



# Intel® Stratix® 10 DX Device Errata



## Contents

---

<b>1. About this Document.....</b>	<b>3</b>
<b>2. Device Errata Summary.....</b>	<b>4</b>
2.1. FPGA.....	5
2.1.1. P-Tile.....	5
2.2. Hard Processor System.....	12
2.2.1. HPS Stops on the First Read Request to SDRAM.....	12
<b>3. Arm* Cortex-A53 MPCore* and CoreSight* Errata.....</b>	<b>13</b>
3.1. ARM Cortex-A53 MPU and CoreSight Errata.....	15
3.1.1. 843819: Memory Locations May be Accessed Speculatively Due to Instruction Fetches When HCR.VM is Set.....	15
3.1.2. 845719: A Load May Read Incorrect Data.....	16
3.1.3. 855871: ETM Does Not Report IDLE State When Disabled Using OSLOCK.....	17
3.1.4. 855872: A Store-Exclusive Instruction May Pass When it Should Fail.....	18
3.1.5. 711668: Configuration Extension Register Has Wrong Value Status.....	20
3.1.6. 720107: Periodic Synchronization Can Be Delayed and Cause Overflow.....	21
3.1.7. 855873: An Eviction Might Overtake a Cache Clean Operation.....	23
3.1.8. 853172: ETM May Assert AFREADY Before All Data Has Been Output.....	24
3.1.9. 836870: Non-Allocating Reads May Prevent a Store Exclusive From Passing.....	25
3.1.10. 836919: Write of the JMCR in EL0 Does Not Generate an UNDEFINED Exception.....	27
3.1.11. 845819: Instruction Sequences Containing AES Instructions May Produce Incorrect Results.....	28
3.1.12. 851672: ETM May Trace an Incorrect Exception Address.....	29
3.1.13. 851871: ETM May Lose Counter Events While Entering WFX Mode.....	30
3.1.14. 852071: Direct Branch Instructions Executed Before a Trace Flush May be Output in an Atom Packet After Flush Acknowledgment.....	31
3.1.15. 852521: A64 Unconditional Branch May Jump to Incorrect Address.....	32
3.1.16. 855827: PMU Counter Values May Be Inaccurate When Monitoring Certain Events.....	33
3.1.17. 855829: Reads of PMEVCNTR<n> are not Masked by HDCR.HPMN.....	35
3.1.18. 855830: Loads of Mismatched Size May not be Single-Copy Atomic.....	36
<b>4. Intel Stratix 10 DX Device Errata Revision History.....</b>	<b>37</b>



## 1. About this Document

---

This document lists known device errata for Intel® Stratix® 10 DX production devices.

## 2. Device Errata Summary

This section lists the errata that apply to the Intel Stratix 10 DX devices. Each listed erratum has an associated status that identifies any planned fixes.

**Table 1. Device Issues**

Issue	Affected Devices	Planned Fix
<b>FPGA</b>		
<b>P-Tile</b>		
<a href="#">Root Port Legacy Interrupt Status register INTx is stuck HIGH on page 5</a>	<ul style="list-style-type: none"> <li>Intel Stratix 10 DX 2800</li> <li>Intel Stratix 10 DX 2100</li> <li>Intel Stratix 10 DX 1100</li> </ul>	No planned fix
<a href="#">TLP Bypass Error Status Register may Report Receiver Errors on page 6</a>	<ul style="list-style-type: none"> <li>Intel Stratix 10 DX 2800</li> <li>Intel Stratix 10 DX 2100</li> <li>Intel Stratix 10 DX 1100</li> </ul>	No planned fix
<a href="#">PCIe CV and PTC Tests in the PCI-SIG Compliance Test Suite may Fail on page 7</a>	<ul style="list-style-type: none"> <li>Intel Stratix 10 DX 2800</li> <li>Intel Stratix 10 DX 2100</li> <li>Intel Stratix 10 DX 1100</li> </ul>	Software fix planned in the future version of the Intel Quartus® Prime Pro Edition
<a href="#">Returning Incorrect Function Number on page 8</a>	<ul style="list-style-type: none"> <li>Intel Stratix 10 DX 2800</li> <li>Intel Stratix 10 DX 2100</li> <li>Intel Stratix 10 DX 1100</li> </ul>	No planned fix
<a href="#">Incorrect Return Value for Power Management Register on page 9</a>	<ul style="list-style-type: none"> <li>Intel Stratix 10 DX 2800</li> <li>Intel Stratix 10 DX 2100</li> <li>Intel Stratix 10 DX 1100</li> </ul>	Software fix planned in the Intel Quartus Prime Pro Edition version 20.1
<a href="#">Register Implementation while using the SR-IOV Feature on page 10</a>	<ul style="list-style-type: none"> <li>Intel Stratix 10 DX 2800</li> <li>Intel Stratix 10 DX 2100</li> <li>Intel Stratix 10 DX 1100</li> </ul>	No planned fix
<a href="#">Register Implementation while using the Multi-function Feature on page 10</a>	<ul style="list-style-type: none"> <li>Intel Stratix 10 DX 2800</li> <li>Intel Stratix 10 DX 2100</li> <li>Intel Stratix 10 DX 1100</li> </ul>	No planned fix
<b>Hard Processor System</b>		
<a href="#">HPS Stops on the First Read Request to SDRAM on page 12</a>	<ul style="list-style-type: none"> <li>Intel Stratix 10 DX 1100</li> </ul>	No planned fix



## 2.1. FPGA

### 2.1.1. P-Tile

#### 2.1.1.1. Root Port Legacy Interrupt Status register INTx is stuck HIGH

##### Description

During root port implementation in the Intel FPGA P-Tile Hard IP for PCIe, the Root Port Legacy Interrupt status register `INT_status` of the port configuration and status register space (address: `0x10414C[3:0]`) is stuck HIGH (it can be any bit of `INT_status`) when the `AssertINTx` message is received but the `DeassertINTx` is not received. This may occur when a warm reset or `PERST` is issued before the endpoint sends the `DeassertINTx` message. As a result, the root port interrupt handling does not operate correctly.

Subsequent warm reset or `PERST` does not clear the stuck interrupt bits.

##### Impacted Modes

- Intel FPGA P-Tile Avalon® Streaming Interface Hard IP for PCIe in Root Port mode

##### Workaround

When using the Intel FPGA P-Tile Hard IP for PCIe in root port mode, Intel recommends you to use the MSI/MSI-X interrupt messages.

To avoid this event when using the Root Port Legacy Interrupt, ensure that an `AssertINTx` message is followed by a `DeassertINTx` message before a warm reset or `PERST` is issued. Otherwise, to clear the stuck interrupt bits, you must re-program the device.

##### Status

Affects:

- Intel Stratix 10 DX 2800
- Intel Stratix 10 DX 2100
- Intel Stratix 10 DX 1100

Status: No planned fix.



### 2.1.1.2. TLP Bypass Error Status Register may Report Receiver Errors

#### Description

During the TLP Bypass implementation using the Intel FPGA P-Tile Avalon Streaming Interface Hard IP for PCIe, the `tlpbypass_err_status` register of port configuration and status registers (address: `0x104190[8]`) may report receiver errors after the `PERST` is released. Therefore, if the user logic implements the Advanced Error Reporting (AER) capability based on the `tlpbypass_err_status` register, the correctable error status register of the AER capability indicates receiver errors.

#### Impacted Modes

- Intel FPGA P-Tile Avalon Streaming Interface Hard IP for PCIe in TLP Bypass mode

#### Workaround

While using the Intel FPGA P-Tile Avalon Streaming Interface Hard IP for PCIe, the user logic must clear the `tlpbypass_err_status` register's receiver error status bit (`0x104190[8]`) of the port configuration and status registers before the `PERST` is released.

#### Status

Affects:

- Intel Stratix 10 DX 2800
- Intel Stratix 10 DX 2100
- Intel Stratix 10 DX 1100

Status: No planned fix.



### 2.1.1.3. PCIe CV and PTC Tests in the PCI-SIG Compliance Test Suite may Fail

#### Description

While using the Intel FPGA P-Tile Hard IP for PCIe, the following PCIe Configuration Verifier (CV) and Protocol Test Card (PTC) tests in the PCI-SIG Compliance Test Suite may fail:

- TD\_1\_4
- TD\_1\_7
- TD\_1\_13
- TD\_1\_27
- TD\_1\_73
- Link Layer Transaction test (Data Link Feature Exchange)

*Note:* The current Intel Quartus Prime Pro Edition software version does not enable every feature that is required to demonstrate compliance.

#### Impacted Modes

- Intel FPGA P-Tile Avalon Streaming Interface Hard IP for PCIe in Endpoint mode
- Intel FPGA P-Tile Avalon Memory-Mapped Interface IP for PCIe in Endpoint mode

#### Workaround

None.

#### Status

Affects:

- Intel Stratix 10 DX 2800
- Intel Stratix 10 DX 2100
- Intel Stratix 10 DX 1100

Status: Software fix planned in the future version of the Intel Quartus Prime Pro Edition.



#### 2.1.1.4. Returning Incorrect Function Number

##### Description

While using the Intel FPGA P-Tile Hard IP for PCIe, the last Physical Function (PF) indicates incorrect next function number if the following conditions are true:

- Alternative Routing-ID Interpretation (ARI) capability is enabled
- Number of Physical Functions (PFs) is less than 8

The next function number incorrectly shows a value of PF+1. As a result, the correctable error status register in the Advanced Error Reporting (AER) capability structure of the root port may report a non-fatal error.

##### Impacted Modes

- Intel FPGA P-Tile Avalon Streaming Interface Hard IP for PCIe in Endpoint mode
- Intel FPGA P-Tile Avalon Memory-Mapped Interface IP for PCIe in Endpoint mode

##### Workaround

To avoid this issue, the user logic must use the Configuration Intercept Interface (CII) to override the read data for the next function number when the above conditions are true.

The user logic must clear the correctable error status register of the AER capability using the Hard IP Reconfiguration interface.

##### Status

Affects:

- Intel Stratix 10 DX 2800
- Intel Stratix 10 DX 2100
- Intel Stratix 10 DX 1100

Status: No planned fix.



### 2.1.1.5. Incorrect Return Value for Power Management Register

#### Description

While using the Intel FPGA P-Tile Hard IP for PCIe, the Power Management Control and Status Register (offset 0x04Ch bit [3]: No\_Soft\_Reset) returns a value of 1'b0 instead of value of 1'b1. As a result, the software expects the Intel FPGA P-Tile Hard IP for PCIe to reset the configuration space during transition from **D3<sub>HOT</sub>** state to **D0** state.

#### Impacted Modes

- Intel FPGA P-Tile Avalon Streaming Interface Hard IP for PCIe in Endpoint, Root Port, and TLP Bypass modes
- Intel FPGA P-Tile Avalon Memory-Mapped Interface IP for PCIe in Endpoint and Root Port modes

#### Workaround

None.

#### Status

Affects:

- Intel Stratix 10 DX 2800
- Intel Stratix 10 DX 2100
- Intel Stratix 10 DX 1100

Status: Fixed in the Intel Quartus Prime Pro Edition software version 20.1.

### 2.1.1.6. Register Implementation while using the SR-IOV Feature

#### Description

While using the Intel FPGA P-Tile Hard IP for PCIe:

- When the SR-IOV feature is enabled:
  - The PCIe capability link status register (offset 0x082h bits [15:0]) returns the parent PF link status register values when accessed by virtual functions. According to the PCIe Base specification revision 4.0 version 1.0, when SR-IOV feature is enabled, this register must be implemented as Reserved and Zero (**RsvdZ**) and accessed by virtual functions.
  - The PCIe device control2 register (offset 0x098h bit [12]: 10-bit tag requested enable bit) returns the parent PF device control2 register values when accessed by virtual functions. According to the PCIe Base specification revision 4.0 version 1.0, when SR-IOV feature is enabled, this register must be implemented as Reserved and Preserved (**RsvdP**) and accessed by virtual functions.
  - The PCIe Address Translation Service (ATS) control register (ATS base address + offset 0x006h bit [4:0]: smallest translation unit (STU) bits) is implemented as read-only (**RO**) register but returns the value of the parent PF when accessed by virtual functions. According to the PCIe Base specification revision 4.0 version 1.0, this register must be implemented as a read-only (**RO**) register and hard-wired to 0 when accessed by virtual functions.

These issues do not cause functional failure.

#### Impacted Modes

- Intel FPGA P-Tile Avalon Streaming Interface Hard IP for PCIe in Endpoint mode
- Intel FPGA P-Tile Avalon Memory-Mapped Interface IP for PCIe in Endpoint mode

#### Workaround

The application logic can use Configuration Intercept Interface (CII) or Direct User Avalon Memory-Mapped Interface to modify the configuration accesses to this register.

#### Status

Affects:

- Intel Stratix 10 DX 2800
- Intel Stratix 10 DX 2100
- Intel Stratix 10 DX 1100

Status: No planned fix.

### 2.1.1.7. Register Implementation while using the Multi-function Feature

#### Description

While using the Intel FPGA P-Tile Hard IP for PCIe:



- When the multi-function feature is enabled, the PCIe device status register (offset 0x07Ah bit [5]: Transactions pending bit) for each of the virtual functions (VF) is implemented as a Write-1-to-Clear status register (**RW1C**). According to the PCIe Base specification revision 4.0 version 1.0, this register must be implemented as read-only (**RO**) when multi-function feature is enabled.

These issues do not cause functional failure.

### Impacted Modes

- Intel FPGA P-Tile Avalon Streaming Interface Hard IP for PCIe in Endpoint mode
- Intel FPGA P-Tile Avalon Memory-Mapped Interface IP for PCIe in Endpoint mode


### Workaround

The application logic can use Configuration Intercept Interface (CII) or Direct User Avalon Memory-Mapped Interface to modify the configuration accesses to this register.

#### Using the Direct User Avalon Memory-Mapped Interface:

The application logic must implement a tracking logic for any pending upstream memory read (MRd) completion. After the last pending MRd completion is received, refer to the [Using Direct User Avalon Memory-Mapped Interface \(Byte Access\)](#) to clear the Transactions pending bit in the Device Status register. The following sequence is an example for VF3 in PF0.

- Application logic programs the User Avalon memory-mapped interface Port Configuration Register (Offset 0x10406A, addressing the third byte of the register) with 0x0A ( $vf\_num[28:18] = 2$ ,  $vf\_select[17] = 1$ ,  $vsec[0] = 0$ ).
- Application logic sets the `hip_reconfig_addr_i[20:0]` with 0x7A which corresponds to the Device Status register within the VF PCI Express Capability Structure<sup>(1)</sup> and performs a write 1 operation to the Transaction pending bit [5] by setting the `hip_reconfig_writedata_i[7:0]` to 0x20.



### Status

Affects:

- Intel Stratix 10 DX 2800
- Intel Stratix 10 DX 2100
- Intel Stratix 10 DX 1100

Status: No planned fix.

(1) For more details, refer to the [PCIe Configuration Registers for Each Virtual Function](#)



## 2.2. Hard Processor System

### 2.2.1. HPS Stops on the First Read Request to SDRAM

#### Description

If the Intel Stratix 10 HPS External Memory Interface (EMIF) is enabled in DDR x32 or x16 mode with or without ECC (IO48 Banks 2M/2N), and the IO48 Bank 2L is being used by the core fabric in specific use cases such as:

- LVDS with Dynamic Phase Alignment (DPA)
- LVDS without DPA and with Clock Phase Alignment (CPA) engaged
- Fabric-EMIF
- PHYLite

then the HPS stops on the first read request to SDRAM, and cannot be recovered.

#### Workaround

When using the Intel Stratix 10 HPS EMIF in DDR x32 or DDR x16 configuration with or without ECC (IO48 Banks 2M, 2N), the IO48 Bank 2L can only be used for LVDS without DPA and without CPA engaged, or GPIO purposes. All other banks have no restrictions and they can be used for LVDS, or Fabric-EMIF, or PHYLite, or GPIO.

#### Status

Affects:

- Intel Stratix 10 DX 1100

Status: No planned fix

#### Related Information

[Intel Stratix 10 High-Speed LVDS I/O User Guide](#)

For more information about LVDS.



### 3. Arm\* Cortex-A53 MPCore\* and CoreSight\* Errata

This section lists the Arm\* Cortex-A53 MPCore\* and CoreSight\* errata.

**Note:** This errata only applies if you are using devices which are enabled with the Hard Processor System (HPS).

Each listed erratum has an associated category number which identifies the degree of the behavior.

The categories are as follows:

- Category 1: Behavior has no workaround and severely restricts the use of the product in all, or the majority of applications, rendering the device unusable.
- Category 2: Behavior contravenes the specified behavior and might limit or severely impair the intended use of the specified features, but does not render the product unusable in all or the majority of applications.
- Category 3: Behavior that was not the originally intended behavior but should not cause any problems in applications.

**Note:** This device only contains category 2 and category 3 errata.

**Table 2. Arm Cortex-A53 MPCore Processor and CoreSight Errata**

Errata Listing	Category Number
<a href="#">843819: Memory Locations May be Accessed Speculatively Due to Instruction Fetches When HCR.VM is Set</a> on page 15	Category 2
<a href="#">845719: A Load May Read Incorrect Data</a> on page 16	Category 2
<a href="#">855871: ETM Does Not Report IDLE State When Disabled Using OSLOCK</a> on page 17	Category 2
<a href="#">855872: A Store-Exclusive Instruction May Pass When it Should Fail</a> on page 18	Category 2
<a href="#">711668: Configuration Extension Register Has Wrong Value Status</a> on page 20	Category 3
<a href="#">720107: Periodic Synchronization Can Be Delayed and Cause Overflow</a> on page 21	Category 3
<a href="#">855873: An Eviction Might Overtake a Cache Clean Operation</a> on page 23	Category 3
<a href="#">853172: ETM May Assert AFREADY Before All Data Has Been Output</a> on page 24	Category 3
<a href="#">836870: Non-Allocating Reads May Prevent a Store Exclusive From Passing</a> on page 25	Category 3
<a href="#">836919: Write of the JMCR in EL0 Does Not Generate an UNDEFINED Exception</a> on page 27	Category 3
<a href="#">845819: Instruction Sequences Containing AES Instructions May Produce Incorrect Results</a> on page 28	Category 3
<a href="#">851672: ETM May Trace an Incorrect Exception Address</a> on page 29	Category 3
<a href="#">851871: ETM May Lose Counter Events While Entering WFX Mode</a> on page 30	Category 3
<i>continued...</i>	

Intel Corporation. All rights reserved. Agilix, Altera, Arria, Cyclone, Enpirion, Intel, the Intel logo, MAX, Nios, Quartus and Stratix words and logos are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Intel warrants performance of its FPGA and semiconductor products to current specifications in accordance with Intel's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

\*Other names and brands may be claimed as the property of others.



### 3. Arm\* Cortex-A53 MPCore\* and CoreSight\* Errata

ES-1075 | 2020.12.23

Errata Listing	Category Number
852071: Direct Branch Instructions Executed Before a Trace Flush May be Output in an Atom Packet After Flush Acknowledgment on page 31	Category 3
852521: A64 Unconditional Branch May Jump to Incorrect Address on page 32	Category 3
855827: PMU Counter Values May Be Inaccurate When Monitoring Certain Events on page 33	Category 3
855829: Reads of PMEVCNTR<n> are not Masked by HDCR.HPMN on page 35	Category 3
855830: Loads of Mismatched Size May not be Single-Copy Atomic on page 36	Category 3



## 3.1. ARM Cortex-A53 MPU and CoreSight Errata

### 3.1.1. 843819: Memory Locations May be Accessed Speculatively Due to Instruction Fetches When HCR.VM is Set

#### Description

The ARMv8 architecture requires that when all associated stages of translation are disabled for the current exception level, software only perform instruction fetches within the same or next translation granule as an instruction which has been or will be fetched due to sequential execution. In the conditions detailed below, this erratum may cause the Cortex-A53 MPCore processor to access other locations speculatively due to instruction fetches.

For this erratum to occur:

- The CPU must be executing at EL3, EL2 or Secure EL1.
- The CPU can be in either AArch32 or AArch64 execution state.
- Address translation is disabled for the current exception level (by clearing the appropriate `SCTLR.M`, `HSCTLR.M` or `SCTLR_ELx.M` bit).
- The `HCR.VM` or `HCR_EL2.VM` bit is set.

#### Impact

If the above conditions are met, then speculative instruction fetches may be made to memory locations not permitted by the architecture.

#### Workaround

Because the `HCR.VM` bit default reset value is low, this situation is most likely to occur in power down code, if EL2 or EL3 software disables address translation before the core is powered down. To work around this erratum, software must clear the `HCR.VM` bit before disabling address translation at EL3, EL2 or Secure EL1.

#### Category

Category 2

### 3.1.2. 845719: A Load May Read Incorrect Data

#### Description

When executing in the AArch32 state at EL0, the Cortex-A53 MPCore processor may read incorrect data if a load is performed to the same offset within a different 4 GB region as a recent previous load in the AArch64 state.

The following sequence must occur for this erratum to be triggered:

**Note:** This sequence requires execution in both AArch32 and AArch64, therefore at least one exception level change is required.

1. At EL0 or EL1, in the AArch32 or AArch64 state, the CPU executes a load, store, preload or data- or instruction-cache maintenance by MVA instruction.
2. At EL0 or EL1 in the AArch64 state, the CPU executes a load with:

```
VA[63:32] != 32'h00000000
```

3. At EL0 in AArch32 state, the CPU executes a load to the same 4 KB region as the instruction in step 1, with VA[31:6] matching the VA from the instruction in step 2.

The erratum does not apply if any of the following occurs between step 1 and step 3 in the sequence:

- A write to any of the following registers:
  - Any SCTLR
  - Any TTBR
  - Any TCR
  - Any MAIR
  - CONTEXTIDR
- A TLB maintenance instruction is executed, followed by a DSB.

This erratum also does not apply if an address translation instruction executes between steps 2 and 3 in the sequence.

#### Impact

If the above conditions are met, data corruption can occur. The load at EL0 can access data written at EL1 for which it does not have read permissions; however, a load in non-secure state cannot access secure data.

#### Workaround

If one of the exception conditions listed after the sequence above is applied between steps 2 and 3, this erratum does not occur. Because there must be an exception return from AArch64 to AArch32 between positions 2 and 3 in the above sequence, the recommended workaround is to insert a write of the CONTEXTIDR\_EL1 register into operating system exception return sequences.

#### Category

Category 2



### 3.1.3. 855871: ETM Does Not Report IDLE State When Disabled Using OSLOCK

#### Description

The OS Lock feature in the Embedded Trace Macrocell (ETM) allows software running on the Cortex-A53 MPCore to disable external debug access and save the register state before powering down the ETM. Software must follow a defined sequence to ensure that the register state is stable and all trace data is output before the system is powered-down. Because of this erratum, when the OS Lock mechanism is used, the ETM never indicates to software that it is safe to be powered-off. This erratum occurs when:

- The ETM is enabled (`TRCPRGCTLR.EN = 1`)
- The OS Lock feature disables the ETM (`TRCOSLAR.OSLK = 1`)

#### Impact

When the OS lock feature is enabled and software polls `TRCSTATR.IDLE`, it remains high even after the ETM is disabled and drained of trace data. If the ETM is already disabled by clearing the `TRCPRGCTLR.EN` bit then a software save and restore sequence behaves correctly.

#### Workaround

To execute the disable sequence properly:

1. Disable the ETM by setting the `TRCOSLAR.OSLK` bit.
2. Save the state of the `TRCPRGCTLR.EN` bit.
3. Clear the `TRCPRGCTLR.EN` bit.

#### Category

Category 2

### 3.1.4. 855872: A Store-Exclusive Instruction May Pass When it Should Fail

#### Description

The Cortex-A53 MPCore processor implements an internal exclusive monitor to manage load-exclusive, store-exclusive, and clear-exclusive instructions. Because of this erratum, a load-exclusive instruction to cacheable memory may set the monitor to the exclusive state when the processor does not have exclusive access to the line. A subsequent store-exclusive instruction may pass when it should fail.

This erratum affects all load-exclusive and store-exclusive instructions, including load-acquire exclusive and store-release exclusive instructions.

The conditions that trigger this erratum are listed below:

- A core executes a store to memory that is marked as both inner-writeback and outer-writeback.
- The store is not a store-exclusive (or a store-release exclusive) or a store-release instruction.
- The store is not followed by a `DMB SY` or `DSB`.
- The store misses in the L1 data cache.
- The store does not trigger a linefill. This requires one or more of the following to be true:
  - The core is in read-allocate mode
  - The memory is marked as no-write-allocate
  - The memory is marked as transient
  - The store is a `STNP` instruction
  - The store is triggered by a `DC ZVA` instruction
- The core starts a linefill to the same address as the store. The linefill is started for one of the following:
  - A `PRFM`, `PLD`, or `PLDW` instruction
  - An automatic data prefetch
  - A pagewalk
- The core executes a load-exclusive (or a load-acquire exclusive) instruction to the same address as the store
- The store data is forwarded to the load-exclusive instruction
- If a core starts a linefill to the same address as the store, the load-exclusive instruction retires before the linefill is serialized.

#### Impact

If the above conditions are met then the Cortex-A53 MPCore processor may set the internal exclusive monitor. This behavior is incorrect because the processor is not guaranteed to have exclusive access to the line. If another master executes a load-exclusive instruction to the same address then both masters may gain access to an exclusive region of code at the same time.



### **Workaround**

The only workaround is to avoid the conditions described above. Disabling read-allocate mode by setting `CPUACTLR.RADIS` to `0x3` degrades write-stream performance. Therefore, the preferred workaround is to use an appropriate store-exclusive (or a store-release exclusive) or a store-release instruction or `DMB` instruction.

### **Category**

Category 2



### 3.1.5. 711668: Configuration Extension Register Has Wrong Value Status

#### Description

The Program Trace Macrocell (PTM) implements several read only registers that provide a mechanism for tools using the PTM to determine which features are present in a specific implementation. The Configuration Code Extension register (0x7A, address 0x1E8) has an incorrect value of 0x000008EA. The correct value is 0x00C019A2.

#### Impact

This erratum has no impact on the generation of trace or the configurations that can be enabled. Tools that read this register in order to determine the capabilities of the PTM detect fewer features than are actually present.

The missing bits include:

- Bit[23] - Return stack implemented.
- Bit[22] - Timestamp Implemented.
- Bit[12] - Reserved, for compatibility with the ETM architecture

Bits[10:3] are incorrect and should indicate 52 external inputs.

#### Workaround

Tools using the PTM-A9 can read the peripheral ID registers (at offsets 0xFE0 to 0xFEC) to determine the version of the PTM-A9 and the features which are implemented, rather than relying on the Configuration Code Extension register.

#### Category

Category 3



### 3.1.6. 720107: Periodic Synchronization Can Be Delayed and Cause Overflow

#### Description

The Program Trace Macrocell (PTM) is required to insert synchronization information into the trace stream at intervals in order to allow a partial trace dump to be decompressed. The synchronization period can be controlled either external to the PTM or by a counter that by default, requests synchronization every 1024 bytes of trace.

Synchronization does not need to be inserted precisely at a regular interval, so allowance is made to delay synchronization if the PTM is required to generate other trace packets, or if there is not sufficient space in the FIFO. To guarantee that some synchronization packets are always inserted regardless of the conditions, a 'forced overflow' mechanism shuts off trace if a new synchronization request occurs before the previous request was satisfied. This forced overflow mechanism is minimally intrusive to the trace stream, but ensures that synchronization is inserted after no more than 2x the requested interval.

Due to this erratum, some specific sequences of instructions prevent the PTM from being able to insert any synchronization into the trace stream while that instruction sequence continues. Typically, there is just a short delay which may not be noticed and the synchronization is inserted once the particular pattern of waypoints changes. It is possible that overflows are generated in the trace regardless of the utilization of the FIFO in the PTM. In this scenario, typically only a single byte of trace (up to 5 waypoints) is lost.

The scenario does not correspond to sustained high rates of trace generation which could genuinely cause the FIFO to become full.

Scenarios that cause this erratum are rare, and are limited to code iterating around a loop many times. The loop would contain several branches and be dependent on memory accesses completing.

This erratum can occur under the following conditions:

1. Tracing is enabled.
2. Branch broadcasting is disabled.
3. Cycle accuracy is disabled.
4. The processor executes tight loops of repeating code, lasting longer than the configured synchronization period.

#### Impact

This erratum reduces the frequency of periodic synchronization and potentially causes trace overflows where some trace is lost. In the case of an overflow, trace following the overflow can be correctly decompressed. This erratum is more noticeable if the periodic synchronization requests are more frequent.



### Workaround

The most appropriate workaround depends on the use-case for the trace:

- If the trace buffer is large enough, consider increasing the synchronization period by setting the `ETMSYNCFR` to a higher value.
- If no trace must be lost (and periodic synchronization does not have to be guaranteed), set `bit[0]` of `ETMAUXCR` to disable the forced overflow function. Synchronization is inserted at the earliest opportunity, dependent on the executed instruction stream.

### Category

Category 3



### 3.1.7. 855873: An Eviction Might Overtake a Cache Clean Operation

#### Description

The Cortex-A53 MPCore processor supports instructions for cache clean operations. To avoid data corruption, the processor must ensure correct ordering between evictions and cache clean operations for the same address.

Because of this erratum, the processor can issue an eviction and an L2 cache clean operation to the interconnect in the wrong order. The processor can also issue transactions that become outstanding in the interconnect at the same time. This situation violates the AXI Coherency Extensions (ACE) protocol specification.

This erratum occurs when the following conditions are met:

1. One or both of the following are true:
  - L2ACTLR[14] is set to 1. This setting allows WriteEvict transactions on the ACE interface when the processor evicts data that it holds in the UniqueClean state.
  - L2ACTLR[3] is set to 0. This setting allows Evict transactions on the ACE interface when the processor evicts clean data.
2. A CPU executes a cache clean-by-address operation for a line that is present and dirty in the L2 cache.
3. A CPU performs any type of memory access to the same set. Memory access types can include a pagewalk, instruction fetch, cache maintenance operation or data access.
4. An instruction fetch to the same set triggers an L2 cache eviction.
5. A present and dirty cache line that is selected for L2 cache eviction at the same time it is targeted for a cache clean operation.

#### Impact

Because of this erratum transactions are erroneously reordered in the interconnect, resulting in data corruption.

#### Workaround

You can workaround this erratum by changing the cache clean-by-address operations to cache clean-and-invalidate operations. To enable these operations, set CPUACTLR.ENDCCASCI to 1.

#### Category

Category 2



### 3.1.8. 853172: ETM May Assert AFREADY Before All Data Has Been Output

#### Description

When the `AFVALID` signal on the Advanced Trace Bus (ATB) interface asserts, the ETM immediately outputs all buffered trace. The ETM must assert the `AFREADY` output one cycle after all trace that was buffered on the cycle in which `AFVALID` was first asserted.

Because of this erratum, the `AFREADY` signal may be asserted before all the necessary trace has been output.

#### Impact

Because of this erratum, the ETM may contain trace data that was generated before the flush request.

#### Workaround

The system can ensure that all trace has been drained from the ETM by disabling it. You can disable the ETM by clearing the `TRCPRGCTLR.EN` bit. Next, system software must poll the `TRCSTATR.IDLE` bit until it reads as 1. This value indicates the ETM is idle and all trace that was generated before the write to `TRCPRGCTLR` has been output.

#### Category

Category 3



### 3.1.9. 836870: Non-Allocating Reads May Prevent a Store Exclusive From Passing

#### Description

A Cortex-A53 MPCore processor executing a load and store exclusive instruction in a loop is allowed to repeatedly fail under certain conditions. An example of an allowed condition is when another processor repeatedly writes to the same cache line.

Because of this erratum, a non-allocating load from another processor may cause the store-exclusive instruction to repeatedly fail.

The failure occurs under these conditions:

- One CPU executes a loop containing a load-exclusive and a store-exclusive instruction. The loop continues until the store-exclusive instruction succeeds.
- Another CPU or master in the system repeatedly executes a load to the same L1 data cache index as the store-exclusive instruction.
- The cache line containing the address of the load is present in the L1 cache of the CPU, and is not present in the L2 cache.
- The load does not cause an allocation into the cache of the CPU or the master executing the load. This response results in a snoop arriving at the CPU holding the cache line every time the load is performed.
- The load is repeated at exactly the same frequency as the load- and store-exclusive loop, such that every time around the loop, the snoop arrives at the same time as the store exclusive instruction is executing.

A load can be non-allocating in Cortex-A53 MPCore processor in the following conditions:

- A load instruction executes on a CPU, while the `CPUACTLR.DTAH` bit is not set, with one of the following conditions:
  - The memory address is marked as writeback cacheable, with no read allocate, in the page tables.
  - The memory address is marked as writeback cacheable, transient, in the page tables.
  - The memory address is marked as writeback cacheable, and an `LDNP` non-temporal load instruction is used.
- A read transaction is made on the ACP interface.

#### Impact

If a CPU or other master is polling a location to determine when the value changes, then it can prevent another CPU from updating that location, causing a software livelock. Most polling routines, however, use memory that can be allocated into their cache. For improved efficiency it is recommended to use a load-exclusive and `WFE` instruction to avoid repeated polling. Because the frequency of the polling loop must match the frequency of the load store-exclusive loop, the conditions that cause this erratum are unlikely. Any other disturbance in the system, such as an interrupt or other bus traffic can easily alter the frequency of either loop sufficiently to break the livelock.



### Workaround

If the repeating load is being executed on a Cortex-A53 CPU, then this erratum can be worked around by setting the `CPUACTLR.DTAH` bit. Note this version of the Cortex-A53 MPCore processor sets this bit by default.

If the repeating load is being executed by another master in the system connected to the ACP interface, then the erratum can be worked around by changing the frequency of the polling so that it no longer aligns with the frequency of the load and store exclusive loop.

If the repeating load is being executed by another master in the system connected through the coherent interconnect, then the erratum can be worked around either by ensuring that the other master allocates the line into its cache, or by changing the frequency of the polling so that it no longer aligns with the frequency of the load and store exclusive loop. Note that the Cortex-A53 MPCore processor always allocates the line into either the cache or a temporary buffer, and so does not require any workaround for typical polling loops.

### Category

Category 3



### 3.1.10. 836919: Write of the JMCR in EL0 Does Not Generate an UNDEFINED Exception

#### Description

When EL0 is using AArch32 register width and a write is performed in EL0 to the Arm Jazelle® Main Configuration Register (JMCR), the write should be UNDEFINED. Because of this erratum, the write is permitted but ignored.

The erratum occurs under the following conditions:

- The processor is executing in AArch32 user mode.
- A write to the JMCR is executed, using the instruction `MCR p14, 7, <Rt>, c2, c0, 0`.

#### Impact

Rather than treating the MCR instruction as UNDEFINED, a write to the JMCR is ignored.

#### Workaround

No available workaround.

#### Category

Category 3

### 3.1.11. 845819: Instruction Sequences Containing AES Instructions May Produce Incorrect Results

#### Description

When the Cortex-A53 MPCore processor is executing in the AArch64 state, certain sequences of instructions that include AES instructions may cause incorrect results.

There are two code sequences that can cause this erratum in the AArch64 execution state:

- Code sequence 1:
  1. The CPU executes an `AESE` instruction.
  2. The CPU executes a `USQADD` instruction.
    - Both the  $V_n$  and  $V_d$  of this instruction must be the same register as  $V_d$  for the `AES` instruction.
    - The size field for this instruction must be '00', indicating byte-sized elements.
    - The `USQADD` instruction can be in either vector or scalar form.
- Code sequence 2:
  1. The CPU executes a `SUQADD` instruction.
    - The size field for this instruction must be 00, indicating byte-sized elements.
    - The `SUQADD` instruction can be in either vector or scalar form.
  2. The CPU executes an `AESMC` or `AESIMC` instruction.
    - Both the  $V_n$  and  $V_d$  of this instruction must be the same register as the  $V_d$  for the `SUQADD` instruction.

For both these sequences, the two instructions listed must be executed consecutively, with no other instructions or exceptions between them. If either of these sequences is met, the final result of the affected instruction sequence may be incorrect.

#### Impact

The sequences of instructions described in the conditions above are not expected to occur in real code because they do not perform useful computation. Therefore, there is no impact expected to real systems.

#### Workaround

Because the code sequences for this erratum are not expected to occur in real code, no workaround is required.

#### Category

Category 3



### 3.1.12. 851672: ETM May Trace an Incorrect Exception Address

#### Description

The address in an exception packet is the preferred exception return address. Because of this erratum, the address may be equal to the target address of the exception. The trace stream is not corrupted, and decompression can continue after the affected packet.

The following sequence is required for this erratum to occur:

1. The ETM must start tracing instructions because of one of the following:
  - a. ViewInst starts.
  - b. The security state changes indicating trace is now permitted.
  - c. The values of the external debug interface (DBGEN, SPIDEN, NIDEN, SPNIDEN) change to indicate trace is now permitted.
  - d. The CPU exits debug mode.
2. Before the CPU executes any other behavior that causes a trace to be generated, it executes a direct branch instruction, which may or may not be taken.
3. The next instruction is a load or store that causes a data abort or watchpoint exception.

After this sequence, provided certain timing specific conditions are met, the address in the exception packet may be incorrect.

#### Impact

The trace decompressor may incorrectly infer execution of many instructions from the branch target to the provided address.

#### Workaround

The trace decompressor can detect that this erratum has occurred by checking if the exception address is in the vector table and identifying if the branch was not expected to be taken to the vector table. A decompressor can infer the correct address of the exception packet. The target of the preceding branch (if the branch was taken), or the next instruction after the branch (if the branch was not-taken) provides the correct address of the exception packet.

#### Category

Category 3

### 3.1.13. 851871: ETM May Lose Counter Events While Entering WfX Mode

#### Description

If the ETM resources become inactive because of a low-power state, there is a one-cycle window during which the counters and the sequencer may ignore counter-at-zero resources.

The following sequence is required for this erratum to occur:

1. The core executes a WFI or WFE instruction.
2. The ETM enters a low-power state.
3. In a one-cycle window around this point, either:
  - a. A counter in self-reload mode generates a counter-at-zero resource.
  - b. A counter in normal mode gets a RLDEVENT on the cycle in which it has just transitioned to zero.
4. A counter or sequencer is sensitive to the counter-at-zero resource.

#### Impact

Counters sensitive to a counter-at-zero resource may not reload or decrement. If the sequencer is sensitive to a counter-at-zero resource, it may not change state, or may change to an incorrect state.

#### Workaround

The ETM can be prevented from entering low-power mode by setting the LPOVERRIDE bit of Trace Event Control 1 (TRCEVENTCTL1R) register. This workaround is only needed if there is a counter or sequencer sensitive to a counter-at-zero resource.

#### Category

Category 3



### 3.1.14. 852071: Direct Branch Instructions Executed Before a Trace Flush May be Output in an Atom Packet After Flush Acknowledgment

#### Description

The Embedded Trace Macrocell (ETMv4) architecture requires that when a trace flush is requested on the AMBA Trace Bus (ATB), a Cortex-A53 MPCore processor must complete any packets that are in the process of being encoded and output them prior to acknowledging the flush request. When trace is enabled, the Cortex-A53 MPCore processor attempts to combine multiple direct branch instructions into a single atom packet. If a direct branch instruction is executed, and an atom packet is in the process of being generated, the processor does not force completion of the packet prior to acknowledging the flush request. This erratum is a violation of the ETMv4 architecture.

The following conditions are required for this erratum to occur:

1. ETM is enabled.
2. Instruction tracing is active.
3. The processor executes one or more direct branch instructions.
4. An atom packet is in the process of being encoded but is not complete.
5. A trace flush is requested on the AMBA ATB.

#### Impact

When the above conditions occur, the ETM must complete encoding and output the atom packet prior to the trace flush request being acknowledged. Because of this erratum, the atom packet is output after the flush is acknowledged. Therefore, it appears to software monitoring the trace that the direct branch was executed after the requested flush.

#### Workaround

Enabling the timestamp by setting the `TS` bit in the Trace Configuration (`TRCCONFIGR`) register resolves the erratum because the atom packets complete when timestamp behavior is enabled.

#### Category

Category 3



### 3.1.15. 852521: A64 Unconditional Branch May Jump to Incorrect Address

#### Description

When executing in AArch64 state with address translation disabled, unconditional immediate branch instructions might jump to an incorrect address.

The following conditions are required for this erratum to occur:

1. The processor is executing in AArch64 state.
2. The `SCTLR_ELx.M` bit for the current exception level is 0.
3. The `HCR_EL2.VM` bit is 0.
4. A `B` or `BL` instruction from the unconditional branch (immediate) encoding class is executed.
5. This branch has an `imm26` field of `0x1FFFFFFF`, encoding a branch target of `{pc} + 0x7FFFFFFC`.

#### Impact

If these conditions are met, then the processor might incorrectly branch to the target `{pc} - 0x8000004` instead of `{pc} + 0x7FFFFFFC`.

#### Workaround

The workaround for this erratum is to avoid the conditions described.

#### Category

Category 3



### 3.1.16. 855827: PMU Counter Values May Be Inaccurate When Monitoring Certain Events

#### Description

The Cortex-A53 MPCore processor implements a Performance Monitor Unit (PMU). The PMU allows programmers to gather statistics on the operation of the processor during runtime. Because of this erratum, software may read inaccurate PMU counter values when monitoring certain events. Specifically:

- The INST\_RETIRE event counts architecturally executed instructions. Because of this erratum, it may count more instructions than were architecturally executed.
- The ST\_RETIRE event does not count store-exclusive instructions that fail. Because of this erratum, it does count these instructions.
- The UNALIGNED\_LDST\_RETIRE event counts loads and stores that fail their alignment check. Because of this erratum, it may also count LDRD and STRD instructions that pass their alignment check.
- The EXC\_TAKEN and EXC\_RETURN events are filtered precisely according to the exception level/security state they were executed in. Because of this erratum, they are filtered according to the destination exception level/security state.

This erratum is present when:

- A performance counter is enabled and configured to count an INST\_RETIRE, ST\_RETIRE, UNALIGNED\_LDST\_RETIRE, EXC\_TAKEN, or EXC\_RETURN event.
- The following conditions occur during the event capture:
  - INST\_RETIRE: An immediate branch instruction is executed.
  - ST\_RETIRE: A store-exclusive instruction is executed and fails.
  - UNALIGNED\_LDST\_RETIRE: A LDRD or STRD instruction is executed that is word aligned but not double word aligned.
  - EXC\_TAKEN: An exception is taken.
  - EXC\_RETURN: An exception return is executed.
- The INST\_RETIRE, ST\_RETIRE, and UNALIGNED\_LDST\_RETIRE filter settings for the performance counter are configured so that an event must be counted if it occurs.
- The EXC\_TAKEN and EXC\_RETURN filter settings are configured to one of the following descriptions:
  - Events must be counted in the original exception level and security state, but must not be counted in the exception level and security state following the exception or the exception return
  - Events must not be counted in the original exception level and security state, but must be counted in the exception level and security state following the exception or the exception return.

#### Impact

If the erratum conditions are met, the performance counter may erroneously increment or fail to increment. Specifically:



- **INST\_RETIRED:** The counter may erroneously increment by two when only one instruction executes.
- **ST\_RETIRED:** The counter may erroneously increment for the failed store-exclusive instruction.
- **UNALIGNED\_LDST\_RETIRED:** The counter erroneously increments for the LDRD or STRD instruction.
- **EXC\_TAKEN:** The counter may erroneously increment or erroneously fail to increment.
- **EXC\_RETURN:** The counter may erroneously increment or erroneously fail to increment.

### Workaround

For the EXC\_TAKEN and EXC\_RETURN events, you can workaround the erratum by changing the filter settings for the performance counter to monitor the destination exception level and security state instead of the exception level and security state in which the exception or exception return are executed.

There are no workarounds for the other PMU events.

### Category

Category 3



### 3.1.17. 855829: Reads of PMEVCNTR<n> are not Masked by HDCR.HPMN

#### Description

Software executing at EL2 can set the `HDCR.HPMN` and `MDCR_EL2.HPMN` field to restrict non-secure EL0 and EL1 software from accessing a subset of the performance counters. If  $n > \text{HPMN}$ , then Performance Counter  $n$  (`PMEVCNTR<n>`) register is not accessible to non-secure EL0 or EL1. Because of this erratum, Cortex-A53 reads of inaccessible registers might not return zero as expected, but instead might read the actual contents of the counter registers.

For this erratum to be present:

1. The processor is executing at non-secure EL1 or EL0.
2. `HDCR.HPMN` and `MDCR_EL2.HPMN` are set to a value  $n$ , where  $n < 6$ .
3. A read of `PMEVCNTR<n>` and `PMEVCNTR<n>_EL0` is executed.

#### Impact

If the conditions described above are met, then non-secure EL0 and EL1 software can read the values of performance counters reserved for EL2.

#### Workaround

Software executing at EL2 can set the `HDCR.TPM` and `MDCR_EL2.TPM` bit. This configuration causes all non-secure EL0 and EL1 accesses to performance monitor registers to trap to EL2. The EL2 software can then emulate accesses as required, and can mask out accesses to reserved registers.

#### Category

Category 3

### 3.1.18. 855830: Loads of Mismatched Size May not be Single-Copy Atomic

#### Description

The Cortex-A53 MPCore processor supports single-copy atomic load and store accesses as described in the ARM architecture documentation. However, in some unusual code sequences, this erratum can cause the CPU executing a store, and later a load, to the same address but with a different access size to load data that does not meet the requirements of a single-copy atomic load.

For this erratum to be present, a store instruction must execute on one CPU. This instruction can be any halfword, word, or double word store instruction that is not a store release and the address must be aligned to the access size. On a second CPU, the following sequence must occur:

1. A store instruction executes. This store must be a smaller access size than the store on the first CPU, and must address bytes accessed by the first CPU. The address must also be aligned to the access size.
2. The store instruction does not allocate into the cache because of any one of the following conditions:
  - The memory address is marked as transient
  - The write allocate hint in the translation table is not set or the memory is marked as non-cacheable
  - The CPU recently executed a stream of stores and has subsequently dynamically switched into a no-write allocate mode
3. A load instruction executes. The load must be a larger access size than the store from the same CPU, and at least some bytes of the load must be to the same address as the store. The address must also be aligned to the access size.

The Arm architecture requires that the load is single-copy atomic. However, in the conditions described, the load may observe a combination of the two stores, indicating that the store on the initial CPU was serialized first. If the load is repeated, it might only see the data from the first CPU's store, indicating that the store on the first CPU was serialized second.

#### Impact

Concurrent, unordered stores are uncommon in multi-threaded code. In the ISO/IEC 9899:2011 (C11) standard, they are restricted to the family of relaxed atomics. Using different size load and store instructions to access the same data is also uncommon. For these reasons, the majority of multi-threaded software does not meet the conditions for this erratum.

#### Workaround

Most multi-threaded software does not satisfy the conditions of this erratum and therefore, does not require a workaround. If a workaround is required, then replace the store on the second CPU with a store-release instruction.

#### Category

Category 3



## 4. Intel Stratix 10 DX Device Errata Revision History

Document Version	Changes
2020.12.23	Erratum: <i>Register Implementation while using the Multi-function and SR-IOV Features</i> is divided into the following two sections: <ul style="list-style-type: none"> <li>• <i>Register Implementation while using the SR-IOV Feature</i></li> <li>• <i>Register Implementation while using the Multi-function Feature</i></li> </ul>
2020.11.16	<ul style="list-style-type: none"> <li>• Added a new erratum: <a href="#">HPS Stops on the First Read Request to SDRAM</a> on page 12.</li> <li>• Added a new section: <a href="#">Arm* Cortex-A53 MPCore* and CoreSight* Errata</a> on page 13.</li> </ul>
2020.05.06	Corrected the offset value for PCIe ATS control register in <a href="#">Register Implementation while using the SR-IOV Feature</a> on page 10.
2020.04.08	Removed the following erratum: <ul style="list-style-type: none"> <li>• <i>Lecroy LTSSM Arc Test may fail</i></li> <li>• <i>Access Control Service Support for Root Port</i></li> </ul>
2020.02.05	Initial release.