Arm DynamIQ Shared Unit-120 MP147

Software Developer Errata Notice

Date of issue: May 16, 2025

Non-Confidential

Document version: 12.0

Document ID: SDEN-2453103

This document contains all known errata since the r0p0 release of the product.

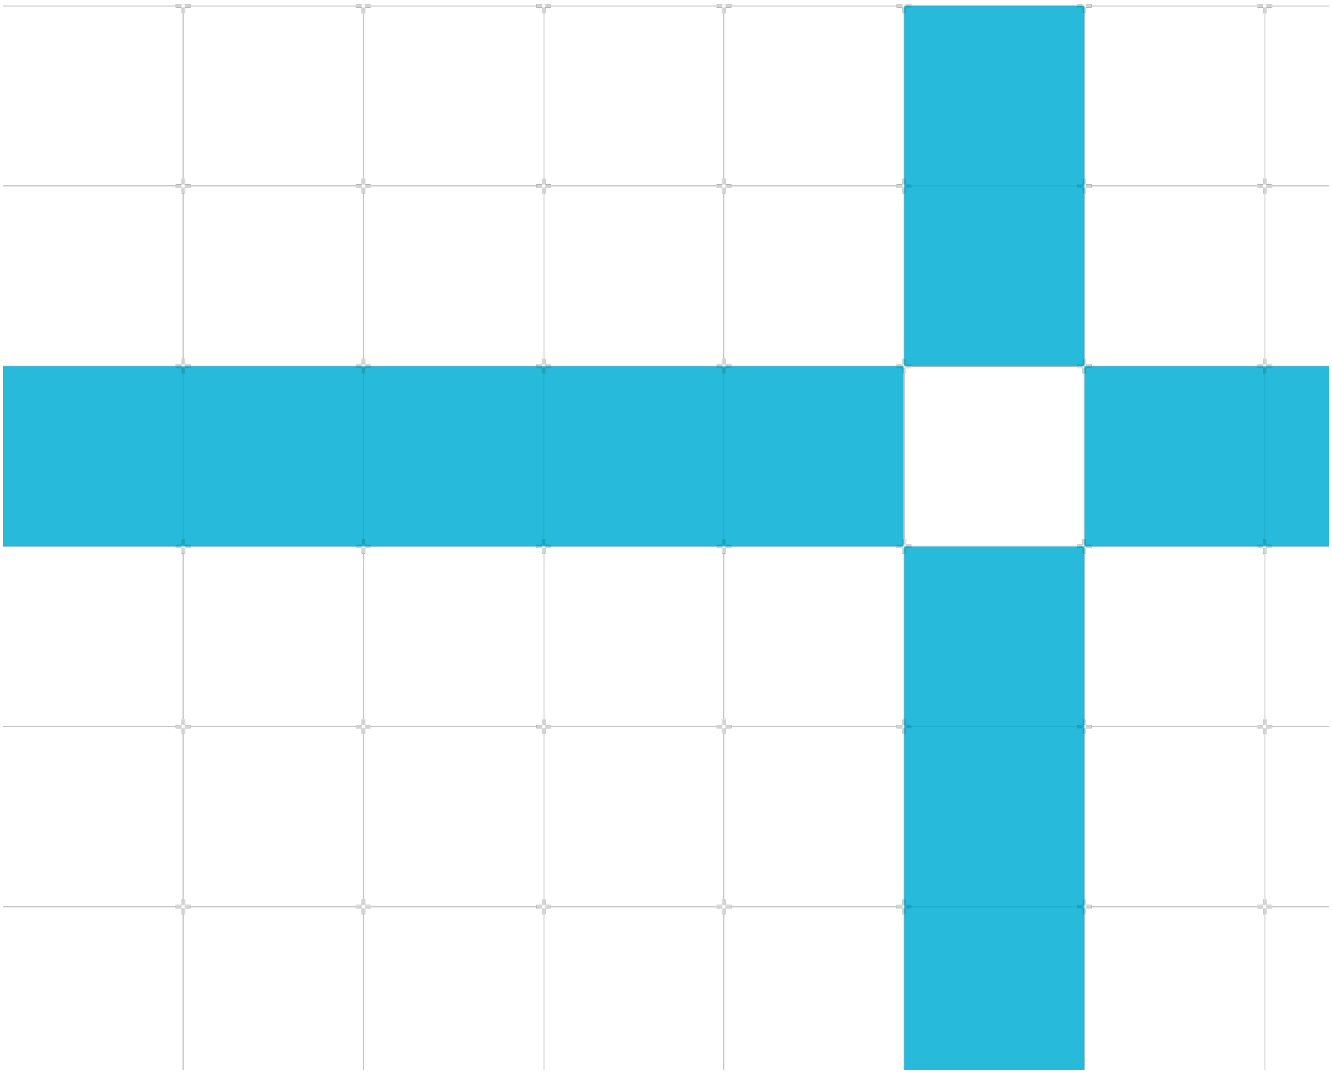

This document (SDEN_2453103_12.0_en) was issued on May 16, 2025.

There might be a later issue at **http://developer.arm.com/documentation/SDEN-2453103**

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.

If you find offensive language in this document, please email **terms@arm.com**.

Feedback

Arm welcomes feedback on this product and its documentation. To provide feedback on Arm DynamIQ Shared Unit-120 MP147, 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.

Contents

r2p0 implementation fixes

Note the following errata might be fixed in some implementations of r2p0. This can be determined by reading the IMP_CLUSTERREVIDR_EL1 register where a set bit indicates that the erratum is fixed in this part.

REVIDR[1]	2900952	CBusy indication causing performance reduction

Note that there is no change to the IMP_CLUSTERIDR_EL1 which remains at r2p0 but the IMP_CLUSTERREVIDR_EL1 is updated to indicate which errata are corrected.
Software will identify this release through the combination of IMP_CLUSTERIDR_EL1 and IMP_CLUSTERREVIDR_EL1.

r1p0 implementation fixes

Note the following errata might be fixed in some implementations of r1p0. This can be determined by reading the IMP_CLUSTERREVIDR_EL1 register where a set bit indicates that the erratum is fixed in this part.

REVIDR[0]	2800804	Power transitions between All slices mode and One slice mode might deadlock

Note that there is no change to the IMP_CLUSTERIDR_EL1 which remains at r1p0 but the IMP_CLUSTERREVIDR_EL1 is updated to indicate which errata are corrected.
Software will identify this release through the combination of IMP_CLUSTERIDR_EL1 and IMP_CLUSTERREVIDR_EL1.

Introduction

Scope

This document describes errata categorized by level of severity. Each description includes:

- The current status of the erratum.
- Where the implementation deviates from the specification and the conditions required for erroneous behavior to occur.
- The implications of the erratum with respect to typical applications.
- The application and limitations of a workaround where possible.

Categorization of errata

Errata are split into three levels of severity and further qualified as common or rare:

Category A	A critical error. No workaround is available or workarounds are impactful. The error is likely to be common for many systems and applications.
Category A (Rare)	A critical error. No workaround is available or workarounds are impactful. The error is likely to be rare for most systems and applications. Rare is determined by analysis, verification and usage.
Category B	A significant error or a critical error with an acceptable workaround. The error is likely to be common for many systems and applications.
Category B (Rare)	A significant error or a critical error with an acceptable workaround. The error is likely to be rare for most systems and applications. Rare is determined by analysis, verification and usage.
Category C	A minor error.

Change Control

Errata are listed in this section if they are new to the document, or marked as "updated" if there has been any change to the erratum text. Fixed errata are not shown as updated unless the erratum text has changed. The **errata summary table** identifies errata that have been fixed in each product revision.

May 16, 2025: Changes in document version v12.0

ID	Status	Area	Category	Summary
3825772	New	Programmer	Category B	Cache maintenance by set/way might operate on wrong L3 index
2777645	Updated	Programmer	Category C	Trace flush may not flush all data in transport
3195828	New	Programmer	Category C	Power transition to ON immediately after warm reset might deadlock
3825768	New	Programmer	Category C	Incorrect Aff0 field in DEVAFF registers
3918694	New	Programmer	Category C	APB access to core, or cluster register might receive incorrect data
3933469	New	Programmer	Category C	Core in a complex transitioning to OFF or OFF_EMU might deadlock
3968468	New	Programmer	Category C	Incorrect fields in MPAMF_IDR register
3968472	New	Programmer	Category C	Incorrect field in PMDEVARCH register
3968477	New	Programmer	Category C	DSU uses default MPAM value with wrong MPAMNS
3968481	New	Programmer	Category C	DSU might not report data loss due to bus error

September 05, 2024: Changes in document version v11.0

ID	Status	Area	Category	Summary
3654550	New	Programmer	Category B	Interrupt outputs remain asserted from cluster after a core is in the OFF_EMU power mode
3654577	New	Programmer	Category C	PPU in static mode with core in FULL_RET might deadlock
3674125	New	Programmer	Category C	Reserved operating modes might cause deadlock

May 28, 2024: Changes in document version v10.0

No new or updated errata in this document version.

April 26, 2024: Changes in document version v9.0

ID	Status	Area	Category	Summary
3207153	New	Programmer	Category B	Denial of service during RetryAcks from CHI interconnect
3203347	New	Programmer	Category C	Incorrect response sent from ACP for poisoned data
3276628	New	Programmer	Category C	MPAM out of range PartID might not be reported

November 22, 2023: Changes in document version v8.0

ID	Status	Area	Category	Summary
2900952	Updated	Programmer	Category B	CBusy indication causing performance reduction
3075552	New	Programmer	Category B	Automated powerdown mechanisms cannot be re-enabled
2873679	New	Programmer	Category C	Error record registers indicate pseudo-fault generation support in configurations without cache protection
2936149	New	Programmer	Category C	ELA not tracing some components in particular configurations
2982213	New	Programmer	Category C	System denying a FULL_RET transition might deadlock
2984621	New	Programmer	Category C	ECC error might cause deadlock if memory attributes mismatched

May 23, 2023: Changes in document version v7.0

No new or updated errata in this document version.

April 04, 2023: Changes in document version v6.0

ID	Status	Area	Category	Summary
2900952	New	Programmer	Category B	CBusy indication causing performance reduction

March 23, 2023: Changes in document version v5.0

No new or updated errata in this document version.

February 24, 2023: Changes in document version v4.0

ID	Status	Area	Category	Summary
2800804	New	Programmer	Category B	Power transitions between All slices mode and One slice mode might deadlock
2777645	New	Programmer	Category C	Trace flush may not flush all data in transport

December 07, 2022: Changes in document version v3.0

ID	Status	Area	Category	Summary
2745150	New	Programmer	Category B	Power transitions between All slices mode and One slice mode might lose coherency

July 26, 2022: Changes in document version v2.0

ID	Status	Area	Category	Summary
2634577	New	Programmer	Category B	ATCLK gating might prevent powerdown
2661093	New	Programmer	Category C	CLUSTERL3HIT and CLUSTERL3MISS registers not saturating on overflow
2667776	New	Programmer	Category C	RAS error during power down might be lost
2679678	New	Programmer	Category C	Cache debug access might deadlock
2700719	New	Programmer	Category C	Clearing RAS interrupt during power down might cause deadlock

February 28, 2022: Changes in document version v1.0

ID	Status	Area	Category	Summary
2457823	New	Programmer	Category C	CLUSTERPMSSRR bits cannot be cleared

Errata summary table

The errata associated with this product affect the product versions described in the following table.

ID	Area	Category	Summary	Found in versions	Fixed in version
2634577	Programmer	Category B	ATCLK gating might prevent powerdown	r0p0	r1p0
2745150	Programmer	Category B	Power transitions between All slices mode and One slice mode might lose coherency	r1p0	r2p0
2800804	Programmer	Category B	Power transitions between All slices mode and One slice mode might deadlock	r1p0	r2p0
2900952	Programmer	Category B	CBusy indication causing performance reduction	r2p0	r2p1
3075552	Programmer	Category B	Automated powerdown mechanisms cannot be re-enabled	r0p0, r1p0, r2p0	r2p1
3207153	Programmer	Category B	Denial of service during RetryAcks from CHI interconnect	r2p0, r2p1	Open
3654550	Programmer	Category B	Interrupt outputs remain asserted from cluster after a core is in the OFF_EMU power mode	r0p0, r1p0, r2p0, r2p1	Open
3825772	Programmer	Category B	Cache maintenance by set/way might operate on wrong L3 index	r0p0, r1p0, r2p0, r2p1	Open
2457823	Programmer	Category C	CLUSTERPMSSRR bits cannot be cleared	r0p0	r1p0
2661093	Programmer	Category C	CLUSTERL3HIT and CLUSTERL3MISS registers not saturating on overflow	r0p0	r1p0
2667776	Programmer	Category C	RAS error during power down might be lost	r0p0	r1p0
2679678	Programmer	Category C	Cache debug access might deadlock	r0p0	r1p0
2700719	Programmer	Category C	Clearing RAS interrupt during power down might cause deadlock	r0p0	r1p0
2777645	Programmer	Category C	Trace flush may not flush all data in transport	r0p0, r1p0, r2p0, r2p1	Open
2873679	Programmer	Category C	Error record registers indicate pseudo-fault generation support in configurations without cache protection	r0p0, r1p0, r2p0	r2p1
2936149	Programmer	Category C	ELA not tracing some components in particular configurations	r2p0	r2p1

ID	Area	Category	Summary	Found in versions	Fixed in version
2982213	Programmer	Category C	System denying a FULL_RET transition might deadlock	r2p0	r2p1
2984621	Programmer	Category C	ECC error might cause deadlock if memory attributes mismatched	r0p0, r1p0, r2p0, r2p1	Open
3195828	Programmer	Category C	Power transition to ON immediately after warm reset might deadlock	r0p0, r1p0, r2p0, r2p1	Open
3203347	Programmer	Category C	Incorrect response sent from ACP for poisoned data	r0p0, r1p0, r2p0, r2p1	Open
3276628	Programmer	Category C	MPAM out of range PartID might not be reported	r0p0, r1p0, r2p0, r2p1	Open
3654577	Programmer	Category C	PPU in static mode with core in FULL_RET might deadlock	r0p0, r1p0, r2p0, r2p1	Open
3674125	Programmer	Category C	Reserved operating modes might cause deadlock	r0p0, r1p0, r2p0, r2p1	Open
3825768	Programmer	Category C	Incorrect Aff0 field in DEVAFF registers	r0p0, r1p0, r2p0, r2p1	Open
3918694	Programmer	Category C	APB access to core, or cluster register might receive incorrect data	r0p0, r1p0, r2p0, r2p1	Open
3933469	Programmer	Category C	Core in a complex transitioning to OFF or OFF_EMU might deadlock	r0p0, r1p0, r2p0, r2p1	Open
3968468	Programmer	Category C	Incorrect fields in MPAMF_IDR register	r0p0, r1p0, r2p0, r2p1	Open
3968472	Programmer	Category C	Incorrect field in PMDEVARCH register	r0p0, r1p0, r2p0, r2p1	Open
3968477	Programmer	Category C	DSU uses default MPAM value with wrong MPAMNS	r0p0, r1p0, r2p0, r2p1	Open
3968481	Programmer	Category C	DSU might not report data loss due to bus error	r0p0, r1p0, r2p0, r2p1	Open

Errata descriptions

Category A

There are no errata in this category.

Category A (rare)

There are no errata in this category.

Category B

2634577
ATCLK gating might prevent powerdown

Status

Fault Type: Programmer Category B
Fault Status: Present in r0p0. Fixed in r1p0.

Description

The DSU provides a Q-Channel to allow the system to gate **ATCLK** when that domain is idle. If the clock is gated during a powerdown sequence, in a system with unusual conditions, then it might prevent the powerdown sequence from completing.

Configurations Affected

This erratum affects all configurations of the DSU. It also requires a system design that can generate the conditions on the **ATCLK** Q-Channel.

Conditions

This erratum occurs when the following sequence of conditions is met:

1. A cluster power mode transition from ON to OFF, or ON to MEM_RET is started.
2. The **ATCLK** Q-Channel is in the Q-Stopped state, but the system is still providing **ATCLK**.
3. The power transition will cause the **ATCLKQACTIVE** signal to be HIGH for more than three **ATCLK** cycles. It might eventually go LOW if **ATCLK** continues to be provided by the system.
4. The system gates **ATCLK** while **ATCLKQACTIVE** is still high, and then does not ungate the clock.

Implications

The system has to behave in an unusual way to satisfy the conditions. Arm does not expect many systems to leave the clock ungated while **ATCLKQACTIVE** is LOW, then gate it when **ATCLKQACTIVE** is HIGH, and then not ungate it. If the clock remains gated, then the power sequence will not complete, which will cause a system deadlock.

Workaround

Most systems are not expected to require a workaround. For those that are affected, firmware should set bit 27 of the IMP_CLUSTERACTLR_EL1 register before the last core in the cluster powers off. This will prevent **ATCLK**, **GICCLK** and **PCLK** from being gated during the power down sequence.

2745150
Power transitions between All slices mode and One slice mode might lose coherency

Status

Fault Type: Programmer Category B
Fault Status: Present in r1p0. Fixed in r2p0.

Description

If the cluster PPU makes an operating mode transition from ALL SLICE to ONE SLICE, or from ONE SLICE to ALL SLICE, under certain conditions the cluster might lose coherency.

Configurations Affected

This erratum affects configurations that meet all of the following conditions:

- More than one L3 slice is configured
- At least one ACP interface is configured
- No Peripheral Port is configured, or a 256-bit Peripheral Port is configured
- The DSU is configured with eight transport nodes. This occurs when:
  - NUM_L3_SLICES is 2 and the total number of complexes and standalone cores is greater than 8
  - NUM_L3_SLICES is 4 and the total number of complexes and standalone cores is greater than 4
  - NUM_L3_SLICES is 8 for any number of cores

Direct connect configurations are not affected.

Conditions

The erratum occurs under the following conditions:

1. The cluster PPU requests a cluster operating mode transition to change from ALL SLICE to ONE SLICE, or from ONE SLICE to ALL SLICE. This will typically be in response to software changing the IMP_CLUSTERPWRCTLR_EL1.SLCRQ bit.
2. During the transition, any of the following conditions is met:
   - There are at least (NUM_LTDBS + 6) ACP transactions outstanding.
   - There are at least 7 ACP transactions outstanding and at least 6 of these are to the same address as outstanding evictions from the DSU to the interconnect, outstanding snoops from the interconnect to the DSU, or other ACP transactions.
   - There are at least 8 Device non-Reorderable ACP transactions outstanding all using different AWIDs or ARIDs. Note that for these purposes, an ARID and AWID with the same value are considered different.
3. During the transition, there is at least one coherent read transaction from a core. This read might be caused by either executing an instruction or a hardware mechanism, such as a translation table walk

or hardware prefetch.

Implications

If the previous conditions and certain microarchitectural timing conditions occur, the DSU might lose coherency for the cacheline that was read by the core.

Workaround

The software should not make use of the ONE SLICE operating mode. It should keep the IMP_CLUSTERPWRCTLR_EL1.SLCRQ bit set, which is the default value. The ONE SLICE power mode is designed to be used when not many cores are active and a lower DSU bandwidth is required. In these situations, the workaround will cause the DSU power to be higher than it would be using the ONE SLICE power mode.

2800804
Power transitions between All slices mode and One slice mode might deadlock

Status

Fault Type: Programmer Category B
Fault Status: Present in r1p0. Fixed in r2p0.

Description

If the cluster PPU makes an operating mode transition from ALL SLICE to ONE SLICE, or from ONE SLICE to ALL SLICE, under certain conditions the power transition might deadlock.

Configurations Affected

This erratum affects configurations that meet all of the following conditions:

- At least one ACP interface is configured.
- A 64-bit AXI Peripheral Port is configured.

Direct connect configurations are not affected.

Conditions

The erratum occurs under the following conditions:

1. The cluster PPU requests a cluster operating mode transition to change from ALL SLICE to ONE SLICE, or from ONE SLICE to ALL SLICE. This will typically be in response to software changing the IMP_CLUSTERPWRCTLR_EL1.SLCRQ bit.
2. During the transition, certain microarchitectural timing conditions occur.

Implications

If certain microarchitectural timing conditions occur, the power transition might not complete. If this occurs, the DSU will not perform new power transitions. It also might not respond to new CPU or ACP transactions. This will likely lead to a system deadlock.

The DSU will continue to respond to snoop transactions. If the DSU receives a non-DVM snoop on the expected CHI port for the physical address, the deadlock might end and the DSU might operate normally again.

Workaround

The software should not make use of the ONE SLICE operating mode. It should keep the IMP_CLUSTERPWRCTLR_EL1.SLCRQ bit set, which is the default value. The ONE SLICE power mode is designed to be used when not many cores are active and a lower DSU bandwidth is required. In these situations, the workaround will cause the DSU power to be higher than it would be using the ONE SLICE power mode.

2900952
CBusy indication causing performance reduction

Status

Fault Type: Programmer Category B
Fault Status: Present in r2p0. Fixed in r2p1.

Description

The CHI protocol includes a CBusy field that allows the system to indicate when it is heavily loaded. The cores can then adjust their behavior, such as reducing the amount of speculative accesses from their prefetchers, to make the most efficient use of the available system bandwidth. In some configurations, the DSU does not correctly pass on this field to the cores, resulting in reduced performance from the cores.

Configurations affected

All configurations except direct connect are affected.

Conditions

The erratum will occur under either of the following conditions:

1. The system interconnect asserts the CBusy field in one of the response or data flits that it sends to a DSU CHI interface, and then deasserts the CBusy field in later flits.
2. One or more cores send heavy traffic to the DSU while other cores in the same cluster do not send much traffic.

Implications

If condition 1 occurs, the DSU will assert CBusy to the cores correctly, but will not deassert it later when the interconnect is no longer busy. The cores will continue to restrict the transactions that they generate which can reduce performance.
If condition 2 occurs, the DSU might assert CBusy to an incorrect core, which can reduce the performance of that core until the condition clears.

Workaround

The DSU can be programmed to ignore the CBusy from the system interconnect by setting the CLUSTERACTLR_EL1 bits [21:20] to 0x3, which will mitigate the performance impact when the system is lightly loaded, however it will prevent the CBusy mechanism from improving the performance in a heavily loaded system. These bits should be programmed as soon as possible because they might not mitigate the performance impact if the erratum occurs before the workaround is applied. CLUSTERACTLR_EL1 bit [8] should be set to 0x1 to assert CBusy to all the cores when the DSU is busy.

3075552
Automated powerdown mechanisms cannot be re-enabled

Status

Fault Type: Programmer Category B
Fault Status: Present in r0p0, r1p0 and r2p0. Fixed in r2p1.

Description

The automated RAM powerdown and slice powerdown mechanisms evaluate whether to change the power mode on a configurable time period. If the mechanism is disabled at the same time as the time period completes, it might prevent re-enabling the mechanism at a later time.

Configurations affected

This erratum affects all configurations except direct connect.

Conditions

The erratum occurs under the following conditions:

1. The IMP_CLUSTERPWRCTLR_EL1.AUTOPRTN field or IMP_CLUSTERPWRCTLR_EL1.AUTOSLC field is set to a non-zero value to enable the automatic RAM powerdown or automatic slice powerdown.
2. The IMP_CLUSTERPWRCTLR_EL1.AUTOPRTN or IMP_CLUSTERPWRCTLR_EL1.AUTOSLC field is later set to zero to disable the functionality.
3. At the same time as it is being disabled, the generic timer value reaches the end of the timeout period.
4. The IMP_CLUSTERPWRCTLR_EL1.AUTOPRTN or IMP_CLUSTERPWRCTLR_EL1.AUTOSLC field is set to a non-zero value to re-enable the automatic RAM powerdown or slice powerdown.

Implications

The automated powerdown mechanism will not make any further transitions, which can lead to either higher power or lower performance if the RAMs or slice remain in a mode that is no longer appropriate for the current workload.
Software can still request manual transitions using the IMP_CLUSTERPWRCTLR_EL1.PRTNRQ or IMP_CLUSTERPWRCTLR_EL1.SLCRQ bits.

Workaround

If software is making use of the automated RAM or slice powerdown, it should avoid disabling it and re-enabling it. If it does need to disable it with the expectation of later enabling again, the following sequence can be used:

1. Read the current value of IMP_CLUSTERPWRCTLR_EL1.AUTOPRTN or IMP_CLUSTERPWRCTLR_EL1.AUTOSLC
2. Proceed with the following steps depending on the current value:
    - If the current value is 0b111 then:
        - Set IMP_CLUSTERPWRCTLR_EL1.AUTOPRTN or IMP_CLUSTERPWRCTLR_EL1.AUTOSLC to 0b001
        - Use the generic timer to wait for the time indicated by the value read by step 1
        - Set IMP_CLUSTERPWRCTLR_EL1.AUTOPRTN or IMP_CLUSTERPWRCTLR_EL1.AUTOSLC to 0b111
        - Use the generic timer to wait for 524288 timer ticks
    - If the current value is not 0b111 then:
        - Set IMP_CLUSTERPWRCTLR_EL1.AUTOPRTN or IMP_CLUSTERPWRCTLR_EL1.AUTOSLC to 0b111
        - Use the generic timer to wait for the time indicated by the value read by step 1
3. Set IMP_CLUSTERPWRCTLR_EL1.AUTOPRTN or IMP_CLUSTERPWRCTLR_EL1.AUTOSLC to 0b000

Note that software could perform other tasks during step 2, while waiting for the time to expire, but if it cannot make use of this time then the delay could have an impact on performance.

3207153
Denial of service during RetryAcks from CHI interconnect

Status

Fault Type: Programmer Category B
Fault Status: Present in r2p0 and r2p1. Open.

Description

If a core or *Accelerator Coherency Port* (ACP) interface sends a stream of memory system transactions to the DSU, and the *Coherent Hub Interface* (CHI) system interconnect repeatedly responds with RetryAck, then one transaction might not make progress.

Configurations Affected

This erratum only affects configurations connected to a CHI interconnect that can generate RetryAck. It does not affect Direct connect configurations.

Conditions

The erratum occurs if all the following conditions apply:

1. At least one of the 2 following conditions applies:
    - One or more cores execute a stream of instructions that causes memory system transactions to the system interconnect.
    - There is a stream of transactions on an ACP interface that causes memory system transactions to the system interconnect.
2. The transactions were generated with different *Memory System Resource Partitioning and Monitoring* (MPAM) PartID values, or from a different security state, which is indicated by MPAM_NS on the ACP interface.
3. The system interconnect repeatedly responds with RetryAck.
4. Certain micro-architectural conditions occur.

Implications

If the erratum occurs, then a stream of transactions with one combination of MPAM PartID and MPAM_NS might prevent the progress of transactions with a different combination of MPAM PartID and MPAM_NS.

Malicious code could be used to prevent instructions from being executed on the processor indefinitely. A Non-secure process or ACP could stall a Secure process on the processor.

Workaround

CLUSTERACTLR_EL1 bit[42] should to be set to 1 to work around this issue. This makes the DSU ignore MPAM PartID and MPAM_NS when arbitrating transactions that have received a RetryAck, so MPAM memory bandwidth regulation will not happen in this situation. This is the behaviour of previous versions of DSU-120.

3654550
Interrupt outputs remain asserted from cluster after a core is in the OFF_EMU power mode

Status

Fault type: Programmer Category B
Fault status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

The interrupt outputs from the core should be deasserted when the core is in the OFF_EMU power mode because the source of the interrupt is no longer active. However, if the source was active and the interrupt was not serviced before the core enters OFF_EMU, then these interrupt pins remain incorrectly asserted.

Configuration affected

This erratum affects all configurations.

Conditions

The erratum occurs if all of the following conditions apply:

1. A core is in the ON power mode.
2. The core has at least one interrupt output request pending. This causes one or more of the below output pins to be asserted:
   - Group 1
     - nTBEIRQ
     - nPMBIRQ
     - nPMUIRQ
     - nVCPUMNTIRQ
   - Group 2
     - nCNTPNSIRQ
     - nCNTPSIRQ
     - nCNTVIRQ
     - nCNTHVIRQ
     - nCNTHPIRQ
     - nCNTHVSIRQ
     - nCNTHPSIRQ
3. The core enters the OFF_EMU power mode before the interrupt is serviced.

Implications

This might result in the interrupt handler receiving spurious interrupts.

The interrupt outputs will remain asserted when the core is in the OFF_EMU power mode and continue to be asserted when the core is powered on again. They will be cleared when either the core is powered off, or when the same core asserts one of these interrupt sources again from the same group.

Workaround

When the core is powered on, software at EL3 should program the source of one of these pins from each group to generate an interrupt before it unmasks interrupts in the PSTATE register.

3825772
Cache maintenance by set/way might operate on wrong L3 index

Status

Fault type: Programmer Category B
Fault status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

In some situations, L3 cache maintenance by set/way might operate on the wrong L3 index.

Configurations Affected

This erratum affects configurations that meet all of the following conditions:

- At least two L3 cache slices
- 3MB or 4MB of L3 cache per slice
- At least one ACP interface
- A 64-bit AXI Peripheral Port

Direct connect configurations are not affected.

Conditions

The erratum occurs under the following conditions:

1. A core executes an L3 set/way cache maintenance operation
2. Before the first cache maintenance operation completes, a different core executes an L3 set/way cache maintenance operation

Implications

If the erratum occurs, cache maintenance operation might be executed on the wrong L3 cache slice, which is part of the L3 index. This means that the intended L3 entry will not be cleaned and invalidated, and instead a different entry will be cleaned and invalidated. The intended entry will remain in the cache, which means its data might not become visible to other observers when expected. If the entry was being invalidated due to being potentially stale, it might cause software to subsequently read the incorrect stale entry.

Cache maintenance by set/way does not provide any guarantees for caches that use hardware coherency unless the caches are disabled. This is because dirty data can migrate between cache levels, avoiding the effect of the cache maintenance. As a result, no general purpose software uses cache maintenance by set/way. Only firmware or software operating with the caches disabled might be affected by this erratum.

Workaround

No workaround is required for most software. Affected software should use one core to perform set/way cache maintenance to the L3 cache, and should ensure it completes by using a `DSB` before allowing another core to perform set/way cache maintenance to the L3 cache.

Category B (rare)

There are no errata in this category.

Category C

2457823
CLUSTERPMSSRR bits cannot be cleared

Status

Fault Type: Programmer Category C
Fault Status: Present in r0p0. Fixed in r1p0.

Description

The CLUSTERPMSSRR register controls whether the PMU counters are reset after a snapshot is taken. If this register is written to set any of the bits, then those bits will be set, but further writes to the register will be unable to clear them.

Configurations Affected

This erratum affects all configurations except Direct connect.

Conditions

The erratum occurs under the following conditions:

1. The CLUSTERPMSSRR register is written with at least one of the RP bits set.
2. The CLUSTERPMSSRR register is written to clear one of the RP bits that is already set.

Implications

The RP bits that are set will not be cleared. The snapshot feature will work with reset of the counters enabled, or disabled, however software will not be able to switch between the two modes more than once, until the cluster is reset.

Workaround

The software should chose which mode it requires and then avoid subsequent switches to the other mode.

2661093
CLUSTERL3HIT and CLUSTERL3MISS registers not saturating on overflow

Status

Fault Type: Programmer Category C
Fault Status: Present in r0p0. Fixed in r1p0.

Description

The CLUSTERL3HIT and CLUSTERL3MISS registers are typically used to decide when to enter or exit from the L3 RAM powerdown modes. These registers are 32-bit wide, and they are defined to saturate when reaching the maximum value. Under certain conditions, the counters might incorrectly wrap rather than saturating.

Configurations Affected

Configurations with an L3 cache that has more than two cache slices are affected.

Conditions

This erratum occurs under the following conditions:

1. The CLUSTERPWRCTLR.AUTOPTRN field is zero and the CLUSTERL3HIT and CLUSTERL3MISS registers are used by software or firmware.
2. The CLUSTERL3HIT or CLUSTERL3MISS register has a value in the range 0xFFFFFFF8 to 0xFFFFFFFD.
3. Between three and eight cache slices detect an L3 hit or an L3 miss on the same cycle.

Implications

When these conditions occur, the CLUSTERL3HIT or CLUSTERL3MISS registers might wrap to a value of 0xFFFFFF00 to 0xFFFFFF05, rather than saturating at 0xFFFFFFFF.
The expected use of these registers is for software or firmware to regularly sample their value, and to reset the count to 0 after each sample. Therefore, in typical use the counter is not expected to reach the maximum value, and if it did then the saturation would have already introduced some inaccuracy. Therefore, this erratum is not expected to significantly affect the accuracy of the values in most systems.

Workaround

The software or firmware should ensure that these registers are sampled and reset frequently enough that they do not reach their maximum value. A sampling period of 100ms or better is sufficient.

2667776
RAS error during power down might be lost

Status

Fault Type: Programmer Category C
Fault Status: Present in r0p0. Fixed in r1p0.

Description

If a RAS error occurs during a power down transition, the transition should be aborted so that the RAS error record can be preserved until the system has been able to deal with the error.
If a RAS error occurs during a small time window in some power transitions, the power transition might continue and the error record can be lost.

Configurations Affected

This erratum affects all configurations of the DSU, except Direct connect configurations.

Conditions

This erratum occurs when the following sequence of conditions is met:

1. The cluster RAS Critical error interrupt, Fault handling interrupt, or Uncorrected error recovery interrupt are enabled.
2. A cluster power mode transition from ON to OFF, OFF_EMU, MEM_RET, or MEM_RET_EMU is started.
3. A RAS error occurs during the power transition. If the transition is to OFF or OFF_EMU, then the error must occur after all the transactions to flush every line have started, but before they have all completed.

Implications

The error will be reported in the RAS registers correctly, however the power sequence will not be aborted. Therefore, the system might not have time to react to the error before the power is removed and the contents of the RAS registers are lost.
The transition to MEM_RET or MEM_RET_EMU is very quick and does not require flushing the L3 cache, therefore it is unlikely that there would be any transactions outstanding during this period, and so the probability of an error being detected is very small.
The transition to OFF or OFF_EMU will flush the L3 cache which can take a long time, however if the error is detected during the majority of this period it will be reported correctly and the power sequence will be aborted. Therefore, the probability of an error occurring on the last few transactions of the flush is very small.
There might be a negligible increase in overall system failure rate because of this erratum.

Workaround

No workaround is required for this erratum.

2679678
Cache debug access might deadlock

Status

Fault Type: Programmer Category C
Fault Status: Present in r0p0. Fixed in r1p0.

Description

The cache Debug register provides a method for software at EL3 to directly access the contents of the RAMs inside the DSU for debug purposes. If the cache Debug register is accessed while there are other memory transactions accessing the DSU, then in rare cases, the system might deadlock.

Configurations Affected

This erratum affects all configurations, except Direct connect.

Conditions

The erratum occurs if all the following conditions apply:

1. Software running at EL3 writes and reads from the IMP_CLUSTERCDBG_EL3 System register to access one of the DSU RAMs.
2. The system interconnect sends an access on the Utility Bus to read a memory mapped register in the DSU.
3. A large number of memory transactions are made to the DSU from cores in the cluster, that miss in L3 and so are sent to the system interconnect. The system interconnect contains a dependency that means it cannot respond to these transactions until the Utility Bus transaction has completed.

Implications

If this erratum occurs, then the system will deadlock. The cache Debug register is an IMPLEMENTATION DEFINED EL3 register, and therefore only custom EL3 debug software is affected.
Typical uses of the cache Debug register would already have other cores idle while reading the cache contents; to avoid disturbing the cache contents. Therefore, other memory transactions are unlikely to occur in such uses.

Workaround

Before executing cache debug instructions, debug software should ensure that other cores and system components that might access the DSU are idle.

2700719
Clearing RAS interrupt during power down might cause deadlock

Status

Fault Type: Programmer Category C
Fault Status: Present in r0p0. Fixed in r1p0.

Description

If a RAS interrupt occurs during a cluster power down transition and the system clears the interrupt then the power transition might deadlock.

Configurations Affected

This erratum affects all configurations of the DSU with an L3 cache. It does not affect Direct Connect configurations.

Conditions

This erratum occurs when the following sequence of conditions occurs:

1. The cluster RAS Critical error interrupt, Fault handling interrupt, or Uncorrected error recovery interrupt are enabled in CLUSTERRAS_ERR0CTLR.
2. All the cores are in OFF or OFF_EMU power states.
3. A cluster power mode transition from ON to OFF or OFF_EMU starts.
4. A RAS error occurs during the power transition and this generates an interrupt.
5. During the power transition and very soon after the interrupt is generated, the system clears the interrupt by writing to the RAS registers using the Utility Bus. This can be a write to CLUSTERRAS_ERR0STATUS to clear the error record, or to CLUSTERRAS_ERR0CTLR to disable RAS interrupt generation or error detection.

Implications

If the RAS registers are cleared very soon after the interrupt is generated, the power transition will deadlock. If this happens, the cluster will process snoop transactions from the interconnect, but it will not process new ACP transactions. It will not be possible to power up the cores or make any new cluster power transitions.

Arm expects that the system will not respond to the interrupt quickly enough to encounter this erratum. There might be a negligible increase in overall system failure rate because of this erratum.

Workaround

No workaround is required.

2777645
Trace flush may not flush all data in transport

Status

Fault Type: Programmer Category C
Fault Status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

If a flush request is made on the *AMBA Trace Bus* (ATB) interface when a trace source has recently generated trace data, the flush might complete before the already generated bytes of the trace data have been output from the DSU transport network.

Configurations affected

This erratum affects all configurations of the DSU except Direct connect.

Conditions

1. The *Embedded Trace Extension* (ETE) in a core or an *Embedded Logic Analyzer* (ELA) in a core or in the cluster is generating trace output.
2. The trace subsystem outside the cluster requests a flush by asserting the **AFVALID** signal before the trace data has naturally drained from the cluster.

The cluster may incorrectly assert the **AFREADY** signal before the already generated trace data has been output on the ATB interface.

Implications

Trace data generated by the ETE in a core or ELA in a core will naturally drain as quickly as it can. If an explicit flush is requested during this time, then the flush may complete before trace data have been output, which could lead to incomplete trace being processed by the debug tools.
If the *TRace Buffer Extension* (TRBE) is in use, then the ATB interface in the DSU is not used and so this erratum has no effect.

Workaround

No workaround is necessary.

2873679
Error record registers indicate pseudo-fault generation support in configurations without cache protection

Status

Fault Type: Programmer Category C
Fault Status: Present in r0p0, r1p0 and r2p0. Fixed in r2p1.

Description

Pseudo-fault generation is a mechanism that software can use to generate an entry in the RAS error record. This makes it possible to test how software would behave if a real fault was reported in the RAS error record registers.
In configurations without cache protection, the DSU does not support pseudo-fault generation, even though the RAS registers for pseudo-fault generation exist.

The following registers indicate the features implemented in the DSU RAS node:

- ERR0FR
- ERR0PFGF

These registers incorrectly indicate that pseudo-fault generation is supported.

The other pseudo-fault generation registers, ERR0PFGCTL and ERR0PFGCDN, can be read and written, but will not generate a pseudo-fault in this configuration.

Configurations Affected

All non-Direct Connect configurations with the SCU_CACHE_PROTECTION parameter set to FALSE are affected.

Conditions

The incorrect behaviour of these registers always occurs in affected configurations. These registers can be accessed by the Utility Bus or by the System Registers.

Implications

In configurations without cache protection, software might incorrectly assume that pseudo-fault generation will work. This might lead to an unexpected error injection test result if a test attempts to inject an error but no error record is created.

Software can determine that a configuration is affected if IMP_CLUSTERCFR_EL1 bits [63:61] are non-zero and bit [14] is zero.

Workaround

There is no workaround.

2936149
ELA not tracing some components in particular configurations

Status

Fault Type: Programmer Category C
Fault Status: Present in r2p0. Fixed in r2p1.

Description

The DSU supports tracing of signals from the design when an ELA-600 is integrated at cluster level. In some configurations data is traced from components that are connected to a different transport node from which the ELA is connected, and the transport network does not correctly transport these signals, resulting in no trace data for these components.

Configurations affected

This erratum affects configurations where DIRECT_CONNECT is FALSE and ELA is TRUE and number of transport nodes is 2, 4, or 8 and at least one of the following configurations:

1. TRANSPORT_HORIZ_REG_SLICE parameter is not equal to 0 and CORNER_PIN_LOCATION is TRUE and
   - Any master or peripheral port is not on node 1 or
   - Any ACP is on node 0, 3, 4, 5, 6, or 7
2. TRANSPORT_HORIZ_REG_SLICE parameter is not equal to 0 and CORNER_PIN_LOCATION is FALSE and
   - Any master port or peripheral port is not on node 0 or
   - Any ACP is on node 1, 2, 6, or 7
3. TRANSPORT_VERT_REG_SLICE is not equal to 0 and CORNER_PIN_LOCATION is TRUE and any ACP is on node 2
4. TRANSPORT_VERT_REG_SLICE is not equal to 0 and CORNER_PIN_LOCATION is FALSE and any ACP is on node 3 or 4

Conditions

1. The **ELADISABLE** pin is LOW.
2. The ELA is programmed to sample a signal group that includes a master port, peripheral port, or ACP interface.

Implications

The affected signal groups will be traced by the ELA as zero, which might impact debug of the affected interfaces using the ELA.

Workaround

There is no workaround for this erratum.

2982213
System denying a FULL_RET transition might deadlock

Status

Fault Type: Programmer Category C
Fault Status: Present in r2p0. Fixed in r2p1.

Description

The FUNC_RET and FULL_RET power modes can be used to save leakage power when the DSU is idle. If these modes are enabled, but the system denies a transition from FUNC_RET to FULL_RET, then it can cause a deadlock if software is disabling the FUNC_RET mode at the same time.

Configurations affected

This erratum affects all configurations of the DSU except Direct connect.

Conditions

1. The IMP_PWRCTRL_EL1.FUNCRET field is set to a non-zero value to enable entry into the FUNC_RET power mode.
2. The IMP_PWRCTRL_EL1.FULLRET field is set to a non-zero value to enable entry into the FULL_RET power mode.
3. The DSU is idle for a long enough period to enter FUNC_RET, but not long enough to enter FULL_RET. This means that no loads, stores, TLB pagewalks, instruction fetches, accesses from the ACP interface, or snoops from the system interconnect access the DSU in this period.
4. While the DSU is in the FUNC_RET mode, software sets the IMP_PWRCTRL_EL1.FUNCRET field to 0b000 to disable FUNC_RET.
5. The DSU remains idle for a long enough period to start an entry into the FULL_RET power mode.
6. As part of the power transition to FULL_RET, the DSU will make a request on the input domain P-Channel, followed by a request on the output domain P-Channel. The system accepts the input domain channel request, but denies the output domain request.

Implications

If the conditions occur, then the power transition will not complete, which can cause a deadlock. However FUNC_RET and FULL_RET are power modes that are internal to the DSU, therefore it is unlikely that a system would have any reason to deny these modes on the output domain P-Channel.

Workaround

Most systems will not deny a transition to FULL_RET and therefore do not require a workaround. For systems that can deny this transition, software should follow this sequence when it wants to disable FUNC_RET:

1. Read the current value of IMP_PWRCTLR_EL1.FULLRET
2. Write 0b000 to IMP_PWRCTLR_EL1.FULLRET
3. Write 0b000 to IMP_PWRCTLR_EL1.FUNCRET
4. Restore the saved value from step 1 back to IMP_PWRCTLR_EL1.FULLRET

2984621
ECC error might cause deadlock if memory attributes mismatched

Status

Fault Type: Programmer Category C
Fault Status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

If an ECC error occurs while there are multiple outstanding transactions to the same address and they use different memory attributes, then the DSU might not respond to a snoop request to that address. This might cause a system deadlock.

Configurations affected

This erratum affects all configurations of the DSU connected to a coherent CHI interconnect except Direct Connect configurations.

Conditions

This erratum occurs when the following sequence of conditions occurs:

1. A core or an ACP requester issues a transaction with any memory type except Normal Inner-WriteBack, Outer-WriteBack, Inner-Shareable, or Outer-Shareable. The transaction is sent to the system interconnect.
2. Before this transaction finishes, a core or an ACP requester issues a transaction to the same address with the memory type Normal Inner-WriteBack, Outer-WriteBack, Inner-Shareable, or Outer-Shareable.
3. That transaction detects an ECC error in the L3 Tag RAM or Snoop Filter RAM.
4. Either:
   - The DEVNRINTERLEAVE input to the DSU is 0b01 and the same core or ACP interface from step 1 issues another transaction to a different address before any of the other transactions have finished. Both this transaction and the transaction from step 1 have the memory type Device non-Reorderable.
   - The CLUSTERECTLR register bit [18] is set and the transaction in step 1 was a write or atomic with any memory type *except* Normal Inner-WriteBack, Outer-WriteBack, or Device non-Reorderable. Before any of the other transactions have finished, the same core or ACP interface from step 1 issues another transaction to the same address with the memory type Normal Inner-WriteBack, Outer-WriteBack.
5. The system interconnect sends a snoop to the DSU for the same address as the transaction in step 1.

Implications

This erratum only occurs when different memory types are used for a single memory location at the same time. This is not expected in most software and systems. The erratum also only occurs if an ECC error occurs. There is still substantial benefit being gained from the ECC logic. This erratum might cause a negligible increase in overall system failure rate.

If the erratum occurs, then the DSU will not respond to the snoop until the transaction from step 1 completes. If the system interconnect waits for the snoop response before processing the transaction from step 1 then the system might deadlock.

Workaround

No workaround is required.

3195828
Power transition to ON immediately after warm reset might deadlock

Status

Fault type: Programmer Category C
Fault status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

The cluster might deadlock if it enters a debug power mode that asserts warm reset, and then immediately transitions to the ON power mode.

Configurations affected

This erratum affects all configurations.

Conditions

The erratum occurs under the following conditions:

1. The cluster enters one of the following debug power modes:
    - OFF_EMU
    - MEM_RET_EMU
    - WARM_RST
    - DBG_RECOV
2. The frequency of SCLK is half or less compared to PPU clock, or SCLK is clock gated externally in the system. This delays the assertion of the warm reset.
3. The cluster transitions to the ON power mode
4. The transition to ON starts soon after the previous power transition, before the warm reset has completed.

Implications

If the conditions occur, the cluster might deadlock during debug.

Workaround

No workaround is expected to be required. Please contact Arm if you believe that a workaround is required for your system.

3203347
Incorrect response sent from ACP for poisoned data

Status

Fault Type: Programmer Category C
Fault Status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

When a read returns data on the ACP interface, it will indicate a SLVERR response for any beats of data that are poisoned due to an uncorrectable ECC error. If only some of the beats of a read transaction are poisoned, then the ACP interface might indicate the SLVERR response on the wrong beats of data.

Configurations affected

Configurations where ACP is TRUE, ACP_WIDTH is 128, and SCU_CACHE_PROTECTION is TRUE are affected. Direct connect configurations are not affected.

Conditions

This erratum occurs under the following conditions:

1. The system sends a read transaction to the ACP interface. This can be either:
   - A read of 32 bytes, to a 64 byte aligned address
   - A read of 64 bytes, to any supported address alignment
2. An uncorrectable ECC error occurs, which causes the affected data to be marked as poisoned. The error could occur in the L3 data RAMs or LTDB RAMs in the DSU, in the L1 or L2 caches in the cores, or in any cache in the rest of the system if a CHI interface is configured and the system supports returning poisoned data on that interface.
3. The poison affects only some of the cache line, the remaining data in the cache line is not poisoned

Implications

If the erratum occurs, then an OKAY response will be sent for read data beats that might contain corrupted data. At least one other beat in the transaction will be sent with a SLVERR response, but this might be a subsequent beat in the transaction. If the requester consumes only part of the transaction, then it might silently consume the poisoned read data beat.

If the source of the poison is an uncorrectable error in the L3, or the L1 or L2 cache of a core, then it would have been reported as Deferred Error in the RAS error record of the source. Any system cache is likely to have a similar record of the source of the error.
There is still substantial benefit being gained from the ECC logic. This erratum might cause a negligible increase in overall system failure rate.

Workaround

No workaround is required.

3276628
MPAM out of range PartID might not be reported

Status

Fault Type: Programmer Category C
Fault Status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

When programming the *Memory System Resource Partitioning and Monitoring* (MPAM) registers in the DSU, register writes with an out of range PartID might not be reported.

Configurations Affected

This erratum affects all configurations, except Direct connect.

Conditions

The erratum occurs if all the following conditions apply:

1. There is a Utility Bus write to the MPAMCFG_PART_SEL register with an out of range PartID.
2. There is a Utility Bus write to either the MPAMCFG_MBW_PROP register or the MPAMCFG_CPBM register.

Implications

During the second write, the DSU should record the out of range PartID as a PARTID_SEL_Range error in the MPAMF_ESR register. For writes to MPAMCFG_MBW_PROP, the error is not recorded. For writes to MPAMCFG_CPBM, only the in-range LSBs of the PartID are recorded in the MPAMF_ESR.PARTID_MON field.

Workaround

No workaround is required because software should determine the number of supported PartIDs before writing to these registers.

3654577
PPU in static mode with core in FULL_RET might deadlock

Status

Fault type: Programmer Category C
Fault status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

A core in the FULL_RET power mode while the *Power Policy Unit* (PPU) is in static mode might deadlock.

Configurations affected

This erratum affects all DSU configurations except Direct connect.

Conditions

The erratum occurs under the following conditions:

1. The IMP_CPUPWRCTLR_EL1.WFI_RET_CTLR or WFE_RET_CTLR field is set to a non-zero value to enable retention.
2. The core executes a WFI or WFE instruction and is idle for long enough so that it is ready to enter the FULL_RET power mode.
3. The core PPU register PPU_PWPR.PWR_DYN_EN is 0, and PPU_PWPR.PWR_POLICY is set to FULL_RET (0b0101).
4. A Utility Bus transaction is sent to a different core in the cluster.
5. A combination of other traffic is sent to the core that is in FULL_RET. This traffic can include debug APB transactions, transactions from the GIC, and must meet configuration dependent microarchitectural timing conditions.

Implications

If the previous conditions occur, the core in FULL_RET will attempt to transition to ON, however the Utility Bus transaction will not complete and this will prevent the system from reacting to the power transition request as it will not be able to program the PPU over the Utility Bus.

The DSU documentation strongly recommends that the retention states are only used when the PPUs are in dynamic mode (PPU_PWPR.PWR_DYN_EN is 1) and therefore it is not expected that systems will meet these conditions.

Workaround

No workaround is needed for typical use cases.

3674125
Reserved operating modes might cause deadlock

Status

Fault type: Programmer Category C
Fault status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

When reserved operating modes are programmed in the cluster PPU_PWPR.OP_POLICY field, power transitions might deadlock.

Configurations affected

This erratum affects all configurations where L3_CACHE is TRUE.

Conditions

This erratum occurs under the following conditions:

1. The PWR_POLICY field in the PPU_PWPR cluster PPU power policy register is programmed for one of the following modes:
   - ON
   - MEM_RET
   - MEM_RET_EMU
   - FUNC_RET
   - FULL_RET
2. The OP_POLICY in the cluster PPU_PWPR register is set to 0xC, 0xD, or 0xF (all reserved values).

Implications

The cluster power transition might deadlock when these reserved operating modes are programmed. These reserved modes should not be used by software.

Workaround

No workaround is necessary.

3825768
Incorrect Aff0 field in DEVAFF registers

Status

Fault type: Programmer Category C
Fault status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

The memory-mapped cluster registers PMDEVAFF and ERRDEVAFF indicate the relationship between core affinities in MPIDR_EL1 and the cluster PMU and RAS registers, respectively. The value of the read-only field Aff0 in these registers is incorrect.

Configurations Affected

This erratum affects all configurations except Direct connect configurations.

Conditions

The erratum occurs under the following conditions:

- There is an APB read to PMDEVAFF, or a Utility bus read to ERRDEVAFF

Implications

The Aff0 field in these registers reads as 0x80, but it should be 0x0.

The value 0x80 indicates that the PE affinity is at level 1, but this contradicts the Aff1 field, which has a value 0x80 and so indicates the PE affinity is at level 2. This might cause software to incorrectly interpret which core affinities the PMU and RAS registers correspond to.

Workaround

There is no workaround.

3918694
APB access to core, or cluster register might receive incorrect data

Status

Fault type: Programmer Category C
Fault status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

An access on the debug APB interface that reads a core, or cluster register might receive incorrect data. This can occur if the previous access from external APB interface was to an ELA register during quiescence of external SCLK Q-Channel.

Configurations affected

This erratum affects configurations when ELA is present in the cluster.

Conditions

The erratum occurs under the following conditions:

1. The ELADISABLE configuration signals is not set.
2. The DSU is idle for a period and so the system starts to request to gate SCLK using the SCLK Q-Channel.
3. During the SCLK Q-Channel quiescence request, there is a debug APB interface access to the cluster ELA.
4. External APB interface accesses core, or cluster registers.

Implications

If the previous conditions occur, the external APB interface accessing core, or cluster register might receive incorrect data due to the previous ELA register access.

Workaround

For production devices ELADISABLE should be set and therefore no workaround is required.

For sample silicon, this erratum can be avoided if external SCLK Q-channel is not quiescent when accessing ELA registers. This can be done by programming IMP_CLUSTERACTLR_EL1[16:15] or memory mapped CLUSTERACTLR[16:15] to 2'b11 while the ELA is in use. This will increase the idle power consumption of the DSU during this time.

3933469
Core in a complex transitioning to OFF or OFF_EMU might deadlock

Status

Fault type: Programmer Category C
Fault status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

A core in a complex that is transitioning to OFF or OFF_EMU might deadlock when a previous power transition to OFF or OFF_EMU is denied. This can happen when the last powered ON core in a complex accepts OFF or OFF_EMU but the shared power transition to OFF or OFF_EMU is denied.

Configurations affected

This erratum affects all configurations where the cluster contains a core that is part of a complex.

Conditions

The erratum occurs under the following conditions:

1. Only one core of the complex is powered ON.
2. The core that is in the ON state has ERRCTLR.ED set. Also at least one of the below controls for handling interrupts is enabled:
   - ERRCTLR.FI
   - ERRCTLR.CFI
   - ERRCTLR.UI
   - ERRCTLR.DUI
3. The core has IMP_CPUPWRCTLR_EL1.CORE_PWRDN_EN set and executes a WFI instruction.
4. The core power transition from ON to OFF or OFF_EMU is denied when a RAS error is seen in the L2 RAMs.
5. Software clears the ERRCTLR register fields that were set in condition 2, or writes to the ERRSTATUS.{V/UE/DE} fields to clear them.
6. The core power transition from ON to OFF or OFF_EMU is started again.

Implications

If the previous conditions occur, the core might deadlock the second power transition to OFF or OFF_EMU. The erratum also only occurs if an ECC error occurs. There is still substantial benefit being gained from the ECC logic. This erratum might cause a negligible increase in overall system failure rate.

Workaround

No workaround is required.

3968468
Incorrect fields in MPAMF_IDR register

Status

Fault type: Programmer Category C
Fault status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

The memory-mapped MPAM register MPAMF_IDR indicates which memory partitioning and monitoring features are present. The value of the reserved offset [37:36] in this register is incorrect: MPAMF_IDR.NO_IMPL_MSMON and MPAMF_IDR.NO_IMPL_PART should both read 0, as MPAMF_IDR.HAS_IMPL_IDR has a value of 0.

Configurations affected

This erratum affects all configurations, except Direct connect configurations.

Conditions

The erratum occurs under the following condition:

- There is an Utility Bus read to MPAMF_IDR register

Implications

Software is expected to read MPAMF_IDR.HAS_IMPL_IDR before reading the NO_IMPL_MSMON or NO_IMPL_PART fields, therefore there are no implications if software uses the value of the HAS_IMPL_IDR field.

Workaround

No workaround is needed.

3968472
Incorrect field in PMDEVARCH register

Status

Fault type: Programmer Category C
Fault status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

The memory-mapped PMU register PMDEVARCH identifies the programmers' model architecture of the Performance Monitor component. The field ARCHID should be 0x0AF0 (CoreSight PMU architecture, revision 0) instead of 0x2A16 (PE PMU architecture, revision 2).

Configurations affected

This erratum affects all configurations except Direct connect configurations.

Conditions

The erratum occurs under the following condition:

- There is an Utility Bus read to PMDEVARCH register

Implications

The PMDEVARCH ARCHID field (bits [15:0]) reads as 0x2A16 but it should be 0x0AF0.

This might cause software to incorrectly assume it is running on a non-architectural version of the PMU.

Workaround

The value of PMIIDR can be used by software to identify the DSU PMU. The ARCHID field can then be ignored.

3968477
DSU uses default MPAM value with wrong MPAMNS

Status

Fault type: Programmer Category C
Fault status: Present in r0p0, r1p0, r2p0 and r2p1. Open.

Description

The DSU uses the default MPAM value in situations where the true MPAM value is not known. The default MPAM value is defined in the architecture as having PartID set to zero and MPAMNS set to match the *Physical Address Space* (PAS) of the transaction. During slice power transitions, the DSU might incorrectly send the default MPAM value with MPAMNS set to Secure for transactions that are to Non-Secure memory.

Configurations affected

This erratum affects all configurations with a CHI port except Direct connect configurations.

Conditions

The erratum occurs under the following conditions:

- The DSU starts a power transition to change the number of slices, that is between any of ALL_SLICES, HALF_SLICE or ONE_SLICE.

Implications

The DSU power transition will cause a flush and some of these transactions will use the default MPAM value.

If the erratum occurs, the DSU might incorrectly set MPAMNS to Secure for transactions that are to Non-Secure memory. This might cause the system to use the wrong performance settings for the transaction, which might have a minor performance impact. MPAM monitors for the Secure world will over-count for PartID zero. MPAM monitors for the Non-Secure world will under-count for PartID zero.

Workaround

No workaround is required.

3968481
DSU might not report data loss due to bus error

Status

Fault type: Programmer Category C
Fault status: Present in r0p0, r1p0, r2p0, and r2p1. Open.

Description

During a read on an external CHI interface, if there is a bus error, then the DSU might silently discard dirty data. The DSU does not report this, even though it might cause data loss and a loss of coherency.

Configurations affected

This erratum affects configurations with at least one CHI interface.

Direct connect configurations are not affected.

Conditions

The erratum occurs under either of the following sets of conditions:
Set of conditions A:

- The DSU sends an allocating read transaction to the CHI system interconnect
- The system responds with CompData_UD_PD, DataSepResp_UD_PD, or Comp_UD_PD
- At least one data packet from the interconnect has RespErr set to Data Error (DERR)

Set of conditions B:

- The line is already cached in the DSU in the *Unique Dirty* (UD) state
- The DSU sends a read transaction with TagOp set to Transfer or Fetch to the CHI system interconnect to get the *Memory Tagging Extension* (MTE) Allocation Tags
- At least one data packet from the interconnect has RespErr set to Non-Data Error (NDERR)

Implications

This erratum can only occur if the DSU receives an error response from the system. If the erratum occurs, then the DSU might drop the dirty data without reporting an Uncontainable error in the RAS error record registers. Dropping the dirty data might cause data loss and a loss of coherency.

For set of conditions A:

- If the DSU is configured with SCU_CACHE_PROTECTION, the DSU CHI data channels have a

Poison field. Arm expects that most systems that support deferring uncorrected ECC errors will use the Poison field rather than DERR. These systems can not cause this set of conditions.
- Arm CMN and CI interconnects cannot generate these conditions when MXP_DEV_POISON_EN is set. Arm SI products do not currently support poisoned data.
- Arm expects that most systems that do not support using Poison for data errors should treat such errors as uncontainable. For systems that do this, this erratum would have no impact when these conditions occur. For systems that use DERR to defer errors, there will be a negligible increase in the *Failures In Time* (FIT) rate.

For set of conditions B:

- These conditions can only occur if software has enabled MTE. Many systems will not generate NDERR if a previous transaction has successfully given the DSU Unique ownership. These systems are not affected by this set of conditions.
- Arm interconnects will not generate these conditions themselves, but might pass through such responses from the rest of the system.

Workaround

No workaround is required.

Proprietary notice

(PRE-1121-V1.0)

Product and document information

Read the information in these sections to understand the release status of the product and documentation, and the conventions used in the Arm documents.

Product status

All products and Services provided by Arm require deliverables to be prepared and made available at different levels of completeness. The information in this document indicates the appropriate level of completeness for the associated deliverables.

Product completeness status

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

Product revision status

The rxpy identifier indicates the revision status of the product described in this manual, where:

rx

Identifies the major revision of the product.

py

Identifies the minor revision or modification status of the product.