Intrinsic Order Dependency from E_1 to E_2 .
 - One of the following applies:
 - E_2 is an Explicit Memory Effect.

- E_2 is an Implicit Tag Memory Read Effect.
- E_2 is an MMU Fault Effect.
- All of the following apply:
 - One of the following applies:
 - E_1 is an Explicit Memory Effect.
 - E_1 is an Implicit Tag Memory Read Effect.
 - E_1 appears in program order before E_2 .
 - E_2 is an Explicit Memory Effect or a Fault Effect and generated by an instruction with Release semantics.
- All of the following apply:
 - One of the following applies:
 - E_1 is an Explicit Memory Effect.
 - E_1 is an Implicit Tag Memory Read Effect.
 - There is an Intrinsic Order Dependency from E_1 to E_2 .
 - E_2 is an Explicit Memory Effect or a Fault Effect and generated by an instruction with Release semantics.

is changed to read:

Barrier-ordered-before

An Effect E_1 is Barrier-ordered-before an Effect E_2 if one of the following applies:

- All of the following apply:
 - E_1 is an Effect with Release semantics.
 - E_1 and E_3 are generated by an atomic instruction.
 - E_3 is an Effect with Acquire semantics.
 - E_1 appears in program order before E_2 .
 - One of the following applies:
 - E_2 is an Explicit Memory Effect.
 - E_2 is an Implicit Tag Memory Read Effect.
 - E_2 is an MMU Fault Effect.
- All of the following apply:
 - E_1 is an Effect with Release semantics.
 - E_1 appears in program order before E_2 .
 - E_2 is an Effect with Acquire semantics.
- All of the following apply:
 - One of the following applies:

- E_1 is an Effect with Acquire semantics.
- E_1 is an Effect with AcquirePC semantics.
- E_1 appears in program order before E_2 .
- One of the following applies:
 - E_2 is an Explicit Memory Effect.
 - E_2 is an Implicit Tag Memory Read Effect.
 - E_2 is an MMU Fault Effect.
- All of the following apply:
 - One of the following applies:
 - E_1 is an Effect with Acquire semantics.
 - E_1 is an Effect with AcquirePC semantics.
 - There is an Intrinsic Order Dependency from E_1 to E_2 .
 - One of the following applies:
 - E_2 is an Explicit Memory Effect.
 - E_2 is an Implicit Tag Memory Read Effect.
 - E_2 is an MMU Fault Effect.
- All of the following apply:
 - One of the following applies:
 - E_1 is an Explicit Memory Effect.
 - E_1 is an Implicit Tag Memory Read Effect.
 - E_1 appears in program order before E_2 .
 - E_2 is an Effect with Release semantics.
- All of the following apply:
 - One of the following applies:
 - E_1 is an Explicit Memory Effect.
 - E_1 is an Implicit Tag Memory Read Effect.
 - There is an Intrinsic Order Dependency from E_1 to E_2 .
 - E_2 is an Effect with Release semantics.

In the same section, the definition of 'Atomic-ordered-before' that reads:

Atomic-ordered-before

An Effect E_1 is Atomic-ordered-before an Effect E_2 if one of the following applies:

- All of the following apply:
 - E_1 is an Explicit Memory Effect.
 - E_1 and E_3 form a successful Read-Modify-Write pair.

- E_2 is a Local read successor of E_3 .
- One of the following applies:
 - E_2 is an Explicit Memory Effect generated by an instruction with Acquire semantics.
 - E_2 is an Explicit Memory Effect generated by an instruction with AcquirePC semantics.

is changed to read:

Atomic-ordered-before

An Effect E_1 is Atomic-ordered-before an Effect E_2 if one of the following applies:

- All of the following apply:
 - E_1 is an Explicit Memory Effect.
 - E_1 and E_3 form a successful Read-Modify-Write pair.
 - E_2 is a Local read successor of E_3 .
 - One of the following applies:
 - E_2 is an Effect with Acquire semantics.
 - E_2 is an Effect with AcquirePC semantics.

2.52 C23587

In section B2.4.1.1 “Peripherals”, the definition of Out-of-band-ordered-before that reads:

A Read Memory or write effect RW_1 is Out-of-band-ordered-before a Read Memory or write effect RW_2 if either of the following cases apply:

- RW_1 appears in program order before a DSB instruction that begins an **IMPLEMENTATION DEFINED** instruction sequence indirectly leading to the generation of RW_2 .
- RW_1 is Ordered-before a Read Memory or write effect RW_3 and RW_3 is Out-of-band-ordered-before RW_2 .

If a Memory effect M_1 is Out-of-band-ordered-before a Read Memory or write effect M_2 , then M_1 is seen to occur before M_2 by all observers.

is changed to read:

An Effect E_1 is Out-of-band-ordered-before an Effect E_2 if and only if all of the following apply:

- E_1 is a Memory Effect.
- Any of the following applies:
 - E_1 is DSB-ordered-before E_3 .
 - E_1 is Instruction-fetch-barrier-ordered-before E_3 .
 - There is an **IMPLEMENTATION DEFINED** instruction sequence from E_3 indirectly leading to the generation of E_2 .

- E_2 is a Memory Effect.

If an Effect E_1 is Out-of-band-ordered-before an Effect E_2 , then E_1 is Ordered-before E_2 .

2.53 D23647

In section C5.5.21 “TLBI RIPAS2E1, TLBI RIPAS2E1NXS, TLB Range Invalidate by Intermediate Physical Address, Stage 2, EL1”, in the description of the field BaseADDR, bits [36:0], the text that reads:

When (FEAT_LPA2 is implemented and TCR_EL1.DS == 1) or (FEAT_D128 is implemented and VTCR_EL2.D128 == 1):

The starting address for the range of the maintenance instructions. This field is BaseADDR[52:16] for all translation granules.

- When using a 4KB translation granule, BaseADDR[15:12] is treated as 0b0000.
- When using a 16KB translation granule, BaseADDR[15:14] is treated as 0b00.

is changed to read:

When (FEAT_LPA2 is implemented and VTCR_EL2.DS == 1) or (FEAT_D128 is implemented and VTCR_EL2.D128 == 1):

The starting address for the range of the maintenance instructions. This field is BaseADDR[52:16] for all translation granules.

- When using a 4KB translation granule, BaseADDR[15:12] is treated as 0b0000.
- When using a 16KB translation granule, BaseADDR[15:14] is treated as 0b00.

The equivalent changes are made in the following sections:

- C5.5.22 “TLBI RIPAS2E1IS, TLBI RIPAS2E1ISNXS”.
- C5.5.23 “TLBI RIPAS2E1OS, TLBI RIPAS2E1OSNXS”.
- C5.5.24 “TLBI RIPAS2LE1, TLBI RIPAS2LE1NXS”.
- C5.5.25 “TLBI RIPAS2LE1IS, TLBI RIPAS2LE1ISNXS”.
- C5.5.26 “TLBI RIPAS2LE1OS, TLBI RIPAS2LE1OSNXS”.

2.54 D23659

In section B2.5.4 “Restrictions on the effects of speculation from Armv8.5”, the text that reads:

If either FEAT_CSV2_1p1 or FEAT_CSV2_3 is implemented, code running in one hardware-defined context (context1) cannot either exploitatively control, or predictively leak to, the

speculative execution of code in a different hardware-defined context (context2) as a result of the behavior of branch target prediction based on the branch history used in context1.

is changed to read:

If either FEAT_CSV2_1p1 or FEAT_CSV2_3 is implemented, code running in one hardware-defined context (context1) cannot either exploitatively control, or predictively leak to, the speculative execution of code in a different hardware-defined context (context2) as a result of the behavior of any prediction mechanism based on the branch history used in context1.

The contents of section D7.5.13 “Branch prediction” are moved into a new section named “Branch history” under B2.5 “Restrictions on the effects of speculation”.

This leaves D7.5.13 “Branch prediction” to read as follows:

If branch prediction is architecturally visible, cache maintenance must also apply to branch prediction.

The exact change will be made available at a later date as C23575.

The new “Branch history” section under B2.5 “Restrictions on the effects of speculation” reads:

If FEAT_CLRBHB is not implemented, then the architecture does not define any branch history maintenance instructions for AArch64 state.

When FEAT_CLRBHB is implemented, the CLRBHB instruction is available. When the CLRBHB instruction is executed, the branch history is cleared for the current context to the extent that branch history information created before the CLRBHB instruction cannot be used by code before the CLRBHB instruction to exploitatively control the execution of any code in the current context appearing in program order after the instruction.

When FEAT_ECBHB is implemented, the branch history information created in a context before an exception entry to a higher Exception level using AArch64 cannot be used by code before that exception entry to exploitatively control the execution of any code in a different context after the exception entry.

Note:

FEAT_CLRBHB and FEAT_ECBHB apply to all forms of prediction that use the branch history and that are protected by FEAT_CSV2.

In section A2.2.10 “The Armv8.9 architecture extension”, the text under ‘FEAT_CLRBHB, Support for Clear Branch History instruction’ that reads:

FEAT_CLRBHB provides a CLRBHB instruction, which clears the branch history for the current context to the extent that branch history information created before the CLRBHB instruction cannot be used by code before the CLRBHB instruction to exploitatively control the execution of any indirect branches in code in the current context that appear in program order after the instruction.

is changed to read:

FEAT_CLRBHB provides a CLRBHB instruction which can be used to clear the branch history.

In section D24.2.83 “ID_AA64MMFR1_EL1, AArch64 Memory Model Feature Register 1”, in the field description of ‘ECBHB, bits [63:60]’, the table that reads:

ECBHB	Meaning
0b0000	The implementation does not disclose whether the branch history information created in a context before an exception to a higher Exception level using AArch64 can be used by code before that exception to exploitatively control the execution of any indirect branches in code in a different context after the exception.
0b0001	The branch history information created in a context before an exception to a higher Exception level using AArch64 cannot be used by code before that exception to exploitatively control the execution of any indirect branches in code in a different context after the exception.

is changed to read

ECBHB	Meaning
0b0000	The implementation does not disclose whether restrictions are imposed on branch history speculation around exceptions.
0b0001	The implementation imposes restrictions on branch history speculation around exceptions.

2.55 D23680

In section D24.10.4 “CNTHP_CVAL_EL2, Counter-timer Physical Timer CompareValue Register (EL2)”, the text that reads:

When CNTHP_CTL_EL2.ENABLE is 1, and TimerConditionMet is TRUE for the EL1 physical timer, the timer condition is met and all of the following are true:

is changed to read:

When CNTHP_CTL_EL2.ENABLE is 1, and TimerConditionMet is TRUE for the EL2 physical timer, the timer condition is met and all of the following are true:

In section D24.10.7 “CNTHPS_CVAL_EL2, Counter-timer Secure Physical Timer CompareValue Register (EL2)”, the text that reads:

When CNTHPS_CTL_EL2.ENABLE is 1, and TimerConditionMet is TRUE for the EL1 physical timer, the timer condition is met and all of the following are true:

is changed to read:

When CNTHPS_CTL_EL2.ENABLE is 1, and TimerConditionMet is TRUE for the Secure EL2 physical timer, the timer condition is met and all of the following are true:

In section D24.10.10 “CNTHV_CVAL_EL2, Counter-timer Virtual Timer CompareValue Register (EL2)”, the text that reads:

When CNTHV_CTL_EL2.ENABLE is 1, and TimerConditionMet is TRUE for the EL1 physical timer, the timer condition is met and all of the following are true:

is changed to read:

When CNTHV_CTL_EL2.ENABLE is 1, and TimerConditionMet is TRUE for the EL2 virtual timer, the timer condition is met and all of the following are true:

In section D24.10.13 “CNTHVS_CVAL_EL2, Counter-timer Secure Virtual Timer CompareValue Register (EL2)”, the text that reads:

When CNTHVS_CTL_EL2.ENABLE is 1, and TimerConditionMet is TRUE for the EL1 physical timer, the timer condition is met and all of the following are true:

is changed to read:

When CNTHVS_CTL_EL2.ENABLE is 1, and TimerConditionMet is TRUE for the Secure EL2 virtual timer, the timer condition is met and all of the following are true:

In section D24.10.23 “CNTPS_CVAL_EL1, Counter-timer Physical Secure Timer CompareValue Register”, the text that reads:

When CNTPS_CTL_EL2.ENABLE is 1, and TimerConditionMet is TRUE for the EL1 physical timer, the timer condition is met and all of the following are true:

is changed to read:

When CNTPS_CTL_EL2.ENABLE is 1, and TimerConditionMet is TRUE for the EL1 secure physical timer, the timer condition is met and all of the following are true:

In section D24.10.26 “CNTV_CVAL_EL0, Counter-timer Virtual Timer CompareValue Register”, the text that reads:

When CNTV_CTL_EL2.ENABLE is 1, and TimerConditionMet is TRUE for the EL1 physical timer, the timer condition is met and all of the following are true:

is changed to read:

When CNTV_CTL_EL2.ENABLE is 1, and TimerConditionMet is TRUE for the EL1 virtual timer, the timer condition is met and all of the following are true:

2.56 D23683

In section D24.10.21 “CNTPOFF_EL2, Counter-timer Physical Offset Register”, under the heading ‘Purpose’, the text that reads:

Holds the 64-bit physical offset. This is the offset for the AArch64 physical timers and counters when Enhanced Counter Virtualization is enabled.

is changed to read:

Holds the 64-bit physical offset. This is the offset for the AArch64 EL1 physical timer and counter when Enhanced Counter Virtualization is enabled.

2.57 D23687

In section J1.1.1.113 “AArch64.WatchpointMatch”, the code that reads:

```
boolean value_match = FALSE;
for byte = 0 to size - 1 do
    value_match = value_match || AArch64.WatchpointByteMatch(n, vaddress +
byte);
```

is changed to read:

```
boolean value_match = FALSE;
watchptinfo.vaddress = vaddress;
for byte = 0 to size - 1 do
    if (!value_match &&
        !AddressInNaturallyAlignedBlock(watchptinfo.vaddress,
                                         vaddress + byte, 64)) then
        // watchpoint should report an address which is in
        // the naturally aligned block of the matched address
        watchptinfo.vaddress = ZeroExtend(vaddress + byte, 64);
    value_match = value_match || AArch64.WatchpointByteMatch(n, vaddress +
byte);
```

An equivalent change is made to function AArch32.WatchpointMatch().

2.58 D23689

In section D14.3.2 “Common microarchitectural events”, the text that reads:

0x802C, FP_MUL_SPEC, Floating-point operation speculatively executed, multiply

The counter counts each speculatively executed floating-point multiply operation counted by FP_SPEC due to any of the following A64 instructions:

- Scalar: FMUL or FMULX.
- Advanced SIMD: FMUL or FMULX.
- SVE: FMUL, FMULX, or FTSMUL.
- SVE2: BFMUL (indexed).

It is **IMPLEMENTATION DEFINED** which floating-point multiply operations are counted in AArch32 state.

is changed to read:

0x802C, FP_MUL_SPEC, Floating-point operation speculatively executed, multiply

The counter counts each speculatively executed floating-point multiply operation counted by FP_SPEC due to any of the following A64 instructions:

- Scalar: FMUL or FMULX.
- Advanced SIMD: FMUL or FMULX.
- SVE: FMUL, FMULX, or FTSMUL.
- SVE2: BFMUL.

Such that BFMUL (vectors, predicated) and BFMUL (vectors, unpredicated) are also counted.

It is **IMPLEMENTATION DEFINED** which floating-point multiply operations are counted in AArch32 state.

A corresponding update adds “SVE2: BFMUL.” to the instruction-list for both SVE_FP_MUL_SPEC (0x802E) and ASE_SVE_FP_MUL_SPEC (0x802F).

2.59 C23701

In section D24.2.205 “VMPIDR_EL2, Virtualization Multiprocessor ID Register”, under the heading ‘Purpose’, the text that reads:

Holds the value of the Virtualization Multiprocessor ID. This is the value returned by EL1 reads of MPIDR_EL1.

is changed to read:

Holds the value of the Virtualization Multiprocessor ID. This is the value returned by EL1 reads of MPIDR_EL1, which in a multiprocessor system, provides an additional PE identification system.

Additionally, the text that reads:

Affinity level 0. This is the affinity level that is most significant for determining PE behavior.

Higher affinity levels are increasingly less significant in determining PE behavior.

The assigned value of the MPIDR.{Aff2, Aff1, Aff0} or MPIDR_EL1.{Aff3, Aff2, Aff1, Aff0} set of fields of each PE must be unique within the system as a whole.

is changed to read:

Affinity level 0. The value of the MPIDR.{Aff2, Aff1, Aff0} or the MPIDR_EL1.{Aff3, Aff2, Aff1, Aff0} set of fields of each PE must be unique within the system as a whole.

The equivalent changes are made in G8.2.170 “VMPIDR, Virtualization Multiprocessor ID Register”.

Furthermore, in section D13.4 “Attributability”, the note that reads:

Note:

- In an implementation containing multiple PEs, each PE is identified by a unique affinity value reported by MPIDR_EL1.{Aff3, Aff2, Aff1, Aff0}, where the value of affinity level 0 is the most significant for determining the PE behavior, and the values of higher affinity levels are less significant. Affinity level 3 is only supported in AArch64 state.
- An implementation is described as multithreaded when the lowest level of affinity consists of logical PEs that are implemented using a multithreading type approach. In this section, when referring to a multithreaded implementation, thread is used to mean processing elements with:
 - MPIDR_EL1.MT or MPIDR.MT set to 1,
 - Different values for affinity level 0.
 - The same values for affinity level 1 and higher.

is changed to read:

Note:

An implementation is described as multithreaded when the lowest level of affinity, as reported by MPIDR_EL1.{Aff3, Aff2, Aff1, Aff0}, consists of logical PEs that are implemented using a multithreading type approach.

In this section, when referring to a multithreaded implementation, thread is used to mean processing elements with:

- MPIDR_EL1.MT or MPIDR.MT set to 1.
- Different values for affinity level 0.
- The same values for affinity level 1 and higher.

2.60 D23713

In section D14.3 “Common event numbers”, the text that reads:

0x8074, SVE_PRED_SPEC, Operation speculatively executed, SVE predicated

The counter counts each speculatively executed SIMD data-processing, load, or store operation due to an instruction with a single Governing predicate operand that determines the Active elements.

When FEAT_SME is implemented, both operations due to SVE instructions and operations due to SME instructions operating on the SVE Z vectors with at least one Governing predicate operand are counted.

Note:

For outer product instructions which are widening, predication is considered with respect to the input element size.

is changed to read:

0x8074, SVE_PRED_SPEC, Operation speculatively executed, SVE or SME predicated

The counter counts each speculatively executed predicated SVE or SME data-processing, load, or store operation.

All of the following SVE operations are counted:

- Data processing or load operations that write to one or more SVE Z vector destination registers under a Governing predicate using either zeroing or merging predication.
- Predicated stores of one or more SVE Z vector registers.

It is **IMPLEMENTATION DEFINED** whether data processing operations due to instructions with a single Governing predicate operand that determines the Active elements which do not write to any SVE Z vector destination registers using either zeroing or merging predication are counted. For example, INCP. When FEAT_SPE is implemented, Arm recommends this event is implemented consistently with the PRED field in the SVE data-processing format Operation Type packet.

All speculatively executed SME data-processing, load, or store operations due to instructions with at least one Governing predicate operand that determines the Active elements are counted.

Note:

For outer product instructions which are widening, predication is considered with respect to the input element size.

In the same section, the text that reads:

0x8394, SME_LDST_TILE_SPEC, SME predicated tile load/store

The counter counts each speculatively executed operation that reads from or writes to memory counted by SME_LDST_REG_SPEC that was due to an instruction with at least one governing predicate which is targeting the ZA array.

is changed to read:

0x8394, SME_LDST_TILE_SPEC, SME predicated tile load/store

The counter counts each speculatively executed operation that reads from or writes to memory counted by SME_INST_SPEC that was due to an instruction with at least one governing predicate which is targeting the ZA array.

Additionally, the text that reads:

0x8074, SVE_PRED_SPEC, Operation speculatively executed, SVE predicated

The counter counts each speculatively executed SIMD data-processing, load, or store operation due to an instruction with a single Governing predicate operand that determines the Active elements.

When FEAT_SME is implemented, both operations due to SVE instructions and operations due to SME instructions operating on the SVE Z vectors with at least one Governing predicate operand are counted.

...

is changed to read:

0x8074, SVE_PRED_SPEC, Operation speculatively executed, SVE or SME predicated

The counter counts each speculatively executed SVE data-processing, load, or store operation due to an instruction with at least one Governing predicate operand that determines the Active elements.

When FEAT_SME is implemented, any speculatively executed SME data-processing, load, or store operation due to an instruction with at least one Governing predicate operand that determines the Active elements is also counted.

...

The equivalent changes are made in D14.3 “Common event numbers” in the following event descriptions:

- '0x8075, SVE_PRED_EMPTY_SPEC, Operation speculatively executed, SVE predicated with no active elements'.
- '0x8076, SVE_PRED_FULL_SPEC, Operation speculatively executed, SVE predicated with all active elements'.
- '0x8077, SVE_PRED_PARTIAL_SPEC, Operation speculatively executed, SVE predicated with partially active elements'.
- '0x8078, SVE_UNPRED_SPEC, Operation speculatively executed, SVE unpredicated'.
- '0x8079, SVE_PRED_NOT_FULL_SPEC, Operation speculatively executed, SVE predicated with at least one inactive element'.

2.61 C23721

In section F5.1.67 “LDM, LDMIA, LDMFD”, the code that reads:

```
if ConditionPassed() then
    ...
    for i = 0 to 14
        if registers<i> == '1' then
            R[i] = MemS[address,4]; address = address + 4;
    ...
    if wback && registers<n> == '0' then R[n] = R[n] + 4*BitCount(registers);
```

```
if wback && registers<n> == '1' then R[n] = bits(32) UNKNOWN;
```

is changed to read:

```
if ConditionPassed() then
    ...
    bits(32) data;
    for i = 0 to 14
        if registers<i> == '1' then
            if i != n then
                R[i] = MemS[address,4];
            else
                data = MemS[address,4];
                address = address + 4;
    ...

    if wback && registers<n> == '1' then
        R[n] = bits(32) UNKNOWN;
    if wback && registers<n> == '0' then
        R[n] = R[n] + 4*BitCount(registers);
    if !wback && registers<n> == '1' then
        R[n] = data;
```

Equivalent changes are made to the following instructions:

- F5.1.68 LDM (exception return).
- F5.1.70 LDMDA, LDMFA.
- F5.1.71 LDMDB, LDMEA.
- F5.1.72 LDMIB, LDMED.

In section F5.1.80 “LDRD (immediate)”, the code that reads:

```
if ConditionPassed() then
    ...
    if IsAligned(address, 8) then
        ...
    else
        R[t] = MemA[address,4];
        R[t2] = MemA[address+4,4];

    if wback then R[n] = offset_addr;
```

is changed to read:

```
if ConditionPassed() then
    ...
    if IsAligned(address, 8) then
        ...
    else
        constant bits(32) data1 = MemA[address,4];
        constant bits(32) data2 = MemA[address+4,4];
        R[t] = data1;
        R[t2] = data2;

    if wback then R[n] = offset_addr;
```

Equivalent changes are made to the following instructions:

- F5.1.81 LDRD (literal).
- F5.1.82 LDRD (register).

In section C6.2.234 “LDXP”, the code that reads:

```

if rt_unknown then
    ...
elseif elsize == 32 then
    ...
else // elsize == 64
    ...
    constant bits(64) address2 = AddressIncrement(address, 8, accdesc);
    X[t, 64] = Mem[address, 8, accdesc];
    X[t2, 64] = Mem[address2, 8, accdesc];

```

is changed to read:

```

if rt_unknown then
    ...
elseif elsize == 32 then
    ...
else // elsize == 64
    ...
    constant bits(64) address2 = AddressIncrement(address, 8, accdesc);
    constant bits(64) data1 = Mem[address, 8, accdesc];
    constant bits(64) data2 = Mem[address2, 8, accdesc];
    X[t, 64] = data1;
    X[t2, 64] = data2;

```

2.62 D23733

In section C6.2.328 “SETGP, SETGM, SETGE”, the pseudocode for SETG when a fault occurs that reads:

```

if memset.implements_option_a then
    while memset.stagesetsize < 0 && !fault do
        (memory_set, memaddrdesc, memstatus) = MemSetBytes(
            memset.toaddress + memset.setsize,
            data, B, accdesc);
        if memory_set != B then
            fault = TRUE;
        else
            tagstep = B DIV TAG_GRANULE;
            tag = AArch64.AllocationTagFromAddress(
                memset.toaddress + memset.setsize);
            while tagstep > 0 do
                tagaddr = memset.toaddress + memset.setsize + (tagstep - 1) *
                TAG_GRANULE;
                AArch64.MemTag[tagaddr, accdesc] = tag;
                tagstep = tagstep - 1;
            end while
            memset.setsize = memset.setsize + B;
        end if
        memset.stagesetsize = memset.stagesetsize + B;
    end while
else
    ... (similar loop for the reverse direction)
end if

```

is changed to read:

```

if memset.implements_option_a then
    (tags_written, memaddrdesc, memstatus) = MemSetTags(
        memset.toaddress + memset.setsize, // start address
        data, // value to write
        B, // block size
        accdesc); // access parameters
    if memstatus.statuscode != Fault_None then
        fault = TRUE;
    else
        memset.setsize = memset.setsize + B;
        memset.stagesetsize = memset.stagesetsize + B;
    end if
else
    (tags_written, memaddrdesc, memstatus) = MemSetTags(
        memset.toaddress,
        data,
        B,
        accdesc);
    if memstatus.statuscode != Fault_None then
        fault = TRUE;
    else
        memset.toaddress = memset.toaddress + B;
        memset.stagesetsize = memset.stagesetsize - B;
    end if
end if

```

2.63 D23744

In section D24.2.206 “VNCR_EL2, Virtual Nested Control Register”, the text that reads:

RESS, bits [63:57]

Reserved, Sign extended. If the bits marked as RESS do not all have the same value, then there is a **CONSTRAINED UNPREDICTABLE** choice between:

- Generating an EL2 translation regime Translation abort on use of the VNCR_EL2 register. If FEAT_D128 is implemented:
- If the virtual address space for EL2 supports 56 bits, bits[63:57] of VNCR_EL2 are treated as the same value as bit[56] for all purposes other than reading back the register.
- If the virtual address space for EL2 supports 56 bits, bits[63:57] of VNCR_EL2 are treated as the same value as bit[56].
- If the virtual address space for EL2 supports 52 bits, bits[63:53] of VNCR_EL2 are treated as the same value as bit[52] for all purposes other than reading back the register.
- If the virtual address space for EL2 supports 52 bits, bits[63:53] of VNCR_EL2 are treated as the same value as bit[52].
- Bits[63:49] of VNCR_EL2 are treated as the same value as bit[48] for all purposes other than reading back the register.
- Bits[63:49] of VNCR_EL2 are treated as the same value as bit[48] for all purposes.

Where the EL2 translation regime has upper and lower address ranges, bit[56] is used to select between those address ranges to determine the number of bits supported by the address space.

BADDR, bits [56:12]

Base Address. If the virtual address space for EL2 does not support more than 48 bits, then bits [56:49] are RESS. If the virtual address space for EL2 does not support more than 52 bits, then bits [56:53] are RESS

When a register read/write is transformed to be a Load or Store, the address of the load/store is to SignOffset(VNCR_EL2.BADDR:Offset<11:0>, 64).

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

is changed to read:

BADDR, bits [63:12]

Base Address. When a register read/write is transformed to be a Load or Store, the address of the load/store is to VNCR_EL2.BADDR:Offset<11:0>.

Bits[63:n] are RESS where n is one of the following:

- If FEAT_LVA3 is implemented, n is 57.
- If FEAT_LVA is implemented, but FEAT_LVA3 is not implemented, n is 53.
- If FEAT_LVA is not implemented, n is 49.

If the bits marked as RESS do not all have the same value as bit[n-1], then one of the following **CONSTRAINED UNPREDICTABLE** choices applies:

- When translating the transformed register read/write, a Translation fault is generated.
- The bits behave as the corresponding RESS for all purposes other than reading back the register.
- The bits behave as the corresponding RESS for all purposes.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

2.64 C23753

In section D1.3.2 “Exception entry”, under rule R_{DQXFW} , the text that reads:

When an exception is taken to an Exception level, ELx, that is using AArch64 state, all the following occur:

- The contents of PSTATE immediately before the exception was taken is written to SPSR_ELx.

- ...

is changed to read:

When an exception is taken to an Exception level, ELx, that is using AArch64 state, all the following occur:

- Fields in SPSR_ELx are written according to their description in SPSR_ELx, based on the values in PSTATE immediately before the exception was taken.
- ...

In section D1.3.4.1 “Legal exception returns from AArch64 state”, under rule R_{BWCFK} , the text that reads:

On a legal exception return from an Exception level, ELx, all of the following occur:

- ...
- The contents of PSTATE are set to the values held in SPSR_ELx.
- ...

is changed to read:

On a legal exception return from an Exception level, ELx, all of the following occur:

- ...
- Fields in PSTATE that have a corresponding field in SPSR_ELx are written according to their description in SPSR_ELx.
- ...

Additionally, in section D1.3.2 “Exception entry”, the informational statement I_{X0001} is added:

I_{X0001}

When an exception is taken to an Exception level that is using AArch64 state, the following PSTATE fields are unchanged:

- If FEAT_DIT is implemented, PSTATE.DIT.
- If FEAT_SME is implemented, PSTATE.{SM,ZA}.

2.65 C23766

In section B2.3.1 “Basic definitions”, the text that reads:

Same Location

An Effect E_1 and an Effect E_2 are to the Same Location if one of the following applies:

...

- All of the following apply:
 - An Implicit TTD Memory Read Effect E_3 is Translation-intrinsically-before E_1 .
 - The output address of the TTD read by E_3 is the same as the output address of the TTD read by an Implicit TTD Memory Read Effect E_4 .
 - An Implicit TTD Memory Read Effect E_4 is Translation-intrinsically-before E_2 .
 - E_1 and E_2 have the Same Low-order Bits.

is changed to read:

Same Location

An Effect E_1 and an Effect E_2 are to the Same Location if one of the following applies:

...

- All of the following apply:
 - E_1 is an MMU Fault Effect.
 - It is not the case that E_1 is an MMU Translation Fault Effect.
 - E_3 is Translation-intrinsically-before E_1 .
 - E_3 is an Implicit TTD Memory Read Effect.
 - E_3 and E_4 have the Same Output Address.
 - E_4 is an Implicit TTD Memory Read Effect.
 - E_4 is Translation-intrinsically-before E_2 .
 - E_1 and E_2 have the Same Low-order Bits.
- All of the following apply:
 - E_3 is Translation-intrinsically-before E_1 .
 - E_3 is an Implicit TTD Memory Read Effect.
 - E_3 and E_4 have the Same Output Address.
 - E_4 is an Implicit TTD Memory Read Effect.
 - E_4 is Translation-intrinsically-before E_2 .
 - E_1 and E_2 have the Same Low-order Bits.
 - E_2 is an MMU Fault Effect.
 - It is not the case that E_2 is an MMU Translation Fault Effect.

2.66 D23768

In section B1.5.6 “About PSTATE.DIT”, the text that reads:

- If SVE2 is not implemented, the data independent timing control introduced by FEAT_DIT does not affect the timing properties of SVE instructions.

- For SVE and SVE2 predicated instructions, it is the programmer's responsibility to use a Governing predicate that does not reflect the values of the data being operated on.

is changed to read:

- If FEAT_SVE2 and FEAT_SME are not implemented, the data independent timing control introduced by FEAT_DIT does not affect the timing properties of SVE instructions.
- If FEAT_SVE2 or FEAT_SME is implemented, the data independent timing control introduced by FEAT_DIT affects the timing properties of all SVE and SME instructions that honor PSTATE.DIT.
- For SVE and SME predicated instructions, it is the programmer's responsibility to use a Governing predicate that does not reflect the values of the data being operated on.

2.67 D23772

In section C6.2.50 “CAS, CASA, CASAL, CASL”, the text that reads:

- CASA and CASAL load from memory with acquire semantics.

is changed to read:

- If the destination register is not one of WZR or XZR, CASA and CASAL load from memory with acquire semantics.

Under the heading “Decode for all variants in this encoding”, the pseudocode that reads:

```
...  
constant boolean acquire = L == '1';  
...
```

is changed to read:

```
...  
constant boolean acquire = L == '1' && t != 31;  
...
```

Equivalent changes are made in the following sections:

- C6.2.51 “CASB, CASAB, CASALB, CASLB”.
- C6.2.52 “CASH, CASAH, CASALH, CASLH”.

In section C6.2.291 “RCWCAS, RCWCASA, RCWCASL, RCWCASAL”, the text that reads:

- RCWCASA and RCWCASAL load from memory with acquire semantics.

is changed to read:

- If the destination register is not XZR, RCWCASA and RCWCASAL load from memory with acquire semantics.

In the same section, under the heading “Decode for all variants in this encoding”, the pseudocode that reads:

```
...
constant boolean acquire = A == '1';
...
```

is changed to read:

```
...
constant boolean acquire = A == '1' && t != 31;
...
```

Equivalent changes are made in the following sections:

- C6.2.293 “RCWCLR, RCWCLRA, RCWCLRL, RCWCLRAL”.
- C6.2.295 “RCWSCAS, RCWSCASA, RCWSCASL, RCWSCASAL”.
- C6.2.297 “RCWSCLR, RCWSCLRA, RCWSCLRL, RCWSCLRAL”.
- C6.2.299 “RCWSET, RCWSETA, RCWSETL, RCWSETAL”.
- C6.2.301 “RCWSSET, RCWSSETA, RCWSSETL, RCWSSETAL”.
- C6.2.303 “RCWSSWP, RCWSSWPA, RCWSSWPL, RCWSSWPAL”.
- C6.2.305 “RCWSWP, RCWSWPA, RCWSWPL, RCWSWPAL”.

In section C6.2.159 “LDAPR”, the text that reads:

The instruction has memory ordering semantics as described in Load-Acquire, Load-AcquirePC, and Store-Release, except that:

- There is no ordering requirement, separate from the requirements of a Load-AcquirePC or a Store-Release, created by having a Store-Release followed by a Load-AcquirePC instruction.
- The reading of a value written by a Store-Release by a Load-AcquirePC instruction by the same observer does not make the write of the Store-Release globally observed.

This difference in memory ordering is not described in the pseudocode.

is changed to read:

If the destination register is not XZR or WZR, LDAPR loads from memory with AcquirePC semantics.

For more information about memory-ordering semantics, see Load-Acquire, Load-AcquirePC and Store-Release.

Under the heading “Post-index”, subheading “Decode for all variants of this encoding”, the pseudocode that reads:

```
...
constant integer offset = 1 << UInt(size);
constant boolean tagchecked = TRUE;
...
```

is changed to read:

```
...
constant integer offset = 1 << UInt(size);
constant boolean acquirepc = t != 31;
constant boolean tagchecked = TRUE;
...
```

Under the heading “No offset”, subheading “Decode for all variants of this encoding”, the pseudocode that reads:

```
...
constant integer datasize = elsize;
constant boolean tagchecked = n != 31;
```

is changed to read:

```
...
constant integer datasize = elsize;
constant boolean acquirepc = t != 31;
constant boolean tagchecked = n != 31;
```

Under the heading “Operation”, the pseudocode that reads:

```
...
constant integer dbytes = datasize DIV 8;
constant AccessDescriptor accdesc = CreateAccDescLDAcqPC(tagchecked);
if n == 31 then
...
```

is changed to read:

```
...
constant integer dbytes = datasize DIV 8;
constant AccessDescriptor accdesc = CreateAccDescLDAcqPC(tagchecked, acquirepc);

if n == 31 then
...
```

Equivalent changes are made in the following sections:

- C6.2.160 “LDAPRB”.
- C6.2.161 “LDAPRH”.

- C6.2.162 “LDAPUR”.
- C6.2.163 “LDAPURB”.
- C6.2.164 “LDAPURH”.
- C6.2.165 “LDAPURSB”.
- C6.2.167 “LDAPURSW”.

In section C6.2.168 “LDAR”, the text that reads:

Load-acquire register

This instruction derives an address from a base register value, loads a 32-bit word or 64-bit doubleword from memory, and writes it to a register. The instruction also has memory ordering semantics as described in Load-Acquire, Store-Release. For information about addressing modes, see Load/Store addressing modes.

Note:

For this instruction, if the destination is WZR/XZR, it is impossible for software to observe the presence of the acquire semantic other than its effect on the arrival at endpoints.

is changed to read:

Load-acquire register

This instruction derives an address from a base register value, loads a 32-bit word or 64-bit doubleword from memory, and writes it to a register. For information about addressing modes, see Load/Store addressing modes.

If the destination register is not WZR or XZR, LDAR loads from memory with Acquire semantics.

For more information about memory ordering semantics, see Load-Acquire, Load-AcquirePC, Store-Release.

Under the heading “Decode for all variants of this encoding”, the pseudocode that reads:

```
...  
constant boolean tagchecked = TRUE;
```

is changed to read:

```
...  
constant boolean acquire = t != 31;  
constant boolean tagchecked = TRUE;
```

Under the heading “Operation”, the pseudocode that reads:

```
bits(64) address;  
constant AccessDescriptor accdesc = CreateAccDescAcqRel(MemOp_LOAD, tagchecked);  
  
if n == 31 then  
...
```

is changed to read:

```
bits(64) address;  
constant AccessDescriptor accdesc = CreateAccDescAcqRel(MemOp_LOAD, tagchecked,  
acquire);  
  
if n == 31 then  
...
```

Equivalent changes are made in the following sections:

- C6.2.169 “LDARB”.
- C6.2.170 “LDARH”.

In section C6.2.171 “LDAXP” the text that reads:

Load-acquire exclusive pair of registers

This instruction derives an address from a base register value, loads two 32-bit words or two 64-bit doublewords from memory, and writes them to two registers. For information on single-copy atomicity and alignment requirements, see Requirements for single-copy atomicity and Alignment of data accesses. The PE marks the physical address being accessed as an exclusive access.

This exclusive access mark is checked by Store Exclusive instructions. See Synchronization and semaphores. The instruction also has memory ordering semantics, as described in Load-Acquire, Store-Release. For information about addressing modes, see Load/Store addressing modes.

is changed to read:

Load-acquire exclusive pair of registers

This instruction derives an address from a base register value, loads two 32-bit words or two 64-bit doublewords from memory, and writes them to two registers. For information on single-copy atomicity and alignment requirements, see Requirements for single-copy atomicity and Alignment of data accesses. The PE marks the physical address being accessed as an exclusive access. This exclusive access mark is checked by Store Exclusive instructions. See Synchronization and semaphores. For information about addressing modes, see Load/Store addressing modes.

If the destination registers are not both WZR or not both XZR, LDAXP loads from memory with Acquire semantics.

For more information about memory ordering semantics, see Load-Acquire, Load-AcquirePC, Store-Release.

Under the heading “Decode for all variants of this encoding”, the pseudocode that reads:

```
...  
constant boolean acqrel = TRUE;  
constant boolean tagchecked = n != 31;  
...
```

is changed to read:

```
...  
constant boolean acqrel = t != 31 && t1 != 31;  
constant boolean tagchecked = n != 31;  
...
```

Equivalent change is made to section C6.2.184 “LDIAPP”.

In section C6.2.172 “LDAXR”, the text that reads:

Load-acquire exclusive register

This instruction derives an address from a base register value, loads a 32-bit word or 64-bit doubleword from memory, and writes it to a register. The memory access is atomic. The PE marks the physical address being accessed as an exclusive access. This exclusive access mark is checked by Store Exclusive instructions. See Synchronization and semaphores. The instruction also has memory ordering semantics as described in Load-Acquire, Store-Release. For information about addressing modes, see Load/Store addressing modes.

is changed to read:

Load-acquire exclusive register

This instruction derives an address from a base register value, loads a 32-bit word or 64-bit doubleword from memory, and writes it to a register. The memory access is atomic. The PE marks the physical address being accessed as an exclusive access. This exclusive access mark is checked by Store Exclusive instructions. See Synchronization and semaphores. For information about addressing modes, see Load/Store addressing modes.

If the destination register is not WZR or XZR, LDAXR loads from memory with Acquire semantics.

For more information about memory ordering semantics, see Load-Acquire, Load-AcquirePC, Store-Release.

Under the heading “Decode for all variants of this encoding”, the pseudocode that reads:

```
...  
constant boolean acqrel = TRUE;  
constant boolean tagchecked = n != 31;
```

is changed to read:

```
...  
constant boolean acqrel = t != 31;  
constant boolean tagchecked = n != 31;
```

Equivalent changes are made in the following sections:

- C6.2.173 “LDAXRB”.
- C6.2.174 “LDAXRH”.

In section C6.2.185 “LDLAR”, the text that reads:

Load LOAcquire register

This instruction loads a 32-bit word or 64-bit doubleword from memory, and writes it to a register. The instruction also has memory ordering semantics as described in Load LOAcquire, Store LORelease. For information about addressing modes, see Load/Store addressing modes.

Note:

For this instruction, if the destination is WZR/XZR, it is impossible for software to observe the presence of the acquire semantic other than its effect on the arrival at endpoints.

is changed to read:

Load LOAcquire register

This instruction loads a 32-bit word or 64-bit doubleword from memory, and writes it to a register. For information about addressing modes, see Load/Store addressing modes.

If the destination register is not WZR or XZR, LDLAR loads from memory with Acquire semantics.

For more information about memory ordering semantics, see Load LOAcquire, Store LORelease and Load-Acquire, Load-AcquirePC, Store-Release.

In addition, under the heading “Decode for all variants of this encoding”, the pseudocode that reads:

```
...  
constant boolean acqrel = TRUE;  
constant boolean tagchecked = n != 31;
```

is changed to read:

```
...  
constant boolean acqrel = t != 31;  
constant boolean tagchecked = n != 31;
```

Under the heading “Operation”, the pseudocode that reads:

```
bits(64) address;  
constant integer dbytes = elsize DIV 8;  
  
constant AccessDescriptor accdesc = CreateAccDescLOR(MemOp_LOAD, tagchecked);  
  
if n == 31 then
```

is changed to read:

```
bits(64) address;  
constant integer dbytes = elsize DIV 8;  
  
constant AccessDescriptor accdesc = CreateAccDescLOR(MemOp_LOAD, tag checked,  
acqrel);  
  
if n == 31 then
```

Equivalent changes are made to the following sections:

- C6.2.186 “LDLARB”.
- C6.2.187 “LDLARH”.

2.68 C23775

The following text is added to B2.3.7 “Ordering relations”:

An Effect E_1 is ETS-ordered-before an Effect E_2 if one of the following applies:

- All of the following apply:
 - FEAT_ETS2 is implemented.
 - E_1 is an Explicit Memory Effect.
 - E_1 appears in program order before E_3 .
 - E_3 is a TLBUncacheable Fault Effect.
 - E_2 is Translation-intrinsically-before E_3 .
 - E_2 is an Implicit TTD Memory Read Effect.
- All of the following apply:
 - FEAT_ETS3 is implemented.
 - E_1 is an Explicit Memory Effect.
 - E_1 appears in program order before E_3 .
 - E_3 is an MMU Fault Effect.
 - E_2 is Translation-intrinsically-before E_3 .
 - E_2 is an Implicit TTD Memory Read Effect.

Note:

The ETS-ordered-before relation does not apply when E_1 is generated by an Advanced SIMD, floating-point, SVE, or SME instruction. It also does not apply when E_2 is an Implicit TTD Memory Read Effect generated for an instruction fetch.

The following text is added to D8.2.6.1 “Ordering of memory accesses from translation table walks”, under the rule R_{NNFPF} :

If FEAT_ETS2 is implemented, then an Effect E_1 is Ordered-before an Effect E_2 when all of the following apply:

- E_1 is not generated by an Advanced SIMD, floating-point, SVE, or SME instruction.
- E_2 is not generated for an instruction fetch.
- E_1 is an Explicit Memory Effect.
- E_1 appears in program order before E_3 .
- E_3 is a TLBUncacheable Fault Effect.
- E_2 is Translation-intrinsically-before E_3 .
- E_2 is an Implicit TTD Memory Read Effect.

The following text is added to the same section under the rule R_{QVWHP} :

If FEAT_ETS3 is implemented, then an Effect E_1 is Ordered-before an Effect E_2 when all of the following apply:

- E_1 is not generated by an Advanced SIMD, floating-point, SVE, or SME instruction.
- E_1 is an Explicit Memory Effect.
- E_2 is not generated for an instruction fetch.
- E_1 appears in program order before E_3 .
- E_3 is an MMU Fault Effect.
- E_2 is Translation-intrinsically-before E_3 .
- E_2 is an Implicit TTD Memory Read Effect.

The following text is added to E2.4 “Ordering of translation table walks”:

If FEAT_ETS2 is implemented, then an Effect E_1 is Ordered-before an Effect E_2 when all of the following apply:

- E_1 is not generated by an Advanced SIMD or floating-point instruction.
- E_2 is not generated for an instruction fetch.
- E_1 is an Explicit Memory Effect.
- E_1 appears in program order before E_3 .
- E_3 is a TLBUncacheable Fault Effect.
- E_2 is Translation-intrinsically-before E_3 .

- E_2 is an Implicit TTD Memory Read Effect.

The following text is added to the same section:

If FEAT_ETS3 is implemented, then an Effect E_1 is Ordered-before an Effect E_2 when all of the following apply:

- E_1 is not generated by an Advanced SIMD, floating-point, SVE, or SME instruction.
- E_2 is not generated for an instruction fetch.
- E_1 is an Explicit Memory Effect.
- E_1 appears in program order before E_3 .
- E_3 is an MMU Fault Effect.
- E_2 is Translation-intrinsically-before E_3 .
- E_2 is an Implicit TTD Memory Read Effect.

2.69 C23782

In section D24.2.167 “SCTLR_EL1, System Control Register (EL1)”, the text in field ‘CMOW, bit[32]’ that reads:

Controls cache maintenance instruction permission for the following instructions executed at ELO.

- IC IVAU and DC CIVAC.
- If FEAT_MTE is implemented, DC CIGDVAC and DC CIGVAC.

CMOW	Meaning
0b0	These instructions executed at ELO with stage 1 read permission, but without stage 1 write permission, do not generate a stage 1 permission fault.
0b1	If enabled as a result of $SCTLR_EL1.UCI=1$, these instructions executed at ELO with stage 1 read permission, but without stage 1 write permission, generate a stage 1 permission fault

When the Effective value of $HCR_EL2.\{E2H, TGE\}$ is $\{1, 1\}$, this bit has no effect on execution at ELO.

For this control, stage 1 has write permission if all of the following apply:

- $AP[2]$ is 0 or DBM is 1 in the stage 1 descriptor.
- Where $APTable$ is in use, $APTable[1]$ is 0 for all levels of the translation table.

is changed to read:

Controls the required permissions for the following cache maintenance instructions executed at ELO:

- Any DC instruction that operates by VA and performs a clean and invalidate operation.

- Any IC instruction that operates by VA.

CMOW	Meaning
0b0	This control does not cause any instruction to generate a stage 1 Permission fault.
0b1	For these instructions, when executed at ELO, the absence of stage 1 unprivileged write permission generates a stage 1 permission fault.

When the Effective value of HCR_EL2.{E2H, TGE} is {1, 1}, this bit has no effect on execution at ELO.

For more information, see:

- Stage 1 permissions
- Implications of enabling the dirty state management mechanism

Equivalent change is made to section D24.2.168 “SCTLR_EL2, System Control Register (EL2)” when ELIsInHost(EL2).

In section D24.2.57 “HCRX_EL2, Extended Hypervisor Configuration Register”, the text in field ‘CMOW, bit[9]’ that reads:

Controls cache maintenance instruction permission for the following instructions executed at EL1 or ELO.

- IC IVAU and DC CIVAC.
- If FEAT_MTE is implemented, DC CIGDVAC and DC CIGVAC.
- ICIMVAU and DCCIMVAC.

CMOW	Meaning
0b0	This control does not cause any instructions to generate a stage 2 Permission fault.
0b1	Unless the instruction generates a higher priority exception, EL1 and ELO execution of the specified instructions without stage 2 write permission generates a stage 2 Permission fault.

For this control, stage 2 has write permission if S2AP[1] is 1 or DBM is 1 in the stage 2 descriptor. The instructions do not cause an update to the dirty state.

is changed to read:

Controls the required permissions for the following cache maintenance instructions executed at EL1 or ELO:

- Any DC instruction that operates by VA and performs a clean and invalidate operation.
- Any IC instruction that operates by VA.

CMOW	Meaning
0b0	This control does not cause any instruction to generate a stage 2 Permission fault.
0b1	For these instructions, when executed at EL1 or ELO, the absence of stage 2 write permission generates a stage 2 permission fault.

For more information, see:

- Stage 2 permissions
- Implications of enabling the dirty state management mechanism

In section D7.5.9.2 “The data cache maintenance instruction (DC)”, the text that reads:

When FEAT_CMOW is implemented, the possible generation of a Permission fault by:

- The execution of a DC CIVAC, DC CIGVAC, or DC CIGDVAC instruction at ELO when the specified address has stage 1 read permission, but does not have stage 1 write permission.
- The execution of a DC CIVAC, DC CIGVAC, or DC CIGDVAC instruction at EL1 or ELO when the specified address does not have stage 2 write permission.

is changed to read:

When FEAT_CMOW is implemented, the possible generation of a Permission fault by:

- The execution of a DC instruction that operates by VA and performs clean and invalidate at ELO when the specified address does not have stage 1 write permission.
- The execution of a DC instruction that operates by VA and performs clean and invalidate at EL1 or ELO when the specified address does not have stage 2 write permission.

In section D7.5.9.3 “ELO accessibility of cache maintenance instructions”, the text that reads:

When FEAT_CMOW is implemented, for the DC CIVAC, DC CIGVAC, DC CIGDVAC, and IC IVAU instructions, both read and write access permission is required:

- When SCTLR_EL1.CMOW is 1, if the instruction executed at ELO has stage 1 read permission, but does not have stage 1 write permission, a Permission fault is generated.
- When HCRX_EL2.CMOW is 1, if the instruction executed at ELO has stage 2 read permission, but does not have stage 2 write permission, a Permission fault is generated.

is changed to read:

When FEAT_CMOW is implemented, for the DC instructions that operate by VA and perform clean and invalidate and IC instructions that operate by VA, write access permission is also required:

- When SCTLR_EL1.CMOW is 1, if the instruction executed at ELO does not have stage 1 write permission, a Permission fault is generated.
- When HCRX_EL2.CMOW is 1, if the instruction executed at ELO does not have stage 2 write permission, a Permission fault is generated.

In section D8.15.3 “MMU faults generated by cache maintenance operations”, the text that reads:

R_{FXTG}

If the Effective value of SCTLR_ELx.CMOW is 1, then when executing an IC IVAU, DC CIVAC, DC CIGVAC, or DC CIGDVAC instruction at ELO that has stage 1 read permission, but does not have stage 1 write permission, a stage 1 Permission fault is generated.

R_{BBLTJ}

If the Effective value of HCRX_EL2.CMOW is 1, then when executing an IC IVAU, DC CIVAC, DC CIGVAC, or DC CIGDVAC instruction at EL1 or ELO that does not have stage 2 write permission, a stage 2 Permission fault is generated.

is changed to read:

R_{FTXTG}

If the Effective value of SCTLR_ELx.CMOW is 1, then when executing any DC instruction that operates by VA and performs clean and invalidate or any IC instruction that operates by VA at ELO that does not have stage 1 write permission, a stage 1 Permission fault is generated.

R_{BBLTJ}

If the Effective value of HCRX_EL2.CMOW is 1, then when executing any DC instruction that operates by VA and performs clean and invalidate or any IC instruction that operates by VA at EL1 or ELO that does not have stage 2 write permission, a stage 2 Permission fault is generated.

2.70 D23794

In section D24.2.164 “SCTLR2_EL1, System Control Register (EL1)”, the text that reads:

If FEAT_SVE is implemented, SCTLR2_EL1.EnADERR is 0, and the access generating the External abort is due to any Active element of an SVE Non-fault vector load instruction or an Active element that is not the first Active element of an SVE First-fault vector load instruction, then no exception is generated and the External abort is reported in the FFR.

is changed to read:

If FEAT_SVE is implemented, SCTLR2_EL1.EnANERR is 0, and the access generating the External abort is due to any Active element of an SVE Non-fault vector load instruction or an Active element that is not the first Active element of an SVE First-fault vector load instruction, then no exception is generated and the External abort is reported in the FFR.

The equivalent changes are made in the following sections:

- D24.2.165 “SCTLR2_EL2, System Control Register (EL2)”
- D24.2.166 “SCTLR2_EL3, System Control Register (EL3)”

2.71 D23796

In section D24.7.9 “PMSFCR_EL1, Sampling Filter Control Register”, in the description of PMSFCR_EL1.FT, the text that reads:

FT, bit [1]

Filter by operation type. The filter is the logical OR of the ST, LD and B bits. For example, if LD and ST are both set, both load and store operations are recorded.

The full description of filtering is provided in the section D17.5.1 “Filtering by operation type,” and does not need to be repeated here.

is changed to read:

FT, bit [1]

Filter by type. The filter is controlled by the {SIMD, FP, ST, LD, B} and {SIMDm, FPm, STm, LDm, Bm} fields.

For more information, see Filtering by operation type.

2.72 D23797

In section D1.3.5.3 “Granule Protection Check faults”, under the rule R_{JXSRX}, the text that reads:

When the PE is in Debug state and EDSCR.SDD is 0, SCR_EL3.GPF is treated as 0, and the following GPC faults are treated as a GPF for the purposes of causing an exception:

- GPT walk faults.
- GPT address size faults.
- Synchronous External abort on GPT fetch.

is changed to read:

When the PE is in Debug state and EDSCR.SDD is 1, SCR_EL3.GPF is treated as 0, and the following GPC faults are treated as a GPF for the purposes of causing an exception:

- GPT walk faults.
- GPT address size faults.
- Synchronous External abort on GPT fetch.

In section J1.3 “Shared pseudocode”, in the function `ReportAsGPCEException()`, the code that reads:

```
boolean ReportAsGPCEException(FaultRecord fault)
...
case fault.gpcf.gpf of
when GPCF_Walk return TRUE;
```

```
when GPCF_AddressSize return TRUE;
when GPCF_EABT return TRUE;
when GPCF_Fail return SCR_EL3.GPF == '1' && PSTATE.EL != EL3;
```

is changed to read:

```
boolean ReportAsGPCEXception(FaultRecord fault)
...
if Halted() && EDSCR.SDD == '1' then
    return FALSE;
case fault.gpcf.gpf of
    when GPCF_Walk return TRUE;
    when GPCF_AddressSize return TRUE;
    when GPCF_EABT return TRUE;
    when GPCF_Fail return SCR_EL3.GPF == '1' && PSTATE.EL != EL3;
```

In the function `AArch64.DataAbort()`, the code that reads:

```
else
    route_to_el2 = (EL2Enabled() && PSTATE.EL IN {EL0, EL1} &&
        (HCR_EL2.TGE == '1' ||
        (IsFeatureImplemented(FEAT_RME) &&
        fault.gpcf.gpf == GPCF_Fail &&
        HCR_EL2.GPF == '1') ||
        ...
```

is changed to read:

```
else
    route_to_el2 = (EL2Enabled() && PSTATE.EL IN {EL0, EL1} &&
        (HCR_EL2.TGE == '1' ||
        (IsFeatureImplemented(FEAT_RME) &&
        fault.gpcf.gpf != GPCF_None &&
        HCR_EL2.GPF == '1') ||
        ...
```

The equivalent change is made in the function `AArch64.InstructionAbort()` for consistency.

In section H2.4.7.1 “Generating exceptions when in Debug state”, the text that reads:

In Debug state:

...

- If FEAT_RME is implemented and EDSCR.SDD is 1, SCR_EL3.GPF is treated as 0, regardless of its actual state, other than for the purpose of a direct read.

is changed to read:

In Debug state:

...

- If FEAT_RME is implemented and EDSCR.SDD is 1:
 - SCR_EL3.GPF is treated as 0, other than for the purpose of a direct read.

- All GPC faults are treated as a GPF for the purposes of causing an exception.

2.73 D23807

In section D20.1 “About the RAS Extension”, in the example D20-1 “Minimal implementation of RAS Extension”, the text that reads:

Any error signaled to the PE generates an asynchronous SError exception, and on taking the SError exception:

- The PE error state is recorded as Uncontainable (UC).
- ESR_ELx.FnV is set to 1, meaning FAR_ELx is **UNKNOWN** and not valid.

FEAT_PFAR is not implemented. See Taking error exceptions.

is changed to read:

Any error signaled to the PE generates an asynchronous SError exception, and on taking the SError exception:

- The PE error state is recorded as Uncontainable (UC).
 - If FEAT_RASv2 is implemented, then ESR_ELx.VFV is set to 0, meaning FAR_ELx is **UNKNOWN** and not valid. If FEAT_RASv2 is not implemented, then FAR_ELx is always **UNKNOWN** and not valid on taking an SError exception.
 - FEAT_PFAR is not implemented, meaning ESR_ELx.PFV is set to 0.

See Taking error exceptions.

2.74 C23808

In section D17.3.1 “Operation sampling”, the text that reads:

1. A sampling interval is written to PMSICR_EL1.COUNT by software. The interval is measured in operations.

is changed to read:

1. Software writes a sampling interval to PMSIRR_EL1.COUNT, and sets PMSICR_EL1 to zero. The interval is measured in operations.

2.75 D23812

In section D1.3.5.7 “Memory Copy and Memory Set exceptions”, under the rule I_{PRXVQ} , the text that reads:

The information in the $ESR_ELx.ISS$ field is sufficient to allow the diagnosis of the reason for a Memory Copy or Memory Set exception being generated and to allow a generic emulation of the memory copy or memory set.

is changed to read:

The information in $SPSR_ELx.C$ (or, equivalently, the $ESR_ELx.ISS.FormatOption$ field), $SPSR_ELx.N$ and the sign of the value in the Xn register is sufficient to determine the current format of the Xd , Xs and Xn registers, and perform any required reformatting, before restarting the memory copy or memory set. However, it is not required to use $FEAT_MOPS$ instructions and the memory copy or memory set could alternatively be restarted using a software implementation.

2.76 E23814

In section C7.2.332 “ $SRSHR$ ”, under ‘Operational Information’, the following text is added:

If $PSTATE.DIT$ is 1, this instruction is required to honor the data-independent timing behaviors described in section B1.5.6 “About $PSTATE.DIT$ ”.

The equivalent changes are made in the following sections:

- C7.2.334 “ $SSHL$ ”.
- C8.2.674 “ $SRSRA$ ”.
- C8.2.868 “ $URSHL$ ”.
- C8.2.872 “ $URSRA$ ”.
- C7.2.406 “ $URSHR$ ”.

These instructions are also added to the list of instructions in C5.2.4 “ DIT , Data Independent Timing”, under ‘Data Processing – Scalar Floating-Point and Advanced SIMD’.

2.77 R23822

In section D18.2.7 “Operation Type packet”, in the description of the $PRED$ field, the text that reads:

PRED	Description
0b0	Not predicated.

PRED	Description
0b1	Predicated SVE operation. The operation is an SVE operation that writes to a vector destination register under a Governing predicate using either zeroing or merging predication.

is changed to read:

PRED	Description
0b0	Not predicated.
0b1	Predicated SVE data-processing operation.

Predicated SVE data-processing operation means an SVE data-processing operation that writes to one or more SVE Z vector destination registers under a Governing predicate using either zeroing or merging predication. It is **IMPLEMENTATION DEFINED** whether this includes data-processing operations due to instructions with a single Governing predicate operand that determines the Active elements which do not write to any SVE Z vector destination registers using either zeroing or merging predication are counted. For example, INCP.

In section D18.2.6 “Events packet” (issue L.a), in the description of the E[18] field, the text that reads:

E[18], byte 2 bit [2], when SZ == 0b10 or SZ == 0b11

Empty predicate.

...

is changed to read:

E[18], byte 2 bit [2], when SZ == 0b10 or SZ == 0b11

Empty predicate.

...

For a sampled SVE operation, it is **IMPLEMENTATION DEFINED** whether this field is valid or reads as zero when the PRED field in the Operation Type field is 0.

The equivalent change is made to the E[17] “Partial or empty predicate” field.

2.78 D23823

In section D18.2.6 “Events packet”, the following text is added to the description of E[19] “Level 2 Data cache access”:

This event is optional. If this event is implemented, then it is further **IMPLEMENTATION DEFINED** and might be **UNPREDICTABLE** whether a store can finish execution before this event is generated, meaning that, for stores, this event is not recorded and this bit reads-as-zero. If this event is not implemented, then this bit reads-as-zero.

Similar text is added to each of the following events:

- E[20] “Level 2 Data cache miss”.
- E[21] “Cached data modified”.
- E[22] “Recently fetched”.
- E[23] “Data snooped”.

2.79 C23830

In section C5.5.1 “TLBI ALLE1, TLBI ALLE1NXS, TLB Invalidate All, EL1”, the text that reads:

If FEAT_XS is implemented, the nXS variant of this System instruction is defined.

Both variants perform the same invalidation, but the TLBI System instruction without the nXS qualifier waits for all memory accesses using in-scope old translation information to complete before it is considered complete.

The TLBI System instruction with the nXS qualifier is considered complete when the subset of these memory accesses with XS attribute set to 0 are complete.

is changed to read:

If FEAT_XS is implemented, the nXS variant of this System instruction is defined.

It is **IMPLEMENTATION SPECIFIC** whether the TLBI System instruction with the nXS qualifier invalidates TLB entries with the XS attribute set to 1.

The TLBI System instruction without the nXS qualifier waits for all memory accesses using in-scope old translation information to complete before it is considered complete.

The TLBI System instruction with the nXS qualifier is considered complete when the subset of these memory accesses with XS attribute set to 0 are complete.

Equivalent changes are made to all TLBI{P} instructions that apply to stage 1 and/or stage 2.

In section D8.17.5.4 “TLB maintenance instruction scope”, the following rule is introduced:

R_{X0001}

It is **IMPLEMENTATION SPECIFIC** whether a TLBI maintenance operation with the nXS qualifier invalidates TLB entries with the XS attribute set to 1.

In section D8.17.8 “Ordering and completion of TLB maintenance instructions” the following rule is added:

I_{X0002}

To guarantee completion of memory accesses with the XS attribute set to 1, software must use a TLBI maintenance instruction without the nXS qualifier, and a DSB with or without the nXS qualifier.

In section J1.1.3.426 “GPTTLBIMatch”, the code that reads:

```
boolean TLBIMatch(TLBIRecord tlbi, TLBRecord tlb_entry)
...
return match;
```

is changed to read:

```
boolean TLBIMatch(TLBIRecord tlbi, TLBRecord tlb_entry)
...
if tlbi.attr == TLBI_ExcludeXS && tlb_entry.context.xs == '1' then
    match = FALSE;
return match;
```

2.80 C23842

In section B2.3.6 “Basic definitions”, the definition of *Local write or MMU Fault successor* that reads:

Local write or MMU Fault successor

An Effect E_2 is a Local write or MMU Fault successor of an Effect E_1 if one of the following applies:

- All of the following apply:
 - One of the following applies:
 - E_1 is an Explicit Memory Effect.
 - E_1 is an Implicit Tag Memory Read Effect.
 - E_1 appears in program order before E_2 .
 - E_1 and E_2 are to the Same Location.
 - E_2 is an Explicit Memory Write Effect.
- All of the following apply:
 - E_1 is an Explicit Memory Effect.
 - E_1 appears in program order before E_2 .
 - E_1 and E_2 are to the Same Location.
 - E_2 is an MMU Fault Effect.

is changed to read:

Local write successor

An Effect E_2 is a Local write successor of an Effect E_1 if one of the following applies:

- All of the following apply:
 - One of the following applies:
 - E_1 is an Explicit Memory Effect.
 - E_1 is an Implicit Tag Memory Read Effect.
 - E_1 appears in program order before E_2 .
 - E_1 and E_2 are to the Same Location.
 - E_2 is an Explicit Memory Write Effect.
- All of the following apply:
 - E_1 is an Implicit TTD Memory Read Effect.
 - E_1 appears in program order before E_2 .
 - E_1 and E_2 are to the Same Location.
 - One of the following applies:
 - E_2 is an Explicit Memory Write Effect.
 - E_2 is a Hardware Update Effect.

Local MMU Fault successor

An Effect E_2 is a Local MMU Fault successor of an Effect E_1 if all of the following apply:

- E_1 is an Explicit Memory Effect.
- E_1 appears in program order before E_2 .
- E_1 and E_2 are to the Same Location.
- E_2 is an MMU Fault Effect.

Additionally, in section B2.3.6, the definition of 'Locally-ordered-before' that reads:

Locally-ordered-before

An Effect E_1 is Locally-ordered-before an Effect E_2 if one of the following applies:

...

- E_2 is a Local write or MMU Fault successor of E_1 .
- All of the following apply:
 - E_3 is a Local write or MMU Fault successor of E_1 .
 - E_3 belongs to the same single-copy-atomic class as E_2

is changed to read:

Locally-ordered-before

An Effect E_1 is Locally-ordered-before an Effect E_2 if one of the following applies:

...

- E_2 is a Local write successor of E_1 .
- All of the following apply:
 - E_3 is a Local write successor of E_1 .
 - E_3 belongs to the same single-copy-atomic class as E_2 .
- E_2 is a Local MMU Fault successor of E_1 .

...

2.81 C23844

In section D13.1 “About the Performance Monitors”, the text that reads:

FEAT_PMUv3 requires that an implementation includes the following common events:

is changed to read:

When at least one event counter is implemented, FEAT_PMUv3 requires that an implementation includes the following common events:

2.82 C23845

In section D24.2.167 “SCTLR_EL1, System Control Register (EL1)”, in the description of the field nAA, bit [6], the text that reads:

Non-aligned access. This bit controls generation of Alignment faults at EL1 and EL0 under certain conditions.

The following instructions generate an Alignment fault if all bytes being accessed are not within a single 16-byte quantity, aligned to 16 bytes for access:

- LDAPR, LDAPRH, LDAPUR, LDAPURH, LDAPURSH, LDAPURSW, LDAR, LDARH, LDLAR, LDLARH.
- STLLR, STLLRH, STLR, STLRH, STLUR, and STLURH.

If FEAT_LRCPC3 is implemented, the following instructions generate an Alignment fault if all bytes being accessed for a single register are not within a single 16-byte quantity, aligned to 16 bytes for access:

- LDIAPP, STILP, the post index versions of LDAPR and the pre index versions of STLR.
- If Advanced SIMD and floating-point instructions are implemented, LDAPUR (SIMD&FP), LDAP1 (SIMD&FP), STLUR (SIMD&FP), and STL1 (SIMD&FP).

is changed to read:

Non-aligned access. This bit controls generation of Alignment faults at EL1 and EL0 under certain conditions.

The following instructions generate an Alignment fault if all bytes being accessed are not within a single naturally aligned 16-byte quantity, for access:

- LDAPR, LDAPRH, LDAPUR, LDAPURH, LDAPURSH, LDAPURSW, LDAR, LDARH, LDLAR, LDLARH.
- STLLR, STLLRH, STLR, STLRH, STLUR, and STLURH.
- If FEAT_LRCPC3 is implemented, the post index versions of LDAPR and the pre index versions of STLR.
- If FEAT_LRCPC3 and Advanced SIMD and floating-point instructions are implemented, LDAPUR (SIMD&FP), LDAP1 (SIMD&FP), STLUR (SIMD&FP), and STL1 (SIMD&FP).

If FEAT_LRCPC3 is implemented, the following instructions generate an Alignment fault if all bytes being accessed for a single register are not within a single naturally aligned 16-byte quantity, for access:

- LDIAPP, STILP

The equivalent changes are made in the following sections:

- D24.2.168 “SCTLR_EL2, System Control Register (EL2)”, field nAA.
- D24.2.169 “SCTLR_EL3, System Control Register (EL3)”, field nAA.

2.83 R23848

In section D24.7.11 “PMSIDR_EL1, Sampling Profiling ID Register” the text that reads:

Recommended minimum sampling interval. This provides guidance from the implementer to the smallest minimum sampling interval, N.

is changed to read:

Minimum sampling interval. A nonzero value provides guidance from the implementer to the smallest sampling interval.

The interval is defined as a number of either operations or instructions, depending on the **IMPLEMENTATION DEFINED** value of PMSIDR_EL1.ArchInst.

...

If this field reads as a nonzero value, then setting the sample interval to a value less than the indicated value is likely to cause a statistically significant number of sample collisions. For example, the smallest interval might indicate the typical size of the speculation window for the processor implementation. That is, the number of operations or instructions that will be executing in parallel, when the processor implementation is not heavily loaded.

Setting the sample interval to a value greater than the indicated value does not guarantee a statistically insignificant number of sample collisions, as this is typically dependent on the program being profiled. For example, if the program has a large number of long latency loads due to cache misses, then this might cause more collisions when these operations or instructions are sampled.

For more information see 'Sample collisions'.

The following text is added in section D17.3.7 "Sample collisions":

Note: Sample collisions are likely to increase with shorter sampling intervals. Samples might not be representative of the overall workload if statistically significant numbers of sample collisions occur.

2.84 D23852

In Section J1.1.1.62 "SPEConstructRecord" the pseudocode that reads:

```

        // Check for overflow
        constant boolean large_counters = boolean
IMPLEMENTATION_DEFINED "SPE 16bit counters";
        if SPESampleCounter\[counter_index\] > 0xFFFF && large\_counters
then
            SPESampleCounter\[counter_index\] = 0xFFFF;
        elseif SPESampleCounter\[counter_index\] > 0xFFF then
            SPESampleCounter\[counter_index\] = 0xFFF;
    
```

is changed to read:

```

        // Check for overflow
        if (boolean IMPLEMENTATION_DEFINED "SPE 16bit counters")
then
            if SPESampleCounter\[counter_index\] > 0xFFFF then
                SPESampleCounter\[counter_index\] = 0xFFFF;
else
            if SPESampleCounter\[counter_index\] > 0xFFF then
                SPESampleCounter\[counter_index\] = 0xFFF;
    
```

2.85 D23854

In section D24.2.56 "HCR_EL2, Hypervisor Configuration Register", in the description of the field FB, bit [9], the text that reads:

Force broadcast. Causes the following instructions to be broadcast within the Inner Shareable domain when executed from EL1:

AArch32: BPIALL, TLBIALL, TLBIMVA, TLBIASID, DTLBIALL, DTLBIMVA, DTLBIASID, ITLBIALL, ITLBIMVA, ITLBIASID, TLBIMVAA, ICIALLU, TLBIMVAL, and TLBIMVAAL.

AArch64: TLBI VMALLE1, TLBI VAE1, TLBI ASIDE1, TLBI VAAE1, TLBI VALE1, TLBI VAALE1, IC IALLU, TLBI RVAE1, TLBI RVAAE1, TLBI RVALE1, and TLBI RVAALE1.

is changed to read:

Force broadcast. Causes the following Non-shareable invalidate instructions to be broadcast within the Inner Shareable domain when executed from EL1:

- In AArch32: BPIALL, TLBIALL, TLBIMVA, TLBIASID, DTLBIALL, DTLBIMVA, DTLBIASID, ITLBIALL, ITLBIMVA, ITLBIASID, TLBIMVAA, ICIALLU, TLBIMVAL, and TLBIMVAAL.
- In AArch64:
 - TLBI instructions that are executable at EL1 and do not specify IS or OS.
 - TLBIP instructions that are executable at EL1 and do not specify IS or OS.
 - IC IALLU.

See also:

- A64 System instructions for TLB maintenance.

2.86 C23859

In section K1.2.13 “Crossing a page boundary with different memory types or Shareability attributes”, the text that reads:

A memory access from a load or store instruction that crosses a page boundary to a memory location that has a different memory type or Shareability attribute results in **CONSTRAINED UNPREDICTABLE** behavior. In this case, the implementation must perform one of the following behaviors:

...

is changed to read:

If a single load or store instruction generates multiple memory accesses, such that the total set of accesses crosses a page boundary to a memory location that has a different memory type, Normal or Device, or Shareability attribute, the result is **CONSTRAINED UNPREDICTABLE** behavior. In this case, the implementation must perform one of the following behaviors:

...

2.87 C23866

In section A2.3.6 “The Armv9.5 architecture extension”, under the description of ‘FEAT_ETS3, Enhanced Translation Synchronization’, the following text is added:

If FEAT_ETS3 is implemented, then FEAT_ETS2 is implemented.

2.88 D23869

In section D24.2.168 “SCTLR_EL2, System Control Register (EL2)”, the text that reads:

The reset behavior of this field is:

- On a Warm reset:
- When the highest implemented Exception level is EL2, this field resets to an architecturally **UNKNOWN** value.

is changed to read:

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

Equivalent changes are made in the following sections:

- D24.2.167 “SCTLR_EL1, System Control Register (EL1)”
- D24.2.169 “SCTLR_EL3, System Control Register (EL3)”
- G8.2.73 “HSCTLR, Hyp System Control Register”
- G8.2.120 “NSACR, Non-Secure Access Control Register”
- G8.2.173 “VTTBR, Virtualization Translation Table Base Register”
- G8.3.32 “HDCR, Hyp Debug Control Register”
- G8.3.35 “SDCR, Secure Debug Control Register”

2.89 C23900

In section D10.4 “Tag checking”, the rule R_{CPKNZ} that reads:

A single-copy atomic explicit memory access can be either:

- A Tag Unchecked memory access.
- A Tag Checked memory access.

is changed to read:

A Memory Read or Write Effect generated by an instruction is one of:

- Tag Unchecked.
- Tag Checked.

In the same section, rule R_{DRGYL} that reads:

A single-copy atomic explicit memory access is Tag Unchecked due to any of the following, otherwise it is Tag Checked:

- The access is:
 - ...
 - An instruction fetch. ...
 - A read of an Allocation Tag due to a Tag Check operation.
 - ...
 - Due to a translation table walk.
 - ...
 - A load or store relative to VNCR_EL2 if FEAT_NV2 is implemented.
 - ...
 - A read generated by an instruction executed at ELx, if FEAT_MTE_STORE_ONLY is enabled at ELx.
 - ...

is changed to read:

A Memory Read or Write Effect is Tag Unchecked due to any of the following, otherwise it is Tag Checked:

- The Memory Effect is an Explicit Memory effect, and is:
 - ...
 - Due to a load or store relative to VNCR_EL2, if FEAT_NV2 is implemented.
 - ...
 - A Memory Read Effect generated by an instruction for which Store-only Tag checking is enabled, if FEAT_MTE_STORE_ONLY is implemented.
 - ...
- The Memory Effect is not an Explicit Memory Effect.

In the same section, rule R_{MNHCZ}, that reads:

If FEAT_MTE_STORE_ONLY is implemented, for a Tag Checked Compare and swap – instruction, if the compare fails, it is **CONSTRAINED UNPREDICTABLE** whether the Tag Check operation is performed.

is changed to read:

For a Compare and swap instruction, if the compare fails, a Memory Write Effect by the instruction is permitted to be treated as Tag Unchecked.

In section D24.2.169 “SCTLR_EL3”, the text that reads:

TCSO	Meaning
0b0	This field has no effect on Tag checking.
0b1	Load instructions executed in EL3 are Tag Unchecked.

is changed to read:

TCSO	Meaning
0b0	This field has no effect on Tag checking.
0b1	Memory Read Effects by instructions executed in EL3 are Tag Unchecked.

Equivalent changes are made to the TCSO, TCSO0, and TCSO1 fields in:

- D24.2.168 “SCTLR_EL2, System Control Register (EL2)”.
- D24.2.167 “SCTLR_EL1, System Control Register (EL1)”.

2.90 R23907

In section D10.3 “Memory region tagging types”, the text that reads:

R_{JRJP}

The memory region tagging type for a memory region is:

- Tagged if all of the following are true:
 - FEAT_MTE2 is implemented.
 - Allocation Tag Access is enabled for the memory region. For more information, see Allocation Tag Access controls.
 - The stage 1 attributes define the Tagged attribute.
 - The combined effects of stage 1 and stage 2 translations define the memory attributes as Normal memory, Inner and Outer Write-Back, Non-Transient, Read-Allocate and Write-Allocate.
- Otherwise it is Canonically Tagged if all of the following are true:
 - FEAT_MTE_CANONICAL_TAGS is implemented.
 - Canonical Tagging is enabled for the memory region.
 - The stage 1 attributes do not define the Tagged attribute.
- Otherwise it is Untagged.

is changed to read:

R_{JRJP}

The memory region tagging type for a memory region is:

- Tagged if all of the following are true:

- FEAT_MTE2 is implemented.
- Allocation Tag Access is enabled for the memory region. For more information, see Allocation Tag Access controls.
- The stage 1 attributes define the Tagged attribute.
- The combined effects of stage 1 and stage 2 translations define the memory attributes as Normal memory, Inner and Outer Write-Back, Non-Transient, Read-Allocate and Write-Allocate.
- Otherwise it is Canonically Tagged if all of the following are true:
 - FEAT_MTE_CANONICAL_TAGS is implemented.
 - Canonical Tagging is enabled for the memory region.
 - The stage 1 attributes do not define the Tagged attribute.
- Otherwise it is **CONSTRAINED UNPREDICTABLE** if all of the following are true:
 - FEAT_MTE_CANONICAL_TAGS is implemented.
 - Canonical Tagging is enabled for the memory region.
 - The stage 1 attributes define the Tagged attribute.
 - The combined effects of stage 1 and stage 2 translations do not define the memory attributes as Normal memory, Inner and Outer Write-Back, Non-Transient, Read-Allocate and Write-Allocate.
- Otherwise it is Untagged.

2.91 R23909

In section D17.6.5 “Additional information for each profiled conditional instruction”, the text that reads:

A conditional compare operation is treated as a conditional operation.

It is **IMPLEMENTATION DEFINED** which conditional select operations (both integer and floating-point), including general-purpose, SIMD&FP, and SVE operations, are treated as conditional:

- Conditional select.
- Conditional select increment.
- Conditional select negate.
- Conditional select invert operations.

Predicated SVE operations are not conditional operations, as the conditionality of the operation is controlled by a predicate.

is changed to read:

Only operations where the conditionality is controlled by $PSTATE.\{N, Z, C, V\}$ are treated as conditional operations. Predicated SVE operations are not treated as conditional operations, because the conditionality of the operation is controlled by a predicate.

For each of the following operations, it is **IMPLEMENTATION DEFINED** whether it is treated as a conditional operation:

- Conditional select, both integer and floating-point.
- Conditional select increment.
- Conditional select negation.
- Conditional select inversion.
- Conditional compare, both integer and floating-point.
- Conditional compare negative.

Arm recommends that all operations where the conditionality is controlled by $PSTATE.\{N, Z, C, V\}$ are treated as conditional operations.

2.92 R23917

In section D24.2.163 “SCR_EL3, Secure Configuration Register”, in the description of the field $PIEn$, bit [45], the table that reads:

$PIEn$	Meaning
0b0	EL0, EL1 and EL2 accesses to the specified registers are disabled, and trapped to EL3. The values in these registers are treated as 0.
0b1	This control does not cause any instructions to be trapped.

is changed to read:

$PIEn$	Meaning
0b0	EL0, EL1 and EL2 accesses to the specified registers are disabled, and trapped to EL3.
0b1	This control does not cause any instructions to be trapped.

If this field is 0, it is **IMPLEMENTATION SPECIFIC** whether the values of the named registers are treated as zero.

2.93 R23928

In section B2.4.1.1 “Peripherals”, the text that reads:

Peripheral coherence order

The Peripheral coherence order of a Memory-mapped peripheral is a total order on all reads and writes to that peripheral.

Note:

The Peripheral coherence order for a Memory-mapped peripheral signifies the order in which accesses arrive at the endpoint.

For a Read Memory or write effect RW1 and a Read Memory or write effect RW2 to the same peripheral, then RW1 will appear in the Peripheral coherence order for the peripheral before RW2 if either of the following cases apply:

- RW1 and RW2 are accesses using Non-cacheable or Device attributes and RW1 is Ordered-before RW2.
- RW1 and RW2 are accesses using Device-nGnRE or Device-nGnRnE attributes, with the same XS attribute value, and RW1 appears in program order before RW2.

Note:

When FEAT_XS is implemented, if accesses marked with the Device-nGnRE or Device-nGnRnE attributes are within the same Memory-mapped peripheral, but the XS attribute is not the same on those accesses, the order of arrival at the endpoint is not defined by the architecture.

is changed to read:

Peripheral arrival order

The Peripheral arrival order of a Memory-mapped peripheral is a total order on all Memory Read Effects and Memory Write Effects to that peripheral.

For a Memory Effect E_1 and a Memory Effect E_2 to the same peripheral if any of the following applies:

- All of the following apply:
 - Any of the following applies:
 - E_1 is to Device-nGnRE Memory.
 - E_1 is to Device-nGnRnE Memory.
 - Any of the following applies:
 - E_2 is to Device-nGnRE Memory.
 - E_2 is to Device-nGnRnE Memory.
 - Any of the following applies:
 - FEAT_XS is not implemented.
 - All of the following apply:
 - FEAT_XS is implemented.
 - E_1 has the XS attribute.
 - E_2 has the XS attribute.
 - All of the following apply:
 - FEAT_XS is implemented.

- E_1 does not have the XS attribute.
- E_2 does not have the XS attribute.
- E_1 is in program-order before E_2 .
- All of the following apply:
 - Any of the following applies:
 - E_1 is to Device-nGRE Memory.
 - E_1 is to Device-nGnRE Memory.
 - E_1 is to Device-nGnRnE Memory.
 - Any of the following applies:
 - E_2 is to Device-nGRE Memory.
 - E_2 is to Device-nGnRE Memory.
 - E_2 is to Device-nGnRnE Memory.
 - E_1 is Ordered-before E_2 .
- All of the following apply:
 - Any of the following applies:
 - E_1 is to Device Memory.
 - E_1 is to Non-cacheable Normal Memory.
 - Any of the following applies:
 - E_2 is to Device Memory.
 - E_2 is to Non-cacheable Normal Memory.
 - E_1 is Barrier-ordered-before E_2 .

Then E_1 is inserted in the Peripheral arrival order for the peripheral before E_2 and it is not permitted that E_1 and E_2 are gathered into a single transaction. Otherwise, exactly one of following applies:

- E_1 and E_2 are gathered into a single transaction and appear as one Effect in Peripheral arrival order for the peripheral.
- E_1 is inserted in the Peripheral arrival order for the peripheral before E_2 .
- E_2 is inserted in the Peripheral arrival order for the peripheral before E_1 .

Note:

- When FEAT_XS is implemented, if E_1 and E_2 are to Device-nGnRE Memory or to Device-nGnRnE Memory and E_1 and E_2 are within the same Memory-mapped peripheral, but E_1 and E_2 use different XS attribute, the order of arrival at the endpoint is not defined by the architecture.

2.94 D23933

In section D24.2.167 “SCTLR_EL1, System Control Register (EL1)”, the description of the field EE, bit [25] is changed to reflect the addition of four new features (FEAT_BigEnd, FEAT_LittleEnd, FEAT_BigEndELO, FEAT_LittleEndELO).

The text that reads:

When FEAT_MixedEnd is implemented: Endianness of data accesses at EL1, and stage-1 translation table walks in the EL1&O translation regime.

EE	Meaning
0b0	Explicit data accesses at EL1, and stage-1 translation table walks in the EL1&O translation regime are little-endian.
0b1	Explicit data accesses at EL1, and stage-1 translation table walks in the EL1&O translation regime are big-endian.

The EE bit is permitted to be cached in a TLB.

When the Effective value of HCR_EL2.{E2H, TGE} is {1, 1}, this bit has no effect on the PE.

The reset behavior of this field is:

- On a Warm reset, this field resets to an **IMPLEMENTATION DEFINED** value.

Otherwise:

Endianness of data accesses at EL1, and stage-1 translation table walks in the EL1&O translation regime.

EE	Meaning
0b0	Explicit data accesses at EL1, and stage-1 translation table walks in the EL1&O translation regime are little-endian.
0b1	Explicit data accesses at EL1, and stage-1 translation table walks in the EL1&O translation regime are big-endian.

If an implementation does not provide Big-endian support at Exception levels higher than ELO, this bit is **RES0**.

If an implementation does not provide Little-endian support at Exception levels higher than ELO, this bit is **RES1**.

The EE bit is permitted to be cached in a TLB.

When the Effective value of HCR_EL2.{E2H, TGE} is {1, 1}, this bit has no effect on the PE.

The reset behavior of this field is:

- On a Warm reset, this field resets to an **IMPLEMENTATION DEFINED** value.

is changed to read:

When FEAT_MixedEnd is implemented:

Endianness of data accesses at EL1, and stage-1 translation table walks in the EL1&0 translation regime.

EE	Meaning
0b0	Explicit data accesses at EL1, and stage-1 translation table walks in the EL1&0 translation regime are little-endian.
0b1	Explicit data accesses at EL1, and stage-1 translation table walks in the EL1&0 translation regime are big-endian.

The EE bit is permitted to be cached in a TLB.

When the Effective value of HCR_EL2.{E2H, TGE} is {1, 1}, this bit has no effect on the PE.

The reset behavior of this field is:

- On a Warm reset, this field resets to an **IMPLEMENTATION DEFINED** value.

When FEAT_BigEnd is implemented

Explicit data accesses at EL1, and stage-1 translation table walks in the EL1&0 translation regime are big-endian.

This field is **RES1**.

Otherwise

Explicit data accesses at EL1, and stage-1 translation table walks in the EL1&0 translation regime are little-endian.

This field is **RES0**.

Equivalent changes are made in the following sections:

- D24.2.168 “SCTLR_EL2, System Control Register (EL2)” – field EE
- D24.2.169 “SCTLR_EL3, System Control Register (EL3)” – field EE

In section D24.2.167 “SCTLR_EL1, System Control Register (EL1)”, in the description of the field EOE, bit [24], the text that reads:

EOE, bit [24]

When FEAT_MixedEndELO is implemented:

Endianness of data accesses at ELO.

EOE	Meaning
0b0	Explicit data accesses at ELO are little-endian.
0b1	Explicit data accesses at ELO are big-endian.

This bit has no effect on the endianness of LDTR, LDTRH, LDTRSH, LDTRSW, STTR, and STTRH instructions executed at EL1.

When the Effective value of HCR_EL2.{E2H, TGE} is {1, 1}, this bit has no effect on execution at ELO.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

Otherwise:

Endianness of data accesses at ELO.

EOE	Meaning
0b0	Explicit data accesses at ELO are little-endian.
0b1	Explicit data accesses at ELO are big-endian.

If an implementation only supports Little-endian accesses at ELO, this bit is **RES0**.

If an implementation only supports Big-endian accesses at ELO, this bit is **RES1**.

This bit has no effect on the endianness of LDTR, LDTRH, LDTRSH, LDTRSW, STTR, and STTRH instructions executed at EL1.

When the Effective value of HCR_EL2.{E2H, TGE} is {1, 1}, this bit has no effect on execution at ELO.

The reset behavior of this field is:

- On a Warm reset, this field resets to an **IMPLEMENTATION DEFINED** value.

is changed to read:

When FEAT_MixedEndELO is implemented: Endianness of data accesses at ELO.

EOE	Meaning
0b0	Explicit data accesses at ELO are little-endian.
0b1	Explicit data accesses at ELO are big-endian.

This bit has no effect on the endianness of LDTR, LDTRH, LDTRSH, LDTRSW, STTR, and STTRH instructions executed at EL1.

When the Effective value of HCR_EL2.{E2H, TGE} is {1, 1}, this bit has no effect on execution at ELO.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

When FEAT_BigEndELO is implemented

Explicit data accesses at ELO are big-endian.

This field is **RES1**.

Otherwise

Explicit data accesses at ELO are little-endian.

This field is **RES0**.

Equivalent changes are made in the following section:

- D24.2.168 “SCTLR_EL2, System Control Register (EL2)” – field EOE.

2.95 D23934

In section D14.3 “Common event numbers”, in the event description of “0x80F1, ASE_FP_DOT_SPEC, Floating-point operation speculatively executed, Advanced SIMD dot-product”, the text that reads:

- When FEAT_BF8DOT2 is implemented, Advanced SIMD: FDOT (2-way, by element) or FDOT (2-way, vector).

is changed to read:

- When FEAT_FP8DOT2 is implemented, Advanced SIMD: FDOT (2-way, by element) or FDOT (2-way, vector).

Equivalent changes are made in the following event description:

- 0x80F3, ASE_SVE_FP_DOT_SPEC, Floating-point operation speculatively executed, Advanced SIMD or SVE dot-product.

In the same section, in the event description of “0x8028, FP_FMA_SPEC, Floating-point operation speculatively executed, FMA”, the text that reads:

- When FEAT_BF8FMA is implemented, Advanced SIMD: FMLALB, FMLALLBB, FMLALLBT, FMLALLTB, FMLALLTT, or FMLALT.

is changed to read:

- When FEAT_FP8FMA is implemented, Advanced SIMD: FMLALB, FMLALLBB, FMLALLBT, FMLALLTB, FMLALLTT, or FMLALT.

Equivalent changes are made in the following event descriptions:

- 0x8029, ASE_FP_FMA_SPEC, Floating-point operation speculatively executed, Advanced SIMD FMA.
- 0x802B, ASE_SVE_FP_FMA_SPEC, Floating-point operation speculatively executed, Advanced SIMD or SVE FMA.

2.96 C23935

In section D24.5.8 “PMCR_ELO, Performance Monitors Control Register”, in the description of the field X, bit [4], the following text is added:

If FEAT_ETE is implemented, this field does not affect the use of PMU events as an External Input by the trace unit. If FEAT_ETM4 is implemented, this field does affect the use of PMU events as an External Input by the trace unit.

The same clarification is added in section G8.4.9 “PMCR, Performance Monitors Control Register”.

2.97 D23936

In section D13.1.4 “Interaction with trace”, the text that reads:

D13.1.4 Interaction with trace

It is **IMPLEMENTATION DEFINED** whether the implementation exports counter events to a trace unit, or other external monitoring agent, to provide triggering information. The form of any exporting is also **IMPLEMENTATION DEFINED**. If implemented, this exporting might be enabled as part of the performance monitoring control functionality.

Arm recommends system designers include a mechanism for importing a set of external events to be counted, but such a feature is **IMPLEMENTATION DEFINED**. When implemented, this feature enables the trace unit to pass in events to be counted.

Exporting PMU events to the ETM is prohibited for some Exception levels when `SelfHostedTraceEnabled() == TRUE`. For more information, see Controls to prohibit trace at Exception levels.

is changed to read:

D13.1.4 Interaction with trace

When FEAT_ETE is implemented, the trace unit uses the PMU events as External Inputs. See External inputs.

When FEAT_ETMv4 is implemented, it is **IMPLEMENTATION DEFINED** whether the implementation exports PMU events to the ETM trace unit, and the form of any exporting is also **IMPLEMENTATION DEFINED**. If implemented, this exporting is controlled by PMCR.X.

For more information about when the trace unit is permitted to observe PMU events, see Controls to prohibit trace at Exception levels and External inputs.

When FEAT_ETE is implemented, the PMU implements a set of events for importing external events from the trace unit and the cross-trigger interface. See Required events. Arm recommends similar functionality is implemented when FEAT_ETMv4 is implemented, and further recommends

re-using the events defined for FEAT_ETE. When implemented, this functionality enables the trace unit to pass in events to be counted.

D13.1.5 Export of PMU events

It is **IMPLEMENTATION DEFINED** whether the implementation exports counter events to any other external monitoring agent to provide triggering information. The form of any exporting is also **IMPLEMENTATION DEFINED**. If implemented, this exporting is controlled by PMCR.X.

2.98 D23947

In section D13.8 “Event threshold and edge counting”, the rule that reads:

R_{HBWB}

The value of $V[n]$ for event counter $\langle n \rangle$ is defined as follows, where TC, TH, and TE are the Effective values of $PMEVTYPER\langle n \rangle_ELO.\{TC, TH, TE\}$ respectively:

$$V[n] = \begin{cases} V_B[n], & \text{if } TC = 0b000 \text{ and } TH = 0 \text{ (disabled)}. \\ V_T[n], & \text{if } TC \neq 0b000 \text{ and } TH \neq 0 \text{ and } TE = 0b0 \text{ (threshold)}. \\ V_E[n], & \text{if } TE = 0b1 \text{ (edge)}. \end{cases}$$

The pseudocode function $PMUCountValue()$ in Armv8 Pseudocode describes this.

is changed to read:

R_{HBWB}

The value of $V[n]$ for event counter $\langle n \rangle$ is defined as follows, where TC, TH, and TE are the Effective values of $PMEVTYPER\langle n \rangle_ELO.\{TC, TH, TE\}$ respectively:

$$V[n] = \begin{cases} V_B[n], & \text{if } TC = 0b000 \text{ and } TH = 0 \text{ (disabled)}. \\ V_T[n], & \text{if } (TC \neq 0b000 \text{ or } TH \neq 0) \text{ and } TE = 0b0 \text{ (threshold)}. \\ V_E[n], & \text{if } TE = 0b1 \text{ (edge)}. \end{cases}$$

The pseudocode function $PMUCountValue()$ in Armv8 Pseudocode describes this.

2.99 C23956

In section B2.6.9 “Data Synchronization Barrier”, the text that reads:

In addition, no instruction that appears in program order after the DSB instruction can alter any state of the system or perform any part of its functionality until the DSB completes other than:

- Being fetched from memory and decoded.
- Reading the general-purpose, SIMD&FP, SVE vector or predicate, Special-purpose, or System registers that are directly or indirectly read without causing side effects.

- If FEAT_ETS2 is not implemented, having any virtual addresses of loads and stores translated.

is changed to read:

In addition, no instruction that appears in program order after the DSB instruction can alter any state of the system or perform any part of its functionality until the DSB completes other than:

- Being fetched from memory and decoded, or any change to architectural or microarchitectural state resulting from being fetched or decoded.
- Indirectly or directly reading any of the following, so long as the reading of these resources does not cause a change in architectural state:
 - General-purpose registers
 - SIMD&FP registers
 - The Stack Pointer register
 - The Program Counter
 - If FEAT_SVE or FEAT_SME is implemented, SVE scalable vector registers or SVE predicate registers
 - If FEAT_SVE is implemented, the First Fault Register
 - If FEAT_SME is implemented, the ZA storage
 - If FEAT_SME2 is implemented, the ZTO register
 - Special-purpose registers
 - System registers other than CNTPCTSS_ELO and CNTVCTSS_ELO
- If FEAT_ETS2 is not implemented, translating any virtual addresses of instructions that generate Explicit Memory Effects.

Equivalent changes are made in section E2.3.5.3 “Data Synchronization Barrier”.

2.100 D23962

In section D14.3.2 “Common microarchitectural events”, the text that reads:

0x8432, SVE_FP_PREDUCE_SPEC, Floating-point operation_speculatively_executed, Advanced SIMD pairwise add step or pairwise reduce step

is changed to read:

0x8432, SVE_FP_PREDUCE_SPEC, Floating-point operation speculatively executed, SVE pairwise add step or pairwise reduce step

2.101 D23966

In section D8.5.2.3 “Hardware dirty state tracking structure”, the text that reads:

R_{YHLRY}

If the HDBSS is full and enabled, then the PE does not update any stage 2 descriptors from writable-clean to writable-dirty.

R_{STJMN}

If the HDBSS is full and enabled, or HDBSSPROD_EL2.FSC is set to any value other than 0b000000, then all of the following apply:

- The PE cannot append entries to the HDBSS.
- Stage 2 dirty state hardware updates are prevented.
- A non-speculative access that requires a stage 2 dirty state hardware update generates an MMU fault that is reported as an Instruction Abort or Data Abort taken to EL2.
 - ESR_EL2 is populated as a stage 2 Permission fault as though VTCR_EL2.HD is 0, indicating the walk level of the stage 2 descriptor that needs to be updated.
 - ESR_EL2.ISS2.HDBSSF is reported as 1, to indicate that the failed update was a result of the HDBSS being full or in an error state.

is changed to read:

R_{YHLRY}

If the HDBSS is full and enabled, then the PE does not update any stage 2 descriptors from writable-clean to writable-dirty.

R_{STJMN}

If the HDBSS is full and enabled, or HDBSSPROD_EL2.FSC is set to any value other than 0b000000, then all of the following apply:

- The PE cannot append entries to the HDBSS.
- Stage 2 dirty state hardware updates are prevented.
- Any access that requires a stage 2 dirty state hardware update generates a stage 2 Permission fault that is reported as it would be if VTCR_EL2.HD were 0.
- If the fault is reported in ESR_EL2, ESR_EL2.ISS2.HDBSSF is reported as 1, to indicate that the failed update was a result of the HDBSS being full or in an error state.

I_{X0001}

MMU faults on speculative accesses are never reported.

If the HDBSS is full and enabled, and this generates a stage 2 Permission fault as part of resolving an AT E1 operation executed at EL2, it is reported in PAR_EL1 as a stage 2 Permission fault on a stage 1 translation table walk, at the level of the stage 2 walk that experienced the fault.

2.102 C23967

In section D24.7.9 “PMSFCR_EL1, Sampling Filter Control Register”, the text that reads:

R_{YGCFK}

When FEAT_SPE_EFT is implemented and PMSFCR_EL1.FT is 1, all of the following apply:

- The sampled operation is discarded if the following is true:

```
(!IsZero(ctrl AND NOT mask) && IsZero(flags AND (ctrl AND NOT mask)))
```

That is, the SPU applies a Boolean-OR filter based on each of the operation type filter controls where the corresponding operation type filter mask bit is 0:

- If all these bits are zero, or all the mask bits are one, then no operations are discarded by this part of the filter.
- Otherwise, operations that match any of these bits are selected, and other operations are discarded.
- ...

is changed to read:

R_{YGCFK}

When FEAT_SPE_EFT is implemented and PMSFCR_EL1.FT is 1, all of the following apply:

- The sampled operation is discarded if the following is true:

```
(!IsZero(ctrl AND NOT mask) && IsZero(flags AND (ctrl AND NOT mask)))
```

That is, the SPU applies a Boolean-OR filter based on each of the operation type filter controls where the corresponding operation type filter mask bit is 0:

- If all the operation type filter mask bits are one, then no operations are discarded by this part of the filter.
- Otherwise, for the operation type bits for which corresponding mask bit is 0:
 - If all these bits are 0, then no operations are discarded by this part of the filter.
 - Otherwise, operations that match any of these bits are selected, and other operations are discarded.
- ...

In the same section, the description of the LD, bit [17], which reads:

LD, bit [17]

Load filter enable.

LD	Meaning
0b0	If FEAT_SPE_EFT is implemented and PMSFCR_EL1.LDm is 1, then record only operations that are not load operations. Otherwise, do not record load operations, unless enabled by another filter.
0b1	If FEAT_SPE_EFT is implemented and PMSFCR_EL1.LDm is 1, then record only operations that are load operations. Otherwise, record all load operations.

This field is ignored by the PE when PMSFCR_EL1.FT == 0.

For filtering purposes, load operations include vector loads and atomic operations that return a value to the PE.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

is changed to read:

LD, bit [17]

When FEAT_SPE_EFT is implemented and PMSFCR_EL1.LDm is 1

Load filter enable.

LD	Meaning
0b0	Record only operations that are not load operations.
0b1	Record only operations that are load operations.

This field is ignored by the PE when PMSFCR_EL1.FT == 0.

For filtering purposes, load operations include vector loads and atomic operations that return a value to the PE.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

Otherwise

Load filter enable.

LD	Meaning
0b0	Do not record load operations, unless enabled by another filter.
0b1	Record all load operations.

This field is ignored by the PE and no records are removed by this filter when any of the following apply:

- PMSFCR_EL1.FT == 0.
- FEAT_SPE_EFT is implemented and the values of all of PMSFCR_EL1.{SIMDm, FPm, STm, LDm, Bm} are zero for which the corresponding PMSFCR_EL1.{SIMD, FP, ST, LD, B} bit is zero.

If FEAT_SPE_EFT is not implemented and the values of PMSFCR_EL1.{ST, LD, B} are all zero, then it is **CONSTRAINED UNPREDICTABLE** whether any records are removed by this filter.

For filtering purposes, load operations include vector loads and atomic operations that return a value to the PE.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

Similar updates are made to each of the other fields {SIMD, FP, ST, B}, noting that PMSFCR_EL1.{SIMD, FP} are only implemented when FEAT_SPE_EFT is implemented.

Additionally, the description of the FT bit with the text that reads:

FT, bit [1]

Filter by operation type. The filter is the logical OR of the ST, LD and B bits. For example, if LD and ST are both set, both load and store operations are recorded

FT	Meaning
0b0	
0b1	Type filtering enabled. Samples not one of the selected operation types will not be recorded

If this field is set to 1 and the PMSFCR_EL1.{ST, LD, B} bits are all set to zero, it is **CONSTRAINED UNPREDICTABLE** whether no samples are recorded or the PE behaves as if PMSFCR_EL1.FT is set to 0

is changed to read:

FT, bit [1]

Filter by operation type. The filter is controlled by the {ST, LD, B} and, if FEAT_SPE_EFT is implemented, {SIMD, FP} and {SIMDm, FPm, STm, LDm, Bm} fields.

FT	Meaning
0b0	Type filtering disabled.
0b1	Type filtering enabled. Samples not one of the selected operation types will not be recorded.

If FEAT_SPE_EFT is not implemented, this field is 1, and the PMSFCR_EL1.{ST, LD, B} bits are all zero, then it is **CONSTRAINED UNPREDICTABLE** whether no samples are recorded or the PE behaves as if PMSFCR_EL1.FT is 0.

2.103 C23982

In section D24.2.197 “TTBRO_EL3, Translation Table Base Register 0 (EL3)”, under the condition “When FEAT_D128 is implemented and TCR_EL3.D128 == 1:”, the description of field CnP, bit [0] the text that reads:

CnP, bit [0]

Common not Private. In an implementation that includes FEAT_TTCNP, indicates whether each entry that is pointed to by VSTTBR_EL2 is a member of a common set that can be used by every PE in the Inner Shareable domain for which the value of VSTTBR_EL2.CnP is 1.

CnP	Meaning
0b0	The translation table entries pointed to by VSTTBR_EL2 are permitted to differ from the entries for VSTTBR_EL2 for other PEs in the Inner Shareable domain. This is not affected by the value of the current VMID.
0b1	The translation table entries pointed to by VSTTBR_EL2 are the same as the translation table entries for every other PE in the Inner Shareable domain for which the value of VSTTBR_EL2.CnP is 1 and the VMID is the same as the current VMID.

This bit is permitted to be cached in a TLB.

Note:

If the value of VSTTBR_EL2.CnP bit is 1 on multiple PEs in the same Inner Shareable domain and those VSTTBR_EL2s do not point to the same translation table entries when using the current VMID, then the results of translations using VSTTBR_EL2 are **CONSTRAINED UNPREDICTABLE**, see ‘**CONSTRAINED UNPREDICTABLE** behaviors due to caching of control or data values’.

When this register has an architecturally-defined reset value, this field resets to a value that is architecturally **UNKNOWN**.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

is changed to read:

CnP, bit [0]

Common not Private. This bit indicates whether each entry that is pointed to by TTBRO_EL3 is a member of a common set that can be used by every PE in the Inner Shareable domain for which the value of TTBRO_EL3.CnP is 1.

CnP	Meaning
0b0	The translation table entries pointed to by TTBRO_EL3, for the current translation regime, are permitted to differ from corresponding entries for TTBRO_EL3 on other PEs in the Inner Shareable domain.
0b1	The translation table entries pointed to by TTBRO_EL3 are the same as the translation table entries for every other PE in the Inner Shareable domain for which the value of TTBRO_EL3.CnP is 1.

This bit is permitted to be cached in a TLB.

Note:

If the value of the TTBR0_EL3.CnP bit is 1 on multiple PEs in the same Inner Shareable domain and those TTBR0_EL3s do not point to the same translation table entries, then the results of translations using TTBR0_EL3 are **CONSTRAINED UNPREDICTABLE**, see '**CONSTRAINED UNPREDICTABLE** behaviors due to caching of control or data values'.

The reset behavior of this field is:

- On a Warm reset, this field resets to an architecturally **UNKNOWN** value.

2.104 D23988

In section D1.3.5.5 “Prioritization of Synchronous exceptions taken to AArch64 state”, under rule R_{ZFGJP}, the following text is added to the priority table between the existing priorities 14 and 15:

Priority	Synchronous exception type
14	...
...	Exceptions taken to EL3 due to the configuration of MPAM3_EL3.TRAPLOWER.

2.105 D23995

In section D24.2.52 “HACDBSBR_EL2, Hardware Accelerator for Cleaning Dirty State Base Register”, the following encoding is added:

SZ, bits [3:0]

Size of the HACDBS.

SZ	Meaning
0b0000	4 KB
...	...

The equivalent change is made in:

- D24.2.58 “HDBSSBR_EL2, Hardware Dirty State Tracking Structure Base Register”.

2.106 C23997

In section B2.2.1 “Requirements for single-copy atomicity”, the statements that read:

For explicit memory effects generated from an Exception level the following rules apply:

...

- Reads to SIMD and floating-point registers of a single 64-bit or smaller quantity that is aligned to the size of the quantity being loaded are treated as single-copy atomic reads.
- Writes from SIMD and floating-point registers of a single 64-bit or smaller quantity that is aligned to the size of the quantity being stored are treated as single-copy atomic writes.
- Element or Structure Reads to SIMD and floating-point registers of 64-bit or smaller elements, where each element is aligned to the size of the element being loaded, have each element treated as a single-copy atomic read.
- Element or Structure Writes from SIMD and floating-point registers of 64-bit or smaller elements, where each element is aligned to the size of the element being stored, have each element treated as a single-copy atomic store.
- Reads to SIMD and floating-point registers of a 128-bit value that is 64-bit aligned in memory are treated as a pair of single-copy atomic 64-bit reads.
- Writes from SIMD and floating-point registers of a 128-bit value that is 64-bit aligned in memory are treated as a pair of single-copy atomic 64-bit writes.
- Atomicity rules for SIMD load and store instructions also apply to SVE load and store instructions.
- SVE predicated load and store instructions are performed as a sequence of memory element accesses.

...

is changed to read:

For explicit memory effects generated from an Exception level the following rules apply:

...

- Reads to SIMD and floating-point registers of a single 64-bit or smaller quantity that is aligned to the size of the quantity being loaded are treated as single-copy atomic reads.
- Writes from SIMD and floating-point registers of a single 64-bit or smaller quantity that is aligned to the size of the quantity being stored are treated as single-copy atomic writes.
- Element or Structure Reads to SIMD and floating-point registers of 64-bit or smaller elements, where each element is aligned to the size of the element being loaded, have each element treated as a single-copy atomic read.
- Element or Structure Writes from SIMD and floating-point registers of 64-bit or smaller elements, where each element is aligned to the size of the element being stored, have each element treated as a single-copy atomic store.
- Reads to SIMD and floating-point registers of a 128-bit value that is 64-bit aligned in memory are treated as a pair of single-copy atomic 64-bit reads.
- Writes from SIMD and floating-point registers of a 128-bit value that is 64-bit aligned in memory are treated as a pair of single-copy atomic 64-bit writes.
- SVE predicated load and store instructions have the same single-copy atomicity guarantees as for SIMD load and store instructions for 64-bit or smaller elements.
- An SVE predicated load or store of a 128-bit element that is 64-bit aligned is treated as a pair of 64-bit single-copy atomic accesses.

- SVE unpredicated load and store instructions are performed as a sequence of byte accesses.
- SVE unpredicated load and store instructions do not guarantee that any access larger than a byte will be performed as a single-copy atomic access.

...

2.107 D24015

In section J1.1.1.34 “ExceptionReturnPPEND”, the code that reads:

```
ExceptionReturnPPEND(bits(64) spsr)
...
    if spsr<33> == '1' then                                // SPSR.PPEND
        bits(2) target_except;
        (enabled_at_source, target_except) = PMUExceptionEnabled();
        masked_at_source = PMUExceptionMasked(target_except, PSTATE.EL, PSTATE.PM);
        ...
        constant boolean masked_at_dest = PMUExceptionMasked(target_except,
target_eret, target_pm);
        ...
```

is changed to read:

```
ExceptionReturnPPEND(bits(64) spsr)
...
    if spsr<33> == '1' then                                // SPSR.PPEND
        bits(2) target_except;
        (enabled_at_source, target_except) = PMUExceptionEnabled();
        if Restarting() then
            masked_at_source = TRUE;
        else
            masked_at_source = PMUExceptionMasked(target_except, PSTATE.EL,
PSTATE.PM);
        ...
        constant boolean masked_at_dest = PMUExceptionMasked(target_except,
target_eret, target_pm);
```

In section J1.1.1.37 “PMUExceptionMasked”, the code that reads:

```
boolean PMUExceptionMasked(bits(2) target_el, bits(2) from_el, bit pm)
    assert IsFeatureImplemented(FEAT_EBEP);

    if Halted() || Restarting() then
        return TRUE;
    elsif UInt(target_el) < UInt(from_el) then
        return TRUE;
    ...
```

is changed to read:

```
boolean PMUExceptionMasked(bits(2) target_el, bits(2) from_el, bit pm)
    assert IsFeatureImplemented(FEAT_EBEP);

    if Halted() then
        return TRUE;
```

```
elseif UInt(target_el) < UInt(from_el) then
    return TRUE;
...
```

2.108 D24019

In section J1.2 “Pseudocode for AArch32 operation”, the pseudocode for function AArch32.S1TTWParamsEL2() that reads:

```
S1TTWParams AArch32.S1TTWParamsEL2 ()
    S1TTWParams walkparams;
    ...
    walkparams.irgn = HTCR.SH0;
    walkparams.orgn = HTCR.IRGNO;
    walkparams.sh   = HTCR.ORGNO;
    ...
    return walkparams;
```

is changed to read:

```
S1TTWParams AArch32.S1TTWParamsEL2 ()
    S1TTWParams walkparams;
    ...
    walkparams.irgn = HTCR.IRGNO;
    walkparams.orgn = HTCR.IRGNO;
    walkparams.sh   = HTCR.SH0;
    ...
    return walkparams;
```

2.109 D24020

In section C5.5.88 “TLBIP IPAS2E1OS, TLBIP IPAS2E1OSNXS, TLB Invalidate Pair by Intermediate Physical Address, Stage 2, EL1, Outer Shareable”, in the field description for the IPA field, the text that reads:

IPA[51:48], bits [107:64]

Bits[55:12] of the intermediate physical address to match.

is changed to read:

IPA[55:12], bits [107:64]

Bits[55:12] of the intermediate physical address to match.

Equivalent changes are made in the same field descriptions in section C5.5.91 “TLBIP IPAS2LE1OS, TLBIP IPAS2LE1OSNXS, TLB Invalidate Pair by Intermediate Physical Address, Stage 2, Last level, EL1, Outer Shareable”.

2.110 C24031

In section D18.2.3 “Counter packet”, the counter definitions that read:

0b00001	Issue latency. Cycle count from the operation being dispatched for issue to the operation being issued for execution. This counts any delay in waiting the operation being ready to issue. Included for all operations.
0b00010	Translation latency. Cycle count from a virtual address being passed to the MMU for translation to the result of the translation being available. Included for all load, store, and atomic operations.

are changed to read:

0b00001	Issue latency. The count of cycles when the operation is waiting to be issued.
0b00010	Translation latency. The count of cycles when at least one part of the operation is waiting for the MMU to complete an address translation. It is IMPLEMENTATION DEFINED whether a cycle is counted if a part of the operation is accessing memory.

Equivalent changes made in section D17.6 “The profiling data” and section D17.6.4 “Additional information for each profiled memory access operation”.

2.111 C24054

In section D8.2.12.1 “Behavior when stage 1 address translation is disabled”, under rule RPPDBS, the rule that reads:

R_{RPPDBS}

If stage 1 address translation is disabled, then all of the following apply to memory accesses that would otherwise be translated at stage 1:

- The stage 1 IA is flat mapped to the OA.
- No Translation faults, Access flag faults, or Permission faults can be generated.
- Address size faults and Alignment faults can be generated.
- ...

is changed to read:

R_{RPPDBS}

If stage 1 address translation is disabled, then all of the following apply to memory accesses that would otherwise be translated at stage 1:

- The stage 1 IA is flat mapped to the OA.
- No Translation faults, Access flag faults, or Permission faults can be generated.
- Address size faults and Alignment faults can be generated.
- An **IMPLEMENTATION DEFINED** fault for an Unsupported Exclusive or Atomic access can be generated if one of the following applies:

- The translation regime is not the EL1&0 regime.
- The translation regime is the EL1&0 regime and the Effective value of $HCR_{EL2}.DC$ is 0.
- ...

The pseudocode function `AArch64.S1DisabledOutput()` with code that reads:

```
if !IsZero(va<addrtop:pamax>) then
    fault.statuscode = Fault_AddressSize;
elseif AArch64.S1HasAlignmentFaultDueToMemType(regime, accdesc, aligned,
    walkparams.ntlsmid, memattrs) then
    fault.statuscode = Fault_Alignment;
```

is changed to read:

```
if !IsZero(va<addrtop:pamax>) then
    fault.statuscode = Fault_AddressSize;
elseif AArch64.S1HasAlignmentFaultDueToMemType(regime, accdesc, aligned,
    walkparams.ntlsmid, memattrs) then
    fault.statuscode = Fault_Alignment;
elseif ((accdesc.exclusive || accdesc.atomicop) &&
    !(regime == Regime_EL10 && EL2Enabled() && HCR_EL2.DC == '1') &&
    ConstrainUnpredictableBool(Unpredictable_Atomic_MMU_IMPDEF_FAULT)) then
    fault.statuscode = Fault_Exclusive;
```

2.112 D24057

In section J1.3.1.114 “PCSRsuspended”, the code that reads:

```
boolean PCSRSuspended()
    if IsFeatureImplemented(FEAT_PCSRv8p9) && PMPCSTL.IMP == '1' then
        return PMPCSTL.EN == '0';
    else
        return boolean IMPLEMENTATION_DEFINED "PCSR is suspended";
```

is changed to read:

```
boolean PCSRSuspended()
    if IsFeatureImplemented(FEAT_PMUv3_SS) && PMPCSTL.SS == '1' then
        return FALSE;
    elseif IsFeatureImplemented(FEAT_PCSRv8p9) && PMPCSTL.IMP == '1' then
```

2.113 C24061

In section D23.3.1 “Instructions for accessing non-debug System registers”, the Note that reads:

Note:

...

All unused encodings in the range $op0 == 0b11$, $op1 == 0b000$, $CRn == 0b0000$, $CRm == \{0b0010-0b0111\}$, $op2 == \{0b000-0b111\}$ are defined to be accessible as Reserved, **RAZ** to ensure correct behavior if the encodings are used for ID registers in the future.

is changed to read:

Direct accesses to some regions of the system register encoding space have a defined behavior even if no register is currently allocated or implemented for the encodings in those regions. The Feature ID space is the System register space in AArch64 with $op0==3$, $op1==\{0, 1, 3\}$, $CRn==0$, $CRm==\{0-7\}$, $op2==\{0-7\}$. The properties for this space are that:

- For unused encodings in the part of the space that has $op1==0$ and $CRm==\{2-7\}$, the following apply:
 - If FEAT_FGT is not implemented, it is **IMPLEMENTATION DEFINED** whether direct reads can be trapped by HCR_EL2.TID3.
 - If FEAT_FGT is implemented, direct reads can be trapped by HCR_EL2.TID3.
 - If FEAT_IDTE3 is implemented, direct reads can be trapped by SCR_EL3.TID3.
 - Direct reads that are not trapped behave as Reserved, **RES0**, to ensure correct behavior if the encodings are used for ID registers in the future.
- For unused encodings in the part of the space that has either $op1==0$ and $CRm==0$, or $op1==\{1, 3\}$, the following apply:
 - If FEAT_IDST is not implemented, direct reads are **UNDEFINED** and therefore generate exceptions reported with EC syndrome value $0x00$.
 - If FEAT_IDST is implemented, direct reads generate exceptions reported with EC syndrome value $0x18$.
 - The behavior for direct reads is consistent with the pseudocode function `UnimplementedIDRegister()`.

2.114 D24062

In section J1.3.1.99 “PMUCaptureEvent”, the code that reads:

```
PMUCaptureEvent ()
...
if IsFeatureImplemented(FEAT_PCSRv8p9) && PMPCSTL.SS == '1' then
    if pc_sample.valid && !debug_state then
        ...
    else
        SetPCSRUnknown ();
...

```

is changed to read:

```
PMUCaptureEvent ()
...
if IsFeatureImplemented(FEAT_PCSRv8p9) && PMPCSTL.SS == '1' then
    if pc_sample.valid && !debug_state then
        ...

```

```
else
    PMPCSR<31:0> = Ones(32);
...
```

2.115 C24064

In section D17.6.9 “Controlling the data that is collected”, the text that reads:

Physical data addresses are collected only if one of the following is true:

- PMSCR_EL1.PA is set to 1 and the Profiling Buffer is owned by Secure EL1, and Secure EL2 is disabled or is not implemented.
- PMSCR_EL2.PA is set to 1 and the Profiling Buffer is owned by Secure or Non-secure EL2.
- PMSCR_EL1.PA is set to 1 and PMSCR_EL2.PA is set to 1 and either the Profiling Buffer is owned by Non-secure EL1, or the Profiling Buffer is owned by Secure EL1 and Secure EL2 is implemented and enabled.

If EL2 is not implemented or is disabled for the current Security state, the PE behaves as if PMSCR_EL2.PA is set to 1, other than for a direct read of the register.

is changed to read:

Physical data addresses are collected only if one of the following is true:

- The Effective value of PMSCR_EL1.PA is 1, the Effective value of PMSCR_EL2.PA is 1, and the Profiling Buffer is owned by EL1.
- The Effective value of PMSCR_EL2.PA is 1 and the Profiling Buffer is owned by EL2.

If EL2 is not implemented or is disabled for the owning Security state, then the Effective value of PMSCR_EL2.PA is 1.

A similar update is made in section D24.7.6 “PMSR_EL2, Statistical Profiling Control Register (EL2)”.

2.116 C24066

In section D14.3.1 “Common architectural events”, in event “0x8000, SIMD_INST_RETIRED, Instruction architecturally executed, SIMD”, the text that reads:

0x8000, SIMD_INST_RETIRED, Instruction architecturally executed, SIMD

The counter counts each architecturally executed SIMD instruction.

It is **IMPLEMENTATION DEFINED** which architecturally executed SIMD instructions are counted in AArch32 state.

When Armv9.5 is not implemented, it is **IMPLEMENTATION DEFINED** whether scalar loads and stores to SIMD&FP registers other than those listed above are counted.

When Armv9.5 is implemented, other scalar loads and stores to SIMD&FP registers are not counted.

is changed to read:

0x8000, SIMD_INST_RETIRED, Instruction architecturally executed, SIMD

The counter counts each architecturally executed instruction counted by INST_RETIRED that is a SIMD instruction.

That is, the counter counts:

- Advanced SIMD data-processing instructions. Advanced SIMD scalar data-processing instructions are not counted.
- SVE and SME SIMD data-processing instructions. Non-SIMD SVE and SME data-processing instructions are not counted.
- Advanced SIMD structure load/store of one or more SIMD&FP registers instructions.
- Advanced SIMD load and replicate to one or more SIMD&FP registers instructions.
- Scalar load/store of a SIMD&FP Q register or pair of Q registers instructions.
- SVE and SME load/store instructions.

It is **IMPLEMENTATION DEFINED** which architecturally executed SIMD instructions are counted in AArch32 state.

When Armv9.5 is not implemented, it is **IMPLEMENTATION DEFINED** whether scalar loads and stores to SIMD&FP registers other than those listed above are counted.

When Armv9.5 is implemented, other scalar loads and stores to SIMD&FP registers are not counted.

2.117 D24086

In section D24.7.11 “PMSIDR_EL1, Sampling Profiling ID Register”, in the description of the field FDS, the text that reads:

From Armv8.9, if FEAT_SPE is implemented, the value 0 is not permitted.

is changed to read:

From Armv8.9, if FEAT_SPE_LDS is implemented, the value 0 is not permitted.

2.118 D24093

FEAT_TME is withdrawn from all future versions of Arm® Architecture Reference Manual for A-profile architecture.

2.119 C24096

In section D8.15.1.6 TLB conflict abort, the rule R_{QWTKL} that reads:

R_{QWTKL}

When a TLB conflict abort is generated, it is **IMPLEMENTATION DEFINED** whether it is a stage 1 abort or a stage 2 abort.

is changed to read:

R_{QWTKL}

For the EL1&0 regime, when a TLB conflict abort is generated and both stages of translation are enabled, then one of the following applies:

- If FEAT_BBML1 is implemented and the TLB conflict abort is generated due to changing the table or block size or Contiguous bit, then the abort is reported to EL2.
- Otherwise, it is **IMPLEMENTATION DEFINED** whether it is reported to EL1 or EL2.

The rule in section D8.17.2 Support levels for changing table or block size, R_{FWRMB} , is removed:

R_{FWRMB}

If all of the following apply, then a TLB conflict abort is reported to EL2:

- Level 1 or level 2 is supported.
- Stage 2 translations are enabled in the current translation regime.
- A TLB conflict abort is generated due to changing the block size or Contiguous bit.

Additionally, the following information statement in the same section, I_{GQTBj} , is removed:

I_{GQTBj}

When a TLB conflict abort is generated, it is **IMPLEMENTATION DEFINED** whether it is a stage 1 abort or a stage 2 abort.

For rule R_{ZMVZT} the text that reads:

R_{ZMVZT}

If EL2 is enabled and stage 2 of the EL1&0 translation regime is disabled, then a stage 2 abort cannot be generated.

is changed to read:

R_{ZMVZT}

A TLB Conflict abort cannot be generated for a disabled stage of translation.

2.120 D24098

In section D24.2.185 “TFSR_EL1, Tag Fault Status Register (EL1)”, the text that reads:

This register is present only when FEAT_MTE2 is implemented. Otherwise, direct accesses to TFSR_EL1 are **UNDEFINED**.

is changed to read:

This register is present only when FEAT_MTE_ASYNC is implemented. Otherwise, direct accesses to TFSR_EL1 are **UNDEFINED**.

Equivalent changes are made in the following registers:

- D24.2.186 “TFSR_EL2, Tag Fault Status Register (EL2)”.
- D24.2.187 “TFSR_EL3, Tag Fault Status Register (EL3)”.
- D24.2.188 “TFSRE0_EL1, Tag Fault Status Register (ELO)”.

2.121 D24099

In section J1.3.5 “shared/translation”, in the pseudocode function `GPCRegistersConsistent()`, the following code segment that reads:

```
boolean GPCRegistersConsistent()
    if IsFeatureImplemented(FEAT_RME_GPC3) then
        ...
        if GPCCR_EL3.GPCBW == '1' then
            ...
            if !GPCBW_EL3BWSTRIDEValid() then
                return FALSE;

            if !IsAligned(GPCBW_EL3.BWADDR, 1 << UInt(GPCBW_EL3.BWSIZE)) then
                return FALSE;
            else
                ...
```

is changed to read:

```
boolean GPCRegistersConsistent()
    if IsFeatureImplemented(FEAT_RME_GPC3) then
        ...
        if GPCCR_EL3.GPCBW == '1' then
            ...
            if !GPCBW_EL3BWSTRIDEValid() then
```

```

        return FALSE;

        constant integer bwstride = 1 << (40 + UInt(GPCBW_EL3.BWSTRIDE));
        constant integer bwaddr = UInt(GPCBW_EL3.BWADDR) << 30;
        if bwstride <= bwaddr then
            return FALSE;

        if !IsAligned(GPCBW_EL3.BWADDR, 1 << UInt(GPCBW_EL3.BWSIZE)) then
            return FALSE;
        else
            ...
    
```

2.122 D24101

In section D24.2.40 “ESR_EL1, Exception Syndrome Register (EL1)”, the description of the ISS2.GCS field that reads:

GCS	Meaning
0b0	The Data Abort is not due to a Guarded control stack data access.
0b1	The Data Abort is due to a Guarded control stack data access.

is changed to read:

GCS	Meaning
0b0	The Data Abort is not due to a Guarded control stack data access.
0b1	The Data Abort is due to a Guarded control stack data access. The ISV field is 0.

A similar update is made in sections D24.2.41 “ESR_EL2, Exception Syndrome Register (EL2)” and D24.2.42 “ESR_EL3, Exception Syndrome Register (EL3)”.

2.123 D24112

In section J1.1.2 “aarch64/exceptions”, in the pseudocode function AArch64.TagCheckFault(), the code that reads:

```

AArch64.TagCheckFault(bits(64) vaddress, AccessDescriptor accdesc)
    ...
    case tcftype of
        when TCFType_Sync
            ...
            fault.write = accdesc.write;
            ...
    
```

is changed to read:

```

AArch64.TagCheckFault(bits(64) vaddress, AccessDescriptor accdesc, boolean
write_fault)
    ...
    case tcftype of
        when TCFType_Sync
    
```

```
...
fault.write = write_fault;
...
```

In section J1.1.3 “aarch64/functions”, in the pseudocode function MemAtomic(), the code that reads:

```
bits(size) MemAtomic(...)
...
// For Store-only Tag checking, the tag check is performed on the store.
if (IsFeatureImplemented(FEAT_MTE2) && accdesc.tagchecked &&
... ) then
    if !AArch64.CheckTag(...) then
        accdesc.write = FALSE; // Tag Check Fault on a read
        AArch64.TagCheckFault(address, accdesc);
...
if IsFeatureImplemented(FEAT_MTE_STORE_ONLY) &&
StoreOnlyTagCheckingEnabled(accdesc.el) then
    ...
    if IsFeatureImplemented(FEAT_MTE2) && accdesc.tagchecked then
        ...
        if !AArch64.CheckTag(...) then
            AArch64.TagCheckFault(address, accdesc);
```

is changed to read:

```
bits(size) MemAtomic(...)
...
// For Store-only Tag checking, the tag check is performed on the store.
if (IsFeatureImplemented(FEAT_MTE2) && accdesc.tagchecked &&
... ) then
    if !AArch64.CheckTag(...) then
        // For a synchronous Tag Check Fault, do not set WnR.
        constant boolean write_fault = FALSE;
        AArch64.TagCheckFault(address, accdesc, write_fault);
...
if IsFeatureImplemented(FEAT_MTE_STORE_ONLY) &&
StoreOnlyTagCheckingEnabled(accdesc.el) then
    ...
    if IsFeatureImplemented(FEAT_MTE2) && accdesc.tagchecked then
        ...
        if !AArch64.CheckTag(...) then
            AArch64.TagCheckFault(address, accdesc, accdesc.write);
```

An equivalent change is made in MemAtomicRCW().

2.124 C24124

In section D20.4.2 PE error state classification, under the rule R_{LFXRD} , the text that reads:

R_{LFXRD}

If and only if all of the following are true, then on taking an Error exception the PE error state is recorded as Corrected (CE):

- The error has been corrected and not silently propagated by the PE.

- The Error exception is taken as an SError exception taken to AArch64 state.
- Software can recover execution from the preferred return address of the exception. Because the error has been corrected, software does not need to take action to locate and repair the error.
- The implementation has not elected to record the PE error state as any other type.

is changed to read:

R_{LFXRD}

If and only if all of the following are true, then on taking an Error exception the PE error state is recorded as Corrected (CE):

- The error has been corrected and not silently propagated by the PE.
- The Error exception is taken as an SError exception taken to AArch64 state.
- Software can recover execution from the preferred return address of the exception. Because the error has been corrected, software does not need to take action to locate and repair the error.
- The implementation has not elected to record the PE error state as any other type.

Arm strongly recommends that a corrected error that is not silently propagated does not generate any error exception at the PE. The producer of the data that corrected the error might still record the error in a RAS error record, and generate a fault handling interrupt, as described in the Arm® Reliability, Availability, and Serviceability (RAS) System Architecture for A-profile architecture.

Additionally, the text that reads:

I_{VKMZB}

The PE error states are summarized by Figure D20-1. Figure D20-1 assumes the type of Error exception supports the resulting PE error state, never elects to record an error as a different PE error state when permitted, and does not show Uncategorized error or **IMPLEMENTATION DEFINED** syndrome.

is changed to read:

I_{VKMZB}

The PE error states for an Error exception generated when the PE consumes an uncorrected error are summarized by Figure D20-1. Figure D20-1 assumes the type of Error exception supports the resulting PE error state, that the error was neither corrected nor deferred, and the PE never elects to record an error as a different PE error state when permitted, an Uncategorized error syndrome, or an **IMPLEMENTATION DEFINED** syndrome.

2.125 D24129

In section B1.5.6 “About PSTATE.DIT” the following text is deleted:

- If FEAT_SVE2 and FEAT_SME are not implemented, the data independent timing control introduced by FEAT_DIT does not affect the timing properties of SVE instructions.
- If FEAT_SVE2 or FEAT_SME is implemented, the data independent timing control introduced by FEAT_DIT affects the timing properties of all SVE and SME instructions that honor PSTATE.DIT.

In section C8.2 “Alphabetical list of SVE instructions”, in the “Operational information” subsection of SVE instructions which were defined prior to FEAT_SVE2, but which are required to honor PSTATE.DIT when FEAT_SVE2 or FEAT_SME is implemented, the text that reads:

If FEAT_SVE2 is implemented or FEAT_SME is implemented, then if PSTATE.DIT is 1:

...

is changed to:

If PSTATE.DIT is 1:

...

2.126 D24146

In section D5.5.2 Common microarchitectural events, the text that describes the PMU event 0x8056, `SVE_SPEC`, Operation speculatively executed, SVE data processing that reads:

The counter counts each operation counted by INST_SPEC that is an SVE data-processing operation.

- Data-processing operations involving SVE scalable vector and predicate registers. This includes operations added by FEAT_SME which involve the SVE registers but do not involve any ZA or ZT registers. This does not include non-SIMD SVE instructions.
- If the Cryptographic Extension and FEAT_SVE2 are implemented, the PMULLB, PMULLT (Q variants) instructions.

This includes all operations that operate on the SVE registers, except those that are counted as one of the following:

- Load or store operations.
- Cryptographic data-processing operations other than those included above.
- When FEAT_SME is implemented, SME data-processing operations.

is changed to read:

The counter counts each operation counted by SE_SPEC that is an SVE data-processing operation.

All operations that operate on the SVE scalable vector and predicate registers are counted as SVE data-processing operations, other than all of the following:

- Load or store operations.
- If the Cryptographic Extension is implemented, operations counted by CRYPTO_SPEC as Cryptographic data-processing operations.
- When FEAT_SME is implemented, operations counted by SME_SPEC as SME data-processing operations.
- Non-SIMD SVE instructions.

If the Cryptographic Extension and FEAT_SVE2 are implemented, operations due to the PMULLB, PMULLT (Q variants) instructions are not counted by CRYPTO_SPEC and are counted as SVE data-processing operations. Operations due to instructions defined by FEAT_SME which involve the SVE registers but do not involve any ZA or ZT registers are counted as SVE data-processing operations.

Additionally, the text that describes the PMU event 0x835C, SME_SPEC, Operation speculatively executed, SME data processing that reads:

The counter counts each operation counted by SE_SPEC that is an SME data-processing operation.

- Data-processing operations involving the ZA and ZT registers.

This includes all operations that operate on the ZA or ZT registers, except those that are counted as load and store operations. Operations due to instructions added by FEAT_SME which involve the SVE registers but do not involve any ZA or ZT registers are counted as SVE data-processing operations. This does not include non-SIMD SME operations.

is changed to read:

The counter counts each operation counted by SE_SPEC that is an SME data-processing operation.

All operations that operate on the ZA or ZT registers are counted as SME data-processing operations, other than all of the following:

- Load and store operations.
- Non-SIMD SME operations.

Operations due to instructions defined by FEAT_SME which involve the SVE registers but do not involve any ZA or ZT registers are counted as SVE data-processing operations.

2.127 C24218

In section C1.2.5 “SVE Condition code aliases” the text that reads:

The SVE assembler syntax defines an alternative set of SVE condition code aliases for use with AArch64 conditional instructions, as follows:

is changed to read:

The SVE and SME assembler syntax defines an alternative set of condition code aliases for use with AArch64 conditional instructions, as follows:

2.128 D24149

In section D1.3.4.2 “Illegal exception returns from AArch64 state”, the text that reads:

R_{VVJHB}

On an illegal exception return from an Exception level, ELx, all of the following occur:

...

- All of the following are **UNKNOWN**:
 - If FEAT_UAO is implemented, then PSTATE.UAO.
 - If FEAT_DIT is implemented, then PSTATE.DIT.
 - If FEAT_MTE is implemented, then PSTATE.TCO.
 - If FEAT_SSBS is implemented, then PSTATE.SSBS.
 - If FEAT_PAuth_LR is implemented, then PSTATE.PACM.

...

is changed to read:

R_{VVJHB}

On an illegal exception return from an Exception level, ELx, all of the following occur:

...

- All of the following are **UNKNOWN**:
 - If FEAT_UAO is implemented, then PSTATE.UAO.
 - If FEAT_DIT is implemented, then PSTATE.DIT.
 - If FEAT_MTE is implemented, then PSTATE.TCO.
 - If FEAT_SSBS is implemented, then PSTATE.SSBS.
 - If FEAT_BTI is implemented, then PSTATE.BTYPE.

- If FEAT_PAAuth_LR is implemented, then PSTATE.PACM.

...

2.129 C24151

In section K14.6.1.2.3 “Resolving by the use of barriers and address dependency”, the following text is deleted:

Where the value returned by a read is used for computation of the virtual address of a subsequent read or write, then these two memory accesses are observed in program order.

Where the value returned by a read is used for computation of the virtual address of a subsequent read or write, this is called an address dependency. An address dependency exists even if the value returned by the first read has no effect on the virtual address. This might occur if the value returned is masked off before it is used, or if it confirms a predicted address value that it might have changed.

This restriction applies only when the data value returned by a read is used as a data value to calculate the address of a subsequent read or write. It does not apply if the data value returned by a read determines the condition flags values, and the values of the flags are used for condition code evaluation to determine the address of a subsequent read, either through conditional execution or the evaluation of a branch. This is called a control dependency.

Where both a control and address dependency exist, the ordering behavior is consistent with the address dependency.

2.130 D24153

In section K7.3 “Recommended authentication interface”, the text that reads:

Arm recommends that any Trace extension has the same authentication interface as the PE it is connected to.

is changed to read:

Any Trace extension has the same authentication interface as the PE it is connected to.

2.131 D24156

In section H9.2.7 “DBGDTRTX_EL0, Debug Data Transfer Register, Transmit” the text that reads:

When DoubleLockStatus(), or !IsCorePowered(), or OSLockStatus(), or !AllowExternalDebugAccess(), accesses to this register return an ERROR.

is changed to read:

When `DoubleLockStatus()`, or `!IsCorePowered()`, or `OSLockStatus()`, accesses to this register return an ERROR.

In section J1.3.1 shared/debug, the pseudocode that reads:

```
// AllowExternalDebugAccess()
// =====
// Returns TRUE if an external debug interface access to the External debug
// registers
// is allowed for the given Security state, FALSE otherwise.
```

is changed to read:

```
// AllowExternalDebugAccess()
// =====
// Returns TRUE if an external debug interface access to the External
// DBGBVR<n>_EL1,
// DBGBCR<n>_EL1, DBGWVR<n>_EL1, DBGWCR<n>_EL1 registers, and, from Armv8.2, the
// OSLAR_EL1 register is allowed for the access. Returns FALSE otherwise.
```

2.132 C24161

In section C8.2.370 “LDFF1B (scalar plus scalar)” the code that reads:

```
// MemNF[] will return fault=TRUE if access is not performed for any reason
(data, fault) = MemNF[addr, mbytes, accdesc];
```

is changed to read:

```
// MemNF[] will return fault=TRUE if access is not performed for any reason
(data, fault) = MemNF[addr, mbytes, accdesc];
faulted = faulted || ConstrainUnpredictableBool(Unpredictable_NONFAULT);
```

Equivalent changes are made to all LDFF* and all LDNF* instructions.

2.133 D24162

In section D8.1.2.3 “Realm EL1&0 translation regime” the rule $R_{\text{BT}\text{V}\text{G}}$ with text that reads:

If all of the following apply, then the Realm EL1&0 translation regime is used to translate addresses:

- The memory accesses occur in Realm state.
- Memory is accessed from EL1 or EL0.
- The Effective value of HCR_EL2.TGE is 0.

is changed to read:

If all of the following apply, then the Realm EL1&O translation regime is used to translate addresses:

- The memory accesses occur in Realm state.
- One of the following applies:
 - Memory is accessed from EL1.
 - Memory is accessed from EL0 and the Effective value of HCR_EL2.{E2H,TGE} is not {1,1}.

Equivalent changes are made to:

- R_{QJSCR} (for the Non-secure EL1&O translation regime).
- R_{DKYXN} (for the Secure EL1&O translation regime).

2.134 D24166

In section J1.1 “Pseudocode for AArch64 operation”, in the function AArch64.S1ComputePermissions(), the code that reads:

```
if slperms.overlay && slperms.wxn == '1' && slperms.ox == '1' then
    slperms.ow = '0';
elseif slperms.wxn == '1' then
    slperms.x = '0';
```

is changed to read:

```
if slperms.overlay then
    // If WXN and the overlay X permission is present, the overlay W permission is
    removed.
    slperms.ow = slperms.ow AND NOT(slperms.wxn AND slperms.ox);
else
    // If WXN and the W and X permissions are present, the X permission is removed.
    // The W and X permissions are factored into the WXN computation.
    slperms.x = slperms.x AND NOT(slperms.wxn);
```

2.135 R24169

In section D24.3.18 “MDCR_EL3, Monitor Debug Configuration Register (EL3)”, the text in field ‘PMEE, bits [41:40]’ that reads:

PMEE	Meaning
0b00	The PMUIRQ signal is asserted on a PMU overflow, and the PMU Profiling exception is disabled.
0b01	The PMUIRQ signal and the PMU Profiling exception are both controlled by MDCR_EL2.PMEE.
0b10	The PMUIRQ signal is deasserted, and the PMU Profiling exception is disabled.
0b11	The PMUIRQ signal is deasserted, and the PMU Profiling exception is enabled.

is changed to read:

PMEE	Meaning
0b00	The PMUIRQ signal is asserted on a PMU overflow, and the PMU Profiling exception is disabled.
0b01	The PMUIRQ signal and the PMU Profiling exception are both controlled by MDCR_EL2.PMEE.
0b11	The PMUIRQ signal is deasserted, and the PMU Profiling exception is enabled.

■ All other values are reserved.

Equivalent changes are made in sections:

- D24.3.17 “MDCR_EL2, Monitor Debug Configuration Register (EL2)”.
- D24.5.9 “PMECR_EL1, Performance Monitors Extended Control Register (EL1)”.

2.136 C24177

In section B1.4.12.2 “ZA multi-slice operands”, the following text is added:

■ In a ZA multi-slice operand that consists of:

- Two ZA tile slices, the lowest numbered slice is a multiple of 2.
- Four ZA tile slices, the lowest numbered slice is a multiple of 4.

Additionally, the text under the declaration statement D_{JFDSB} that reads:

■ In instructions operating on ZA multi-slice operands the lowest-numbered slice is a multiple of 2 for a two-slice ZA operand, and a multiple of 4 for a four-slice ZA operand, selected by the sum of a 32-bit general-purpose register (slice index register Ws) and an offset (slice index offset $offs$).

is changed to read:

■ In instructions operating on ZA multi-slice operands, the lowest-numbered slice is selected by rounding down the sum of a 32-bit general-purpose register (slice index register Ws) and an offset (slice index offset $offs$) to the nearest lower multiple of the number of the slices in the ZA multi-slice operand, modulo the number of slices in the named tile.

In section C9.2.148 “MOV (tile to vector, two registers)”, the text that reads:

■ The consecutive slice numbers within the tile are selected starting from the sum of the slice index register and immediate offset, modulo the number of such elements in a vector.

is changed to read:

■ The first slice of the two-slice operand is selected by rounding down the sum of the slice index register and the immediate offset to the nearest lower multiple of two, modulo the number of slices in the tile.

Equivalent changes are made to the following instruction descriptions:

- C9.2.146 “MOV (tile to vector, four registers)”.
- C9.2.151 “MOV (vector to tile, four registers)”.
- C9.2.153 “MOV (vector to tile, two registers)”.
- C9.2.156 “MOVA (tile to vector, four registers)”.
- C9.2.158 “MOVA (tile to vector, two registers)”.
- C9.2.161 “MOVA (vector to tile, four registers)”.
- C9.2.163 “MOVA (vector to tile, two registers)”.
- C9.2.166 “MOVAZ (tile to vector, four registers)”.
- C9.2.168 “MOVAZ (tile to vector, two registers)”.

2.137 D24180

In section D1.3.6.4.2 “Virtual interrupt masking”, the text which reads:

B When the interrupt is pending, it might be subject to masking, as defined in R_{SNLJH} and R_{XSPSG} . If the interrupt is masked, it is not taken. If the interrupt is not masked, it is taken.

is changed to read:

B When the interrupt is pending, it might be subject to masking, as defined in R_{SNLJH} , R_{XSPSG} , and R_{MMYBS} . If the interrupt is masked, it is not taken. If the interrupt is not masked, it is taken.

In the same section, the text that reads:

R_{SNLJH}

If the target Exception level of a virtual IRQ interrupt is the current Exception level, ELx, the following controls determine whether the interrupt is masked:

PSTATE.I	SCTLR_ELx.NMI	AllIntMask	vIRQ	vIRQ with Superpriority
0	0	x	Not Masked	Not Masked
0	1	0	Not Masked	Not Masked
0	1	1	Masked	Masked
1	0	x	Masked	Masked
1	1	0	Masked	Not Masked
1	1	1	Masked	Masked

R_{XSPSG}

If the target Exception level of a virtual FIQ interrupt is the current Exception level, ELx, the following controls determine whether the interrupt is masked:

PSTATE.F	SCTLR_ELx.NMI	AllIntMask	vFIQ	vFIQ with Superpriority
0	0	x	Not Masked	Not Masked
0	1	0	Not Masked	Not Masked
0	1	1	Masked	Masked
1	0	x	Masked	Masked
1	1	0	Masked	Not Masked
1	1	1	Masked	Masked

R_{MMYBS}

If the target Exception level of a virtual SError exception is the current Exception level, ELx, the PSTATE.A control determines whether the interrupt is masked. However, if all of the following are true, the PSTATE.A control is ignored, and the interrupt is taken:

- The target Exception level is the current Exception level.
- FEAT_DoubleFault is implemented.
- SCR_EL3.NMEA is 1.

is changed to read:

R_{SNUH}

The following controls determine whether a virtual IRQ interrupt is masked:

PSTATE.I	SCTLR_EL1.NMI	PSTATE.EL	AllIntMask	vIRQ	vIRQ with Superpriority
0	0	xx	x	Not Masked	Not Masked
0	1	ELO	x	Not Masked	Not Masked
0	1	EL1	0	Not Masked	Not Masked
0	1	EL1	1	Masked	Masked
1	0	xx	x	Masked	Masked
1	1	ELO	x	Masked	Not Masked
1	1	EL1	0	Masked	Not Masked
1	1	EL1	1	Masked	Masked

R_{XSPSG}

The following controls determine whether a virtual FIQ interrupt is masked:

PSTATE.F	SCTLR_EL1.NMI	PSTATE.EL	AllIntMask	vFIQ	vFIQ with Superpriority
0	0	xx	x	Not Masked	Not Masked
0	1	ELO	x	Not Masked	Not Masked
0	1	EL1	0	Not Masked	Not Masked
0	1	EL1	1	Masked	Masked
1	0	xx	x	Masked	Masked
1	1	ELO	x	Masked	Not Masked

PSTATE.F	SCTLR_EL1.NMI	PSTATE.EL	AllIntMask	vFIQ	vFIQ with Superpriority
1	1	EL1	0	Masked	Not Masked
1	1	EL1	1	Masked	Masked

R_{MMYBS}

The following controls determine whether a virtual SError exception is masked:

PSTATE.A	PSTATE.EL	vSError
0	xx	Not Masked
1	xx	Masked

Additionally, and info statement is added to the same section:

I_{x0001}

The ability to execute at EL0 with virtual interrupts taken to EL1 masked is required by some user level driver code.

In section D1.3.6.4.3 “Delegated SError exception masking”, the text that reads:

B When the exception is pending, it might be subject to masking. If the exception is masked, it is not taken. If the exception is not masked, it is taken.

is changed to read:

B When the exception is pending, it might be subject to masking. If PSTATE.A is 1, then the exception is masked and is not taken. If PSTATE.A is 0, then the exception is not masked and is taken.

2.138 C24186

In section D24.2.164 “SCTLR2_EL1, System Control Register (EL1)”, the text that reads:

NMEA, bit [2] When FEAT_DoubleFault2 is implemented: Non-maskable External Aborts.
Controls whether PSTATE.A masks SError exceptions at EL1.

NMEA	Meaning
0b0	SError exceptions are not taken at EL1 if PSTATE.A == 1, unless routed to a higher Exception level.
0b1	SError exceptions are taken at EL1 regardless of the value of PSTATE.A, unless routed to a higher Exception level.

is changed to read:

NMEA, bit [2] When FEAT_DoubleFault2 is implemented: Non-maskable External Aborts.
Controls whether PSTATE.A masks physical SError exceptions at EL1.

NMEA	Meaning
0b0	Physical SError exceptions are not taken at EL1 if PSTATE.A == 1, unless routed to a higher Exception level.
0b1	Physical SError exceptions are taken at EL1 regardless of the value of PSTATE.A, unless routed to a higher Exception level.

This field has no effect on the masking of virtual and delegated SError interrupts.

Equivalent changes are made in section D24.2.165 “SCTLR2_EL2, System Control Register (EL2)”, but without mention of virtual SError interrupts.

2.139 C24190

In section A1.5.9.9 “Combinations of floating-point exceptions” the following text is added:

If FPCR.AH is 1 and the result of a floating-point operation is a denormalized number that is flushed to zero, both an Inexact floating-point exception and an untrapped Underflow floating-point exception are generated, which sets the cumulative Underflow exception flag FPSR.UFC to 1 regardless of whether or not the Inexact exception is trapped.

2.140 C24218

In section J1.1.1 “aarch/debug”, in the pseudocode function SPECCollectRecord(), the code that reads:

```
constant boolean is_lat = (total_latency < UInt(PMSLATFR_EL1.MINLAT));
```

is changed to read:

```
constant boolean is_lat = (total_latency >= UInt(PMSLATFR_EL1.MINLAT));
```

2.141 D24224

In section D4.5.12.4 “Return stack address matching”, the following text is added:

When a Branch with Link instruction is speculatively executed and generates an E Atom element, if that Atom element is subsequently cancelled by a Cancel element then any corresponding entry on the trace unit return stack must be popped from the trace unit return stack.

Note: Discarding the entire contents of the trace unit return stack meets this obligation.

2.142 C24236

In section D8.7.1.3 “Implementation options when the Contiguous bit is misprogrammed”, the text that reads:

R_{JHQPP}

An implementation is permitted to ignore the Contiguous bit to ensure that OA range checking is done.

is changed to read:

I_{JHQPP}

An implementation is permitted to ignore the Contiguous bit and treat it as 0 to ensure that OA range checking is done.

This text is then moved to section D8.7.1.2 “Architectural guarantees when the Contiguous bit is misprogrammed” to sit alongside the existing Rule R_{XFDVL}.

In section J1.1.4.65 “AArch64.ContiguousBit”, the code that reads:

```
bit AArch64.ContiguousBit(TGx tgx, bit d128, integer level, bits(N) descriptor)
    if d128 == '1' then
        if (tgx == TGx_64KB && level == 1) || (tgx == TGx_4KB && level == 0) then
            return '0'; // RESO
        else
            return descriptor<111>;
        // When using TGx 64KB and FEAT_LPA is implemented,
        // the Contiguous bit is RESO for Block descriptors at level 1

    if tgx == TGx_64KB && level == 1 then
        return '0'; // RESO

    // When the effective value of TCR_ELx.DS is '1',
    // the Contiguous bit is RESO for all the following:
    //     * For TGx 4KB, Block descriptors at level 0
    //     * For TGx 16KB, Block descriptors at level 1

    if tgx == TGx_16KB && level == 1 then
        return '0'; // RESO

    if tgx == TGx_4KB && level == 0 then
        return '0'; // RESO

    return descriptor<52>;
```

is changed to read:

```
bit AArch64.ContiguousBit(TGx tgx, bit d128, integer level, bits(N) descriptor)

    if boolean IMPLEMENTATION_DEFINED "Treat Contiguous bit as 0" then
        return '0';

    if d128 == '1' then
        if (tgx == TGx_64KB && level == 1) || (tgx == TGx_4KB && level == 0) then
            return '0'; // RESO
        else
```

```

        return descriptor<111>;

// When using TGx 64KB then the Contiguous bit is RES0 for Block
// descriptors at level 1. For VMSAv8-64, level 1 Block descriptors
// can exist only if FEAT_LPA is implemented.

if tgx == TGx_64KB && level == 1 then
    return '0'; // RES0

// When the Effective value of TCR_ELx.DS is '1',
// the Contiguous bit is RES0 for all the following:
//     * For TGx 4KB, Block descriptors at level 0
//     * For TGx 16KB, Block descriptors at level 1
// For VMSAv8-64, the above Block descriptors at the specified
// levels can exist only if the Effective value of TCR_ELx.DS is 1.

if tgx == TGx_16KB && level == 1 then
    return '0'; // RES0

if tgx == TGx_4KB && level == 0 then
    return '0'; // RES0

return descriptor<52>;

```

2.143 C24254

In section D8.17 “TLB maintenance”, the text that reads:

R_GNJPZ

When a TLB maintenance instruction applies to both stage 1 and stage 2 entries, all of the following apply:

- The maintenance applies to any entries in caching structures that include stage 1 information used to translate the address or range of addresses being invalidated.
- The maintenance applies to any entries in caching structures that include stage 2 information used to translate the address or range of addresses being invalidated.
- The maintenance applies to structures combining stage 1 and stage 2 information used to translate the address or range of addresses being invalidated.

is changed to read:

R_GNJPZ

When a TLB maintenance instruction applies to both stage 1 and stage 2 entries, all of the following apply:

- The maintenance applies to any entries in caching structures that include stage 1 information.
- The maintenance applies to any entries in caching structures that include stage 2 information.
- The maintenance applies to structures combining stage 1 and stage 2 information.

2.144 D24257

In section D4.4.1.4 “Meaning of Atom elements”, the following entries are added to the table in Rule R_{HGDNB}:

Instruction	E	N
RETAASPPC	The branch was taken and the PE returned from the subroutine	Reserved
RETAASPPCR	The branch was taken and the PE returned from the subroutine	Reserved
RETABSPPC	The branch was taken and the PE returned from the subroutine	Reserved
RETABSPPCR	The branch was taken and the PE returned from the subroutine	Reserved

2.145 R24624

In section D24.2.41 “ESR_EL2, Exception Syndrome Register (EL2)”, in the description for the ISS encoding for an exception from a Data Abort, the description of the AR has the following text added:

For a load-acquire instruction that does not have acquire semantics as the result of the destination register being ZR, it is **IMPLEMENTATION SPECIFIC** whether this field is reported as 0 or 1.

Equivalent changes are made to ESR_EL3 and ESR_EL1.

In section D24.2.167 “SCTLR_EL1, System Control Register (EL1)”, the description of the nAA bit has the following text added:

For a load-acquire instruction that does not have acquire semantics as the result of the destination register, or registers, being ZR, it is **IMPLEMENTATION SPECIFIC** whether this field behaves as 1 or the programmed value.

Equivalent changes are made to SCTLR_EL2 and SCTLR_EL3.

2.146 D24274

In section D24.9.7 “ERXGSR_EL1, Selected Error Record Group Status Register”, the text that reads:

If ERRIDR_EL1.NUM is 0x0000 or ERRSELR_EL1.SEL[15:6] is greater than or equal to ERRIDR_EL1.NUM, then one of the following occurs:

- An **UNKNOWN** group of error records are selected.
- ERXGSR_EL1 is **RAZ**.
- Direct reads of ERXGSR_EL1 are NOPs.
- Direct reads of ERXGSR_EL1 are **UNDEFINED**.

is changed to read:

If `ERRIDR_EL1.NUM` is `0x0000` or `ERRSELR_EL1.SEL` is greater than or equal to `ERRIDR_EL1.NUM`, then one of the following occurs:

- An **UNKNOWN** group of error records are selected.

In section D24.9.3 “`ERRSELR_EL1`, Error Record Select Register”, the text that reads:

If `ERRSELR_EL1.SEL` is greater than or equal to `ERRIDR_EL1.NUM`, then all of the following apply:

- The value read back from `ERRSELR_EL1.SEL` is **UNKNOWN**.
- One of the following occurs:
 - An **UNKNOWN** error record is selected.
 - The `ERX*_EL1` registers are **RAZ/WI**.
 - `ERX*_EL1` register reads and writes are NOPs.
 - `ERX*_EL1` register reads and writes are **UNDEFINED**.

is changed to read:

If `ERRSELR_EL1.SEL` is greater than or equal to `ERRIDR_EL1.NUM`, then all of the following apply:

- The value read back from `ERRSELR_EL1.SEL` is **UNKNOWN**.
- Unless the access to an `ERX*_EL1` register generates a higher priority exception, one of the following occurs:
 - An **UNKNOWN** error record is selected.
 - The `ERX*_EL1` registers are **RAZ/WI**.
 - `ERX*_EL1` register reads and writes are NOPs.
 - `ERX*_EL1` register reads and writes are **UNDEFINED**.

2.147 D24282

In section D8.5.2.3 “Hardware dirty state tracking structure”, the text that reads:

`RSTJMN`

If the HDBSS is full and enabled, or `HDBSSPROD_EL2.FSC` is set to any value other than `0b000000`, then all of the following apply:

[...]

is changed to read:

`RSTJMN`

If the HDBSS is enabled, and either the HDBSS is full or HDBSSPROD_EL2.FSC is set to any value other than 0b000000, then all of the following apply:

[...]

2.148 D24321

In section A1.5.10 “Alternate BFloat16 behaviors”, the text that reads:

The behaviors in this section apply to the following instructions:

- Advanced SIMD&FP: BFCVT, BFCVTN, BFCVTN2, BFMLALB, BFMLALT.
- SVE: BFCVT, BFCVTNT, BFMLALB, BFMLALT, BFMLSLB, BFMLSLT.
- SME: BFCVT, BFCVTN, BFMLAL, BFMLSL.

is changed to read:

The behaviors in this section apply to the following instructions:

- Advanced SIMD&FP: BFCVT, BFCVTN, BFCVTN2, BFMLALB, BFMLALT.
- SVE: BFCVT, BFCVTNT, BFMLALB, BFMLALT, BFMLSLB, BFMLSLT.
- SME: BFCVT, BFCVTN.

2.149 C24337

In section D14.1.5 “Counting exceptions taken locally or not taken locally”, the following entries are added to Table D14-1:

ESR.EC	Description		Taken locally	Not Taken locally
0x2D	GCS Exception	GCS Data Check exception:	0x0084, EXC_DABORT	0x008C, EXC_TRAP_DABORT
		EXLOCK exception:	0x0081, EXC_UNDEF	0x008D, EXC_TRAP_OTHER
		Trap on GCSSTR or GCSSTTR:	0x0081, EXC_UNDEF	0x008D, EXC_TRAP_OTHER

2.150 D24341

In section J1.1.2.22 “AArch64.TakeVirtualSErrorException”, the code that reads:

```
AArch64.TakeVirtualSErrorException()
    assert PSTATE.EL IN {EL0, EL1} && EL2Enabled();
    assert HCR_EL2.TGE == '0' && HCR_EL2.AMO == '1'; // Virtual SError enabled if
    TGE==0 and AMO==1
    ...
```

is changed to read,

```
AArch64.TakeVirtualErrorException()  
  
    assert PSTATE.EL IN {EL0, EL1} && EL2Enabled();  
    // Virtual Error enabled if TGE==0 and AMO==1 or TMEA==1  
    assert HCR_EL2.TGE == '0' && (HCR_EL2.AMO == '1' ||  
                                   (IsHCRXEL2Enabled() && HCRX_EL2.TMEA == '1'));  
    ...
```

2.151 D24352

In section D4.6.6.2 “Address Range Comparators”, the equation in rule R_{LLQPL} that reads:

■ $ARC_ADDR_i[n/2] = (ThisInstrAddr)_i \geq start_address) \# (ThisInstrAddr)_i$

is changed to read:

■ $ARC_ADDR_i[n/2] = (ThisInstrAddr)_i \geq start_address) \# (ThisInstrAddr)_i \leq end_address$