Arm DynamIQ Shared Unit 120 MP147

Software Developer Errata Notice

Date of issue: 09-Jun-2023

Non-Confidential

Document version: 4.1

Document ID: SDEN-2453103

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

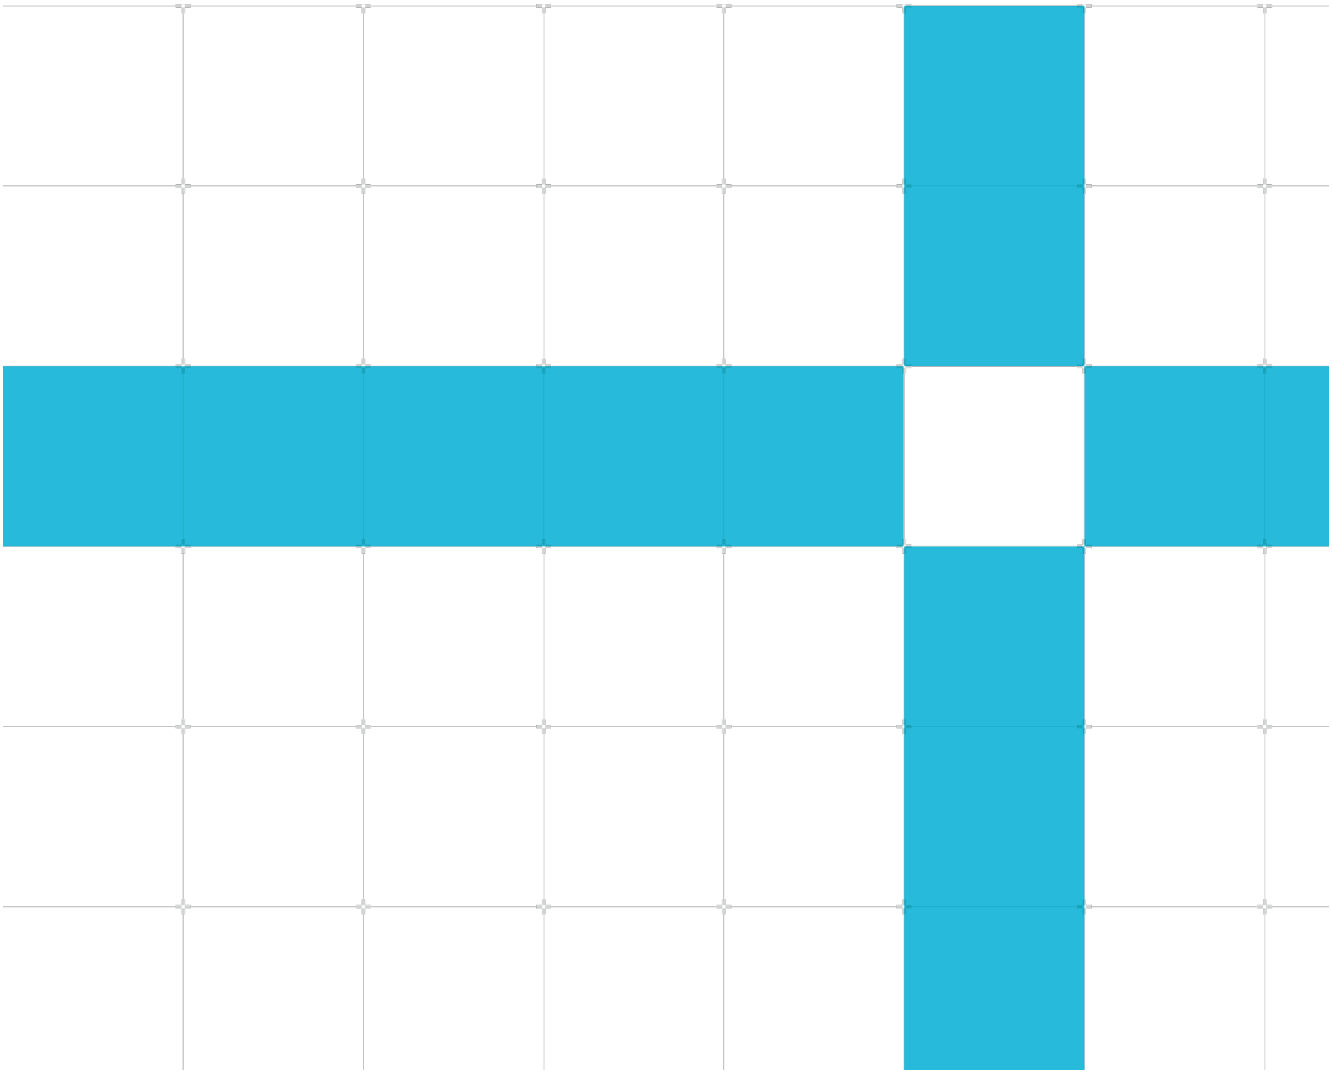

Non-confidential proprietary notice

Arm Limited. Company 02557590 registered in England.

110 Fulbourn Road, Cambridge, England CB1 9NJ.

(LES-PRE-20349)

Confidentiality status

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

Unrestricted Access is an Arm internal classification.

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**.

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**.

Contents

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.

09-Jun-2023: Changes in document version v4.1

No new or updated errata in this document version.

24-Feb-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

07-Dec-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

26-Jul-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

28-Feb-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
2800804	Programmer	Category B	Power transitions between All slices mode and One slice mode might deadlock	r1p0	Open
2745150	Programmer	Category B	Power transitions between All slices mode and One slice mode might lose coherency	r1p0	Open
2634577	Programmer	Category B	ATCLK gating might prevent powerdown	r0p0	r1p0
2777645	Programmer	Category C	Trace flush may not flush all data in transport	r0p0, r1p0	Open
2700719	Programmer	Category C	Clearing RAS interrupt during power down might cause deadlock	r0p0	r1p0
2679678	Programmer	Category C	Cache debug access might deadlock	r0p0	r1p0
2667776	Programmer	Category C	RAS error during power down might be lost	r0p0	r1p0
2661093	Programmer	Category C	CLUSTERL3HIT and CLUSTERL3MISS registers not saturating on overflow	r0p0	r1p0
2457823	Programmer	Category C	CLUSTERPMSSRR bits cannot be cleared	r0p0	r1p0

Errata descriptions

Category A

There are no errata in this category.

Category A (rare)

There are no errata in this category.

Category B

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

Status

Fault Type: Programmer Category B
Fault Status: Present in r1p0. Open.

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.

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. Open.

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.

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.

Category B (rare)

There are no errata in this category.

Category C

2777645
Trace flush may not flush all data in transport

Status

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

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.

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.

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.

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.

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.