

Intel[®] Itanium[®] Processor 9300 Series

Specification Update

December 2010

Notice: The Intel[®] Itanium[®] Processor 9300 Series Processor may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Document number: 323169-006



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® Itanium® processor 9300 series may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.


Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature may be obtained by calling 1-800-548-4725 or by visiting Intel's website at <http://www.intel.com>.

Intel, Itanium, Intel Interconnect BIST (Intel IBIST), Intel Scalable Memory Interconnect (Intel SMI), and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

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

Copyright © 2010, Intel Corporation. All Rights Reserved.




Contents

Revision History	5
Preface	6
Identification Information	7
Summary Table of Changes	8
Errata	13
Documentation Changes and Clarifications	29


§





Revision History

Revision Number	Description	Date
006	Updated Intel® Itanium® Processor 9300 Series Stepping Summary Table	December 2010
005	Added PAL 4.30 and associated errata fixes. Added new erratum 53	December 2010
004	Added new errata 51 and 52	September 2010
003	Added new errata 48 - 50	July 2010
002	Added new errata 40 - 47. PAL 4.29 released fixing errata 38, and 40 - 45. QDF numbers replaced with S-SPEC numbers. Added two datasheet documentation changes	April 2010
001	Initial release of the document	February 2010



Preface

This document is an update to the specifications contained in the Affected/Related Documents table below. This document is a compilation of device and documentation errata, specification clarifications, and specification changes. It is intended for hardware system manufacturers and software developers of applications, operating systems, or tools.

This document may contain information that was not previously published.

Affected/Related Documents


Title	Document #
<i>Intel® Itanium® Processor 9300 Series Datasheet</i>	322821-001
<i>Intel® Itanium® Processor 9300 Series Based Platform Design Guide</i>	404697
<i>Intel® Itanium® Processor 9300 Series Reference Manual for Software Development and Optimization</i>	410879

Nomenclature

Errata are design defects or errors. These may cause the processor's behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.

Documentation Changes and Clarifications are modifications to the current published specifications. These changes will be incorporated in the next release of the specification. They can describe a specification in greater detail or further highlight a specification's impact to a complex design situation. These clarifications will be incorporated in the next release of the specification.

Errata remain in the Specification Update throughout the product's lifecycle, or until a particular stepping is no longer commercially available. Under these circumstances, errata removed from the errata report are archived and available upon request. Specification changes, specification clarifications and documentation changes are removed from the specification update when the appropriate changes are made to the appropriate product specification or user documentation (datasheets, manuals, and so forth).



Identification Information

Intel® Itanium® Processor 9300 Series Stepping Summary

Intel® Itanium® Processor 9300 Series Stepping Summary

S-Spec Number	Processor Number	Processor Stepping Revision	CPUID	Cores	TDP	Core Freq	L3 (MB)	Notes
LBMX	9350	EO	0020020404	4	185	1.73 GHz with turbo up to 1.86 GHz	24	
LBMW	9340	EO	0020020404	4	185	1.60 GHz with turbo up to 1.73 GHz	20	
LBMU	9330	EO	0020020404	4	155	1.46 GHz with turbo up to 1.60 GHz	20	
LBN2	9320	EO	0020020404	4	155	1.33 GHz with turbo up to 1.46 GHz	16	
LBMV	9310	EO	0020020404	2	130	1.60 GHz	10	
LC3A	9350	EO	0020020404	4	185	1.73 GHz with turbo up to 1.86 GHz	24	
LC39	9340	EO	0020020404	4	185	1.60 GHz with turbo up to 1.73 GHz	20	
LC38	9330	EO	0020020404	4	155	1.46 GHz with turbo up to 1.60 GHz	20	
LC3B	9320	EO	0020020404	4	155	1.33 GHz with turbo up to 1.46 GHz	16	
LC37	9310	EO	0020020404	2	130	1.60 GHz	10	



Summary Table of Changes

The following tables indicate the errata, specification changes, and specification clarifications which apply to the processor. Intel may fix some of the errata in a future stepping of the component, and account for the other outstanding issues through documentation or specification changes as noted. These tables use the following notations:

Codes Used in Summary Tables

Stepping


- X: Errata exists in the stepping indicated. Specification Change or Clarification that applies to this stepping.
- (Blank box): This errata is fixed in listed stepping or specification change does not apply to listed stepping.

Status

- Doc: Document change or update will be implemented.
- Plan Fix: This errata may be fixed in a future stepping of the product.
- Fixed: This errata has been previously fixed.
- No Fix: There are no plans to fix this errata.

Row

| Change bar to left of table row indicates this errata is either new or modified from the previous version of the document.



Errata Summary (Sheet 1 of 3)

	Processor Stepping	PAL Version					Status	Errata Title
	E0	4.15	4.25	4.28	4.29	4.30		
1	X						No Fix	Configuration Agent Responds with Poison Error to Non-Aligned Writes with Poison Set
2	X						No Fix	Hold of Incoming PTC.G Pending During PAL-based IA-32 Execution Can Cause Deadlock
3	X						No Fix	Error During Intel QPI Link Initialization Can Cause Hang
4	X						No Fix	Bbox Violates Message Class Dependency
5	X						No Fix	CDEF Memory Region Coherency
6	X						No Fix	Px[n]_CSIWCI Register for Half Link Widths is not Correct After Overwriting
7	X						No Fix	Two Writes Required to Clear CSITTLLECR.[Error Overflow]
8	X						No Fix	CRC Errors With 16 Bit CRC
9	X						No Fix	Reset While In Calibration State Can Cause Hang On Intel SMI
10	X						No Fix	Intel SMI PZ[n]_PBOXFS2.CALIB_DONE Can Only Be Cleared In Reset
11	X						No Fix	Memory Buffer Must Not Be Reset Directly by FPGA or Other System Logic On Warm-State Reset
12	X						No Fix	R_CSR_OPER0.PHYTRAINLIMIT Is Not Asserted When All Clocks Fail
13	X						No Fix	Frame Alert Logged After Warm Logic Reset
14	X						No Fix	Bad Parity In Route Table Can Cause Unexpected Error
15	X						No Fix	Lower 2 Bits Of IHA Have Read/Write Behavior
16	X						No Fix	Read or Write of TAD CSRs While System Is Not Quiesced Can Cause Data Corruption
17	X						No Fix	After An Intel QPI Link Soft Reset CRC Errors May Be Observed
18	X						No Fix	In CLL A Warm-Logic Reset During LOL Event Causes Unexpected Results
19	X						No Fix	In CLL A State Reset May Reenter CLL
20	X						No Fix	Physical Damage To Intel SMI Lane Can Cause Training To Fail
21	X						No Fix	Northbound Intel SMI CRC Persistent Error Can Cause South Bound CRCs Resulting in Fast Reset Loop
22	X						No Fix	CRC Errors Occur on Intel QPI After Physical Layer Reset When Scrambling And Periodic Retraining Are Enabled
23	X						No Fix	Transmitter Parameter Values
24	X						No Fix	L2i Fills In 1 Way Of Set When All Other Ways Are Valid And Way Is Disabled In Set
25	X						No Fix	If A Global Fatal MCA Occurs While A chk.a, chk.s Or fchkf Is Executing A Logical Processors May Not Enter Machine Check Abort and Min-State Save Area May Be Invalid
26	X						No Fix	DC Common Mode Clock At Rx Input For Intel SMI and Intel QPI



Errata Summary (Sheet 2 of 3)

	Processor Stepping	PAL Version					Status	Errata Title
	EO	4.15	4.25	4.28	4.29	4.30		
27	X						No Fix	North Bound Link Errors Followed Immediately By Erasure+1 Errors Can Cause Correctable Error To Be Uncorrectable
28	X						No Fix	Corrupted ALERT Frame Not Detected By Zbox
29	X						No Fix	Z_CSR_ECC_LLE_SUCCESS_[7-0] Does Not Freeze On Error
30	X						No Fix	In Core Level Lockstep A Ubox Request Can Be Dropped Causing Hang
31		X					Fixed	PAL_THREAD_CONTROL Does Not Properly Deconfigure Cores
32		X					Fixed	PAL_SHUTDOWN Has Incorrect Index
33		X					Fixed	PAL_ECC_ERROR_MASK Does Not Check Mask For Determining No Error
34	X						No Fix	Rbox Blocking And Intel QPI Min Credit Occurring At The Same Time With Some Traffic Patterns Can Cause Hang
35	X						No Fix	Persistent CRC Error During Aggressive Pin Throttling Can Cause Hang
36		X	X				Fixed	In A Virtualized Environment Guest OS IIB0/IIB1 Reads/Writes May Not Behave As Expected If Virtualization Acceleration is Enabled
37		X	X				Fixed	A register file error may be signaled on wrong thread
38		X	X	X			Fixed	Neither thread on a core entering SAL_CHECK during a SYSINIT MCA
39		X	X				Fixed	A fault during RSE (Register Stack Engine) load in PAL_CHECK may cause system hang
40		X	X	X			Fixed	Register file error is reported incorrectly by PAL_MC_ERROR_INFO with every ERROR1 assertion
41		X	X	X			Fixed	PAL_REGISTER_INFO incorrectly reports CR.iib0 and CR.iib1 as unimplemented
42		X	X	X			Fixed	Nested MCA may occur during recoverable MCA testing
43		X	X	X			Fixed	Infinite Viral MCA's may occur when testing viral MCA's
44		X	X	X			Fixed	INIT flow will hang if invoked prior to MINSTATE initialization
45		X	X	X			Fixed	An unexpected MCA may occur if multiple thermal error injections are applied to a socket
46		X	X	X	X	X	No Fix	A pended QR CMCI may not be cleared during the OS boot sequence
47	X						No Fix	Firmware Setting Z_CSR_FBD_FAIL_CHIP_STATE_[3-0]. chipFailEnable[3-0] During Erasure Can Cause Uncorrectable ECC Error
48	X						No Fix	CPE Outbound Request Buffer Timeout (CPB_ORB_TIMEOUT_CFG) counters are not deterministically assigned to the correct transaction types
49	X						No Fix	Incorrect PIROM value for THERMALERT_N Hot Deassertion Hysteresis
50		X	X	X	X	X	No Fix	PAL_VP_CREATE does not perform reserved bit checking on the virtualization acceleration control (VAC) and virtualization disable control (VDC) fields in the virtual processor descriptor (VPD)




Errata Summary (Sheet 3 of 3)


	Processor Stepping	PAL Version					Status	Errata Title
	EO	4.15	4.25	4.28	4.29	4.30		
51	X						Fixed ¹	False ERROR1_N signal may be asserted during power cycle
52		X	X	X	X		Fixed	PAL_MC_ERROR_INFO incorrectly returns target address valid when originating from ptc.g with RspStatus=Fail
53		X	X	X	X	X	No Fix	Intel QPI State Machine Does Not Increment On Every Initialization Failure Instance

1. Fixed in S-spec parts numbered LCxx.

Documentation Changes and Clarifications

§





Errata

Note: In this document the Intel® Itanium® processor 9300 series will be referred to as “the processor”.

1. Configuration Agent Responds With Poison Error To Non-Aligned Writes with Poison Set

Problem: When a zero byte request (zero byte error) happens in the same packet with poison for flit 0 set, the configuration agent (Ubox) responds with a poisoned data error. This poisoned data error can be identified by U_CSR_SESR.e[9] set to 1.

Implication: Intel® QuickPath Interconnect agents must never issue non-aligned or zero length writes with poison to the configuration agent (Ubox).

Workaround: None at this time

Status: No Fix

2. Hold of Incoming PTC.G Pending During PAL-based IA-32 Execution Can Cause Deadlock

Problem: A hold of an incoming PTC.G pending during a PAL-based IA-32 execution can cause a deadlock.

Implication: Deadlock conditions in some code sequences involving PAL-based IA-32 execution and PTC.G can occur if workaround is not applied.

Workaround: Contact your Intel technical representative for details on workaround.

Status: No Fix

3. Error During Intel QPI Link Initialization Can Cause Hang

Problem: An error during link initialization can cause an Intel® QuickPath Interconnect link to timeout and lead to a system hang.

Implication: If this condition is encountered a hang during Intel QuickPath Interconnect initialization can occur.

Workaround: Firmware can poll Px[n]_PBOXMSCCTL.NO_RESPONSE_FROM_LL. If firmware finds this bit is asserted it can trigger a reset in the phy layer.

Status: No Fix

4. Bbox Violates Message Class Dependency

Problem: The *Intel® QuickPath Interconnect Specification* does not allow an NCB to DRS dependency. The Ubox has this dependency since the NCB and NCS share the same DRS credits. There is also a DRS to NCB dependency at the home agent in mirroring mode. Logic has been added to the Rbox to prevent this deadlock. Since this logic does not apply across a Node Controller, mirroring across a Node Controller is not supported. This is documented in the *Intel® Itanium® Processor 9300 Series Datasheet*.


Implication: Mirroring across a Node Controller is not supported.

Workaround: None at this time

Status: No Fix

5. CDEF Memory Region Coherency

Problem: Due to a coherency issue with the CDEF region, if this region is enabled SAL must ensure the CDEF region has RdEn = 0 and WrEn = 0 in the SAD I/O Decoder if multiple agents will access this memory region. This will treat the CDEF memory region as coherent.



If SAL sets RdEN = 1 (non-coherent) for the CDEF region, then SAL must ensure only one core ever accesses the CDEF region. SAL is responsible for maintaining cache coherency since the CDEF region is now mapped to MMIO. As no practical use is known for setting WrEN = 1, it should always be left as 0.

Implication: SAL must ensure CDEF region is coherent or ensure coherency of CDEF region.

Workaround: None at this time

Status: No Fix

6. P_x[n]_CSIWCI Register for Half Link Widths is not Correct After Overwriting

Problem: The default value on the P_x[n]_CSIWCI register for Half Link widths is not correct after overwriting. The P_x[n]_CSIWCI register lists the LMs (Lane Maps) supported by a particular QPI link. Half width links (port 4 and 5) have a default value of 0x182. This is correct at startup. When writing an incorrect value to the P_x[n]_CSIWCI register to these ports, the value is not accepted and overwritten by the default value (which is expected). This time it is being written as 0x183. This implies that the full width combination is allowed, and this is not correct. It should always read a 0x182 when over written with a value which is not supported.

Implication: This behavior should be taken into account when accessing the P_x[n]_CSIWCI register.

Workaround: None at this time

Status: No Fix

7. Two Writes Required to Clear CSITTLECR.[Error Overflow]

Problem: The CSITTLECR.[Error Overflow] register field is set when an error counter has overflowed in a selected lane. Writing a 1 to this field should clear this register field, but it does not. It takes 2 writes of a 1 to properly clear this register field.

Implication: This behavior should be taken into account when accessing the P_x[n]_CSIWCI register.

Workaround: None at this time


Status: No Fix

8. CRC Errors With 16 Bit CRC

Problem: A Link layer and Phy Layer reset can both cause multiple CRC errors when 16 bit rolling CRC is enabled in at least one direction in the link. This can also occur on a reset where 16 bit CRC is changed to 8 bit CRC, or 8 bit CRC is changed to 16bit CRC. This can possibly triggering ERROR# 301, "Rbox Intel QuickPath Interconnect CRC Error" if enabled. Use the following flow when performing a link layer reset where the CRC mode is set or changed to 16 bit rolling CRC.

If doing a Link Layer Reset, use the following flow on the port doing a link layer reset.

1. Clear the RBOX error IPER0.CRCEC by writing 0x0.
2. Disable errors 300, 301, 304 by setting IPER0[26:24] = '000
3. Perform the following steps to link layer reset while 16 bit CRC is enabled:
 - Disable PhyInitBegin by setting P_x[n]_CSIPHCTR.PHY_INIT_BEGIN = 0 and P_x[n]_CSIPHCTR.PHY_RESET = 1
 - Set the R_CSR_CSILCL.CRCMODE= '01 for 16 bit CRC mode
 - Set R_CSR_CSILCL.LNKRSTC = 1 for a soft reset
 - Set P_x[n]_CSIPHCTR.PHY_INIT_BEGIN = 1 to start the initialization
4. Again, clear the RBOX error IPER0.CRCEC by writing 0x0
5. Enable promotion of Rbox errors to MCA by setting IPER0[26:24] = '111



If performing just a PHY Reset (Px[n]_CSIPHCTR.PHY_RESET = 1) clear R_CSR_IPER0.CRCERRD, R_CSR_IPER0.CRCESH, and R_CSR_IPER0.CRCEC following the PHY reset.

Implication: Intel QPI CRC Errors can occur if 16 bit rolling CRC is enabled and a link layer or PHY layer reset occurs.

Status: No Fix

9. **Reset While In Calibration State Can Cause Hang On Intel® Scalable Memory Interconnect (Intel® SMI)**

Problem: A reset while in the calibration state can cause a hang on Intel SMI. The state machine is in calibration if either PZ[n]_PBOXFS3.INIT_RX_STATE is in the calibrate state, or PZ[n]_PBOXFS2.INIT_TX_STATE is in the calibrate state. Reset is asserted via PZ[n]_PBOXFS1.FBD_PHY_RESET.

Implication: A reset must not be performed while in this state.

Workaround: Before issuing a CSR\PAL reset via a CSR write, firmware must ensure that the state machine is in not in calibrate.

Status: No Fix

10. **Intel SMI PZ[n]_PBOXFS2.CALIB_DONE Can Only Be Cleared In Reset**

Problem: PZ[n]_PBOXFS2.CALIB_DONE must not be cleared if the link is not in reset (PZ[n]_PBOXFS2.INIT_TX_STATE = Reset).

PZ[n]_PBOXFS2.CALIB_DONE is set to one once calibration is complete. To perform a calibration in the next re-initialization, this bit must be cleared. In order to clear this bit with the link in reset, the following flow can be used:

```
Px[n]_CSIPHCTR.PHY_INIT_BEGIN=0  \\Keep link in reset following reset
PZ[n]_PBOXFCTL1.FBD_PHY_RESET=1  \\Perform reset
PZ[n]_PBOXFCTL1.FBD_PHY_RESET=0
```

```
PZ[n]_PBOXFS2.CALIB_DONE = 0  \\ok to set calib_done=0
Px[n]_CSIPHCTR.PHY_INIT_BEGIN = 1  \\Allow state to advance past reset
PZ[n]_PBOXFCTL1.FBD_PHY_RESET=1  \\Perform reset to do recalibration
PZ[n]_PBOXFCTL1.FBD_PHY_RESET=0
```

Implication: Calibration of the Intel SMI can fail if PZ[n]_PBOXFS2.CALIB_DONE is written to a 0 when in reset.

Workaround: None at this time

Status: No Fix

11. **Memory Buffer Must Not Be Reset Directly by FPGA or Other System Logic On Warm Reset**

Problem: On a warm-state or warm-logic reset an Intel® 7500 Scalable Memory Buffer can be reset sooner than the processor. When this occurs the Zbox detects an error on Intel SMI. The processor is unable to signal an MCA in response to this Zbox error due to the pending reset. A cold reset is required to recover from this state.

Implication: FPGA or other system logic should not directly reset the Intel 7500 Scalable Memory Buffer on a warm-state or warm-logic reset.

Workaround: Firmware should provide some capabilities so that it can control the Mill Brook reset on warm reset.

Status: No Fix



12. R_CSR_OPER0.PHYTRAINLIMIT Is Not Asserted When All Clocks Fail

Problem: R_CSR_OPER0.PHYTRAINLIMIT should assert anytime the Pbox fails to train after multiple attempts. This is reflected in Error# 303. In the case that all clocks are dead on the respective link, R_CSR_OPER0.PHYTRAINLIMIT will not assert. A subsequent error will assert and will be dependent on activity of the system at time of failure.

If alternate clocks are enabled, all alternate clocks would be required to fail to hit this condition.

Implication: Error#303 will not assert on failed clock(s).

Workaround: None at this time

Status: No Fix

13. Frame Alert Logged After Warm-Logic Reset

Problem: After a warm-logic reset where the Intel 7500 Scalable Memory Buffer is not reset, the Zbox will log a "Zbox Memory Alert Error" Error# 605. Error# 605 does not cause any response. Its only response is to set Z_CSR_ERR_LOG.status_frm_alert. In addition, Z_CSR_CHNL_ERR_LOG.sts_frm_alert_ch[1-0] will be set to '11 which notifies the system frame alert channel 0 and channel 1 fired will be logged. Also Z_CSR_ERR_LOG.errflw_success = 1 will be set to 1 after a warm-logic reset, but Z_CSR_ERF_CTL_STS_0.err_1st will not show any persistent errors.

Implication: SAL needs be aware of this expected behavior. SAL should clear these errors fields after a warm-logic reset. A warm-logic reset is performed on a fatal error which causes an MCA.

Workaround: None at this time

Status: No Fix

14. Bad Parity In Route Table Can Cause Unexpected Error

Problem: SAL has access to the R_CSR_RTWR CSR to write to the route table. This CSR has three even parity bits for each of the 3 VN* fields in R_CSR_RTWR. Bad parity in the route table can cause an unexpected recoverable or fatal error. This error can occur even if the route table entry is not accessed by an Intel QPI transaction. There is no parity protection for the RTA Index (RTAINDX).

Implication: SAL may try to use bad parity in route entries to detect if unsupported NodeIDs are accessed. This is not supported on the processor.

Workaround: Ensure good parity is used in the route table.

Status: No Fix


15. Lower 2 Bits Of IHA Have Read/Write Behavior

Problem: Volume 2 of the 2.2 *Intel® Itanium® Architecture Software Developer's Manual*, Section 3.3.5.9 "Interruption Hash Address (IHA – CR25)" states that the lower 2-bits of the IHA register are ignored. Ignored fields have the property that writes are ignored and reads return zeros. The processor implements the lower 2 bits as Read/Write.

Implication: When a VMM runs an operating system virtualized using the virtualization acceleration control (a_to_int_cr and a_from_int_cr), the lower 2-bits of IHA will not be ignored but will have regular read/write behavior.

Workaround: None at this time

Status: No Fix



16. Read or Write of TAD CSRs While System Is Not Quiesced Can Cause Data Corruption

Problem: The *Intel® Itanium® Processor 9300 Series Datasheet* defines the TAD Structure. This section describes the 4 TAD CSRs B_CSR_TAD_CNTL, B_CSR_TAD_ADDR, B_CSR_TAD_DATA0 and B_CSR_TAD_DATA1.

Implication: Reads and Writes to these CSRs while the system is not quiesced can cause data corruption to inflight traffic in the Bbox.

Workaround: None at this time

Status: No Fix

17. After An Intel QPI Link Soft Reset CRC Errors May Be Observed

Problem: Following a soft reset (R_CSR_CSILCL.LNKRSTC=1) to the Intel QPI Link, a stream of CRC errors will be detected by the Rbox. This stream of errors will cause R_CSR_IPERO.CRCEC to be set to 0x7F indicating multiple CRCs occurred with an overflow (Error#301 and ERROR# 302). The stream of CRC errors will also trigger a physical layer inband reset.

Implication: When a soft reset occurs, an unexpected MCA may occur if the MCA is enabled.

Workaround: Firmware needs to disable MCA's for these errors when performing a Link Layer soft reset (R_CSR_IPERO.CRCERRCOINTEN=0 and R_CSR_OPERO.RETRYABORTEN=0). Following the soft reset clear any CRC errors and enable the respective MCAs.

Status: No Fix

18. In CLL A Warm-Logic Reset During LoL Event Causes Unexpected Results

Problem: If a fatal error occurs in the middle of handling a LoL (Loss of Lockstep) event, then the socket already in CLL (Core Level Lockstep) recovery may do one of the following after the subsequent warm-logic reset caused by the fatal error:

- Reenter CLL unexpectedly without going through PAL_HR_ENTER (which is not supported)
- Reenable the slave core as an active core during MCA handling. Slave core is expected to be disabled in this situation.

Both responses could result in a warm-logic reset loop preventing PAL to SAL handoff.

Implication: If this condition occurs the system could enter a reset loop.

Workaround: A Service Processor can poll for this condition, and reinitialize the CLL.

Status: No Fix

19. In CLL A State Reset May Reenter CLL

Problem: The only proper way to enter CLL is through PAL_HR_ENTER. Under some reset scenarios, the processor will come up in CLL in an unsupported manner. This occurs if the warm-state reset occurs before PAL reinitializes the processor.

- If a warm-state reset occurs during CLL, but before PAL has initialized CLL, CLL will be entered following the warm-state reset.
- A warm-logic reset followed by a subsequent warm-logic reset will result in the second warm-logic reset to take the warm-state reset flow. This scenario will also cause the processor to come up in CLL in an unsupported manner.

Both responses could result in a warm-logic reset loop preventing PAL to SAL handoff.

Implication: If this condition occurs, CLL will be brought up in an unsupported manner.

Workaround: A Service Processor can poll for this condition, and reinitialize the CLL. PAL release ensures that CLL is disabled on boot. This change will mitigate the potential of this issue from occurring, but does not eliminate it.



Status: No Fix

20. Physical Damage To Intel SMI Lane Can Cause Training To Fail

Problem: Noise on an Intel SMI lane can appear as a TSO training header to the receiver causing training to fail. This is seen on Intel SMI lanes which have some physical damage (generally connector, trace, or pad connectivity issues). For this issue to occur there must be physical damage to the lane to cause reflection, and the noise at the receiver needs to match the TSO header during Intel SMI training. If this condition occurs, the processor will trigger a fast reset on the Intel SMI channel, and the link will attempt training again. In a rare case if this condition is hit consecutively and the retry threshold is met, this will cause the Intel SMI link to fail with a fatal error.

The damage to the lane will log an Error# 617 or 618 (Zbox Correctable Memory Link CRC Error, Northbound or Southbound with lane mapout) on the processor.

Implication: A retraining can be caused by memory initialization, RAS event, or firmware generated fast reset. Anytime one of these events occurs and a lane with physical damage is present, this condition can occur. When this condition occurs, the Intel SMI link will attempt to retrain (up to the retry threshold) to recover. If the number of retrains exceeds the retry threshold, a fatal Error# 608\610 (Zbox Uncorrectable Memory Link CRC Error) will occur which will cause a fatal MCA.

Workaround: Set PZ[n]_PBOXFCTL1.RETRY_THR to '0111 to enable multiple retraining attempts. With this value, achieving the retry threshold due to this issues is expected to be a very rare event.

Status: No Fix

21. Northbound Intel SMI CRC Persistent Error Can Cause South Bound CRCs Resulting In Fast Reset Loop

Problem: Injection of a persistent Error #618 (Zbox Correctable Memory Link CRC Error - Northbound with Lane Mapout) can cause the processor to log Error #617 (Zbox Correctable Memory Link CRC Error - Southbound with Lane Mapout). In this case, the Intel SMI lane is properly mapped out on subsequent fast resets, but the Error #617 is not expected. The Error #617 is persistent enough to cause a fast reset due to the CRC threshold being hit. The Error #617 will occur on subsequent fast resets causing the Intel SMI link to loop on fast resets. This does not affect Intel QPI links.

Implication: Error #617 is incorrectly logged.

Workaround: The PAL call PAL_CRC_ERR_THRESHOLD_CONFIG provides the ability to change the window and threshold by which South Bound CRC errors are counted and promoted to persistent.

Status: No Fix


22. CRC Errors Occur on Intel QPI After Physical Layer Reset When Scrambling And Periodic Retraining Are Enabled

Problem: During an Intel QPI physical layer reset at the local socket, packets with CRC errors can be issued to the remote socket when scrambling is enabled. This will cause Error# 300 "Rbox Intel QPI CRC Error" to be logged at the remote socket. The link layer needs to be in the "normal operation" state, and a physical layer reset needs to hit a two cycle window just before the Physical layer initialization state machine is reaching the active state (LO) to encounter this issue. This is due to an issue with the scrambling logic when these condition occur.

The physical layer reset observed by the local socket can be initiated by Px[n]_CSIPHCTR.PHY_RESET=1, inband reset, or by RAS event.

Two scenarios can occur:

- An Intel QPI physical layer reset occurs on the local socket. The local socket encounters this issue and packets with CRC errors are sent to the remote socket



which logs Error# 300 "Rbox Intel QPI CRC Error". The remote socket encounters a "Rbox Intel QPI CRC Error - CRC Counter Overflow" and logs Error#301. This in turn causes the remote socket to generate a "Rbox Intel QPI Retry Abort Error" which is logged as Error# 302. The condition self heals, but respective errors are logged. Error# 302 is recoverable and an MCA will be issued. On the Intel® 7500 Chipset, a B1 Error will be logged (versus an Error#302 on the processor).

- An Intel QPI physical layer reset occurs on the local socket. The local socket encounters this issue and packets with CRC errors are sent to the remote socket which logs Error# 300 "Rbox Intel QPI CRC Error". Before the remote socket can encounter a CRC overflow, the local socket times out on its link level retry request (R_CSR_CSILCL.LLRTLNRST) waiting for a link level Retry Ack and issues a link reset which is logged as Error# 302 on the local socket. The condition self heals, but respective errors are logged. Error# 302 is recoverable and an MCA will be issued.

Implication: During a physical layer reset due to a CSR based local or remote socket phy layer reset, inband reset, or RAS event an unexpected Error# 302 "Rbox Intel QPI Retry Abort Error" may be logged if scrambling is enabled. Disabling Intel QPI periodic retraining decreases the probability of hitting this issue when scrambling is enabled.

In the extremely rare case where both the local and remote sockets of a link observe a physical layer reset at the same time, and this issue occurs on both sides of the link, a fatal Error# 303 "Rbox Uncorrectable Intel QPI CRC Error" could be logged by one or both of the processors. This would occur if the links fails to train after 3 attempts.

Workaround: Disabling Intel QPI scrambling on the link will prevent this issue from occurring. If scrambling is enabled (during the system link bring up), enable scrambling after all links have been brought up to mitigate this issue on boot. Not enabling Intel QPI Periodic Retraining will greatly reduce the probability of this issue from occurring when scrambling is enabled. The 2 cycle window occurs at the end of each Periodic Retraining Sequence (10mS) when scrambling is enabled (as well as on all physical layer resets). If scrambling is disabled, this 2 cycle window only occurs on a physical layer reset.

Status: No Fix

23. Transmitter Parameter Values


Problem: The Intel® Itanium® Processor 9300 Series Datasheet defines the Transmitter Parameter Values for Intel® QuickPath Interconnect and Intel SMI Channels at 4.8 GT/s. The processor does not follow the Intel® QuickPath Interconnect, Version 1.0, Electrical specifications for three values and are reflected in the table below.

ITEM	SYMBOL	Intel® Itanium® Processor 9300 Series Datasheet			1.0 Intel® QuickPath Interconnect Specification Electrical specifications		
		Min	Max	Unit	Min	Max	Unit
1	T _{SLEW-RISE-FALL-PIN}	6	12	V/ns	9	20	V/ns
2	V _{TX-CM-RIPPLE-PIN}	0	14	%	-0.0375	0.0375	ratio
3	TX _{DUTY-CYCLE-PIN}	-0.076	0.076	UI-UI jitter	-0.055	0.055	UI-UI jitter

Implication: Min/Max measured values may not comply with the Intel® QuickPath Interconnect, Version 1.0, Electrical specifications.

Workaround: None at this time

Status: No Fix



24. L2i Fills In 1 Way Of Set When All Other Ways Are Valid And A Way Is Disabled In Set

Problem: When Intel Cache Safe Technology disables any of the 8 ways in a set, the ways which are not disabled are used for the fill. For the L2i one preferred way will always be selected for a set if all the enabled ways are valid. If any of the valid ways in an affected set are invalidated, there will be no preferred way and the replacement algorithm for L2i will fill as expected.

This does not affect L2i sets where all ways are enabled, only sets which have ways where cache lines are disabled via Intel Cache Safe Technology.

Implication: The expected LRU algorithm is not followed for a L2i set which has one or more disabled cache lines (via Intel Cache Safe Technology) and all other enabled ways set to valid.

Workaround: None at this time

Status: No Fix

25. If A Global Fatal MCA Occurs While A chk.a, chk.s Or fchkf Is Executing A Logical Processors May Not Enter Machine Check Abort and Min-State Save Area May Be Invalid

Problem: If a global fatal MCA occurs while a chk.a, chk.s or fchkf instruction is in the pipeline, some threads may not restateer to PALE_CHECK, preventing the MCA from being initiated. This issue may also result in the min-state save area IIP, XIP, IPSR, and XPSR being invalid for the respective logical processor.

Implication: IIP, XIP, IPSR, and XPSR will be invalid in the min-state save area.

Workaround: In PAL the first thread in a socket to enter PALE_CHECK on a global fatal error will monitor all active threads to determine whether they enter PALE_CHECK. If any threads do not arrive at PALE_CHECK within 100 mS, PAL will force the missing thread(s) to restateer to PALE_CHECK. When this occurs, PAL will write IIP and IPSR to -1 in minstate for the threads that initially failed to enter MCA.

With the PAL workaround there are global fatal MCA cases where the IIP/IPSR will be invalid, but all threads will arrive at SAL_CHECK. In this condition the IIP, XIP, IPSR, and XPSR will be invalid, but PAL will not write -1 to IIP and IPSR, and the MCA will proceed.


In a rare global fatal MCA corner case this issue could affect all threads. In this case PAL will not be able to monitor the arrival of all threads at SAL_CHECK. For PAL to monitor arrival of all threads, at least one thread must arrive at the PALE_CHECK vector. An XPN timeout is likely to occur if the corner case condition happens due to hang resulting from all threads failing to enter PALE_CHECK.

Status: No Fix

26. DC Common Mode Clock At Rx Input For Intel SMI and Intel QPI

Problem: The *Intel® Itanium® Processor 9300 Series Datasheet* defines the Receiver Parameter Values for Intel QuickPath Interconnect and Intel SMI Channels @ 4.8 GTs. The processor does not follow the minimum value of 125mV documented in the Intel® QuickPath Interconnect, Version 1.0, Electrical specifications for the parameter in the table below. Instead it uses a Min value of 175mV.

Symbol	Parameter	Min	Nom	Max	Units	Notes	Symbol
V _{Rx-clock-cm-pin}	DC common mode ranges at the Rx input for any clock channel	175		350	mV	2	V _{Rx-clock-cm-pin}



Implication: Min measured value does not comply with the Intel® QuickPath Interconnect, Version 1.0, Electrical specifications and Intel Scalable Memory Interconnect Electrical specifications.

Workaround: None at this time

Status: No Fix

27. North Bound Link Errors Followed Immediately By Erasure+1 Errors Can Cause Correctable Error To Be Uncorrectable

Problem: The following sequence will cause the 2nd correctable ECC error (4th bullet) to be logged as uncorrectable.

- Northbound Link error on rank A
- Followed immediately by Erasure +1 error on rank B (not rank A)
- Followed fairly quickly by another Northbound link error on rank C (C could be same as A)
- Followed immediately by Erasure +1 error on a different rank (not rank C)

Erasure is where a DRAM is mapped out by the memory controller as it has exceeded the number of correctable errors.

Implication: This issue is only expected to be seen during error injection testing.

Workaround: None at this time

Status: No Fix

28. Corrupted ALERT Frame Not Detected By Zbox


Problem: Any corrupted ALERT frame will not be detected by the processor. Since the Intel® 7500 Scalable Memory Buffer issues a series of ALERT frames, in most cases missing an ALERT frame is not an issue. There is no CRC protection on ALERT frames, and the processor does not have a mechanism to detect corrupt Alert Frames.

Implication: A rare corner case condition could occur due to this issue. This corner case has not been seen on a real system and is expected to be a very rare event. When the memory controller issues a write to the Intel 7500 Scalable Memory Buffer, the memory controller waits a round trip latency to confirm the southbound command arrived at the Intel 7500 Scalable Memory Buffer intact. In the event of a southbound Intel SMI CRC error (transient or persistent) at the Intel 7500 Scalable Memory Buffer, the error is detected and sends a series of ALERT frames on the northbound link to notify the host that it received a southbound CRC error. The processor detects the ALERT frames, resets the link, and reissues the packet with the southbound error. If the initial Alert frame for the respective south bound error is corrupted by a secondary error on the northbound link due to this issue, the processor will complete the write without reissuing. The memory controller uses two Intel SMI channels in lockstep for each cache line access, thus on a future read if one channel was affected by this issue the other would likely return valid data.

- If a Link error occurs on southbound write data, the Intel 7500 Scalable Memory Buffer will write data with BAD ECC to DRAM. The memory controller ECC would catch this, and should fix resulting in correctable error.
- If a Link error occurs on southbound write command, the Intel 7500 Scalable Memory Buffer will likely drop the write to memory. The memory controller ECC would catch this, but likely not be able to fix resulting in uncorrectable error.

Workaround: None at this time

Status: No Fix



29. **Z_CSR_ECC_LLE_SUCCESS_[7-0] Does Not Freeze On Error**

Problem: Z_CSR_ECC_LLE_SUCCESS_[7-0] does not freeze when an error is detected, while Z_CSR_PLD_CORR_LOG_[1-0] does freeze when an error is detected. Thus, it is not always possible to identify a bad DIMM in DIMM pair. Z_CSR_PLD_CORR_LOG_[1-0] identifies a correctable error to a DIMM pair. Z_CSR_ECC_LLE_SUCCESS_[7-0] identifies the bad DIMM in a DIMM pair.

Implication: If multiple errors have occurred Z_CSR_ECC_LLE_SUCCESS_[7-0] may not identify a bad DIMM in a DIMM pair.

Workaround: As a workaround, firmware can configure Z_CSR_PMU_CNT_[5] and Z_CSR_PMU_CNT_CTL_[5] PMU counters to count ECC errors. If only one error has occurred, then Z_CSR_ECC_LLE_SUCCESS_[7-0] is valid.

To initialize:

- Set Z_CSR_PMU_CNT_CTL_[5].incr_sel = 0xD (Event as selected by Z_CSR_PMU_ZDP_CTL_FVC.evnt1).
- Set Z_CSR_PMU_ZDP_CTL_FVC.evnt1 = 0b001 (mem_ecc_err - Memory ECC error detected)
- Set Z_CSR_PMU_CNT_[5] = 0

When an error is detected by assertion of Z_CSR_PLD_CORR_LOG_1.Valid:

- Read Z_CSR_ECC_LLE_SUCCESS_[7-0]
- Immediately read Z_CSR_PMU_CNT_[5] to find the number of ECC errors
 - If Z_CSR_PMU_CNT_[5] = 1, ECC_LLE_SUCCESS is valid
 - If Z_CSR_PMU_CNT_[5] > 1, ECC_LLE_SUCCESS is not valid
- Immediately set Z_CSR_PMU_CNT_[5] = 0

Status: No Fix

30. **In Core Level Lockstep A Ubox Request Can Be Dropped Causing Hang**

Problem: The processor's Ubox does not have a dedicated port. The Ubox has two shared ports, one shared with CPE2 and one shared with CPE3. Disabling either core 2 or core 3 (connected to CPE2 and CPE3 respectively) can cause an issue with the respective Ubox port. With this issue, the tracking of the Intel QPI DRS (Data Response) and NDR (No Data Response) credits is done incorrectly and the Ubox may drop a response. This issue only occurs with specific corner case traffic. If this issue occurs, it will cause a hang which will subsequently cause an ORB timeout from the requesting agent.

Implication: This issue affects the Core Level Lockstep SKU.

Workaround: None at this time

Status: Under Investigation


31. **PAL_THREAD_CONTROL Does Not Properly Deconfigure Cores**

Problem: The processor only supports 1 core deconfiguration as documented in the *Itanium® 2-Based Platform Compatible Processors Firmware Guide* via PAL_THREAD_CONTROL. The only deconfiguration supported is where only C1 and C0 are disabled on the 4 core SKU. The current version of PAL does not support this 4 core to 2 core configuration option.

Implication: Affected PAL versions do not support core configuration.

Workaround: SAL must not use PAL_THREAD_CONTROL to disabled cores with affected PAL releases. PAL_THREAD_CONTROL thd_control bits [11:8] must be 0.

Status: Plan Fix in PAL



32. PAL_SHUTDOWN Has Incorrect Index

Problem: Volume 2 of the 2.2 *Intel® Itanium® Architecture Software Developer's Manual*, Table 11-39 lists the PAL Power Information and Management Procedures. The Idx for PAL_SHUTDOWN is 45 in the table, but PAL implements this as 518.

In addition, this table in the *Intel® Itanium® Architecture Software Developer's Manual* states there is a dependency on PAL_MEMORY_BUFFER. PAL is removing this dependency. Also, this change required PAL_A_SPEC and PAL_B to be modified.

Implication: The index of PAL_SHUTDOWN is incorrect for affected PAL versions.

Workaround: None at this time

Status: Plan Fix in PAL

33. PAL_ECC_ERROR_MASK Does Not Check Mask For Determining No Error

Problem: PAL_ECC_ERROR_MASK checks the Z_CSR_ECC_LLE_SUCCESS_[7-0] CSRs to determine if there is an error. For the even CSRs, if position and mask are both 0, PAL_ECC_ERROR_MASK should return no error. Else it should return error.

PAL_ECC_ERROR_MASK incorrectly only checks Z_CSR_ECC_LLE_SUCCESS_[n].position when determining no error. It does not check Z_CSR_ECC_LLE_SUCCESS_[n].mask.

Implication: It is possible that an error is logged in the Z_CSR_ECC_LLE_SUCCESS_[7-0] CSRs, but PAL_ECC_ERROR_MASK could miss it if the position field is 0 and mask is not 0 on affected PAL versions.

Workaround: Service processor can access the Z_CSR_ECC_LLE_SUCCESS_[7-0] CSRs directly.

Status: Plan Fix in PAL

34. Rbox Blocking And Intel QPI Min Credit Occurring At The Same Time With Some Traffic Patterns Can Cause Hang

Problem: If the Rbox is blocked and unable to make forward progress for some period of time, and at the same time the Intel QPI credits are used up (both VNA credits on an Intel QPI link and VNO DRS credits on the caching agent), a deadlock can occur. For this condition to occur there must also be local CSR traffic, and 2 sockets must be accessing each others Ubox. The cross Ubox traffic can be either from CSR accesses or interrupts. The cross Ubox traffic can also be caused by 2 cores on the same socket accessing the local Ubox in a specific traffic corner case.

An ORB timeout or Snoop timeout will be logged if this rare condition occurs. Usually at least one ORB timeout will be logged and the address will show CSR or interrupt addresses to another Ubox.

Implication: If this condition occurs, a transaction will fail to make forward progress resulting in a hang.

Workaround: Intel QPI Periodic Retraining is an event which exposes this issue. Intel QPI Periodic retraining must be disabled (Px[n]_CSIPHPRT.RETRAIN_INTERVAL = 0). OEMs are responsible for thermal and environmental testing of their platform, and this testing needs to take into account the impact of Intel QPI Periodic Retraining disabled.

Contact your Intel Technical representative for questions on or analysis of your thermal and environmental testing around Intel QPI Periodic Retraining.

Status: No Fix

35. Persistent CRC Error During Aggressive Pin Throttling Can Cause Hang

Problem: If a refresh is held up and pending when it is time for the next refresh to that same rank, the processor prevents the second refresh from dispatching (which is expected behavior). Due to a logic issue in this state, if an Intel SMI northbound or southbound



persistent CRC link error occurs, or an Intel SMI southbound transient CRC link error occurs while in this state the processor will hang.

Implication: Hang if the condition occurs.

Workaround: If aggressive pin throttling is used via the MEM_THROTTLE_L pin, firmware must enforce the following max throttling. More aggressive throttling is not supported. The supported max throttling rates vary based on the setting of Z_CSR_REFRESH_CTL.rfr_accelerate_ena.

- Z_CSR_REFRESH_CTL.rfr_accelerate_ena = 0
 - Z_CSR_PT_CTL.window = 0x100 (256 frames per window)
 - Z_CSR_PT_CTL.max = 0x80 (max throttling of 50% - 128 will be NOPs in window)
- Z_CSR_REFRESH_CTL.rfr_accelerate_ena = 1
 - Z_CSR_PT_CTL.window = 0x100 (256 frames per window)
 - Z_CSR_PT_CTL.max = 0x4C (max throttling of 30% - 76 will be NOPs in window)

Status: No Fix.

36. In A Virtualized Environment Guest OS IIB0/IIB1 Reads/Writes May Not Behave As Expected If Virtualization Acceleration is Enabled

Problem: The *Intel® Itanium® Architecture Software Developer's Manual Specification Update October 2009*, Specification Change 16 adds "Interruption Instruction Bundle Registers" IIB0/1 control registers:

- CR26 IIB0 Interruption Instruction Bundle 0
- CR27 IIB1 Interruption Instruction Bundle 1

If the Virtualization Acceleration Control a_to_int_cr optimization is enabled, writes to IIB0/1 should complete without an intercept. With the affected PAL releases, an unexpected intercept to the VMM will occur on writes to IIB0/IIB1 if the a_to_int_cr optimization is enabled.

If the Virtualization Acceleration Control a_from_int_cr optimization is enabled, reads from IIB0/1 may not return the correct Virtual Processor Descriptor values.

Exposure to this erratum only occurs when the guest operating system is running (PSR.vm = 1).

Implication: With affected PAL releases, guest operating systems that access IIB0/IIB1 in a virtual machine monitor that enables a_to_int_cr or a_from_int_cr virtualization acceleration may not behave as expected or may cause an unexpected intercept to the VMM.

Workaround: With affected PAL releases, avoid virtualizing operating systems that access IIB0/IIB1 with VMMs that enable a_to_int_cr or a_from_int_cr accelerations.

Status: Plan Fix in PAL


37. A register file error may be signaled on wrong thread

Problem: When register parity file detection creates a MCA, a register file parity error could be signaled to the wrong thread. If this condition occurs the register file parity error will not be recoverable.

Implication: The wrong thread receiving the register file error could lead to inappropriate MCA handling.

Workaround: None at this time.

Status: Fixed in PAL 4.28. For affected versions, see the *Summary Table of Changes*.



38. Neither thread on a core entering SAL_CHECK during a SYSINIT MCA

Problem: When the SYSINIT MCA PAL thread semaphore is owned and the Intel QPI viral bit is set, the semaphore is not released causing the SYSINIT MCA to be continuously signaled.

Implication: The SYSINIT MCA handoff to SAL may not complete.

Workaround: None at this time.

Status: Fixed in PAL 4.29. For affected versions, see the *Summary Table of Changes*.

39. A fault during RSE (Register Stack Engine) load in PAL_CHECK may cause system hang

Problem: During PAL_CHECK, the second load to the RSE could result in a General Exception Fault.

Implication: A General Exception Fault may cause system to hang.

Workaround: None at this time.

Status: Fixed in PAL 4.28. For affected versions, see the *Summary Table of Changes*.

40. Register file error is reported incorrectly by PAL_MC_ERROR_INFO with every ERROR1 assertion

Problem: A register file error may be incorrectly reported by PAL_MC_ERROR_INFO.

Implication: When retrieving info from PAL_MC_ERROR_INFO for an ERROR1 assertion (resulting in a global fatal MCA), a register file error will be incorrectly detected.

Workaround: None at this time. For affected versions, see the *Summary Table of Changes*.

Status: Fixed in PAL 4.29.

41. PAL_REGISTER_INFO incorrectly reports CR.iib0 and CR.iib1 as unimplemented

Problem: The PAL_REGISTER_INFO call reports CR.iib0 and CR.iib1 as unimplemented.

Implication: When making PAL_REGISTER_INFO call, it returns incorrect information on the implementation status for CR.iib0 and CR.iib1 as not implemented.

Workaround: None at this time.

Status: Fixed in PAL 4.29. For affected versions, see the *Summary Table of Changes*.

42. Nested MCA may occur during recoverable MCA testing

Problem: Recoverable MCA testing may result in a PAL hang.

Implication: When testing recoverable MCA's, a nested MCA condition may occur which may cause a hang in PAL MCA handling.

Workaround: None at this time.

Status: Fixed in PAL 4.29. For affected versions, see the *Summary Table of Changes*.

43. Infinite Viral MCA's may occur when testing viral MCA's

Problem: Viral MCA's may not reach SAL_CHECK.

Implication: When testing Viral MCA's, the MCA flow may not successfully reach SAL_CHECK.

Workaround: None at this Time.

Status: Fixed in PAL 4.29. For affected versions, see the *Summary Table of Changes*.

44. INIT flow will hang if invoked prior to MINSTATE initialization

Problem: Processor initialization flow interrupt (INIT) will hang if the MINSTATE area is not available.



Implication: When testing INIT's prior to initialization of the MINSTATE area, it will result in a hang or an illegal address access.

Workaround: None at this time.

Status: Fixed in PAL 4.29. For affected versions, see the *Summary Table of Changes*.

45. An unexpected MCA may occur if multiple thermal error injections are applied to a socket

Problem: An unexpected MCA may occur on a QR ERROR (charge rationing) assertion.

Implication: When testing thermal error injections, a move of the QR master during the same time that a QR ERROR is signaled may cause an unexpected MCA.

Workaround: None at this time.

Status: No Fix. Affects all PAL versions.

46. A pended QR CMCI may not be cleared during the OS boot sequence

Problem: A pended QR (charge rationing) CMCI prior to booting to the OS needs to be cleared, or handled by the OS in order to handle subsequent QR CMCI's.

Implication: If a pended QR CMCI is not handled in the OS at a reasonable time during boot, additional QR CMCI's will not be triggered until the pended QR CMCI is handled or cleared.

Workaround: Implement polling in the OS to clear or process any logged CMCI's during the OS boot sequence. Another option is to call PAL_MC_ERROR_INFO and PAL_MC_CLEAR_LOGS early in the OS boot sequence to handle a pended QR CMCI.

Status: No Fix. Affects all PAL versions.

47. Firmware Setting Z_CSR_FBD_FAIL_CHIP_STATE_[3-0]. chipFailEnable[3-0] During Erasure Can Cause Uncorrectable ECC Error

Problem: Z_CSR_FBD_FAIL_CHIP_STATE_[3-0] are programmed by processor ECC microcode when an ECC threshold has occurred. In this mode all cacheline reads are ECC corrected on the fly. Firmware is responsible for clearing the erasure as processor ECC microcode will not do so. Clearing erasure is accomplished by writing 0x01 to the respective Z_CSR_FBD_FAIL_CHIP_STATE_[3-0].chipFail*[n] CSR fields. When chipFailEnable[n] is set there is a very small window which can cause an unexpected uncorrectable ECC error. This is expected to be a very rare event.

Implication: Unexpected ECC error when setting chipFailEnable[n] to 1 to clear erasure.

Workaround: To mitigate this very rare event, firmware can Quiesce the socket before writing to Z_CSR_FBD_FAIL_CHIP_STATE_[3-0]. chipFailEnable[3-0].

Status: No Fix. Affects all steppings.


48. CPE Outbound Request Buffer Timeout (CPB_ORB_TIMEOUT_CFG) counters are not deterministically assigned to the correct transaction types

Problem: The timeout values configured by programming the CPB_ORB_TIMEOUT_CFG CSR fields (level 1, 3, 4, and 5) are not guaranteed to be applied to their corresponding transaction types. Any transaction type may be randomly assigned to any level of timeout. Similarly, the timeout hierarchy described in the QPI specification is not preserved on Intel® Itanium® Processor 9300 Series.

Implication: Miscorrelation of timeout level to transaction type can lead to false timeout failures. Also this can lead to erroneous fault isolation.

Workaround: Program all timeout counters to the same value or disable all counters.

Status: No fix.



49. Incorrect PIROM value for THERMALERT_N Hot Deassertion Hysteresis

Problem: The THERMALERT_N hot deassertion hysteresis value programmed in the PIROM (Thermal Reference Sec #108, Offset 0x6Ch) has an incorrect value of 0x02.

Implication: None due to the implemented workaround.

Workaround: The Intel® Itanium® Processor 9300 Series microcode programs THERMALERT_N hot deassertion hysteresis to the correct value of 0x04 during initialization.

Status: No Fix.

50. PAL_VP_CREATE does not perform reserved bit checking on the virtualization acceleration control (VAC) and virtualization disable control (VDC) fields in the virtual processor descriptor (VPD)

Problem: The software developers manual describes Bits [63:7] of the VAC and VDC fields in the VPD as reserved. Architecturally reserved bits passed into PAL procedures return an error indication in the return argument. PAL_VP_CREATE does not verify that the reserved bits are all zero and does not return an error indication.

Implication: If a customer is setting a reserved bit in these fields, no indication will be returned and it could have an impact on the behavior of a virtual machine if the reserved bits get defined in future processors.

Workaround: Callers of the PAL_VP_CREATE procedure should verify that they are not setting any reserved bits in the VAC/VDC fields of the VPD passed to this procedure.

Status: No Fix.

51. False ERROR1_N signal may be asserted during power cycle

Problem: An uninitialized circuit node within the processor may lead to a false assertion of the ERROR1_N signal only during a power cycle. This false error indication is visible at the deassertion of RESET_N.

Implication: During a power cycle, the ERROR1_N signal may erroneously assert before RESET_N deassertion. The false error indication is not cleared by a RESET_N / PWRGOOD sequence.

Workaround: The workaround for this issue consists of two components: Incorporate PAL version 4.30 into the platform firmware. Assert a PWRGOOD reset 2ms after detection of the false ERROR1_N condition.

Status: Fixed in S-spec parts numbered LCxx.

52. PAL_MC_ERROR_INFO incorrectly returns target address valid when originating from ptc.g with RspStatus=Fail

Problem: PAL_MC_ERROR_INFO incorrectly sets the target address valid bit for ptc.g transactions when RspStatus=Fail.


Implication: ptc.g transaction failures are incorrectly reported by PAL_MC_ERROR_INFO as having a valid transaction address.

Workaround: None at this time.

Status: Fixed in PAL 4.30.

53. Intel QPI State Machine Does Not Increment On Every Initialization Failure Instance


Problem: When there is a link layer initialization failure, per the Intel QPI specification the Intel QPI logic should track such a failure in order to issue the appropriate phy reset when an internal threshold is met. If a link layer initialization failure occurs while the Intel QPI logic in the processor is in some internal states, some failed Intel QPI link layer inits would not be recorded in the counter. These missed counts can cause a delay in when the count limit would be reached.



Implication: This will normally cause the processor to take longer to issue a phy reset. In most cases the Intel QPI link will recover, or eventually assert the phy reset encountering enough counted Initialization failures. In other cases the agent on the link would trigger the phy reset first. The conditions to create a livelock condition where both side of the link never increment their counter is extremely rare.

Workaround: None at this time.

Status: No Fix. Affects all PAL versions.



Documentation Changes and Clarifications

§

