Intel® E8501 Chipset North Bridge (NB)

Specification Update

May 2006

Notice: The Intel® E8501 Chipset North Bridge (NB) may contain design defects or errors known as errata that may cause the product to deviate from published specifications. Current characterized errata are documented in this specification update.

Document Number: 309622-001

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, or life sustaining 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® E8501 Chipset North Bridge (NB) 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.

I^{2}C is a two-wire communications bus/protocol developed by Philips. SMBus is a subset of the I^{2}C bus/protocol and was developed by Intel. Implementations of the I^{2}C bus/protocol may require licenses from various entities, including Philips Electronics N.V. and North American Philips Corporation.

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

Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Contents

§

Revision History

Version	Description	Date
-001	• Initial Release	May 2006

Preface

This is an update to the specifications in the documents listed in the "Affected Documents" and "Nomenclature" tables. It is a compilation of device and document errata and specification clarifications/changes, and is intended for hardware system manufacturers and software developers.

Information types defined in the Nomenclature section of this document are consolidated into this document and are no longer published in other documents. This document may also contain previously unpublished information.

Affected Documents

Document Title	Document Number
Intel® E8501 Chipset North Bridge (NB) Datasheet	309620
Intel® E8501 Chipset North Bridge External Design Specification	19548
64-bit Intel® Xeon® Processor MP, Dual-Core Intel® Xeon® Processor 7000 Sequence, Tulsa and Intel® E8501 Chipset Platform Design Guide (PDG)	19552
RS - Twin Castle Chipset BIOS Specification	Contact your Intel representative

Nomenclature

S-Spec Number is a five-digit code used to identify products. Products are differentiated by their unique characteristics,e.g., core speed, L3 cache size, package type, etc. as described in the processor identification information table. Read all notes associated with each S-Spec number.

QDF Number is a several digit code used to distinguish between engineering samples. These samples are used for qualification and early design validation. The functionality of these parts can range from mechanical only to fully functional. This document has a processor identification information table that lists these QDF numbers and the corresponding product sample details.

Errata are design defects or errors. These may cause the NB 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.

Specification Changes/Clarifications are modifications to the current published specifications. These changes will be incorporated in the next release of the specification.

Documentation Changes include typos, errors, or omissions from the current published specifications. These will be incorporated in the next release of the specification.

Note: Errata remain in the specification update throughout the product's life cycle, or until a particular stepping is no longer commercially available. Under these circumstances, errata removed from the specification update 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.

Summary Table of Changes

The tables included in this section indicate the , errata, specification changes/ clarifications, or documentation changes that apply to the Intel® E8501 Chipset North Bridge (NB). 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.

Codes Used in Summary Table

Stepping/Version

X: Applies to this stepping.

Blank: Fixed in listed stepping or does not exist in listed stepping.

Status

No Fix - Root caused to a silicon issue that will not be fixed.

Plan Fix - Root cased to a silicon issue and will be fixed in a future stepping.

Fixed - Root cased to a silicon issue and has been fixed in a subsequent stepping.

Row

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

Errata

Number	Affected Stepping		Status	ERRATA
	B-0	B-1		
1	X	X	No Fix	Hub interface data parity error may not generate poison
2	X	X	No Fix	FSB ECC error followed by a 0-length transfer logs two errors
3	X	X	No Fix	Uncorrectable errors are logged twice
4	X	X	No Fix	VPP ports may report errors across system resets
5	X	X	No Fix	PCI Express Errors should not be routed to H_{A,B}_MCERR#
6	X	X	No Fix	PCI Express Cards Intermittently Train x4 Instead of x8
7	X	X	No Fix	PCI Express Port A0 will not Train x1
8	X	X	No Fix	NB may hang under heavy stress
9	X	X	No Fix	EXP_SLOTSTS.CMDCOMP cannot be cleared
10	X	X	No Fix	Replay Timer errors logged when PCI Express ports exit L1 power manangement state

Errata (Continued)

Number	Affected Stepping		Status	ERRATA
	B-0	B-1		
11	X	X	No Fix	NB hang during stress and PCI Express link events
12	X	X	No Fix	NB may cause IERR# when using SYRE.CPURST with the Intel® Xeon® Processor MP with 1MB L2 Cache
13	X	X	No Fix	NB may trigger ECC errors during warm reset
14	X	X	No Fix	NB PCI Express ports get stuck in reset state
15	X	X	No Fix	Receiver errors reported on NB when link is disabled
16	X	X	No Fix	Receiver errors reported on PortB when link is disabled in Port A in the x8 configuration
17	X	X	No Fix	PCI Express Max Payload Size does not default to 128B
18	X	X	No Fix	NB Coherency issue with Locked Sequence
19	X	X	No Fix	Hot Plug Memory Board powered down after hard reset
20	X	X	No Fix	SMI Escalation via ERR[2:0]# Pins May Result in IERR#
21	X	X	No Fix	Can not read across MI link after XMB fatal error
22	X	X	No Fix	PCIe hang on L1 entry when a x4 link has downshifted to x2

Specification Changes

Number	SPECIFICATION CHANGES
	There are no new Specification Changes in this revision of the Specification Update.

Specification Clarifications

Number	SPECIFICATION CHANGES
1	PCI Express Configuration completion time-out enable needs non-configuration completion time-out to be enabled in order to work

Documentation Changes

Number	DOCUMENTATION CHANGES
	There are no new Documentation Changes in this revision of the Specification Update.

Identification Information

Component Identification via Programming Interface

The Intel® E8501 Chipset North Bridge (NB) can be identified by the following register contents:

Stepping	Vendor ID[a]	Device ID[b]	Revision Number[c]
B-0	8086h	2600h	10h
B-1	8086h	2600h	11h

NOTES:

a. The Vendor ID corresponds to bits 15:0 of the Vendor ID Register located at offset 00 - 01h in the PCI function 0 configuration space.
b. The Device ID corresponds to bits 15:0 of the Device ID Register located at offset 02 - 03h in the PCI function 0 configuration space.
c. The Revision Number corresponds to bits 7:0 of the Revision ID Register located at offset 08h in the PCI function 0 configuration space.

Component Marking Information

The Intel® E8501 Chipset North Bridge (NB) stepping can be identified by the following component markings:

Stepping	S-Spec	Marking	Notes
B-0	QG84010TNBS	876226	Pb-Free Package
B-0	NQ84010TNBS	876225	
B-1	SL8XF	876226	Pb-Free Package
B-1	SL8XE	876225	

Figure 1. Top-Side Marking Example (optional)

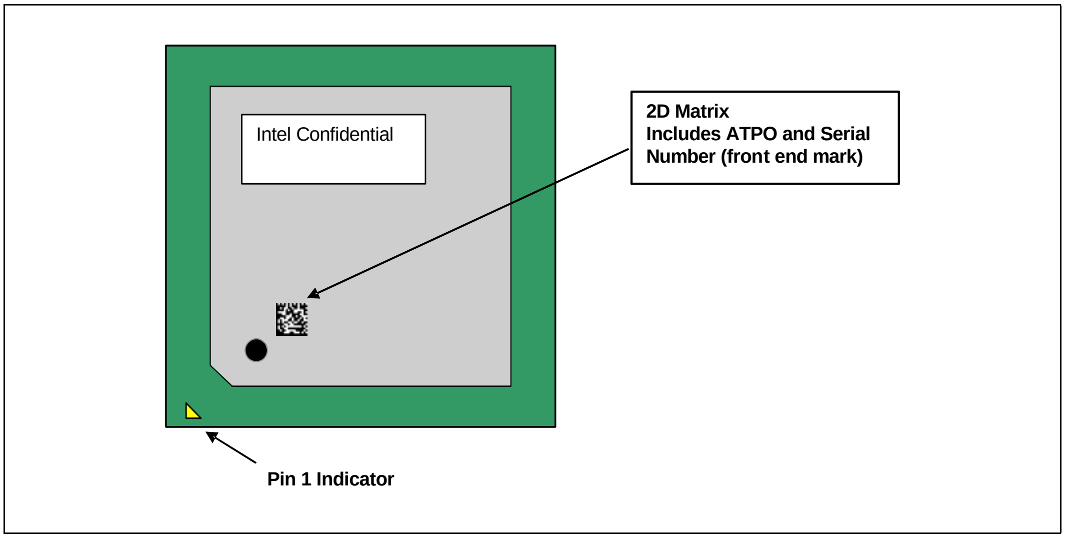

Figure 2. Bottom-Side Marking Example (optional)

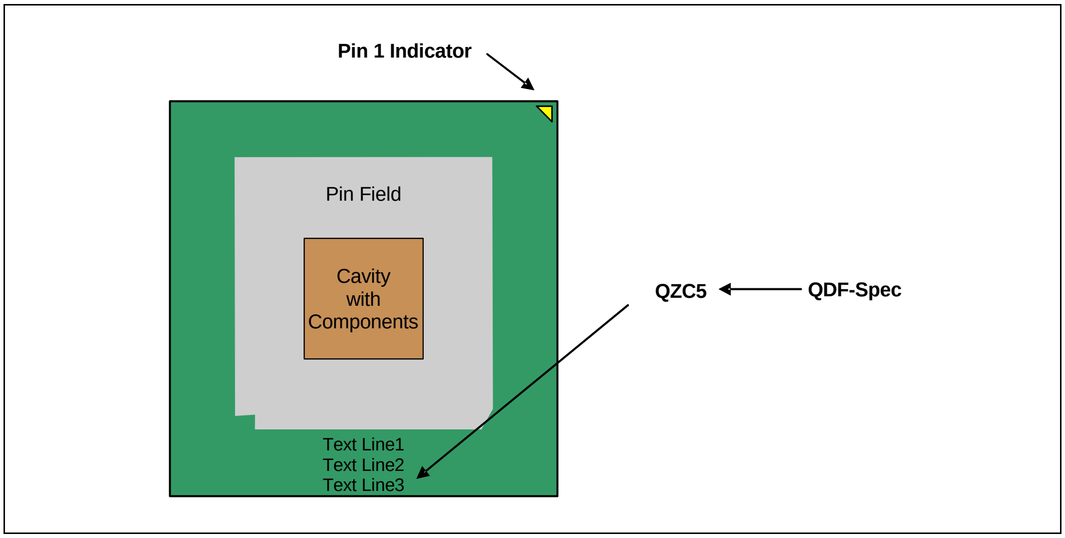

Errata

1. Hub interface data parity error may not generate poison

Problem: If a data parity error occurs in the very last data double word of an incoming hub interface transaction and the transaction size is a multiple of 4 double words, the poison information does not get propagated to CDC.

Implication: If this error is masked, silent data corruption may occur.

Workaround: Hub interface parity errors must be treated as fatal to contain this possible data corruption.

Status: No Planned Fix.

2. FSB ECC error followed by a 0-length transfer logs two errors

Problem: The NB does an ECC check on the data from the previous transaction on a 0-length transfer.

Implication: If a data ECC error occurs on the FSB subsequently followed by a 0 length transfer, the NB will log the ECC error twice.

Workaround: None.

Status: No Planned Fix.

3. Uncorrectable errors are logged twice

Problem: If an uncorrectable error occurs, this error will be logged in both the NB’s FERR and NERR registers. This issue does not affect errors on IMI.

Implication: Details on errors occurring before an uncorrectable error will be lost.

Workaround: None.

Status: No Planned Fix.

4. VPP ports may report errors across system resets

Problem: VPP Ports may report errors in the EXP_unitERR.T9Err bit across system resets. Per the I^{2}C* specification a data line can change when the clock is low. The only exceptions to this are Start and Stop commands in which is a master drives a data falling edge when clock is high (for Start) or drives a data rising edge when clock is high (for Stop).

The NB VPP state machine is constantly cycling through each VPP port and reading its state. If the NB gets reset when the clock is high and one of the VPP ports is driving the data line low (as part of the read or an acknowledge) then, upon reset, the VPP will not recognize the initial Start command and will not properly initialize. The NB recognizes the VPP did not respond properly and sets the T9Err bit.

Implication: Hot plug VPPs may non-deterministically be unresponsive after a reboot.

Workaround: This issue can be avoided by reinitializing the VPP ports when this error occurs as follows.

Disable Error Logging (EXP_unitDMASK.T9DetMsk) for this port

Enable this VPP

Delay 50 ms

Disable this VPP

Wait 1 second

Clear the Error bit (EXP_unitERR.T9Err)

Enable Error Logging (EXP_unitDMASK.T9DetMsk)

Enable VPP

Status: No Planned Fix.

5. PCI Express Errors should not be routed to H_{A,B}_MCERR#

Problem: If PCI Express errors are routed to the FSB H_{A,B}_MCERR# pins via settings in the EXP_ERR_DOCMD register, when an error occurs, MCERR# will be continually asserted (3 cycles on, 3 cycles off) until the error is cleared.

Implication: MCERR# should only be asserted for 2 cycles per error. Continuos assertion of MCERR# will result in an interrupt storm and system hang.

Workaround: PCI Express errors routed via the EXP_ERR_DOCMD register (D1-7, F0, 0x148) should be routed through one of the ERR[2:0] signals rather than the MCERR# pins.

Status: No Planned Fix.

6. PCI Express Cards Intermittently Train x4 Instead of x8

Problem: Intel validation boards which support x8 links will intermittently train as x4 links coming out of an AC power cycle only. This behavior seems to follow cards rather than systems. Some cards show this behavior up to 25% of the time.

Implication: This behavior has not been reproduced on other cards. Cards will not experience full bandwidth if they train as a x4 rather than x8 link.

Workaround: All ports should be strapped to their maximum widths (i.e. x8 slots should be strapped as x8, regardless of the poplulated card width) via the EXP_WIDTH strapping options.

Status: No Planned Fix.

7. PCI Express Port A0 will not Train x1

Problem: A logic error in NB B0 occasionally prevents port A0 lane 0 from gaining symbol lock.

Implication: During a x1 negotiation (i.e. with a x1 device or after x4 link training has failed because of a bad lane), link training will fail.

Workaround: A workaround for this bug is described in the BIOS Spec Update. This workaround must be run each time the link trains with a x1 device.

Status: No Planned Fix.

8. NB may hang under heavy stress

Problem: Under heavy memory write and PCI configuration space stress, internal NB’s queues may become completely full and enter an arbitration deadlock

Implication: Under these conditions, the NB may hang the system.

Workaround: A workaround for this issue which adjusts internal buffer aribitration policy is described in the BIOS Spec Update.

Status: No Planned Fix.

9. EXP_SLOTSTS.CMDCOMP cannot be cleared

Problem: Due to the width of the internal NB bus, a write to clear the EXP_SLOTSTS.CMDCOMP bit causes the NB to also register a write to the EXP_SLOTCTRL register. As a result, the NB will set the EXP_SLOTSTS.CMDCOMP within a few cycles.

Implication: The EXP_SLOTSTS.CMDCOMP bit is effectively stuck high

Workaround: BIOS and OS's should not rely on the value of this bit and disable command complete interrupts by clearing setting EXP_SLOTCTRL.CCIEN

Status: No Planned Fix.

10. Replay Timer errors logged when PCI Express ports exit L1 power manangement state

Problem: When a PCI Express link is in L1, the link is in eletrical idle, but the NB physical layer incorrectly indicates that it is ready to accept transactions to the NB link layer. If the NB issues a transaction during this time, it is dropped by downstream device, but does cause an exit from L1. The link goes into recovery and then L0. Since the initial transaction was dropped, there is no ack sent by the downstream device, which eventually causes the replay timer to expire.

Implication: A reply timer error may be logged whenever a PCI Express link exits L1. This is a correctable error. The dropped transaction will be resent and the link will operate normally after this error.

Workaround: COREDMASK.IO11DETMSK should be set to disable logging of replay timer errors.

Status: No Planned Fix

11. NB hang during stress and PCI Express link events

Problem: An arbitration deadlock can occur when a PCI express device issues unaligned DMA writes to memory and on the same I/O unit, either a link goes down or a PCI Express-related internal MSI is generated.

Implication: Under these conditions, the NB may hang the system or silent data corruption may occur.

Workaround: A workaround for this issue which disables internally generated MSIs and adjusts arbiter timing is described in the BIOS Spec Update.

Status: No Planned Fix

12. NB may cause IERR# when using SYRE.CPURST with the Intel® Xeon® Processor MP with 1MB L2 Cache

Problem: During CPU-only resets with Intel® Xeon® Processor MP with 1MB L2 Cache invoked via the NB's SYRE.CPURST bit, there is a window of 1/2 cycle where the NB will not be able to stop the CPU from driving a new FSB transaction before the NB is able to drive reset on the bus. If a transaction occurs during this window, the NB will queue a response, and drive it after the reset has been de-asserted.

Implication: The processor will assert IERR# if such a spurious transaction is issued after reset.

Workaround: Programmers should not use the SYRE.CPURST bit with Intel® Xeon® Processor MP with 1MB L2 Cache.

Status: No Planned Fix

13. NB may trigger ECC errors during warm reset

Problem: The NB treats reset assertion as an asychronous event. If the bus is active during a system reset, NB FSB I/O flip-flops may capture a cycle of data from the NB core as the core flip-flops are transitioning to the reset state. Because setup time is not met, some of the I/O flip-flops capture the pre-reset value and others capture the reset value.

Implication: The NB may drive spurious FSB data or exhibit a FSB strobe glitch during the reset sequence. The processor may detect an ECC or parity error in this case.

Workaround: Programmers should clear and ignore processor ECC, glitch, or parity errors logged immediately after a NB reset

Status: No Planned Fix

14. NB PCI Express ports get stuck in reset state

Problem: The NB x8 PCI Express I/O units have a unique, narrow band of susceptibility when connected to an un-powered PCI Express end agents that have low impedances to ground. This low impedance causes the IOU to oscillate rapidly between Detect and Polling state.

Implication: If this oscillation continues for a long period (a few seconds) the unit may eventually become stuck in a reset state until the NB is power cycled or reset. Only x8 hot-plug slots with unpowered end agents are susceptible to this issue.

Workaround: A workaround for this issue which disables the NB PCIe transmitters when a slot is unpowered is described in the BIOS Spec Update.

Status: No Planned Fix

15. Receiver errors reported on NB when link is disabled

Problem: The NB is flagging the correctable PCI Express error “receiver error” when the link is disabled.

Disabling any of the PCI Express ports by setting the LNKDIS bit, bit 4, in the EXP_LNKCTRL[7:1] register at offset 0x7C, device 1-7, function 0, will cause correctable receiver errors to be erroneously reported. The link errors are logged in bit 0 of register CORERRSTS[7:1] at offset 0x110, device 1-7, function 0. Bit 0 is the IO7ERR bit which denotes a correctable receiver error. This error cannot be cleared as long as the PCI Express port is held in the disabled state.

Implication: This issue affects all of the PCI Express ports regardless of link width configuration with exception of port B when it is configured as x8.

Workaround: Receiver error detection should be masked prior commanding a port to disable its link.

Workaround Implementation

Use the following procedure to disable a link:

1. Mask the correctable error bit by setting IO7MSK bit, bit 0, in the COREDMASK[7:1] register. The COREDMASK[7:1] register is at offset 0x150, device 4-7, function 0.
2. Set the disable link bit, LNKDIS bit, bit 4, in register EXP_LNKCTRL[7:1] at offset 0x7C, device 1-7, function 0.

Use the following procedure to enable the link:

1. Clear the disable link, LNKDIS bit, bit 4, in register EXP_LNKCTRL[7:1] at offset 0x7C, device 1-7, function 0.
2. If the port is A0, and this port meets the conditions for #65587, then apply the workaround for #65587
3. Wait 100 ms.
4. Clear the receiver error detect mask bit, IO7MSK, bit 0 bit in register COREDMASK[7:1] at offset 0x150, device 1-7, function 0.

Status: No planned fix

16. Receiver errors reported on PortB when link is disabled in Port A in the x8 configuration

Problem: The NB hub is reporting the correctable PCI Express error “receiver error” on port B when ports A0 and or A1 are disabled. This issue manifests itself only when port B is configured as a x8 link (ports B0 and B1 combined). This error can not be cleared as long as the link is held in the disabled state.

Disabling either PCI Express port A0 or A1 by setting the LNKDIS bit, bit 4, in the EXP_LNKCTRL[7:1] register at offset 0x7C, device 6-7, function 0, will cause correctable receiver errors to be erroneously reported on port B. The link errors are logged in bit 0 of register CORERRSTS[7:1] at offset 0x110, device 4, function 0. Bit 0 is the IO7ERR bit which denotes a correctable receiver error. This error cannot be cleared as long as the PCI Express port is held in the disabled state.

Implication: This issue affects only port B when it is configured as a x8 link (ports B0 and B1 combined). This issue manifests its self regardless of whether ports A0 and A1 are combined as a x8 link or are separate x4 links.

Workaround: Receiver error detection should be masked prior commanding a port A0 or A1 to disable its link, and should be left disabled as long as the link in port A0 or port A1 is disabled

Workaround Implementation

Use the following procedure to disable a link:

1. If port B is configured as a x8 link and either ports A0 and A1 are to be disabled.
2. Mask the correctable error bit by setting IO7MSK bit, bit 0, in the COREDMASK[7:1] register. The COREDMASK[7:1] register is at offset 0x150, device 6 or 7, function 0.
3. Mask the correctable error bit by setting IO7MSK bit, bit 0, in the COREDMASK[7:1] register. The COREDMASK[7:1] register is at offset 0x150, device 4, function 0.
4. Set the disable link bit, LNKDIS bit, bit 4, in register EXP_LNKCTRL[7:1] at offset 0x7C, device 6 or 7, function 0.

Use the following procedure to enable the link:

1. If port B is configured as a x8 link and either ports A0 and A1 are to be enabled.
2. Clear the disable link, LNKDIS bit, bit 4, in register EXP_LNKCTRL[7:1] at offset 0x7C, device 6 or 7, function 0.
3. If the port is A0, and this port meets the conditions for #65587, then apply the workaround for #65587
4. Wait 100 ms.
5. Clear the receiver error detect mask bit, IO7MSK, bit 0 bit in register COREDMASK[7:1] at offset 0x150, device 6 or 7, function 0.
6. If links on ports A0 and A1 are enabled; clear the receiver error detect mask bit, IO7MSK, bit 0 bit in register COREDMASK[7:1] at offset 0x150, device 4, function 0.

Status: No planned fix

17. PCI Express Max Payload Size does not default to 128B

Problem: The NB does not set the PCI Express Max Payload Size to 128B by default.

Implication: This is a PCI Express Base Specificiation 1.0a compliance issue.

Workaround: BIOS should set Max Payload Size to 128B by writing 000b to bits 2:0 of EXP_DEVCTRL[7:1]

Status: No Planned Fix

18. NB Coherency issue with Locked Sequence

Problem: When a locked request conflicts with an earlier request to the same address and the NB is at a specific point in its lock flow, the NB may not preoperly order implicit writeback data that results from the earlier request.

Implication: Under a very specific set of timing conditions, stale data is returned for the first request of a locked sequence and cohernecy is violated. Dual core processors are most susceptible to this issue. These conditions are extremely unlikely to occur on the Cranford processor and are impossible with the Potomac processor.

Workaround: BIOS should set bit 21 of NB Device 17, Function 2, Offset C8

Status: No Planned Fix

19. Hot Plug Memory Board powered down after hard reset

Problem: The IMIHPC and EXPSLTCRTL PWRCTRL bit values are not sticky through hard reset as specified. After a hard reset, when a VPP is reprogrammed by BIOS, the other port of the 9555 component will be reinitialized to all inputs since it is disabled by default at reset.

Implication: It is impossible for PWREN# to be continuously asserted for both ports of a 9555 during the VPP reprogramming process. XMB error information will be lost across resets without implementing a workaround. If an errors occur on two separate IMIs which use the same 9555 component, only one set of error information may be recovered after the NB is reset. In addition, hot-plug PCI Express ports may experience a power glitch after NB resets, but there is no further impact to the port (or system) due to this problem.

Workaround: For memory boards, System BIOS must identify the failing IMI from NB Error Status Registers, and enable the corresponding VPP via the IMIHPC register for that IMI before another VPP which utilizes the same 9555 component. This workaround insures the memory board with error information will not be inadvertently powered down. No workaround is necessary for hot-plug PCI Express cards.

Status: No Planned Fix

20. SMI Escalation via ERR[2:0]# Pins May Result in IERR#

Problem: Any condition routed to MCH ERR[2:0]# output pins for SMI# escalation that would result in data poisoning on FSB will result in a CPU IERR# assertion due to race condition. The race condition exists between poisoned data presented to the FSB by the MCH and the uncorrectable error escalated to SMI# via ERR[2:0]# pins. If the CPU observes an SMI# assertion before the poisoned data is presented to the FSB, the CPU asserts MCERR# followed by IERR#. The CPU expects the SMI# assertion to occur at least 10 BCLKs after the associated MCERR# assertion. Conditions that result in data poisoning on FSB are uncorrectable memory ECC errors or poisoned TLPs received from any PCI Express* port.

Implication: CPU may assert IERR# and cause a system hang before the error may be logged in SMM.

Workaround: The following workarounds have been identified:

1. Log errors upon reboot
   — System components (TNB, CPU, Southbridge) preserve errors across warm resets via sticky error registers

   — Requires a mechanism to reset system upon IERR# (server management controller, depress system reset button, etc.)
2. Delay SMI# assertion by inserting a 60µs-150µs delay on the ERR0# pin
   — Two recommended circuit solutions (other solutions may be possible):

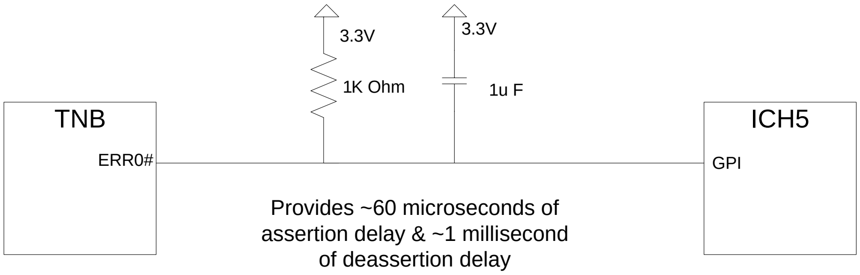

Provides ~60 microseconds of assertion delay & ~1 millisecond of deassertion delay

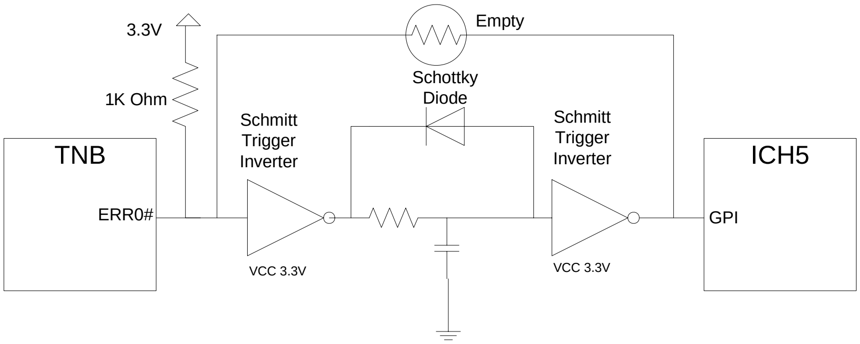

RC filter network can be designed for 60us – 150us assertion delay. Diode minimizes deassertion delay.

   — Route errors that result in FSB data poisoning to the ERR0# pin (IO4, M2, M4-13, M15)
3. Disable FSB data parity (ECC on Potomac/Tulsa) checking on CPU
   — Prevents CPU from asserting MCERR# on poisoned data

Note: As in all cases, error-logging on reboot in required.

Status: No Fix

21. Can not read across MI link after XMB fatal error

Problem: XMB can not be accessed vi the MI link after an unrecoverable error. has occured on the XMB.

Implication: XMB CSR’s cannot be read after an unrecoverable error even after the system is reset.

Workaround: When a fatal error occurs on a memory board and the associated MI link goes down, the system will eventually reset. Upon reset the following should be done to re-establish the MI link.

1. Clear bit 28 on D:13, F: 0-3, O:19C for the MI link that failed.

2. Initialize that MI link.
3. Read and clear the XMB errors.
4. Read and clear the associated TNB errors.
5. Set the bit for the MI link that failed (D:13, F:0-3, O:19C, B:28).
6. Continue with memory initialization.

Note: An attempt can be made to re-initialize the DIMM that caused the fatal error or it can be taken offline.

Status: No Planned Fix

22. PCIe hang on L1 entry when a x4 link has downshifted to x2

Problem: If a PCIe link downshifts from x4 width to x2, due to a HW failure on the PCIe lanes, AND subsequently enters L1 a system hang may occur.

Implication: PCIe cards that support degradation to x2 are susceptible; only a few have been identified.

Workaround: Given susceptibility is limited to only PCIe cards that support degradation to x2 and given opportunity for hardware failures vary per platform design and usage model. An optional workaround has been created. **Optional Workaround**: There is an SMI based workaround that will detect the downshift scenario and further degrade the link to a x1. Refer to the Intel E8501 chipset BIOS spec update for details.

Status: No planned fix.

Specification Changes

There are no new Specification Changes in this revision of the Specification Update.

Specification Clarifications

1. PCI Express Configuration completion time-out enable needs non-configuration completion time-out to be enabled in order to work

The PCI Express configuration completion time-out timer requires both of the TIMEOUT_ENABLE and TIMEOUT_ENABLE_CFG bits to be set. Simply setting the TIMEOUT_ENABLE_CFG bit will not enable the completion time-out timer. These bits are found in the EXP_CTRL[7:1] register.

This affects all of the PCI Express ports. There are only three choices of completion time-outs:

1. No completion time-outs.
2. Non-configuration completion time-outs.
3. Both configuration and non-configuration time-outs.

There is no way to enable just the configuration time-out counter.

Use of configuration completion time-outs is not recommended (the PCIE specification requres that all non-posted transations be timed out excepting configures cycles which should not time out), if you require them, use the following process. When time-out for configuration transactions is needed, both of the TIMEOUT_ENABLE and TIMEOUT_ENABLE_CFG bits in the EXP_CTRL[7:1] register need to be set.

To enable both configuration and non-configuration transaction time-outs:

1. Set the TIMEOUT_ENABLE_CFG bit, bit 23 in register EXP_CTRL[7:1] at offset 0x48, device: 1-7, function 0.
2. Set the TIMEOUT_ENABLE bit, bit 22 in register EXP_CTRL[7:1] at offset 0x48, device: 1-7, function 0.

To enable just non- configuration transactions time-out:

1. Set the TIMEOUT_ENABLE bit, bit 22 in register EXP_CTRL[7:1] at offset 0x48, device: 1-7, function 0.

To disable both configuration and non- configuration transactions time-outs:

1. Clear the TIMEOUT_ENABLE_CFG bit, bit 23 in register EXP_CTRL[7:1] at offset 0x48, device: 1-7, function 0.
2. Clear the TIMEOUT_ENABLE bit, bit 22 in register EXP_CTRL[7:1] at offset 0x48, device: 1-7, function 0.

Documentation Changes

There are no new Documentation Changes in this revision of the Specification Update.

§